How In-Kind Transfer Design Shapes Supply Incentives: Evidence from U.S. Housing Vouchers*

Hector Blanco
Rutgers University

Jaehee Song
University of Colorado Boulder

August 5, 2026

Abstract

Many safety-net programs deliver in-kind benefits through private providers, and program design determines who participates and at what price. We study landlord participation and pricing decisions in the U.S. Housing Choice Voucher program, linking voucher administrative data to national rental listings and to a new panel of the program's rent ceilings. We show that landlords are more likely to participate and to explicitly target voucher holders in high-poverty neighborhoods and that they adjust rents substantially toward the program's rent ceiling between listing and contracting. We estimate a model of landlord behavior and show that these price adjustments sustain participation. Participation is costly, markups below the ceiling offset part of the cost, and concessions above it reveal landlords' willingness to pay for the program's payment guarantee. Reforms paying landlords only where the ceiling binds deliver more voucher leases per dollar than unconditional transfers, and targeted ceiling increases expand neighborhood access most cost-effectively.

Keywords: in-kind transfers, housing vouchers, landlord participation, rent ceilings, bunching, neighborhood access.
JEL codes: H42, H53, I38, R21, R31, R38.

---

*Blanco, H.: Department of Economics, Rutgers University. Contact: hector.blanco@rutgers.edu. Song, J.: Leeds School of Business, University of Colorado Boulder. Contact: jaehee.song@colorado.edu. We thank seminar participants at Columbia-NYU-Yale Housing Day, NYU Furman Center, and Rutgers for helpful comments. Oliver Kim and Eric Wenz provided excellent research assistance. We gratefully acknowledge financial support from the Michael A. Klump Center for Real Estate.

1 Introduction

Several of the largest U.S. safety-net programs, including the Housing Choice Voucher (HCV) program, the Supplemental Nutrition Assistance Program (SNAP), Medicaid managed care, and the Women, Infants, and Children program, share a common structure. The government finances these in-kind benefits, but private providers, such as landlords, grocers, and insurers, deliver them, and the government cannot compel their participation. Program outcomes therefore depend not only on who is eligible and what benefits they receive, but also on whether providers serve program participants and at what price. Providers weigh what the program offers against what participation costs them, and program design shapes both sides of that trade-off. In turn, these provider decisions determine where recipients can redeem benefits and who captures program spending.

We study landlords' participation and pricing decisions in the Housing Choice Voucher program, the largest rental-assistance program in the United States, which served 2.4 million households in 2025 with more than \$30 billion in annual assistance payments. Under the program, voucher holders lease units in the private market and pay roughly 30 percent of their income toward rent, with the government covering the rest up to a rent ceiling, known as the *payment standard*, set by each local Public Housing Agency (PHA). Though the program aims both to provide decent housing for low-income households and to expand their neighborhood choice, it has fallen short of both. Only about 60 percent of households issued a voucher succeed in leasing a unit (Ellen et al., 2024), and those who do remain concentrated in high-poverty neighborhoods (Galvez, 2010; Lens, 2013; Horn et al., 2014; Collinson et al., 2019). Because a voucher pays out only when a private landlord takes its holder as a tenant, landlords' decisions about whether to lease to voucher holders and at what rent are a natural place to look for the source of these shortfalls.

The program's payment rule ties landlords' pricing to their participation decision. Because the tenant's payment is a fixed share of income for any rent up to the payment standard, the tenant bears none of a rent increase below the standard, while rent in excess of the standard falls on the tenant dollar for dollar.[1] A landlord listing below it can therefore raise the rent toward the ceiling at no cost to the tenant, which we term *upcharging*, and a landlord listing above the ceiling can attract voucher holders by cutting the rent, which we term *downcharging*. Because the ceiling is largely uniform within a market and bedroom category while rental units are heterogeneous, nearly every listing faces one of these incentives, so the rent a landlord would set under the program is a critical part of the decision to lease to voucher holders at all. The two adjustments also carry opposite budgetary implications, since downcharging lowers what the program pays relative to the listed rent while upcharging raises it.

[1] This distinguishes HCV from programs such as SNAP, where prices vary across stores and promotions but the program imposes no administratively set, beneficiary-specific price schedule.

To examine these landlord decisions, we assemble and link several datasets that follow units from the moment a landlord lists them to the moment a contract is signed with a voucher holder, together with the applicable payment standard. For the listing stage, we use online rental listings from RentHub, which aggregates postings across property-management and rental platforms. These data let us classify landlords by whether they target voucher holders. We label units posted on affordable housing platforms specialized in HCV as *hard-targeted*, and units posted elsewhere but explicitly soliciting voucher holders in their descriptions as *soft-targeted*. For the contracting stage, we use restricted-access, address-level HUD administrative data on all HCV contracts, including household, unit, and program characteristics and the applicable payment standard. We link the two sources at the address level to track participation by targeting status and the change in rent from listing to contract. To our knowledge, this is the first such linkage at a national scale. In addition, we construct the first nationwide panel of payment standards, combining inference from HUD administrative records with manual collection, allowing us to study how landlords interact with the ceiling at both stages. Together, these data cover more than 900 CBSAs and roughly 2,000 PHAs administering vouchers between 2015 and 2023, with payment standard information for nearly 500 of the largest PHAs.

We document several facts about landlord participation and price-setting in the HCV program. First, landlord targeting of voucher holders is concentrated in high-poverty neighborhoods. On average, 1.1 percent of listings hard- or soft-target voucher holders, rising to 2.6 percent in the highest-poverty quartile. Targeting also rises with the shares of Black and Hispanic residents, with lower education levels, and with higher vacancy rates in the neighborhood. Targeting strongly predicts participation, as hard-targeted units are roughly twenty times as likely as non-targeted units to lease under the program, and soft-targeted units nearly nine times as likely. These facts suggest that landlords who find participation profitable are concentrated in high-poverty, high-vacancy areas, which may steer voucher holders toward such neighborhoods.

Second, both participating and targeting landlords set rents with the payment standard in mind. Contract rents for new voucher contracts bunch sharply at the payment standard, as do listed rents, but only among hard- and soft-targeted listings, indicating that targeting landlords act on the ceiling before any voucher holder makes contact. Contract rent bunching alone, however, cannot separate landlord price-setting from tenant search, since voucher holders have their own incentive to lease units priced at the payment standard, where their subsidy is maximized. Our HCV–listings linked sample, which contains both the listed and the contract rent for the same unit, allows us to separate the two.

Third, voucher rents adjust toward the payment standard from both directions. In the linked sample, bunching is far more pronounced in contract rents than in listed rents, indicating that most of the adjustment arises at the contracting stage rather than through tenants' search. The listed rent distribution is also far wider than the contract rent distribution, pointing to two coexisting adjustments, downcharging among listings above the standard and upcharging among those below it. To quantify them, we implement a differences-in-bunching design that exploits variation over time in the real value of rents relative to nominal payment standards to recover the counterfactual rent a landlord would have posted absent the program. We estimate that 3.66 percent of linked voucher units bunch at the payment standard, with excess mass extending to units priced up to 20 percent below it, so upcharging combines bunching with smoother, partial adjustment. Only about a quarter of this bunching arises at the listing stage, with three-quarters emerging between listing and contracting. The offsetting missing mass appears on both sides, among units 20 to 50 percent below the payment standard, reflecting upcharging, and among units up to 40 percent above it, reflecting downcharging.

Fourth, these listing-to-contract rent adjustments are large, and they are specific to the HCV program. Among linked units, 35.9 percent of landlords raise rents between listing and contract and 35.7 percent lower them, each by about 8 percent on average. The adjustments grow with the distance between the listed rent and the payment standard in both upcharging and downcharging alike. In Multiple Listing Service transactions outside the program, for example, the same listing-to-contract comparison yields changes near zero over a comparable range of listed rents around the payment standard.

Finally, we exploit quasi-experimental variation in the rent ceiling from the staggered adoption of ZIP code-level payment standards by a subset of PHAs during our sample period. The reform, intended to expand voucher holders' access to lower-poverty neighborhoods, replaces the metro-wide ceiling with ZIP-level ones, raising the ceiling in higher-rent ZIP codes and lowering it in lower-rent ones within the same metro. Landlords respond on both margins identified above. Targeting of voucher holders rises, and bunching among voucher contract rents at the payment standard increases. Both responses are consistent with the ZIP-level ceiling sitting closer to local market rents, which shortens the distance rents must travel to reach it. Consistent with prior work (Collinson and Ganong, 2018; Ellen et al., 2025; Eriksen et al., 2025), the average exposed voucher holder nonetheless ends up in a somewhat lower-poverty neighborhood, potentially because the higher ceiling in higher-rent ZIP codes brings additional units within reach.

Guided by these facts, we develop and estimate a three-stage equilibrium model of landlord behavior in the voucher program. In the first stage, landlords choose whether to advertise to voucher holders, at a cost, and at what rent to list. The second stage allocates the metro-level flow of newly admitted vouchers across listings, where each listing's attractiveness to voucher holders varies with its targeting status and rent, and, because the flow is fixed, one listing's gain comes partly at the expense of others. Given that most rent adjustment arises after listing, the third stage lets matched landlords move the contract rent toward the payment standard, upward below the ceiling and downward above it. The key parameters are the landlord's cost of leasing to a voucher holder, the value of the program's guaranteed payment relative to the landlord's outside option, an uncertain market tenancy, and the scale of the two rent adjustments. Together, they determine how many landlords find a voucher match worthwhile, where those landlords are, and how participation and rents respond to program design.

We estimate the model sequentially backwards, combining maximum likelihood for the contracting and allocating stages with a moments-based estimator for the listing stage that matches the targeting shares, participation rates, and rent distributions constituting the facts above. The participation cost is identified from the fact that most landlords forgo targeting options that are relatively cheap yet would sharply raise their odds of a voucher match. The voucher match value is identified from how participation varies with local leasing prospects, and the adjustment scales from the listing-to-contract rent changes in the linked sample. We further discipline the model's policy elasticities with the quasi-experimental responses to ZIP-level payment standards documented above, so the counterfactuals inherit the credibility of the event studies rather than the cross-section alone.

The estimates put a price on program participation. Leasing to a voucher holder costs the typical landlord \$1,245 per contract, equivalent to about one month of the typical rent, a cost that captures not only the program's administrative burden but also any aversion to voucher holders themselves (Cunningham et al., 2018; Aliprantis et al., 2022). Against this cost, a voucher match resolves market uncertainty worth roughly \$780 per month for the typical listing, expected upcharging adds \$63, and downcharging landlords surrender \$414 in rent, about half the value of the payment guarantee they secure. Accordingly, a voucher match is worthwhile for only 3 percent of listings, and the landlords who pay to target voucher holders are precisely those for whom it pays, which rationalizes why targeting is rare in the aggregate yet decisive for where vouchers land.

We then use the model to evaluate counterfactual program designs. Three takeaways emerge, robust to how much of the unobserved rental market responds to policy. First, how a dollar reaches landlords matters more than how many dollars do. A uniform increase in payment standards, which raises payments only on contracts priced at the ceiling, delivers roughly two and a half times more matches per dollar than comparable unconditional cash transfers to landlords, which mostly subsidize landlords who would have participated anyway. Second, aggregate lease-up and neighborhood access call for different instruments. Uniform generosity expands matches roughly evenly across neighborhoods, rent-indexed payment standards reallocate matches toward low-poverty tracts at the cost of total matches, and payment standard increases targeted to low-poverty tracts buy access most cheaply, at a monthly cost of roughly \$1,500 per match moved into the lowest-poverty quartile.

Third, the program can fund access out of its own inefficiency, since tightening rent-reasonableness enforcement claws back the upcharging transfer and recycling the savings into targeted standards raises low-poverty matches by nearly a fifth at no net budget cost.

Related literature. We contribute to several strands of literature. We first add to the economics of in-kind transfers and the private provision of public programs. A large literature shows that private providers respond to contract design on margins that are difficult to contract on (Hart et al., 1997; Galiani et al., 2005; Knutsson and Tyrefors, 2022; Duggan et al., 2023; Agafiev Macambira et al., 2026; Coube et al., 2026), and studies of administered payment schedules find that providers adjust participation and treatment when fees change (Clemens and Gottlieb, 2014; Alexander and Schnell, 2024). Closest to our setting, providers may price-discriminate between program and non-program customers, as documented for WIC (Meckel, 2020) though not for SNAP (Goldin et al., 2022). A complementary public finance literature attributes incomplete take-up to informational barriers and transaction costs on the recipient side (Bettinger et al., 2012; Bhargava and Manoli, 2015; Alatas et al., 2016; Deshpande and Li, 2019; Finkelstein and Notowidigdo, 2019). We expand this literature in the housing voucher context, one of the primary programs relying on private providers where the provider side remains relatively underexplored despite its evident relevance, given low take-up rates (Ellen et al., 2023), the concentration of recipients in low-income neighborhoods (Galvez, 2010; Lens, 2013; Horn et al., 2014; Collinson et al., 2019), and the well-documented discrimination against voucher holders (Phillips, 2017; Cunningham et al., 2018; Aliprantis et al., 2022; Faber and Mercier, 2022). In particular, we show how landlord behavior at the listing and contracting stages shapes both the quality of the housing the program delivers, in terms of neighborhood access, and its price.

We also contribute to a growing body of work on landlord incentives in low-rent markets. Recent studies show that low-rent properties earn high investment returns (Damen et al., 2025), learning about tenants' nonpayment risk is an important driver of eviction decisions (Humphries et al., 2025), and landlords' demand for rent insurance (Bézy et al., 2025; Abramson and Nieuwerburgh, 2026), together indicating that insuring rental income is valuable, especially in the low-rent market segment. We bring together this body of work by showing that some landlords use the voucher program itself as insurance, prioritizing voucher holders at the listing stage and, at times, forgoing part of the rent to attract them, and by quantifying their willingness to pay for the payment guarantee the program provides.

Finally, we build on prior work on housing vouchers. One strand examines their effect on private market rents, generally finding that vouchers raise rents through demand (Susin, 2002; Eriksen and Ross, 2015; Park, 2026; Fack, 2006), with case study evidence that landlords charge voucher holders more than comparable tenants (Desmond and Perkins, 2016). Another studies how program design shapes outcomes through the demand side of the market, documenting bunching at payment standards in particular metropolitan areas (McMillen and Singh, 2020; Aliprantis et al., 2022, for LA and DC, respectively) and showing nationally that raising payment standards in low-poverty areas improves neighborhood outcomes for voucher holders without reducing lease-up (Collinson and Ganong, 2018; Ellen et al., 2025). Bergman et al. (2024) further show that customized support for voucher holders can substantially improve these outcomes as well. We complement this demand-side evidence by modeling the supply side, studying participation and pricing jointly from the listing stage onward, and moving the analysis of landlord behavior from localized settings to national scale.

2 Institutional Background

2.1 The Housing Choice Voucher Program

The Housing Choice Voucher (HCV) program, established in 1974 under the Housing and Community Development Act and formerly known as Section 8, is the largest rental-assistance program in the United States, serving 2.4 million households as of 2025.[2] The program is federally funded by the U.S. Department of Housing and Urban Development (HUD) and administered locally by approximately 2,200 Public Housing Authorities (PHAs). Unlike project-based subsidies, which tie assistance to specific developments, most HCV assistance is tenant-based, meaning that eligible households may lease any unit in the private market that meets the program's quality and pricing rules.

A defining feature of the program is that voucher holders must themselves find a willing landlord in the private rental market. Once a voucher is issued, the household has a fixed window, typically 60 to 120 days, depending on the PHA, to find a unit, have it pass inspection, and execute a lease. PHAs may grant extensions at their discretion, but these are not guaranteed and vary widely in practice. If the household fails to sign a lease within the allotted window, the voucher expires and reverts to the PHA, and the household is returned to the waiting list.[3] This search-on-deadline structure places voucher holders under substantial time pressure. Approximately 40 percent of households who received a voucher fail to lease up within the search window, and failure rates are higher in tight rental markets, among larger households, and among Black and Hispanic households (Ellen et al., 2024). These search frictions imply that the effective reach of the program depends directly on landlord participation, since a voucher that no landlord will accept provides no assistance. Notably, participation is voluntary in most of the country. Outside jurisdictions with source-of-income (SOI)

[2]Source: HUD Picture of Subsidized Households (PSH).

[3]Average waiting times for HCV vouchers exceeded two years in most large metropolitan areas as of the late 2010s, and several major PHAs closed their waitlists entirely during this period (Collinson et al., 2019).

anti-discrimination laws, landlords may refuse to accept vouchers for any reason, and even where SOI protections exist, enforcement tends to be weak.[4] This dependence on voluntary landlord participation motivates our focus on the supply side of the market.

Even a willing landlord must clear several further requirements before a lease can be signed, and these can shape both whether the tenancy is completed and the rent ultimately agreed upon. The unit must pass a Housing Quality Standards (HQS) inspection before lease-up and periodically thereafter, and the contracted rent must pass a rent reasonableness test, under which the PHA compares it to comparable unassisted units to ensure the program does not subsidize inflated rents. The landlord then signs a Housing Assistance Payments (HAP) contract with the associated PHA, and the tenant signs a standard lease with the landlord. Once the contract is executed, the PHA pays the subsidy directly to the landlord each month, and the tenant pays the landlord their share of the rent.

2.2 Key Program Features

Under the HCV program, voucher holders generally pay 30 percent of their income toward rent, with the government covering the remainder, as long as gross rent does not exceed a rent ceiling, the payment standard, and passes a rent reasonableness test. We now turn to each of these features, which are central to our setting.

Payment standards. The payment standard, set by each PHA, determines both the maximum subsidy a voucher holder can receive in the PHA's jurisdiction and the maximum gross rent, the sum of contract rent and a utility allowance, that the program will subsidize in full. Each PHA may set its payment standards within a permitted range, usually between 90 and 110 percent of the metro-by-bedroom count-level Fair Market Rent (FMR), published by HUD, which is based on the 40th percentile of recent-mover gross rents and the 50th percentile in some tight markets. Since 2016, HUD has required a subset of metropolitan areas to set payment standards at the ZIP-code level using Small Area FMRs (SAFMRs), which track sub-market rent differences more finely than metro-wide FMRs.[5]

Payment standards introduce a kink into the rent schedule that both landlords and voucher holders face. For a voucher holder with monthly income I_i, their out-of-pocket payment stays

[4] As of 2023, roughly half of U.S. voucher holders live in jurisdictions with some form of source-of-income protection. Source: Poverty & Race Research Action Council. (January 2023). *Appendix B: State, local, and federal laws barring source-of-income discrimination. Expanding Choice: Practical Strategies for Building a Successful Housing Mobility Program.* See Cunningham et al. (2018) and Aliprantis et al. (2022) for evidence on the prevalence of voucher denial under both regimes.

[5] SAFMRs were introduced in 24 metros in 2018 and in 41 metros in 2025. A larger set of PHAs adopted ZIP-code payment standards voluntarily before the mandatory rule took effect. See Collinson and Ganong (2018) and Ellen et al. (2025) for the policy history and evaluations of uniform and SAFMR-based ceiling changes.

at $\tau I_i = 0.3\, I_i$, 30 percent of their income, as long as gross rent R^g does not exceed the payment standard S. Gross rent above S is passed on to the tenant dollar-for-dollar, $\tau I_i + (R^g - S)$, subject to an overall out-of-pocket cap of $\bar{\tau} I_i = 0.40\, I_i$. On the flip side, the government's housing assistance payment (HAP) rises one-for-one with rent up to $R^g = S$ and is flat beyond, because each additional dollar above the ceiling must come from the tenant. Appendix Figure A.8 illustrates the government's and the tenant's shares of the rent as a function of the rent relative to the payment standard.

Unit inspection and paperwork. Participation carries administrative requirements with both explicit and implicit costs. Before a lease can begin, the unit must pass a Housing Quality Standards (HQS) inspection, and it must be reinspected periodically thereafter. A failed inspection requires repairs and a follow-up visit, and the PHA pays no subsidy until the unit passes, so inspection lags translate into forgone rent on a unit the landlord holds off the market. Beyond inspections, the landlord must execute the HAP contract and the program's lease addendum, document ownership and tax information, and obtain PHA approval for subsequent rent increases. Interview evidence suggests that scheduling, paperwork, and payment setup often add weeks between tenant selection and lease-up (Greenlee, 2014; Garboden et al., 2018; Aliprantis et al., 2022).

Rent reasonableness. Another friction facing landlords is the rent reasonableness test. Before executing a HAP contract, the PHA must document that the proposed rent is reasonable relative to comparable unassisted units and the unit's own characteristics. In principle, the test can also bind below the payment standard, since a rent well under the ceiling may still be deemed unreasonable if comparable units rent for less. In practice, implementation varies across PHAs, and the test's effective stringency has been the subject of debate. Qualitative work suggests landlords retain scope to secure higher rents by upgrading or reclassifying unit features, for example, counting extra rooms or making cosmetic renovations, to justify them under the program (Rosen, 2014; Garboden et al., 2018). We return to this channel in our counterfactual analysis (Section 8), where we model tightened enforcement as a proportional scaling down of the expected upward rent adjustment at the contracting stage.

2.3 Landlord Participation and Price-Setting Incentives

Landlords face a trade-off when leasing to voucher holders, well documented in prior qualitative studies. On the benefit side, the program offers guaranteed rent payments, since the government directly covers a portion of the rent, along with reduced non-payment and vacancy risk and access to a broader tenant pool (Rosen, 2014; Greenlee, 2014). On the cost side, participating landlords face unit inspections, administrative paperwork, and the soft rent ceilings imposed by payment standards (Greenlee, 2014; Garboden et al., 2018; Aliprantis et al., 2022). The balance between these benefits and costs varies across landlords and markets. Participation is more attractive where rents are low, landlords' outside options are thin, and the guaranteed payment is more valuable, and less attractive where market rents substantially exceed the payment standard.

On the price-setting margin, two forces motivate our analysis. First, landlords whose market rents would exceed the payment standard may lower their rents to the payment standard, where the tenant pays only the base income share and the unit becomes fully affordable to voucher holders. In exchange for the forgone rent, these landlords secure the guaranteed government payments and the program's insurance against nonpayment and vacancy, protection worth most in low-rent markets (Humphries et al., 2025). We call this *downcharging* behavior. Second, landlords whose market rents would fall below the payment standard may instead raise them toward it to capture a larger government transfer. They can do so because voucher holders' demand is inelastic below the payment standard, where the tenant's payment is a fixed share of income no matter the rent. We call this *upcharging*, consistent with qualitative accounts of such behavior under the program (Rosen, 2014; Desmond and Perkins, 2016; Garboden et al., 2018). Bunching exactly at the standard maximizes the landlord's subsidy, but frictions such as the rent reasonableness test may stop some landlords partway.

The two forces have opposite budgetary implications. Downcharging transfers rent from landlords to voucher holders and can open markets that would otherwise be out of reach, while upcharging transfers taxpayer dollars to landlords. They are also inseparable from the participation decision itself, since the rent a landlord expects to collect under the program, whether above the listed rent through upcharging or below it through downcharging, enters the return to leasing to voucher holders in the first place. A program design that curbs upcharging therefore saves money but may discourage participation, while one that accommodates downcharging expands access at landlords' expense. Distinguishing the two forces, and quantifying their relative importance, is the central task of this paper.

3 Data

Our analysis combines three data sources: restricted-access administrative data on all HCV participants from HUD, rental listings from RentHub, and a nationwide panel of payment standards that we construct by combining information inferred from HUD data with manual collection.

3.1 HUD voucher administrative data

We use HUD's restricted-access contract-level administrative data on the HCV program, covering all voucher holders from 2015 through 2023. The data provide individual-level records on household characteristics (composition, race, ethnicity, age, disability status, and income), unit characteristics (precise address, bedroom count, contract rent, gross rent, tenant out-of-pocket payment, HAP, and utility allowance), and program parameters (the payment standard applied to each contract, and indicators for new admissions).

Our main analyses focus on the 1.2 million new admissions observed between 2015 and 2023, corresponding to roughly 110,000 to 170,000 admissions per year. New admissions are the natural sample for studying landlord price-setting because the contract rent is freshly negotiated under the payment standard in force at the time of lease-up, while continuing contracts may inherit the rent amount from prior periods. We impose two restrictions in our analyses. First, we exclude project-based voucher participants, since our focus is on landlords who are not bound by development-level contracts. Second, we restrict to units with 0 to 4 bedrooms, the range for which payment standards are consistently published. Column (1) of Table 1 reports summary statistics for this new admissions sample; the median contract rent is $800, the median payment standard is $936, and the distribution spans 923 CBSAs.

3.2 RentHub rental listings data

RentHub compiles rental listings from property management websites, rental platform aggregators, and other public sources, and re-scrapes active listings on a weekly to bi-weekly cadence, so a given unit appears in many scrapes over the time it is on the market. Each scrape records the asking rent, bedroom count, bathroom count, square footage, property type, amenities, the full-text description, and the street address, which we geocode to a standardized address, Census tract, and county. Each scrape also carries platform and unit identifiers assigned by RentHub. We use the platform identifier and the description text to classify voucher targeting and the unit identifier together with the geocoded address to chain scrapes into listing spells, described below.

We organize the raw scrapes into listing spells, each meant to capture a single unit's continuous time on the market. We identify a unit by the RentHub unit identifier where present, which maps to a single geocoded address for the large majority of listings, and otherwise by its geocoded address and any extracted unit number. Within each unit we order scrapes by date and begin a new spell whenever the gap between consecutive scrapes exceeds 60 days, treating a longer gap as evidence that the unit went off the market between spells, typically because it was leased.[6] We chain scrapes

[6] The 60-day threshold sits in a valley of the within-unit gap distribution, between the weekly to bi-weekly scrape

only within a constant bedroom count, so a unit whose recorded bedroom count changes is split into separate spells. From each spell we retain the first and last listed rent, the minimum and maximum listed rent, the first and last scrape dates, the spell duration, the number of scrapes, bedroom count, bathroom count, square footage, building type, the voucher-targeting classification described below, an indicator for whether the spell links to a HUD voucher admission at the same address, and the payment standard that applies given the unit's ZIP code, bedroom count, and month.

The spell is our primary unit of observation and anchors the listing-to-contract analysis in Section 4 and the model estimation in Section 7. For our analyses we restrict to spells anchored between 2015 and 2023, to studio and one- to four-bedroom units, and to units we can place within a CBSA, and we trim listed rents below $100 or above the 99th percentile within each state to remove data-entry errors and an implausible upper tail.[7]

Landlords who participate in the HCV program do so through distinct channels that plausibly reflect different levels of voucher-targeting effort. Recognizing this, we classify each listing spell into one of three mutually exclusive categories based on platform and listing text, with hard targeting taking precedence over soft and soft over non-targeted.

- **Non-targeted.** Listings that are not on an affordable housing platform and whose description does not explicitly welcome vouchers. This includes listings on mainstream platforms (e.g., Zillow, Realtor.com, Apartments.com), listings with no description, and listings whose only voucher language is a refusal.

- **Soft-targeted.** Listings that are not on an affordable housing platform but whose description explicitly welcomes vouchers, for example "Section 8 welcome," "vouchers accepted," or "all vouchers welcome." Acceptance language is screened against a refusal list, so phrases such as "no section 8" do not qualify. See Appendix B.3 for further details.

- **Hard-targeted.** Listings on an affordable housing platform or with an affordable housing management organization, channels that exist specifically to serve voucher holders. Within our data these are dominated by AffordableHousing.com (formerly GoSection8.com), the most widely used voucher-specific platform. See Appendix B.3 for further details.[8]

This three-tier classification plays two roles in the paper. First, it underlies Fact 1 in Section 4,

cadence of an active listing and the much longer gaps that precede genuine re-listings, so the exact cutoff is not consequential. See Appendix B.2.

[7] RentHub coverage is uneven across states and years. In particular, the District of Columbia is missing between 2015 and 2022, and Hawaii is thin between 2017 and 2022, with 2019 missing entirely. See Appendix B.1 for coverage by state and year.

[8] Over 97 percent of our hard-targeted listings originate from AffordableHousing.com. Because the platform and company metadata that flag hard-targeted listings fill in sharply from 2019, all specifications include year fixed effects, and we re-estimate the model parameters using 2023 data alone as a robustness check.

which documents how targeting intensity correlates with voucher match rates and with neighborhood characteristics. Second, it enters the structural model of Section 6 as the landlord's choice variable, whose composition across listings identifies the Stage 1 parameters governing the cost and the match-rate benefit of active voucher targeting.

We report statistics at two levels of observation, the individual scrape and the listing spell, in Table 1. This yields a RentHub full sample of 287,613,176 scrapes across 928 CBSAs, which collapse to 22,693,226 listing spells, an average of about 13 scrapes per spell. Rental units in the RentHub data are considerably more expensive than the voucher stock. The median listed rent is \$1,707 at the scrape level and \$1,730 at the spell level, roughly twice the \$800 median contract rent among all new voucher admissions and well above the \$950 median among the admissions we link to a listing. Part of this gap reflects size, since listed units have a median of about 1,090 to 1,180 square feet against roughly 858 for linked voucher units, and part reflects that listings concentrate in higher-rent areas, where the applicable payment standard is also higher, with a median between \$1,342 and \$1,497. Voucher targeting is rare among listings as a whole, at 0.12 to 0.35 percent hard-targeted and 0.51 to 0.79 percent soft-targeted across the two levels, but it is an order of magnitude more common among the voucher units we link to a listing, at 6.29 and 6.66 percent, which we describe in Section 3.4.

3.3 Payment standards

Payment standards are central to our analysis of landlord behavior; thus, measuring them accurately is essential. Yet no centralized or consistently maintained nationwide dataset of payment standards exists. PHAs vary widely in how they publish this information, with some posting schedules in administrative plans or on their websites and others providing little accessible documentation. HUD administrative data record the payment standard applied to each contract, but these records do not indicate whether it is set at the metro or ZIP-code level and, in raw form, cannot be used to build a panel of payment standards applicable to units that did not actually enter the HCV program. In addition, HUD administrative data report payment standards only for PHAs that do not participate in the Moving to Work (MTW) Demonstration program.[9]

We construct a harmonized monthly panel of payment standards by PHA, geographic unit, and bedroom size, combining three complementary sources. First, we infer each PHA's schedule from HUD administrative data using the modal payment standard in each PHA-period-bedroom cell. This procedure yields valid metro-level payment standard inferences for 454 PHAs over all or part of

[9]The Moving-To-Work Demonstration program grants participating PHAs flexibility to waive or adapt standard public housing and HCV rules, including how payment standards are set. MTW began with 39 initial agencies, which retain MTW status throughout most of our sample period. Under the expansion authorized by the 2016 Consolidated Appropriations Act, HUD began designating up to 100 additional agencies starting in 2021.

the 2015–2023 sample period, including PHAs that adopted ZIP-level payment standards partway through the sample (for which the metro-level payment standards applies in pre-adoption years). Second, we exploit the regularity that PHAs typically set payment standards as a fixed multiplier of the FMR or SAFMR (e.g., 100 or 110 percent) and recover the applicable multipliers from the same administrative records to extend coverage to finer geographic levels, particularly for PHAs applying ZIP-level payment standards. We obtain data for 401 PHAs using this method. Third, for 76 of the largest PHAs affected or almost affected by SAFMR mandates in 2016, we manually collect payment standards from administrative plans, websites, and direct correspondence. When multiple sources cover the same PHA-period-bedroom cell, we apply a priority hierarchy that favors manually collected data and, within the two inference methods, prioritizes the modal method over the multiplier-based one. In total, our dataset contains payment standard information for at least some period in 2015–2023 for 474 PHAs. Appendix A details the construction of this dataset.

Utility allowances, which enter the tenant's gross rent calculation but not the landlord's contract rent, vary across PHAs and bedroom sizes with some PHAs reporting zero allowances for all contracts and others consistently reporting positive values. Although the payment standard formally applies to gross rent, our analyses focus on contract rent, which is the object landlords choose and the one that maps directly to listed rents. As a robustness check, we replicate the main findings in the subsample of contracts with zero utility allowance, for which contract and gross rent coincide.

3.4 Linking rental listings to HUD contracts

A contribution of our data construction is the linkage between HUD administrative records and RentHub rental listings, which lets us observe both the listed rent and the signed contract rent for the same voucher unit. This linkage identifies the listing-to-contract rent change studied in Fact 3 (Section 4) and disciplines the Stage 3 contracting parameters in the structural model.

We link each HCV new admission to any RentHub listing spell posted within the six months prior to the admission date, matching on cleaned street address, city, state, and bedroom count. Address cleaning proceeds in several steps. We first extract unit identifiers (e.g., "apt," "unit," "suite," "#") from the full address field, using a dictionary of prefixes followed by alphanumeric sequences, and strip these identifiers from the address string. We geocode the cleaned addresses with the Census Geocoder to obtain standardized street numbers, street names, and street types; for addresses without a successful Census match, we fall back to the ArcGIS World Geocoder and retain the original address when both attempts fail. We then standardize all key address components in both the HUD and RentHub files.

Because unit identifiers (apartment or suite numbers) are missing or unreliable in RentHub for most multi-unit properties, our main linkage is at the address level rather than the unit level. The correspondence between admissions and listing spells is therefore many-to-many rather than one-to-one. A single admission may be linked to several spells, either because the address was posted more than once or because a multi-family building contains units beyond the one the admission refers to, and a single spell may in turn link to several admissions in the same building. In multi-family housing, then, a linked spell may correspond to a different unit than the voucher holder's. This procedure links 12 percent of new admissions over 2015–2023 (17 percent among single-unit addresses).

For analyses conducted at the voucher-contract level, we address this concern in two ways. First, when multiple listings are linked to an address, we retain the single listing whose posting date is closest to the admission date and use it as our main sample. Second, we repeat all analyses on a robustness subsample restricted to single-unit addresses, for which the concern is absent by construction.[10]

For analyses conducted at the listing spell level, where the unit of observation is the spell itself, we retain every linked spell and attach a linkage weight in place of a binary linked indicator. Each linked admission receives a total weight of one, divided equally among the spells linked to it, so an admission linked to k spells contributes $1/k$ to each. A spell's linkage weight is the sum of these fractions over the admissions linked to it, and is zero for an unlinked spell. Because each admission's weight sums to one regardless of how many spells it touches, the linkage weights sum across spells to the number of distinct linked admissions, so no admission is counted more than once.

The HCV units we link to listing spells are not a random sample of HCV admissions. Table 1, Column (2), reports summary statistics for the linked sample. Relative to the full sample of new admissions, linked units have higher contract rents and payment standards, though they have lower rents and are smaller than the average RentHub listing. Appendix Figure A.9 further plots coefficients from a linear probability model of linkage on tenant and unit observables (standardized within PHA). Linkage is more likely not only for units with higher contract rents but also for units in neighborhoods with higher renter shares and in multi-family structures, consistent with online listing platforms over-representing rental housing in denser, more heavily rented submarkets. We correct for this selection in Section 7 using an inverse-probability weight that reweights linked units to match the full HUD admission distribution on observables.

[10]We classify an address as single-unit if none of the following hold in either the HUD or the RentHub data. The unit is recorded as being in a multi-family building; the record carries an apartment number; or the same address appears elsewhere in that source with an apartment number.

4 Stylized Facts

This section presents four facts about landlords' participation and price-setting under the HCV program. Fact 1 uses the listings universe to show that voucher targeting concentrates in high-poverty neighborhoods and strongly predicts participation. The remaining facts turn to pricing. Fact 2 shows that contract rents bunch sharply at the payment standard while listed rents do so only among targeting landlords. Fact 3 uses the linked listings-to-contracts sample and a differences-in-bunching design to quantify the adjustment, finding that rents move toward the standard from both directions and that three-quarters of the bunching arises between listing and contracting. Fact 4 shows that these within-unit rent changes are large, about 8 percent in either direction, with no counterpart in a market-rate benchmark. Together, the facts indicate that participating supply is geographically selected and that pricing responds to the payment rule at both the listing and contracting stages, patterns that the model of Section 6 incorporates.

Fact 1: Voucher targeting is concentrated in high-poverty neighborhoods and strongly predicts participation

Landlords are more likely to seek voucher tenants in high-poverty neighborhoods within a metropolitan area. Figure 1a plots the share of listing spells that are hard- and soft-targeted, defined as in Section 3.2, by the poverty quartile of the unit's Census tract, computed within the CBSA. The share of hard-targeted listings, those advertised on affordable housing platforms, rises from near zero in the lowest-poverty quartile to roughly 0.9 percent in the highest. The share of soft-targeted listings, those whose text explicitly seeks out voucher holders, rises from 0.25 percent to almost 1.7 percent across the same quartiles, a nearly seven-fold increase. Soft targeting exceeds hard targeting throughout the poverty distribution, indicating that landlords more often signal openness to vouchers within standard listings than post on voucher-specific platforms. A plausible explanation is that landlords seek to limit the cost of posting across multiple platforms while preserving the appeal to a broader tenant pool that includes voucher holders.

Several other tract characteristics also predict targeting. Figure 1b reports coefficients from separate listing-level regressions of a hard- or soft-targeted indicator on standardized Census tract characteristics, each including CBSA-by-year and CBSA-by-bedroom fixed effects, so the coefficients measure how each characteristic relates to the probability of targeting across neighborhoods within the same metropolitan area. Both forms of targeting are more likely in tracts with larger Black and Hispanic populations, lower education levels, higher unemployment, and lower median household income. Among housing-market characteristics, targeting is more common in higher-vacancy markets with larger renter shares and lower median rents.

Landlords who target vouchers are also more likely to enter the program. Figure 2a shows, by targeting status and within-CBSA poverty quartile, the share of listing spells linked to a new voucher contract. Among hard-targeted listings, the linkage rate rises from about 6 percent in the lowest-poverty quartile to 10 percent in the highest. Among soft-targeted listings, it rises from 2.5 to 5 percent. Non-targeted listings link at far lower rates throughout, from near zero in the lowest-poverty quartile to about 1 percent in the highest. Pooling across neighborhoods, hard-targeted listings are roughly twenty times as likely as non-targeted listings to enter the voucher program, and soft-targeted listings nearly nine times as likely.

These differences in program participation also appear across other tract characteristics. Figure 2b reproduces the pairwise regressions of Figure 1b using an HCV-linked indicator as the outcome, a proxy for the program participation, estimated separately for hard-, soft-, and non-targeted listings. The pattern mirrors that of targeting. Listed units are more likely to enter the program in tracts with larger Black populations, lower education levels, lower median household income, larger renter shares, and lower median rents. Most associations hold across the three samples, with a few exceptions. Among hard-targeted listings, we cannot reject that linkage rates are unrelated to tract Black share, education, and median rent. That is, once hard targeting, which is already concentrated in these areas, is accounted for, these characteristics carry no additional weight for program entry. Likewise, landlords in higher-vacancy areas are no more likely to enter the program in any targeting category, once their greater propensity to seek voucher holders is taken into account.

Together, these patterns are consistent with landlords seeking voucher tenants where the benefits of participation presumably outweigh the costs. That targeting is concentrated in lower-income, less-educated areas suggests landlords there value the reduced nonpayment risk vouchers provide, since part of the rent is guaranteed by the government. Higher targeting in higher-vacancy areas is consistent with lower vacancy risk, as landlords open to voucher holders draw from an expanded tenant pool. That several of these characteristics lose predictive power once targeting is held fixed indicates that they operate through the decision to seek voucher holders in the first place.

Fact 2: Voucher contract rents bunch sharply at the payment standard, and listed rents only among targeting landlords

We next examine how rents are distributed around payment standards for landlords participating in, or seeking to participate in, the HCV program. Figure 3 plots the distribution of rents around payment standards in 0.5-percentage-point bins for four samples, contract rents for new admissions in the administrative data (Panel a), and listed rents for hard-, soft-, and non-targeted listing spells (Panels b–d). Voucher contract rents bunch sharply at payment standards. For listings, bunching declines with targeting intensity. Hard-targeted listings bunch by about as much as contract rents, soft-targeted listings show less bunching, and non-targeted listings show almost none. Listed rents thus bunch only where landlords have signaled interest in the program, and they do so before any tenant is involved, indicating that landlords themselves drive part of the price adjustment.

The extent of contract rent bunching also varies across neighborhoods, though not along a single gradient. Appendix Figure A.10 reports pairwise regressions of an indicator for the contract rent equaling the applicable payment standard on standardized tract characteristics, with PHA-by-year and PHA-by-bedroom fixed effects. Bunching is more common in high-poverty, high-vacancy tracts with more renter households, where listed rents tend to sit below the standard and upcharging can carry rents to it, but also in tracts with lower Black population shares and more educated residents, where rents more often start above the standard and bunching is more plausibly reached by downcharging. Because these regressions condition on a signed contract, they cannot separate the two forces, a distinction we take up directly in Fact 3.

Fact 3: Both upcharging and downcharging contribute to the bunching, mostly at the contracting stage

To decompose the excess mass at the payment standard by direction and timing, we restrict attention to newly admitted HCV units linked to our listings sample, constructed as in Section 3.4. This sample follows rents from the listing stage to the contracting stage while holding unit composition fixed. Figure 4 overlays histograms of contract and listed rents for these linked units. Bunching is far more prominent in contract rents, indicating substantial rent changes between listing and contracting, but it is present in listed rents as well. This listing-stage bunching may reflect voucher holders seeking out units priced at the payment standard, so that rents at the standard are overrepresented among linked listings, or it may reflect landlords already manipulating rents when they list, as the targeting patterns of Fact 2 suggest. Either way, the listed rent distribution cannot serve directly as the no-manipulation benchmark, and quantifying the full extent of rent adjustment under the program requires estimating one. Intuitively, once that benchmark is in hand, missing mass above the payment standard points to downcharging, while missing mass below it points to upcharging.

We recover the counterfactual distribution of listed rents under no manipulation with a differences-in-bunching design that exploits the interaction of nominal payment standards with rent inflation (Dube et al., 2025).[11] Because payment standards are set in nominal terms and do not generally

[11] The traditional bunching estimator (Saez, 2010; Chetty et al., 2011; Kleven, 2016) fits a high-order polynomial to the focal distribution outside an excluded manipulation range around the threshold. It relies on strong parametric assumptions to identify the counterfactual (Blomquist et al., 2021) and detects only sharp adjustments near the threshold,

adjust with inflation, a focal year's payment standard falls at arbitrary, non-salient points of other years' rent distributions once those rents are expressed in focal-year dollars. The deflated distributions from other years are therefore free of bunching at the focal year's payment standard and can be used to infer its counterfactual.

We implement the design by stacking the data by focal year, pairing each focal year's observations with those from years more than two apart, since standards can be sticky across adjacent years, deflating rents into focal-year dollars, and binning them in 0.5-percentage-point intervals.[12] Let $\tilde{r}_{i,t,\tau}$ denote the deviation of unit i's year-t rent, deflated into year-τ dollars, from the year-τ payment standard, in percent of the standard, and let $p_{\tilde{r},t,\tau}$ denote the share of stack-τ observations from year t in bin $\tilde{r}$. Within each stack, we estimate

$$p_{\tilde{r},t,\tau} = \sum_{\tau'} \sum_{i=0}^{k=20} \alpha_{\tau,i}\, \mathbb{1}(\tau = \tau') \times \tilde{r}^i + \sum_{\tau'} \sum_{j=-w}^{w} \beta_{\tau,j}\, \mathbb{1}(\tau = \tau') \times \mathbb{1}(\tilde{r} = j) \times \mathbb{1}(t = \tau) + \varepsilon_{\tilde{r},t,\tau} \quad (1)$$

where the first term fits a 20th-order polynomial as the counterfactual from the non-focal years and the second term recovers the focal year's deviation from it in every bin, so excess and missing mass can be estimated anywhere in the distribution rather than only near the threshold. We estimate the model on a sample of 441 PHAs with payment standard data available consistently over 2015 to 2023, computing standard errors with a 1,000 iterations of a residual bootstrap procedure following the literature (Kleven and Waseem, 2013; Dube et al., 2025).[13][14]

Figure 5 shows the results, where Panel (a) overlays the contract rent histogram and the estimated counterfactual listed rent distribution, and Panel (b) plots their difference. We estimate that 3.66 percent of voucher new admissions bunch at the payment standard, with excess mass extending to units priced up to 20 percent below it rather than confined to the threshold. The excess mass alone cannot reveal how much comes from units priced above versus below their counterfactual rent, but the offsetting missing mass can. The difference plot in Panel (b) reveals significant missing mass among units with contract rents 50 to 20 percent below the payment standard, reflecting upcharging landlords, and significant missing mass among units 0 to 40 percent above the payment standard, reflecting downcharging landlords.

whereas the differences-in-bunching design also captures smoother manipulation far from it.

[12] We deflate rents using the smoothed monthly Zillow Observed Rent Index for the United States, applying the December value of each year.

[13] The sample restriction is necessary because the design applies a given year's payment standards to contract rents from other years. We therefore exclude PHAs whose payment standards we could neither infer from HUD administrative records (mostly MTW and small PHAs) nor collect directly.

[14] In each bootstrap iteration b, we resample residuals from Equation (1) with replacement, add them to the fitted values, and re-estimate the model; we repeat this 1,000 times and compute standard errors from the resulting empirical distribution.

This approach also characterizes the timing of manipulation. Comparing the listed rent distribution with its counterfactual, we find that only about 1 percent of participating landlords bunch at the payment standard at the listing stage (see Appendix Figure A.12). That is, roughly one-quarter of voucher contract bunching occurs at the listing stage and the remaining three-quarters at the contracting stage.

These results survive a range of robustness checks, including alternative polynomial orders and bin widths, excluding the post-COVID period of unusually fast rent growth, and using absolute rather than normalized deviations from the payment standard (see Appendix Table A.9). One exception is the sample of single-unit addresses, where we estimate smaller though still statistically significant bunching of 1.4 percent (see column 8 of Appendix Table A.9 and Appendix Figure A.13). A plausible explanation is that single-unit addresses over-represent smaller landlords, who may be less attentive to the rent ceiling, hold less bargaining power, or be more willing to accommodate utility allowances so that the payment standard binds on gross rather than contract rent. At the same time, the missing mass from above the payment standard is considerably larger in this sample, rising from 17 percent in the main sample to 26 percent. This pattern suggests that smaller landlords value the program's payment guarantee, and its access to a broader tenant pool, more than larger landlords.

Fact 4: Rent adjustments far exceed the market-rate benchmark and are larger for units listed farther from the payment standard

We now evaluate the size of listing-to-contract rent adjustments. Figure 6a shows a binned scatter plot of the relative change from listed to contract rent against the listed rent among linked units. Units with listed rents at least roughly 10 percent below the payment standard end up with higher contract rents, and the gap grows with distance from the ceiling, reaching increases of up to 50 percent, though few units listed that far below. This is consistent with the upcharging incentive being stronger well below the payment standard where there is ample room to raise rents while the tenant's payment stays fixed and demand from voucher holders is perfectly price-inelastic. Above the payment standard the pattern reverses as the downcharging incentive takes over. Rents fall from listing to contract by up to 20 percent, indicating that some landlords find the program valuable enough to accept lower rents than they listed in the private market.

Overall, we estimate that 35.9 percent of landlords adjust rents upward by approximately 8.1 percent, and 35.7 percent adjust rents downward by the same proportion. Scaled by baseline rents, the first group ends up with roughly $49 more per month under the program, and the second foregoes roughly $63 per month on average in 2015 dollars.

The magnitude of these listing-to-contract rent changes are far larger than those in private rental markets. Figure 6b repeats the exercise for market-rate MLS rental transactions outside the program, which record both the listed and the signed contract rent. Over the same range of listed rents relative to the payment standard, listing-to-contract changes are close to zero, with only a small discount that grows modestly with the rent level. A rent schedule that bends toward the ceiling, from both sides and increasingly with distance, is therefore specific to the voucher program rather than a general feature of rental transactions.

We also relate the magnitude of these rent changes to voucher holder and neighborhood characteristics. Figure 7 plots coefficients from pairwise regressions of the rent change on standardized voucher-holder and Census tract characteristics, estimated separately for linked units listed above the payment standard (left panel), which on average lower their rents under the program, and units listed below it (right panel), which on average raise them. To ease interpretation, the left panel uses the negative of the rent change as the outcome, so that coefficients in both panels measure the size of the adjustment toward the payment standard.

For units initially listed above the payment standard, one pattern stands out. Discounts at contracting are largest when the program's guarantee is presumably worth the most, for elderly tenants, whose Social Security income is stable, and for lower-income tenants, for whom the government covers a larger share of the rent, as well as in lower-income, less-educated areas, where market-rate tenants tend to carry larger nonpayment risks and outside market tends to be thin. The pattern reverses for units initially listed below the payment standard. These units raise rents more when the voucher holder's income is higher and in higher-education, higher-income, and higher-rent neighborhoods, where the gap between the payment standard and the listed rent leaves more room to upcharge.

5 Quasi-Experimental Evidence from Shifts in the Rent Ceiling

The previous section documents how landlords interact with the HCV program through their participation and price-setting decisions, and how these patterns vary with the rent ceiling and neighborhood characteristics. This section uses a differences-in-differences design to show that landlords respond along both margins to quasi-exogenous changes in the rent ceiling. The same variation later serves to discipline the model estimation in Section 6.

5.1 Background and Empirical Strategy

Complementing the descriptive evidence in Section 4 requires variation in the rent ceiling, payment standards, that is plausibly orthogonal to local demand and supply conditions. Following prior work (Collinson and Ganong, 2018; Ellen et al., 2025; Eriksen et al., 2025; Park, 2026), we exploit the staggered adoption of Small Area Fair Market Rents (SAFMRs) across PHAs. Conventional payment standards are anchored to the metro-level Fair Market Rent (FMR) and are therefore largely uniform within a metropolitan area. SAFMRs instead set the FMR benchmark, and hence the payment standard, at the ZIP-code level. SAFMR adoption therefore raises the rent ceiling in high-rent ZIP codes and lowers it in low-rent ZIP codes within the same metro. The rollout has been driven largely by HUD, which mandated ZIP-level payment standards for PHAs in 24 metropolitan areas by 2018 and in 41 by 2025, while other PHAs have been adopting them voluntarily. Our payment standard dataset captures this gradual rollout, with roughly 99 PHAs switching from metro- to ZIP-level payment standards between 2015 and 2025 (see Appendix Figure A.14).

We adopt a staggered difference-in-differences design by comparing outcomes in PHAs that adopt ZIP code-level payment standards earlier to those that adopt them later in the sample period. Identification rests on the assumption that, absent adoption, payment standards and landlord behavior in early-adopting PHAs would have evolved in parallel with those in later-adopting PHAs. Appendix Table A.10 compares early-adopting, later-adopting, and remaining PHAs in our nationwide payment standard dataset along several dimensions. Although later-adopting PHAs tend to serve more voucher holders, the three groups of PHAs are located in comparable areas in terms of county characteristics, suggesting that later-adopting PHAs provide a plausible comparison group.

For observation i (listing or voucher contract) in year t, let g_i denote the year in which the PHA serving i adopts ZIP-level payment standards and $e = t - g_i$ denote event time. We estimate

$$Y_{it} = \sum_{e \neq -1} \beta_e \, \mathbb{1}\{t - g_i = e\} + \mu_{c(i)} + \delta_t + \varepsilon_{it}, \tag{2}$$

normalizing $\beta_{-1} = 0$, where $\mu_{c(i)}$ and δ_t are geography and year fixed effects, with geography measured at the county level for listings and the PHA level for contracts, and standard errors are clustered at the corresponding geography. The coefficients β_e trace the evolution of the outcome Y around adoption relative to the control group. We study four outcomes: the log payment standard, as a first-stage check; an indicator for either hard- and soft-targeting at the listing stage; an indicator for contract rent equal to the payment standard; and the poverty rate of the tract in which a voucher holder leases. Because adoption raises the subsidy limit in some ZIP codes and lowers it in others, we estimate Equation (2) separately for ZIP codes where adoption is expected to raise versus lower the payment standard (higher- and lower-PS ZIP codes). We classify each ZIP code by taking the difference between its 40th-percentile rent and the 40th-percentile rent of its CBSA, both measured at baseline in the 2011–2015 American Community Survey (5-year estimates), assigning ZIP codes with a positive difference to the higher-PS group and the rest to the lower-PS group.

We adopt the interaction-weighted estimator of Sun and Abraham (2021), which is robust to heterogeneity in treatment effects across adoption cohorts. Because our voucher administrative data span 2015–2023, we use the PHAs that adopt ZIP-level payment standards in 2023–2025, and are therefore not yet treated within the sample window, as the control group.[15]

5.2 Results

Figure 8a reports first-stage event studies for the log payment standard separately for higher- and lower-PS ZIP codes. SAFMR adoption shifts the payment standard as expected, in opposite directions across the two groups. Before adoption, both groups evolve in parallel with later adopters, with no significant pre-trends. After adoption, the payment standard rises by roughly 4 percent in higher-PS ZIP codes and falls by about 6 percent in lower-PS ZIP codes, with most effects emerging immediately and persisting over time.

Landlords respond to the change in the rent ceiling already at the listing stage. Figure 8b reports event studies for an indicator that a listing is hard- or soft-targeted. In ZIP codes where the payment standard rose, we find no significant change in voucher targeting, especially right after the adoption. This is consistent with the descriptive evidence, since targeting is rare in these higher-rent areas to begin with (Figure 1a) and a higher rent ceiling need not offset the costs of participation. In ZIP codes where the payment standard fell, by contrast, targeting rises by roughly 0.4 to 0.8 percentage points, imprecisely estimated but sizable relative to baseline targeting shares of around 1 percent (Section 4). One interpretation is that landlords in these lower-income areas intensify their efforts to retain voucher tenants, even though the subsidy itself has decreased.

The clearest pricing response appears in the probability of contract rents bunching at payment standards (see Figure 8c). Pre-adoption coefficients are flat and statistically indistinguishable from zero. After adoption, the probability that contract rent equals the payment standard rises steadily for both higher- and lower-PS ZIP codes and stabilizes at around 2 percentage points. This is direct causal evidence that the rent ceiling shapes price-setting decisions, reinforcing the bunching documented in Section 4.

[15] Ellen et al. (2025) and Eriksen et al. (2025) instead define the control group by exploiting the cutoff rules HUD used to select the first 24 mandated metropolitan areas, based mainly on the number of voucher holders served, the concentration of voucher holders in low-income Census tracts, and the share of rental units in ZIP codes where the SAFMR exceeds 110 percent of the FMR, using the metropolitan areas just below these thresholds. While we do not adopt this design explicitly, our comparison group largely overlaps with theirs, since such PHAs not mandated by 2018 were largely mandated by 2025 and enter our controls.

Finally, we ask whether these shifts in the payment standard translate into where voucher holders live, an outcome plausibly mediated by the landlord responses just documented. Figure 8d reports effects on the poverty rate of the tract in which a voucher holder leases. After adoption, we observe a downward level shift of about 1 percentage point in the poverty rate of voucher holders' tracts, consistent with SAFMR's intended effect of expanding access to lower-poverty neighborhoods and with prior evidence (Ellen et al., 2025; Eriksen et al., 2025).

Taken together, the event studies show that quasi-exogenous shifts in the rent ceiling move landlord behavior along precisely the margins identified in Section 4: targeting, contract rent setting, and the neighborhoods voucher holders reach. These reduced-form responses motivate and discipline the structural model in Section 6, which we use to quantify landlords' valuation of program participation and to evaluate counterfactual payment standard designs.

6 A Model of Landlord Participation and Price-Setting

The evidence in Sections 4 and 5 establishes that landlords target voucher tenants where the program is most valuable to them, that contract rents gravitate toward the payment standard from both directions, and that both margins respond to quasi-exogenous shifts in the rent ceiling. Several questions remain that this evidence cannot answer. First, how large a rent concession are landlords willing to make in exchange for the program's guaranteed payment? Second, how costly is program participation for landlords? Third, how would landlord participation and pricing respond to alternative subsidy designs, such as uniform payment standard increases or tighter rent-reasonableness enforcement? To address these questions, we develop and estimate a three-stage model of landlord behavior.

6.1 Setup

Listing j in year t is characterized by observed attributes Y_{jt}, comprising its bedroom count and the characteristics of its Census tract (the poverty rate, the rental vacancy rate, and log median rent), its metro area c_j, a voucher-targeting intensity $V_{jt} \in \{\text{none}, \text{soft}, \text{hard}\}$, and the payment standard S_{jt} that applies to the unit given its bedroom count and PHA jurisdiction. A tenant i is characterized by voucher status $v_i \in \{0, 1\}$.

Under voucher rules, we restrict attention to contract rents at or below the payment standard, $r_{jt} \leq S_{jt}$. The tenant's monthly payment and the government's housing assistance payment are

$$p_{it}(r, S_{jt}) = \tau I_i, \qquad \text{HAP}_{it}(r, S_{jt}) = r - \tau I_i,$$

where the tenant pays a fixed $\tau = 30\%$ of income regardless of the contract rent. Appendix Figure A.15 supports this restriction, showing that the tenant out-of-pocket payment in the data is 30 percent of income in the majority of cases.[16]

The market unfolds in three stages.

1. **Stage 1 (Listing).** The payment standard S_{jt} is exogenously given. The landlord chooses a listing rent R_{jt} and a voucher-targeting intensity V_{jt}.

2. **Stage 2 (Match).** The metro-year flow of new voucher admissions is allocated across listings. Listing j matches a voucher tenant with probability π_{jt}^{HCV}, which depends on $(R_{jt}, S_{jt}, Y_{jt}, V_{jt})$.

3. **Stage 3 (Contracting).** Conditional on a voucher match, the landlord and tenant contract on a rent r_{jt} subject to the voucher rules. Absent a voucher match, the unit leases on the private market with a prospect that depends on the listing rent and local market conditions.

6.2 Stage 1: Listing choice

In Stage 1, the landlord chooses the listing rent R and the voucher-targeting intensity $V \in \{\text{none}, \text{soft}, \text{hard}\}$ to maximize expected revenue net of a targeting cost,

$$\begin{aligned} \max_{R,V} \quad & \underbrace{\pi^{\text{HCV}}(R, S_{jt}, Y_{jt}, V) \cdot \Pi^{\text{HCV}}(R, S_{jt}, Y_{jt}, V)}_{\text{expected revenue if a voucher match}} \\ & + \underbrace{\left[1 - \pi^{\text{HCV}}(R, S_{jt}, Y_{jt}, V)\right] \cdot \Pi^{\text{MKT}}(R, Y_{jt}, q_{jt})}_{\text{expected revenue if no voucher match}} \\ & - \underbrace{c^V(V)}_{\text{targeting cost}} \end{aligned} \tag{3}$$

where π^{HCV} is the voucher match probability (see Section 6.3) and $\Pi^{\text{HCV}}(R, S_{jt}, Y_{jt}, V)$ is expected revenue with a voucher match (see Section 6.4). Expected revenue without a voucher match is $\Pi^{\text{MKT}} = \rho(\cdot)\, R$, the listing rent times a market prospect $\rho_{jt} \in (0, 1)$ given by

$$\rho(R, Y_{jt}, q_{jt}) = \Lambda\big(\gamma_0 + \gamma_R \log\big(R/R_{jt}^{\text{mkt}}\big) + Y_{jt}'\gamma_Y + q_{jt}\big) \tag{4}$$

[16] More generally, the tenant's monthly payment and the government's housing assistance payment are $p_{it}(r, S_{jt}) = \tau I_i + (r - S_{jt})_+$ and $\text{HAP}_{it}(r, S_{jt}) = \min\{r, S_{jt}\} - \tau I_i$. The subsidy is capped at the payment standard, so a contract rent above S_{jt} is borne dollar for dollar by the tenant, and this kink at S_{jt} makes the payment standard a focal point for pricing. Contract rents above the payment standard are nonetheless bounded above. At initial lease-up, program rules cap the tenant's total share at $\bar{\phi} = 0.40$ of monthly income, which implies a feasibility ceiling $\bar{r}_{it} = S_{jt} + (\bar{\phi} - \tau)\, I_i$ on the contract rent. The two program parameters τ and $\bar{\phi}$ set the 30% base share and the 40% tenant-pay cap, and the latter is slack whenever $r_{jt} \leq S_{jt}$.

where $\Lambda(z) = e^z/(1+e^z)$ is the logistic function. The benchmark R^{mkt}_{jt} is the typical rent of comparable local units, constructed at the tract-by-bedroom-by-year level (see Appendix C.3), and with $\gamma_R < 0$ the prospect falls as the rent rises above it. Covariates Y_{jt} let the prospect vary with local conditions. The landlord observes a listing-level idiosyncratic shock $q_{jt} \sim N(0, \sigma_q^2)$.[17] The targeting cost $c^V(V)$ equals 0 when V = none, c^V_{soft} when V = soft, and c^V_{hard} when V = hard, where c^V_{soft} and c^V_{hard} denote the costs of advertising a voucher-friendly listing and of listing on an affordable housing platform, respectively.

6.3 Stage 2: Allocation of new voucher tenants

Stage 2 allocates the flow of new vouchers across listings. In metro area c and year t, a flow M_{ct} of new voucher admissions searches for housing. Listing j matches a voucher holder with probability

$$\pi^{\text{HCV}}_{jt} = \frac{M_{ct}\,\exp(x'_{jt}\theta)}{s^0_{ct} + \sum_{k\in ct}\exp(x'_{kt}\theta)}, \qquad x'_{jt}\theta = \theta_1(S_{jt}-R_{jt})_+ + \theta_2(R_{jt}-S_{jt})_+ + Y'_{jt}\theta_Y + V'_{jt}\theta_V, \quad (5)$$

where the index $x'_{jt}\theta$ measures a listing's attractiveness to the voucher flow and $s^0_{ct} > 0$ is an outside mass absorbing new vouchers that lease units outside the listing platforms we observe. Here $(\cdot)_+$ denotes the positive part, and the piecewise-linear form in $R - S$ reflects that listing attractiveness responds asymmetrically below versus above the payment standard.

6.4 Stage 3: Contract rent adjustments

Conditional on a voucher match, the landlord sets the contract rent r_{jt} subject to the program ceiling $r_{jt} \le S_{jt}$. We model the contract rent as a two-part adjustment of the listing rent, specified separately in the upcharging region ($R_{jt} \le S_{jt}$), where the landlord may raise the rent toward the payment standard, and the downcharging region ($R_{jt} > S_{jt}$), where the landlord may lower it to the payment standard to secure the voucher tenant.

Expected revenue with a voucher match $\Pi^{\text{HCV}}(R, S_{jt}, Y_{jt}, V)$ is then

$$\Pi^{\text{HCV}}(R, S_{jt}, Y_{jt}, V) = \begin{cases} R + m^+_{jt}(R, S_{jt}, Y_{jt}, V) - c^{\text{HCV}} & \text{if } R \le S_{jt}, \\ S_{jt} - m^-_{jt}(R, S_{jt}, Y_{jt}, V) - c^{\text{HCV}} & \text{if } R > S_{jt}, \end{cases} \quad (6)$$

where $m^+_{jt} \ge 0$ and $m^-_{jt} \ge 0$ are the realized upward and downward rent adjustments and c^{HCV} is a

[17] Absent q_{jt}, units with identical observables and payment standards would all list at the same optimal rent. The idiosyncratic shock exists only so the model matches the rent dispersion observed among such units.

program participation cost.

Case A (Upcharging, $R_{jt} \leq S_{jt}$). The contract rent is

$$r_{jt} = \begin{cases} R_{jt} & \text{w.p. } \pi^{0+} \quad \text{(no markup)}, \\ R_{jt} + \min\{\tilde{m}_{jt}^{+},\ S_{jt} - R_{jt}\} & \text{w.p. } 1 - \pi^{0+}, \end{cases} \tag{7}$$

where $\tilde{m}_{jt}^{+} > 0$ is a desired dollar adjustment. The point mass at $r_{jt} = R_{jt}$ occurs with probability π^{0+} and captures landlords who post no markup, leaving the contract rent at the listing rent. The point mass at $r_{jt} = S_{jt}$ is the censoring mass $(1 - \pi^{0+})\,\mathbb{P}(\tilde{m}_{jt}^{+} \geq S_{jt} - R_{jt})$. Between these two masses, the rent varies continuously with the truncated latent adjustment.

Case B (Downcharging, $R_{jt} > S_{jt}$). The contract rent is

$$r_{jt} = \begin{cases} S_{jt} & \text{w.p. } \pi^{0-} \quad \text{(no reduction below the payment standard)}, \\ S_{jt} - \tilde{m}_{jt}^{-} & \text{w.p. } 1 - \pi^{0-}, \end{cases} \tag{8}$$

where $\tilde{m}_{jt}^{-} > 0$ is a desired dollar adjustment. The point mass at $r_{jt} = S_{jt}$ is the downcharging-to-ceiling bunching mass and captures landlords who reduce the rent from the listing rent to exactly the payment standard and no further. Larger reductions place the rent continuously below the payment standard with probability $1 - \pi^{0-}$.

Latent adjustment. $s \in \{+, -\}$ indexes the upcharging $(+)$ and downcharging $(-)$ regions. The desired adjustment $\tilde{m}_{jt}^{s} \sim \text{Weibull}(\alpha^{s}, \lambda_{jt}^{s})$, with

$$\log \lambda_{jt}^{s} = \mu_0^{s} + \mu_1^{s}\, d_{jt}^{s} + Y_{jt}'\mu_Y^{s} + V_{jt}'\mu_V^{s}, \qquad s \in \{+, -\}, \tag{9}$$

where $d_{jt}^{+} = S_{jt} - R_{jt}$ is the headroom in the upcharging region, and $d_{jt}^{-} = R_{jt} - S_{jt}$ is the gap in the downcharging region. The shape α^{s} governs how concentrated the adjustment is near zero, where $\alpha^{s} = 1$ is the exponential case and $\alpha^{s} > 1$ places the mode away from zero.[18]

[18]Under the Weibull parameterization used here, with shape α^{s} and scale λ_{jt}^{s}, the latent adjustment has mean $\mathbb{E}[\tilde{m}_{jt}^{s}] = \lambda_{jt}^{s}\,\Gamma\big(1 + 1/\alpha^{s}\big)$ and standard deviation $\text{SD}[\tilde{m}_{jt}^{s}] = \lambda_{jt}^{s}\sqrt{\Gamma\big(1 + 2/\alpha^{s}\big) - \Gamma\big(1 + 1/\alpha^{s}\big)^{2}}$.

6.5 Discussion

The model deliberately focuses on the landlord side. Three scope choices follow from this design and deserve further discussion. First, we do not model tenant demand. Voucher holders enter only through the matching stage and through one program fact, that their out-of-pocket payment is flat in the rent up to the payment standard, so their demand is price-inelastic on the region where landlords price. This is what allows Stage 3 to treat contract rent setting as a landlord problem and makes the estimates robust to unmodeled tenant preferences, and it is also why we restrict attention to contracts at or below the standard, since above it the tenant pays the excess and income begins to matter. The flat schedule is preserved by all the counterfactuals we consider, which move the level or geography of the ceiling, the participation cost, or the expected markup while leaving the schedule's shape intact. Policies that reshape the schedule itself, such as a proportional subsidy or tenant out-of-pocket payments rising in rent, would activate the demand margin the model deliberately omits and lie outside its domain.

Second, we do not model landlord acceptance as a separate stage. The match indicator is a reduced form folding together platform coverage, tenant search, landlord screening, and unit inspection, and the participation cost bundles their burdens into a single net cost, which is what its revealed-preference identification requires and what the cost-offset counterfactual exploits. The limitation is that policies aimed at a single component, such as source-of-income enforcement, search assistance, or inspection reform, enter only through the net cost or the match surface rather than being identified separately.

Third, our listings cover part of each metro's rental stock, and matches formed outside the listed market are not observed. Rather than assuming this incompleteness away, we carry it into the policy analysis explicitly. Section 8 introduces a parameter μ for the share of the unobserved pool that responds to policy symmetrically with observed listings, reports results under a central value bracketed by the polar cases, and shows which conclusions the choice does and does not affect.

7 Estimation

We estimate the model sequentially, solving it backwards. This section summarizes each estimator and the variation that identifies it. Appendix C presents the optimization formulas and computational details.

7.1 Estimation procedure

Stage 3 estimation. The contracting stage is estimated on the linked sample of new-admission HCV contracts described in Section 3.4, covering admissions from 2015 through 2023.[19] For each contract we observe the listing rent R, the payment standard S, and the contract rent r, so the data directly reveal whether the landlord left the rent alone, adjusted it partway toward the standard, or moved it exactly to the standard. Each region $s \in \{+, -\}$ is estimated by maximum likelihood on its subsample, $\hat{\xi}^s = \arg\max_{\xi^s} \sum_{jt \in s} \ell^s_{jt}(\xi^s)$. The Stage 3 parameters $\xi^s = (\mu^s, \pi^{0s}, \alpha^s)$ collect the scale-index coefficients, with bedroom, year, and CBSA fixed effects, the no-adjustment probability, and the shape. The per-contract log-likelihood ℓ^s_{jt} assigns one piece to each of the three observed outcomes (see Appendix C.1).

Stage 2 estimation. The allocation stage characterizes which listings the flow of new voucher tenants lands on. Taking logs of Equation (5) shows that a listing's log match probability is additive in a metro-year scarcity term and the listing's attractiveness index, $\log \pi^{\mathrm{HCV}}_{jt} = \log A_{ct} + x'_{jt}\theta$, where $A_{ct} \equiv M_{ct} / \big(s^0_{ct} + \sum_k \exp(x'_{kt}\theta)\big)$. We first estimate the slopes $\theta = (\theta_1, \theta_2, \theta_Y, \theta_V)$ as a fixed-effects logit, using within-metro-year variation in the voucher match rate, and then calibrate the outside mass s^0_{ct} so that the model reproduces each metro-year's flow M_{ct} of new voucher admissions (see Appendix C.2). The estimation panel is restricted to the CBSAs on which Stage 3 is estimated.[20]

Stage 1 estimation. We estimate the listing-stage parameters $\psi = (c^{\mathrm{HCV}}, c^V_{\mathrm{soft}}, c^V_{\mathrm{hard}}, \gamma)$ by minimum distance, restricting the costs to be positive and $\gamma_R < 0$, and calibrating the shock scale σ_q to the within tract-by-bedroom-by-payment standard-by-year dispersion of log listing rents. The estimation cells are tract-by-bedroom-by-payment standard-by-year combinations weighted by their listing counts, restricted to the CBSA universe on which Stage 3 and Stage 2 are estimated, resulting in 1.64 million cells. We discretize the listing rent choice into the rent grid $\mathcal{G}_{jt} = \{\, \nu\, R^{\mathrm{mkt}}_{jt}$ for $\nu = 0.50, 0.51, \ldots, 2.00 \,\} \cup \{S_{jt}\}$, where the payment standard is included so that bunching is representable.[21] For a candidate ψ, every cell solves the listing problem of Equation (3) on this grid, integrating over the market shock q (see Appendix C.3).

[19]We exclude Moving-to-Work agencies, keep units with zero to four bedrooms, and restrict to 181 CBSAs with at least 100 linked contracts so the fixed effects are well populated, dropping about 14,000 linked contracts. Links whose listing and contract rents differ by more than a factor of two in either direction are treated as unreliable and removed. Consistent with the $r \leq S$ restriction, a handful of voucher contracts observed above the payment standard are censored at S and enter through the mass points at $r = S$, rather than being dropped.

[20]We show robustness of the parameter estimates without the CBSA restriction in Appendix Table A.13.

[21]We construct R^{mkt}_{jt} as the median rent among untargeted listings within each tract-by-bedroom-by-year cell, shrinking thin cells toward the corresponding CBSA median (with state and national medians as fallbacks) using an empirical Bayes weight $n/(n + 10)$, where n is the cell's listing count.

The estimator matches two kinds of moments. The first is a vector of 36 cross-sectional moments. These include the shares of listings that match a voucher tenant, target vouchers, and price at or above the payment standard, together with the splits of those shares across high-poverty, high-vacancy, and high-rent tracts. They also include the mean, the dispersion, the payment standard position bins, and the neighborhood splits of the rent premium $\log(R/R^{\text{mkt}})$, matching and bunching conditional on targeting, and the level and neighborhood splits of the tract-level vacancy-implied leasing prospect. The second collects the quasi-experimental responses to SAFMR adoption from Section 5, the event study effects on the log listing rent and on the shares of listings at and above the payment standard. Appendix Table A.2 states each moment's exact definition and the parameter it primarily disciplines. Stacking the two into a single moment vector $\hat{h}$ with model counterpart $h(\psi)$, the estimator solves

$$\hat{\psi} = \arg\min_{\psi} \ \big(h(\psi) - \hat{h}\big)' \, W \, \big(h(\psi) - \hat{h}\big), \tag{10}$$

where W is a diagonal inverse-variance weighting matrix.[22]

7.2 Identification

The Stage 3 parameters are identified by the shape of the linked contract rent distribution within each adjustment region. The shares of contracts written exactly at the listing rent (in the upcharging region) and exactly at the payment standard (in the downcharging region) pin the probabilities π^{0+} and π^{0-}, and the interior adjustments identify the scale index μ^s and the shape α^s. The covariates in Equation (9) shift where the adjustment distribution is centered, while α^s governs how tightly it concentrates near zero. The share of upcharging contracts bunched at the payment standard must equal the probability that the latent markup exceeds the available headroom $d^+_{jt} = S_{jt} - R_{jt}$. The decline of the bunching share in the headroom therefore disciplines the tail of the Weibull and, with it, the distance slope μ^+_1.

The Stage 2 slopes θ are identified by within-metro-year covariation between matching and a listing's position relative to the payment standard, its neighborhood attributes, and its targeting status, while the outside mass s^0_{ct} is identified by metro-year linkage rates given the administrative count of new admissions. The two-step design absorbs differential linkage across metro-years into the intercepts, but differential linkage correlated with listing characteristics within a metro-year would bias the slopes, and the robustness of the estimates across CBSA universes and sample

[22] The diagonal entries of W are proportional to the inverse sampling variances of the corresponding data moments, normalized so that the cross-sectional and SAFMR blocks each have a unit mean diagonal. The block-wise normalization is needed because the raw inverse variances are not comparable across the two blocks, and without it the quasi-experimental moments would carry essentially no weight in the objective. Appendix C.3 gives further details.

windows (Appendix Table A.13) is reassuring on this margin.

The central identification challenge arises in Stage 1, where the participation cost, the targeting costs, and the leasing preferences can partially substitute for one another. A landlord who rarely targets may do so because targeting is costly, because voucher tenants are costly once matched, or because pricing for the private market is attractive. Aggregate shares alone do not separate these parameters, so the moment vector adds conditional and joint moments (Appendix Table A.2). One key input is the Stage 2 estimate that targeting raises the match probability substantially. Targeting is nonetheless nearly universally absent, so the levels of the targeting shares identify the targeting costs c^V_{soft} and c^V_{hard}. The per-contract cost c^{HCV} enters the landlord's payoff in expectation, multiplied by the match probability. That probability is under one percent for the typical listing, so most moments are nearly flat in c^{HCV}, but it is roughly ten times larger among targeted listings, where the expected cost is material. Conditional on the targeting costs pinned above, the choice to forgo a cheap targeting option with a large match benefit reveals that a voucher contract itself must be costly, and matching conditional on targeting pins the level of c^{HCV}. Finally, the leasing-prospect coefficients (γ_0, γ_R) are identified by the location and dispersion of the rent premium $\log(R/R^{\text{mkt}})$, and the covariate coefficients γ_Y by how that premium and the vacancy-implied leasing prospect vary across neighborhoods.

The cross-sectional moments identify these parameters only through the model's structure, and the SAFMR moments add quasi-experimental discipline. The estimator requires the model's pricing response to a shift in the rent ceiling to match the response observed when SAFMR adoption moved payment standards, so a parameter vector that fits the cross-section but implies the wrong sensitivity of listing rents and bunching to the standard is rejected. This is also the margin most relevant for the counterfactual analyses in Section 8, which shift the payment standard directly.

7.3 Parameter Estimates

Table 2 reports the Stage 2 and Stage 3 estimates. The contracting estimates (Stage 3) quantify the two adjustment margins. In the upcharging region, which includes voucher contracts listed at or below the payment standard, 56 percent of listings leave the rent at the listing. Conditional on adjusting, desired markups scale strongly with the headroom to the ceiling, as the scale of the latent markup rises by 25 percent for every \$100 of headroom. Markups are larger in higher-poverty, higher-vacancy, and higher-rent tracts, though the median rent coefficient is smaller. Hard-targeted listings seek markups 15 percent larger than observably similar non-targeted listings, while soft targeted listings seek comparable markups to the latter. In the downcharging region, which includes voucher contracts listing above the standard, 37 percent of matched listings reduce the rent exactly to the payment standard and no further. Conditional on reducing below the standard, the desired reduction is nearly invariant to how far above the standard the unit was listed, so conceding landlords price from the payment standard rather than from their listing rent. Reductions are larger in higher-rent tracts and smaller among hard-targeted listings. In both regions, the estimated Weibull shape exceeds one ($\hat{\alpha}^+ = 1.36$ and $\hat{\alpha}^- = 1.24$, each precisely estimated), placing the mode of the desired adjustment away from zero.

The allocation (Stage 2) estimates quantify how new voucher holders select among listings. Pricing \$100 further above the payment standard lowers a listing's match odds by about 20 percent, while pricing \$100 further below it raises them by about 4 percent, so attractiveness bends sharply at the standard and the penalty for exceeding it is roughly six times the reward for undercutting it, dollar for dollar. Hard-targeted listings match at roughly eleven times the rate of observably similar non-targeted listings and soft-targeted listings at roughly five times ($\hat{\theta}_V = 2.40$ and 1.66). Matching is concentrated in higher-poverty, lower-rent neighborhoods, as a ten percentage point higher tract poverty rate raises match odds by 22 percent. The match probability falls roughly one for one, in proportional terms, with the tract's median rent. Combining the slopes with the calibrated scarcity levels, the implied probability that a listing matches a new voucher admission is 0.5 percent for the typical non-targeted listing, about 5 percent for soft-targeted listings, and about 11 percent for hard-targeted listings.

The Stage 2 and 3 estimates are robust across specifications. Appendix Tables A.11 and A.12 show that the contracting estimates are stable under extended tract covariates, a binary targeting indicator, principal-component controls, dropping all fixed effects, and the 2023 cross-section, with the hurdle probabilities and shapes moving little throughout. Appendix Table A.13 shows that the above-standard penalty θ_2, the targeting effects, and the neighborhood gradients are likewise stable across CBSA universes and sample windows, while the below-standard slope θ_1 is identified by the pooled panel's payment standard variation and is small and insignificant in single-year cross-sections.

Table 3 reports the Stage 1 estimates. Participating in the program costs the landlord $\hat{c}^{\text{HCV}} =$ $\$1{,}245$ per voucher contract, equivalent to about one month of the typical rent.[23] This cost estimate reflects not only the administrative cost of program participation but also the non-pecuniary aversion of landlords to voucher tenants. At the typical listing, the estimate leaves expected revenue conditional on a voucher match near zero, against roughly \$660 in expected market revenue. The magnitude follows from three features of the data. Stage 2 match probabilities rise steeply with targeting, such targeting is nonetheless nearly universally absent, and per-month voucher rents after Stage 3 adjustments weakly exceed market rents. It indicates strong aversion to voucher tenants, reflected in the participation cost parameter. Targeting is far cheaper on the same scale, at \$5.4 for

[23] Stage 1 payoffs are denominated in one month of rent revenue per listing episode ($\Pi^{\text{MKT}} = \rho R$ with R the monthly rent), so the dollar costs are identified relative to that scale rather than as out-of-pocket lump-sum costs.

advertising a voucher-friendly listing ($\hat{c}^V_{\text{soft}}$) and \$26.0 for listing on an affordable housing platform ($\hat{c}^V_{\text{hard}}$). These costs reflect the combined direct expense and any market-side stigma of voucher targeting, since for example some market-rate tenants may avoid rental listings that signal voucher targeting. The leasing prospect declines steeply in the rent premium ($\hat{\gamma}_R = -2.21$), so listing 10 percent above the local benchmark lowers the odds of leasing by about 19 percent. Prospects are weaker in higher-poverty tracts, where a ten percentage point higher poverty rate lowers leasing odds by 10 percent, and essentially flat in the tract vacancy rate and median rent.

7.4 Model Fit

We evaluate how the estimated model matches (1) the distribution of contract rent adjustments within the linked sample, (2) the cross-sectional gradients of voucher matching across neighborhoods, and (3) untargeted variation in targeting, bunching, and matching across neighborhoods and metropolitan areas. We also document where the fit is imperfect and discuss the implications for the counterfactual analyses.

Panel (a) of Figure 9 compares the mean adjustment among contracts with interior adjustments, by distance to the payment standard. The model closely aligns with the data in both regions. Among upcharging contracts, mean markups rise with the headroom below the ceiling in both the model and the data, except for the highest bin in the data, because adjustment is worth more when the scope for it is larger. Among downcharging contracts, the mean reduction varies little with how far the listing sits above the standard in the data. The model aligns with this because conceding landlords price from the standard rather than from their listed rent. The estimated mixture also reproduces the point masses, by construction, with 56 percent of upcharging contracts and 37 percent of downcharging contracts showing no adjustment. Two features fit less well in the upcharging region. The model over-predicts ceiling bunching among contracts with little headroom, and it holds the no-adjustment share flat while in the data it declines with headroom. Both arise from the constant hurdle probability assumption.

Panel (b) of Figure 9 compares the share of listings matched to a new voucher admission by deciles of tract poverty, tract vacancy, and tract median rent, against the Stage 2 allocation prediction. The model follows the data along all three gradients. Match rates rise steeply and log-linearly in poverty and fall in median rent, in the model and in the data, decile by decile including the extremes. Vacancy moves match rates little in either the model or the data, so the model earns the flat panel just as it earns the steep ones. These decile-level comparisons discipline the functional form well beyond the binary splits used in estimation, since the logit index is linear in the covariates while the decile patterns are free to bend.

Stage 1 targets 36 cross-sectional moments and 3 quasi-experimental responses, reported in

Appendix Tables A.2 and A.3. The pricing block fits tightly, including the share of listings priced at the payment standard, the share priced above it, the dispersion of rent premia, and the premium's gradients across neighborhood types. It also matches the quasi-experimental response of bunching to SAFMR adoption within the 95% confidence band of the event study estimate. Panel (a) of Figure 10 shows that the fit extends to untargeted within-metro variation. Bunching at the standard rises with within-CBSA tract poverty in the model as in the data, and the targeting share rises steeply across poverty quartiles in both, though the model understates the concentration of targeting in the poorest quartile. Panel (b) plots the fit metro by metro, none of which is targeted. The above-standard and at-standard shares line up along the diagonal across the 181 estimation CBSAs, and match shares are strongly rank-correlated with a level gap we return to below.

Three misfits deserve discussion. First, the model under-predicts matching and targeting levels, with the shortfall nearly uniform across neighborhood types, as shown in the poverty-quartile profiles in Figure 10. Counterfactual aggregates that scale with the number of matches can be understated in proportion, while variations in voucher activity across neighborhoods remain unaffected. Second, the model makes hard-targeted listings too textbook, concentrating their prices at the payment standard, whereas hard-targeted listings in the data price more idiosyncratically. Counterfactual results that condition on targeting status inherit this, so the model likely overstates how tightly voucher surplus concentrates among hard-targeted landlords. Third, the model transmits almost none of a payment standard increase into listed rents. The SAFMR event study, however, implies listed rents rise 8.7 percent (standard error 4.3 percent) with adoption. In the model, the payment standard enters profits only through a match probability of around 1 percent, so shifting the standard barely moves the optimal listing price. The bunching response to the same reform is matched, however, so the allocation and pricing-location margins in the counterfactuals remain disciplined by policy variation. Rent-level responses to changes in the payment standard, by contrast, should be read as conservative.

Overall, the model accounts for the joint distribution of listing prices, targeting, and matching across neighborhoods and for the contract rent adjustments that determine what landlords ultimately receive, while understating levels of voucher activity and the pass-through of payment standards into listing rents.

7.5 Model-Based Decompositions

The estimated model offers a decomposition of what a voucher match is worth to a landlord. For any matched listing, the value of the match relative to the market alternative, in the same monthly-revenue scale of the Stage 1 estimates, is

$$\Pi^{\mathrm{HCV}} - \Pi^{\mathrm{MKT}} \;=\; \underbrace{(1-\rho)\,R}_{\text{market uncertainty}} \;+\; \underbrace{\mathbb{E}[r \mid R] - R}_{\text{rent adjustment}} \;-\; \underbrace{c^{\mathrm{HCV}}}_{\text{participation cost}}\,, \tag{11}$$

where $\mathbb{E}[r \mid R]$ is the expected contract rent implied by the Stage 3 adjustment process, conditional on the listing rent. This is evaluated at the solution of Equation (20), so the listing rent entering each term is the endogenous choice $R^*_{jt}(q)$. Whether a matched listing sits in the upcharging region ($R^* \leq S$) or the downcharging region ($R^* > S$) is therefore itself an outcome of the listing stage, and the split of matches between the two regions below is model implied.

The first term is the value of resolving market uncertainty. Under a voucher match revenue is assured, while on the private market the listing leases with prospect ρ, so the term equals the shortfall of expected market revenue relative to assured revenue, and it is largest where the market prospect is weak and rents are high. Landlords are risk neutral in the model, so the term is an expected-value gap rather than a price of risk. The second term is the expected rent adjustment of the contracting stage. In the upcharging region, it is the expected markup toward the payment standard, reflecting the subsidy extraction. In the downcharging region, it is negative, and its absolute value is the average rent concession, reflecting the surrendered listing premium plus any further reduction below the standard, thus interpretable as what downcharging landlords pay to resolve market uncertainty.

The parameter estimates imply that the market-uncertainty term amounts to \$780 per month for the typical voucher-matched listing. Upcharging matches add an expected markup of \$63 on top of it. Downcharging matches surrender $\$414$ in rent against a market-uncertainty value of \$752, so the concession prices the assured voucher tenancy and remains well below its value. Net of the participation cost, however, a voucher match is worth $-\$387$ to the typical upcharging match, $-\$906$ to the typical downcharging match, and $-\$575$ overall, and the net value is positive for only 3.3 percent of listings. In the upcharging region, the net-positive listings are those with large headroom below a high payment standard, which gives the widest scope for the expected markup. In the downcharging region, they are listings with weak market prospects and priced above but still close to the standard, so the surrendered premium is small and the expected market revenue given up is low. The targeting margin makes the selection explicit. A voucher match is worth $-\$726$ to the typical non-targeted listing but $\$326$ to the typical soft-targeted listing and $\$507$ to the typical hard-targeted listing, so the landlords who pay the targeting cost are precisely those for whom a match pays.

8 Counterfactual Analysis

We use the estimated model to evaluate alternative program designs. Each counterfactual changes one policy instrument and re-solves the full equilibrium, where landlords re-optimize targeting and rents, market-level attractiveness adjusts, and voucher holders re-sort across listings, holding market primitives fixed.

8.1 Policy Experiments

We consider four families of policies, each in a global version and in targeted versions where they are implemented only in the lowest-poverty quartile (Q1) or the two lowest-poverty quartiles (Q1–Q2) of tracts. Additional details are explained in Appendix D.1.

- *Uniform increase in payment standards.* Payment standards rise by 10 and 20 percent, either in all tracts or only in targeted tracts, $S \to \chi S$ for $\chi \in \{1.10, 1.20\}$. This is the natural benchmark of greater program generosity, akin to a PHA moving from the FMR to the top of the exception payment standard range.
- *Rent-indexed payment standards.* Each tract's standard is set to ωR^{mkt}, a fixed proportion of its local market rent, with the factor ω chosen within each CBSA, year, and bedroom size so that the listing-weighted average standard is unchanged. This is a stylized version of nationwide SAFMR adoption that raises standards in expensive neighborhoods and lowers them in cheap ones.[24] A hold-harmless variant never lowers a standard below its baseline, and additional variants layer the uniform or targeted increases above on top of the indexed schedule.
- *Cash transfer to participating landlords.* The net participation cost is reduced by 12.5, 18, and 25 percent, mimicking a cash transfer to landlords who lease to a voucher holder, such as a signing bonus or an insurance or damage-mitigation fund.[25]
- *Rent-reasonableness enforcement.* The expected upward rent adjustment m^+ at the contracting stage is scaled to 75, 50, 25, and 0 percent of its estimated size, mimicking stricter enforcement of rent-reasonableness tests. Downcharging concessions are untouched. A final set of

[24]SAFMRs are implemented at the ZIP code level, while our counterfactual sets standards at the Census tract level, a typically finer geography and the one at which the market-rent benchmark R^{mkt} varies in our setting.

[25]The doses are chosen so that the transfers are dose-matched to the payment standard increases per contract at the mean standard. The 12.5 percent offset is worth about \$156 per contract, close to the roughly \$155 increase in the mean standard under the 10 percent uplift, and the 25 percent offset (\$311) matches the 20 percent uplift (\$310). The two families therefore deliver a similar transfer to a matched landlord and differ mainly in who receives it.

budget-neutral experiments recycles the enforcement savings into higher payment standards, uniformly or in targeted tracts, with the uplift factor chosen so that the payment standard increase cost exhausts the enforcement saving.

Our listings data cover only part of each metro's rental stock, and matches formed outside the listed market are not observed. We therefore report results under three alternative assumptions on the outside market, indexed by $\mu \in \{0, 0.5, 1\}$, the share of the unobserved outside pool that responds to policy symmetrically with observed listings. Under $\mu = 0$, the outside pool is fixed at its baseline attractiveness, so a policy that makes listed units more attractive pulls voucher holders into the listing platforms observed in our data. This assumption is more relevant if the outside pool consists largely of units and landlords beyond the reach of the policy levers, such as informal or long-tenured arrangements. Under $\mu = 1$, the outside pool responds one-for-one, aggregate matches are held at their baseline level by construction, and only the composition of matches across tracts is affected. This assumption is more relevant if the outside pool consists of units similar to those on the online rental platforms our data cover, whose listings the sample nonetheless fails to capture. The intermediate case $\mu = 0.5$ lets half the outside mass respond. We take $\mu = 0.5$ as the baseline and use the polar cases as bounds, with $\mu = 0$ providing an upper bound on aggregate match effects and $\mu = 1$ a composition-only lower bound, reported in Appendix D.4. Section 8.3 summarizes which results the choice of μ does and does not affect.

8.2 Results

Table 4 reports the baseline equilibrium and seven headline scenarios under the assumption of $\mu = 0.5$. Each column describes a full counterfactual equilibrium through five outcome panels. The first two report matches in the listed market, in the aggregate and by baseline within-metro poverty quartile, so the aggregate panel measures how many voucher holders lease through observed listings and the quartile panel measures the neighborhood access those matches deliver. The third panel reports rent manipulation among matched listings, the mean expected markup among upcharging matches and the mean concession among downcharging matches. The final two panels translate each policy into fiscal terms, reporting the change in government rent payments and transfer outlays per matched voucher per month together with the landlord-surplus incidence, and the implied cost per marginal match and per match moved into the lowest-poverty quartile. Appendix D reports the full dosage grids under all three values of μ in Tables A.5–A.7. In the status-quo baseline, shown in Column (1), the listed market absorbs 74,000 of the 825,000 voucher-holder flow, or 9.0 percent.

A uniform 20 percent increase in payment standards (Column 2) raises listed-market matches by 14.6 percent, at an additional monthly cost of $3,270 per marginal match. The mechanism runs through both participation and pricing. The share of listings with positive net match value nearly triples, targeting roughly triples, and among matched upcharging listings the mean markup rises from \$63 to \$137, so part of the added spending passes through to inframarginal rents rather than new matches. Critically, the gains are spread broadly across the poverty distribution, between $+14$ and $+18$ percent in every quartile, leaving the concentration of voucher holders in high-poverty neighborhoods largely unchanged.

Restricting the same 20 percent increase to the lowest-poverty quartile (Column 3) raises matches in those tracts by 57.8 percent while total matches rise 2.2 percent, with modest displacement elsewhere (-3.6 to -3.8 percent) as market-level congestion redirects voucher holders toward the newly attractive tracts, since the total voucher flow is fixed. At \$1,540 per match moved into Q1, this is the cheapest mobility price among the spending policies we consider.

Generosity-neutral rent indexing trades total matches for neighborhood access. Replacing standards with rent-indexed ones at constant average generosity (Column 4) lowers total matches by 5.5 percent, since standards fall in the cheap, high-poverty tracts where most baseline matches form, so Q3 and Q4 lose 7 to 8 percent, while Q1 matches rise 5.2 percent. Indexing also compresses rent manipulation. The mean markup falls from \$63 to \$39, the upcharging share of matches falls from 63 to 53 percent, and the program saves \$94 per matched voucher. Adding a hold-harmless floor (Column 5) converts the reform into a targeted generosity increase, with matches rising 3.5 percent at \$3,689 per marginal match and the largest gain in Q1, but by construction it gives up the budget savings. As Section 8.3 shows, the terms of this trade-off improve further the more the outside market responds.

Cash transfers to landlords are dominated by dose-matched payment standard increases. The 25 percent cost offset (Column 6), dose-matched to the 20 percent standard increase, raises matches by 5.7 percent, well under half the effect of its twin, at \$8,103 per marginal match, roughly 2.5 times the price of payment standard increases. The targeted version (Column 7) costs \$9,286 per marginal match and \$3,306 per match moved into Q1, about twice the price of targeted payment standards. The incidence rows explain why. Because the transfer is paid on every match, most of it accrues to landlords who would have participated anyway, and landlord surplus rises by \$299 per matched voucher, against \$91 under the uniform standard increase. The offset also has the opposite mobility effect. Its gains concentrate in the highest-poverty quartile ($+10.2$ percent) while Q1 matches are essentially flat.

Halving the expected contracting-stage markup (Column 8) cuts the mean markup on matched upcharging listings from \$63 to \$26 and saves \$29 per matched voucher, at the cost of a 0.5 percent decline in matches. Expected markups are part of what makes participation worthwhile, so removing them discourages some marginal landlords, and full elimination cuts matches by 0.7 percent. Landlord surplus also falls \$17.9 per matched voucher, while contract rents fall mechanically with the suppressed markup, since listed rents are essentially unchanged and the markup is the wedge between listed and contract rents among upcharging matches. Recycling the enforcement savings into targeted payment standards turns this into a budget-neutral mobility policy. When we cut the expected upcharging by 50 percent and use the savings to fund higher payment standards in Q1 tracts, Q1 matches rise 17.6 percent and total matches rise 0.3 percent, with a residual budget change below \$5 per matched voucher (see Appendix D.3 for details).

8.3 Sensitivity to the Outside-Market Response

We discuss how the choice of μ affects the counterfactual outcomes, with full results across μ reported in Appendix D.4. First, aggregate match effects shrink as the outside pool responds more, while costs per match rise. Fixing the outside pool under $\mu = 0$ roughly doubles to triples the aggregate effects, as the uniform standard increase rises from 14.6 to 40.4 percent and the economy-wide offset from 5.7 to 17.4 percent, and lowers the cost per marginal match, from \$3,270 to \$2,655 for the uniform increase. Full outside symmetry under $\mu = 1$ instead holds aggregate matches at baseline by construction, so only composition remains.

Second, the ranking of policy instruments is robust to μ. Payment standard increases dominate dose-matched direct transfers to landlords under the fixed pool (\$2,655 against \$4,486 per marginal match) and by more at the baseline (\$3,270 against \$8,103), since cost offsets lose their effects fastest as the outside market responds. Likewise, the compositional effects of low-poverty tract targeting remain essentially intact, with the Q1-targeted payment standard increase raising Q1 matches by 63.8, 57.8, and 53.0 percent under $\mu = 0$, 0.5, and 1.

Third, the performance of rent-indexed payment standards improves with μ. The reform's Q1 effect rises from $+0.2$ percent under $\mu = 0$ to $+5.2$ at the baseline and $+11.1$ under $\mu = 1$, while the losses in the high-poverty quartiles shrink. When the outside market re-prices along with listed units, the reform stops draining the listed market in cheap tracts and operates as intended, as a reallocation toward lower-poverty neighborhoods. Whether SAFMR-style indexing is a good policy therefore hinges on how broadly the rental stock responds to the payment standard schedule, precisely the margin the quasi-experimental evidence in Section 5 speaks to.

In summary, the counterfactual analysis yields three policy implications, robust to the choice of μ. First, payment standard increases pay only where the ceiling binds, so they deliver roughly two and a half times more matches per dollar than dose-matched cash transfers, which mostly subsidize landlords who would have participated anyway. Second, aggregate lease-up and neighborhood access call for different instruments. Uniform generosity expands matches roughly evenly across neighborhoods, rent indexing reallocates matches toward low-poverty tracts at the cost of total matches, and targeted payment standard increases buy access most cheaply, at \$1,540 per match moved into the lowest-poverty quartile. Third, the program can fund access out of its own inefficiency. Tightening rent-reasonableness enforcement claws back the upcharging transfer at a cost in participation, but recycling the savings into targeted payment standards raises low-poverty matches substantially at no net budget cost.

9 Conclusion

We assemble the first national linkage of online rental listings to administrative voucher contracts, together with a newly constructed panel of the program's rent ceilings, to study how landlords' participation and pricing decisions shape the Housing Choice Voucher program. The data reveal that landlords act on the program's payment rule at every stage. Targeting of voucher holders concentrates in high-poverty neighborhoods and predicts participation by an order of magnitude. Rents bunch at the program's rent ceiling, do so among listings only where landlords target vouchers, and mostly reach the ceiling between listing and contracting. The adjustment runs in both directions, with 3.66 percent of linked voucher units bunching at the standard and rents rising below the ceiling and falling above it by magnitudes with no counterpart in market transactions over the same range of rents. Quasi-experimental variation from ZIP-level payment standards confirms that both the participation and the pricing margins respond when the ceiling moves.

Motivated by this evidence, we develop and estimate a three-stage equilibrium model of landlord targeting, pricing, and matching. The estimates put a price on program participation. Leasing to a voucher holder costs the typical landlord \$1,245 per contract, about one month of rent, a cost that spans administrative burden and any aversion to voucher holders. Against it, the payment guarantee is worth roughly \$780 per month to the typical listing, expected upcharging adds \$63, and downcharging landlords surrender \$414 in rent, about half the value of the guarantee they secure. A voucher match is accordingly worthwhile for only 3 percent of listings, and the landlords who pay to target voucher holders are those for whom it pays, which is why participating supply concentrates in the neighborhoods the program hopes families can leave.

The counterfactuals turn these estimates into policy implications. Instruments that pay landlords only where the ceiling binds deliver roughly two and a half times more matches per dollar than unconditional transfers, most of which accrue to landlords who would have participated regardless. Aggregate lease-up and neighborhood access call for different instruments, and access is cheapest to buy directly, with rent ceiling increases targeted to low-poverty tracts moving a voucher holder into the lowest-poverty quartile at a monthly cost of roughly \$1,500. Lastly, the program can fund access out of its own inefficiency, since enforcing rent reasonableness to reduce upcharging and recycling the savings into targeted rent ceiling increases raises voucher matches in low-poverty tracts by nearly a fifth at no net budget cost.

These lessons speak to in-kind programs beyond housing. Wherever private providers choose whether to serve program participants and payment rules cap what the program pays, a similar logic applies, transfers to all participants subsidize the inframarginal, payments that bind only at the margin buy more participation per dollar, and the geography of willing providers, not only of eligible recipients, determines what the program delivers.

References

Abramson, B. and Nieuwerburgh, S. V. (2026). Rent guarantee insurance. *Journal of Political Economy*.

Agafiev Macambira, D., Geruso, M., Lollo, A., Ndumele, C. D., and Wallace, J. (2026). The private provision of public services: Evidence from random assignment in medicaid. *American Economic Review*, 116(6):2038–84.

Alatas, V., Banerjee, A., Hanna, R., Olken, B. A., Purnamasari, R., and Wai-Poi, M. (2016). Self-Targeting: Evidence from a Field Experiment in Indonesia. *Journal of Political Economy*, 124(2):371–427. Publisher: The University of Chicago Press.

Alexander, D. and Schnell, M. (2024). The impacts of physician payments on patient access, use, and health. *American Economic Journal: Applied Economics*, 16(3):142–177.

Aliprantis, D., Martin, H., and Phillips, D. (2022). Landlords and access to opportunity. *Journal of Urban Economics*, 129:103420.

Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L. F., and Palmer, C. (2024). Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice. *American Economic Review*, 114(5):1281–1337.

Bettinger, E. P., Long, B. T., Oreopoulos, P., and Sanbonmatsu, L. (2012). The Role of Application Assistance and Information in College Decisions: Results from the H&R Block Fafsa Experiment*. *The Quarterly Journal of Economics*, 127(3):1205–1242.

Bézy, T., Levy, A., and McQuade, T. (2025). Insuring landlords.

Bhargava, S. and Manoli, D. (2015). Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment. *American Economic Review*, 105(11):3489–3529.

Blomquist, S., Newey, W. K., Kumar, A., and Liang, C.-Y. (2021). On bunching and identification of the taxable income elasticity. *Journal of Political Economy*, 129(8):2320–2343.

Chetty, R., Friedman, J. N., Olsen, T., and Pistaferri, L. (2011). Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records *. *The Quarterly Journal of Economics*, 126(2):749–804.

Clemens, J. and Gottlieb, J. D. (2014). Do physicians' financial incentives affect medical treatment and patient health? *American Economic Review*, 104(4):1320–1349.

Collinson, R., Ellen, I. G., and Ludwig, J. (2019). Reforming Housing Assistance. *The ANNALS of the American Academy of Political and Social Science*, 686(1):250–285. Publisher: SAGE Publications Inc.

Collinson, R. and Ganong, P. (2018). How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality? *American Economic Journal: Economic Policy*, 10(2):62–89.

Coube, M., Felipe Fontes, L., and Rocha, R. (2026). Public services under private management. *The Quarterly Journal of Economics*, page qjag028.

Cunningham, M. K., Galvez, M. M., Aranda, C., Santos, R., Wissoker, D. A., Oneto, A. D., Pitingolo, R., and Crawford, J. (2018). A Pilot Study of Landlord Acceptance of Housing Choice Vouchers.

Damen, S., Korevaar, M., and Nieuwerburgh, S. V. (2025). An alpha in affordable housing?

Deshpande, M. and Li, Y. (2019). Who Is Screened Out? Application Costs and the Targeting of Disability Programs. *American Economic Journal: Economic Policy*, 11(4):213–248.

Desmond, M. and Perkins, K. L. (2016). Are Landlords Overcharging Housing Voucher Holders? *City & Community*, 15(2):137–162. Publisher: SAGE Publications.

Dube, A., Manning, A., and Naidu, S. (2025). Monopsony and employer misoptimization explain why wages bunch at round numbers. *American Economic Review*, 115(8):2689–2721.

Duggan, M., Gupta, A., Jackson, E., and Templeton, Z. S. (2023). The impact of privatization: Evidence from the hospital sector. Working Paper 30824, National Bureau of Economic Research.

Ellen, I. G., O'Regan, K., and Strochak, S. (2024). Race, Space, and Take Up: Explaining housing voucher lease-up rates. *Journal of Housing Economics*, 63:101980.

Ellen, I. G., O'Regan, K., and Strochak, S. (2025). Pricing for opportunity: The impact of spatially varying rent subsidies on housing voucher neighborhoods and take-up. *Journal of Public Economics*, 250:105465.

Ellen, I. G., O'Regan, K. M., and Harwood, K. W. H. (2023). Advancing Choice in the Housing Choice Voucher Program: Source of Income Protections and Locational Outcomes. *Housing Policy Debate*, 33(4):941–962. Publisher: Routledge _eprint: https://doi.org/10.1080/10511482.2022.2089196.

Eriksen, M. D., Kwon, E., and Yang, G. (2025). Scaling of small-area fair market rents: Evidence on neighborhood choices of voucher recipients.

Eriksen, M. D. and Ross, A. (2015). Housing Vouchers and the Price of Rental Housing. *American Economic Journal: Economic Policy*, 7(3):154–176.

Faber, J. W. and Mercier, M.-D. (2022). Multidimensional Discrimination in the Online Rental Housing Market: Implications for Families With Young Children. *Housing Policy Debate*, 0(0):1–24. Publisher: Routledge _eprint: https://doi.org/10.1080/10511482.2021.2010118.

Fack, G. (2006). Are housing benefit an effective way to redistribute income? evidence from a natural experiment in france. *Labour Economics*, 13(6):747–771. Evaluating Social Policy Changes in EU Countries (realized in collaboration with the National Poverty Centre and the European Union Centre at the University of Michigan).

Finkelstein, A. and Notowidigdo, M. J. (2019). Take-Up and Targeting: Experimental Evidence from SNAP*. *The Quarterly Journal of Economics*, 134(3):1505–1556.

Galiani, S., Gertler, P., and Schargrodsky, E. (2005). Water for life: The impact of the privatization of water services on child mortality. *Journal of Political Economy*, 113(1):83–120.

Galvez, M. (2010). What Do We Know About Housing Choice Voucher Program Location Outcomes? A Review of Recent Literature. Technical report, Urban Institute.

Garboden, P. M. E., Rosen, E., DeLuca, S., and Edin, K. (2018). Taking Stock: What Drives Landlord Participation in the Housing Choice Voucher Program. *Housing Policy Debate*, 28(6):979–1003. Publisher: Routledge _eprint: https://doi.org/10.1080/10511482.2018.1502202.

Goldin, J., Homonoff, T., and Meckel, K. (2022). Issuance and incidence: Snap benefit cycles and grocery prices. *American Economic Journal: Economic Policy*, 14(1):152–78.

Greenlee, A. J. (2014). More Than Meets the Market? Landlord Agency in the Illinois Housing Choice Voucher Program. *Housing Policy Debate*, 24(3):500–524. Publisher: Routledge _eprint: https://doi.org/10.1080/10511482.2014.913649.

Hart, O., Shleifer, A., and Vishny, R. W. (1997). The proper scope of government: Theory and an application to prisons*. *The Quarterly Journal of Economics*, 112(4):1127–1161.

Horn, K. M., Ellen, I. G., and Schwartz, A. E. (2014). Do Housing Choice Voucher holders live near good schools? *Journal of Housing Economics*, 23:28–40.

Humphries, J. E., Nelson, S., Nguyen, D. L., van Dijk, W., and Waldinger, D. (2025). Nonpayment and eviction in the rental housing market.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, 8(Volume 8, 2016):435–464.

Kleven, H. J. and Waseem, M. (2013). Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from pakistan *. *The Quarterly Journal of Economics*, 128(2):669–723.

Knutsson, D. and Tyrefors, B. (2022). The quality and efficiency of public and private firms: Evidence from ambulance services*. *The Quarterly Journal of Economics*, 137(4):2213–2262.

Lens, M. C. (2013). Safe, but Could Be Safer: Why Do HCVP Households Live in Higher Crime Neighborhoods? *Cityscape*, 15(3):131–152. Publisher: US Department of Housing and Urban Development.

McMillen, D. and Singh, R. (2020). Fair market rent and the distribution of rents in los angeles. *Regional Science and Urban Economics*, 80:103397. Special Issue on Housing Affordability.

Meckel, K. (2020). Is the cure worse than the disease? unintended effects of payment reform in a quantity-based transfer program. *American Economic Review*, 110(6):1821–65.

Park, J. (2026). Improving access to opportunity: Housing vouchers and residential equilibrium.

Phillips, D. C. (2017). Landlords avoid tenants who pay with vouchers. *Economics Letters*, 151:48–52.

Rosen, E. (2014). Rigging the Rules of the Game: How Landlords Geographically Sort Low–Income Renters. *City & Community*, 13(4):310–340. Publisher: SAGE Publications.

Saez, E. (2010). Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2(3):180–212.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199.

Susin, S. (2002). Rent vouchers and the price of low-income housing. *Journal of Public Economics*, 83(1):109–152.

Figures and Tables

Table 1: Summary Statistics

	HCV New Admissions		**RentHub Listings**	
	All	Linked	Scrape level	Spell level
	(1)	(2)	(3)	(4)
# Observations	1,256,063	140,245	287,613,176	22,693,226
# CBSAs	923	831	928	928
Rent (USD)				
25th Percentile	615	710	1,220	1,200
Median	800	950	1,707	1,730
75th Percentile	1,120	1,415	2,250	2,400
Mean	948.1	1,149.6	1,879.9	1,959.1
Standard Deviation	497.6	612.7	1,025.8	1,150.8
Payment Standard (USD)				
25th Percentile	726	805	1,017	1,139
Median	936	1,082	1,342	1,497
75th Percentile	1,279	1,583	1,721	1,950
Mean	1,086.3	1,284.0	1,434.6	1,620.6
Standard Deviation	534.0	669.9	561.9	669.6
Number of Bedrooms				
25th Percentile	1	1	1	2
Median	2	2	2	2
75th Percentile	3	2	3	3
Mean	1.88	1.85	2.12	2.27
Standard Deviation	0.95	0.96	1.07	1.05
Square Footage				
25th Percentile	—	672	800	858
Median	—	858	1,087	1,180
75th Percentile	—	1,084	1,456	1,605
Mean	—	934.4	1,204.5	1,310.8
Standard Deviation	—	477.1	576.4	694.8
Targeting composition (%)				
Hard-targeted	—	6.29	0.12	0.35
Soft-targeted	—	6.66	0.51	0.79

Note. The table reports summary statistics of rental unit characteristics, including rents (contract for HCV, listed for RentHub), number of bedrooms, payment standards, and square footage. Columns (1) and (2) use HUD voucher contract data restricted to new admissions, with Column (2) further restricted to admissions linked to a RentHub listing. Columns (3) and (4) use RentHub listings restricted to spells first posted in 2015–2023. Column (4) reports statistics at the spell level, and Column (3) at the level of the underlying raw scrapes that compose those spells. Payment standard statistics are computed on the subsample with non-missing payment standards. Summary statistics for that subsample are reported in Appendix Table A.8. Square footage information in Column (2) is obtained from RentHub.

Figure 1: Landlord targeting of voucher holders by neighborhood characteristics

(a) By neighborhood poverty

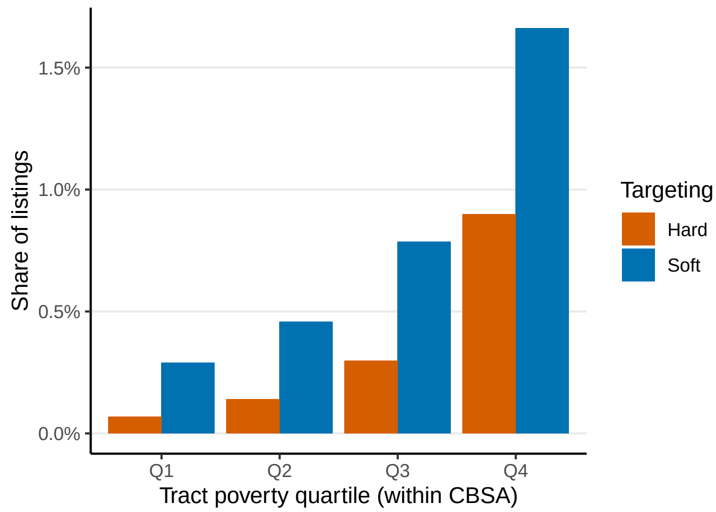

(b) Pairwise regressions

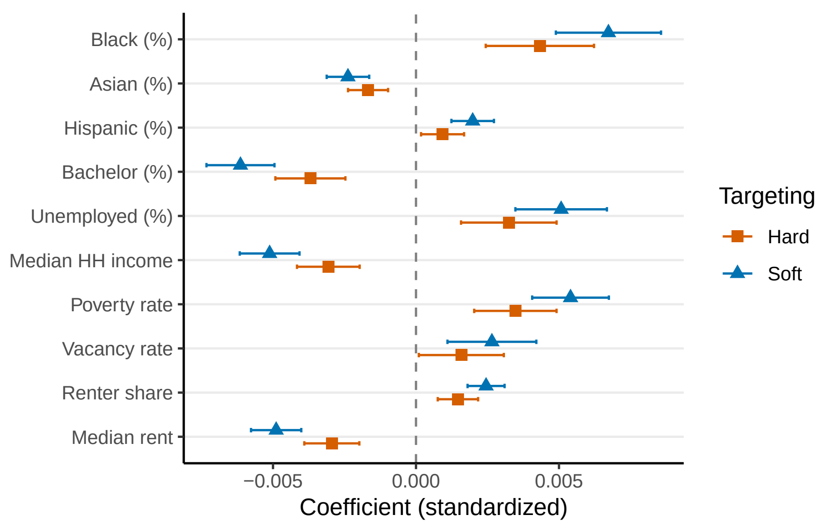

Note. Panel (a) shows the share of RentHub listings classified as hard-targeted (orange) or soft-targeted (blue) within each poverty quartile of the listing's Census tract, where quartiles are defined within the CBSA. Panel (b) plots coefficients from separate listing-level regressions in which the outcome is an indicator for hard-targeted (orange) or soft-targeted (blue) status and the regressor is a tract-level characteristic standardized within the CBSA. All regressions include CBSA-by-year and CBSA-by-bedrooms fixed effects; standard errors are clustered at the tract level. Hard-targeted listings appear on affordable housing platforms; soft-targeted listings appear on general platforms and contain voucher-seeking language, as defined in Section 3.2. Census characteristics are obtained from the 2016–2020 American Community Survey (5-year estimates).

Figure 2: Voucher participation across listing spells by targeting and neighborhood characteristics

(a) Voucher participation by neighborhood poverty

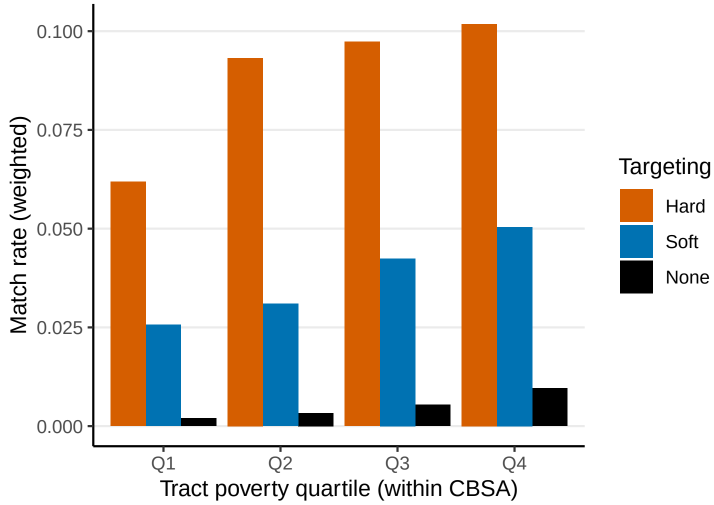

(b) Pairwise regressions

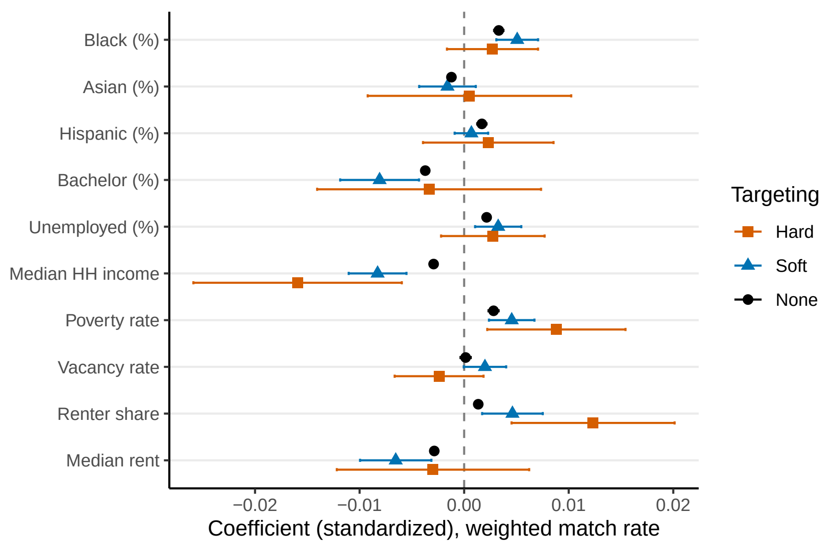

Note. Panel (a) shows the share of RentHub listing spells linked to an HCV unit within each poverty quartile of the listing's Census tract, where quartiles are defined within the CBSA, separately for hard-targeted (orange), soft-targeted (blue), and non-targeted (black) listings. Panel (b) plots coefficients from separate listing-level regressions in which the outcome is an indicator for whether the listing spell is linked to an HCV unit and the regressor is a tract-level characteristic standardized within the CBSA, estimated separately by targeting status. All regressions include CBSA-by-year and CBSA-by-bedrooms fixed effects, and standard errors are clustered at the tract level. Hard-targeted listings appear on affordable housing platforms while soft-targeted listings appear on general platforms and contain voucher-soliciting language, as defined in Section 3.2. Census characteristics are obtained from the 2016–2020 American Community Survey (5-year estimates). The linkage weight assigns a listing spell, for each HCV unit linked to it, a fraction equal to one over the number of spells that unit is linked to, so every unit contributes a total weight of one and unlinked spells contribute zero, as defined in Section 3.4. The estimated intercepts, which reflect the level differences in linkage rates across targeting types, are 0.0982 (s.e. 0.0085) for hard-targeted listings, 0.0434 (s.e. 0.0026) for soft-targeted, and 0.0050 (s.e. 0.0003) for non-targeted.

Figure 3: Histogram of contract and listed rents relative to the payment standard

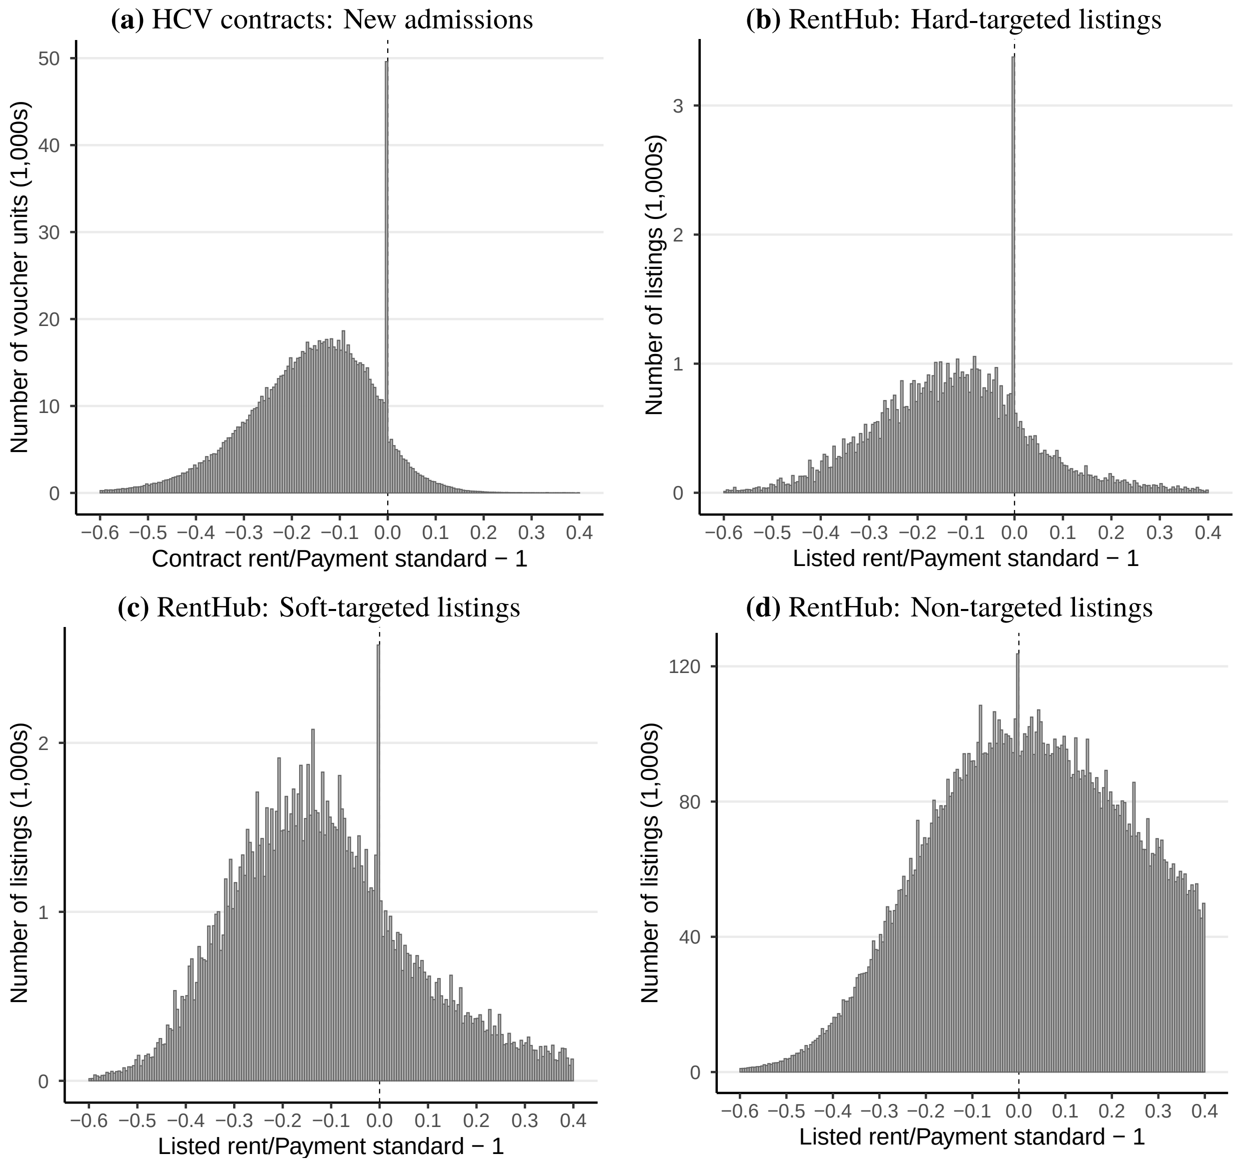

Note. The panels show the histogram of contract (panel a) and listed (panels b–d) rents relative to the payment standard, measured as (rent/payment standard) − 1. Panel (a) uses new admissions from HCV administrative data, 2015–2023. Panels (b)–(d) use RentHub listings over the same period, restricted to the geographic areas and years for which payment standards are available, as described in Section 3.3. Hard-targeted listings appear on affordable housing platforms; soft-targeted listings appear on general platforms and contain voucher-seeking language, as defined in Section 3.2. Bins are 0.5 percentage points wide. The vertical dashed line marks zero deviation from the payment standard.

Figure 4: Histogram of contract and listed rents for HCV units linked to listings

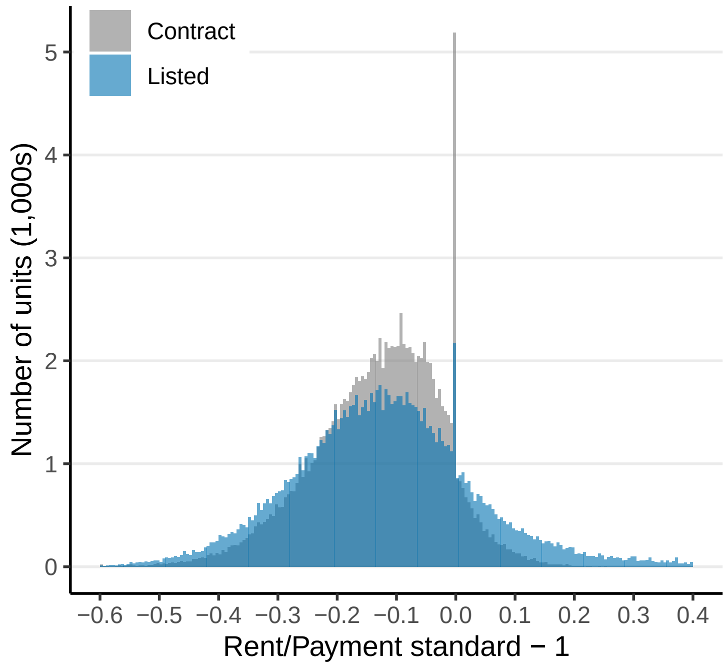

Note. The plot shows the histogram of contract (gray) and listed (blue) rents relative to the payment standard, measured as (rent/payment standard) − 1. The sample includes new admissions in HCV administrative data from 2015 to 2023 that are linked to a RentHub listing spell, as described in Section 3.4. Bins are 0.5 percentage points wide.

Figure 5: Differences-in-bunching: Contract vs counterfactual rents for linked units

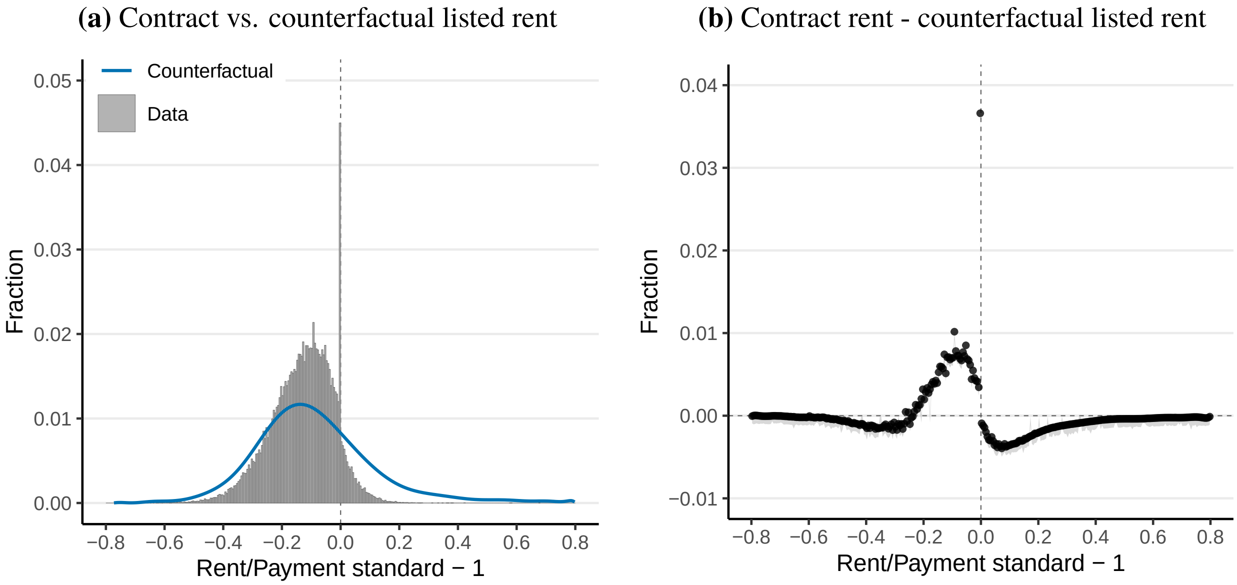

Note. Panel (a) plots the histogram of contract rents for new admissions in HCV administrative data, 2015–2023, that are linked to a RentHub listing spell (gray), together with the counterfactual distribution of listed rents (black), computed from linked listed rents using the differences-in-bunching estimator described in Section 4. Panel (b) plots the difference between the two, with 95% confidence bands (gray) from the iterative bootstrap procedure described there. Bins are 0.5 percentage points wide.

Figure 6: Listed-to-contract rent changes around payment standards, by market

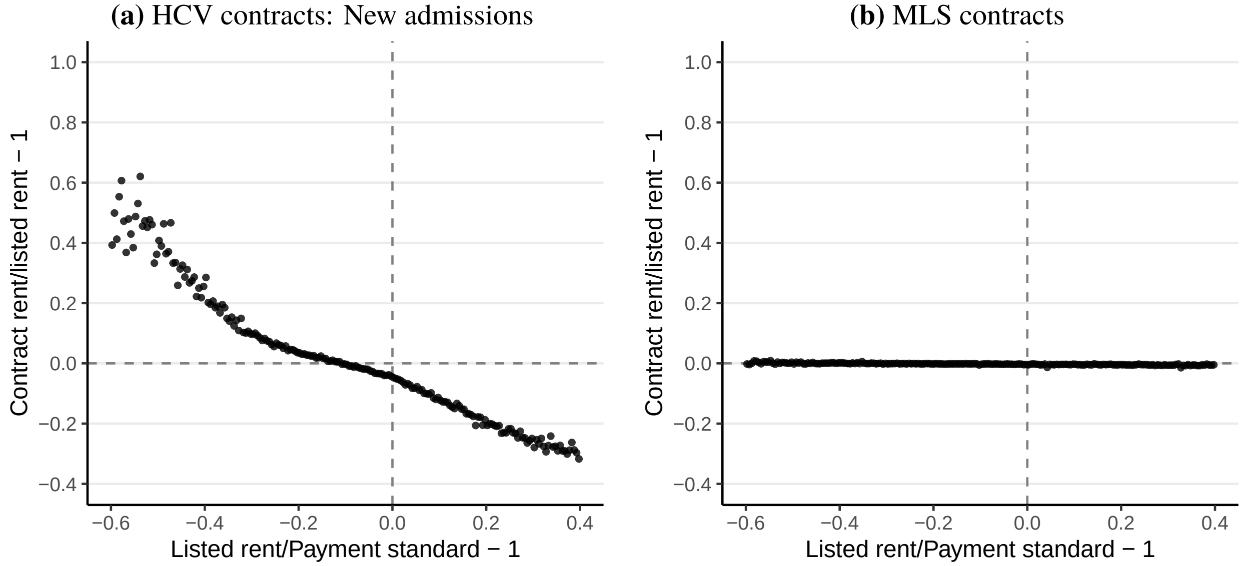

Note. The figure shows binscatter plots of the mean rent change from the listing to the contract stage, as a function of the listed rent relative to the payment standard, measured as (listed rent/payment standard) − 1. Panel (a) uses 2015-2023 HCV new admissions linked to RentHub listing spells; panel (b) uses contracts for units listed in Multiple Listing Service (MLS) as a comparison market. Bins are 0.5 percentage points wide.

Figure 7: Pairwise regressions of listing-to-contract rent changes on voucher-holder and tract characteristics for linked units, by listed rent position relative to payment standards

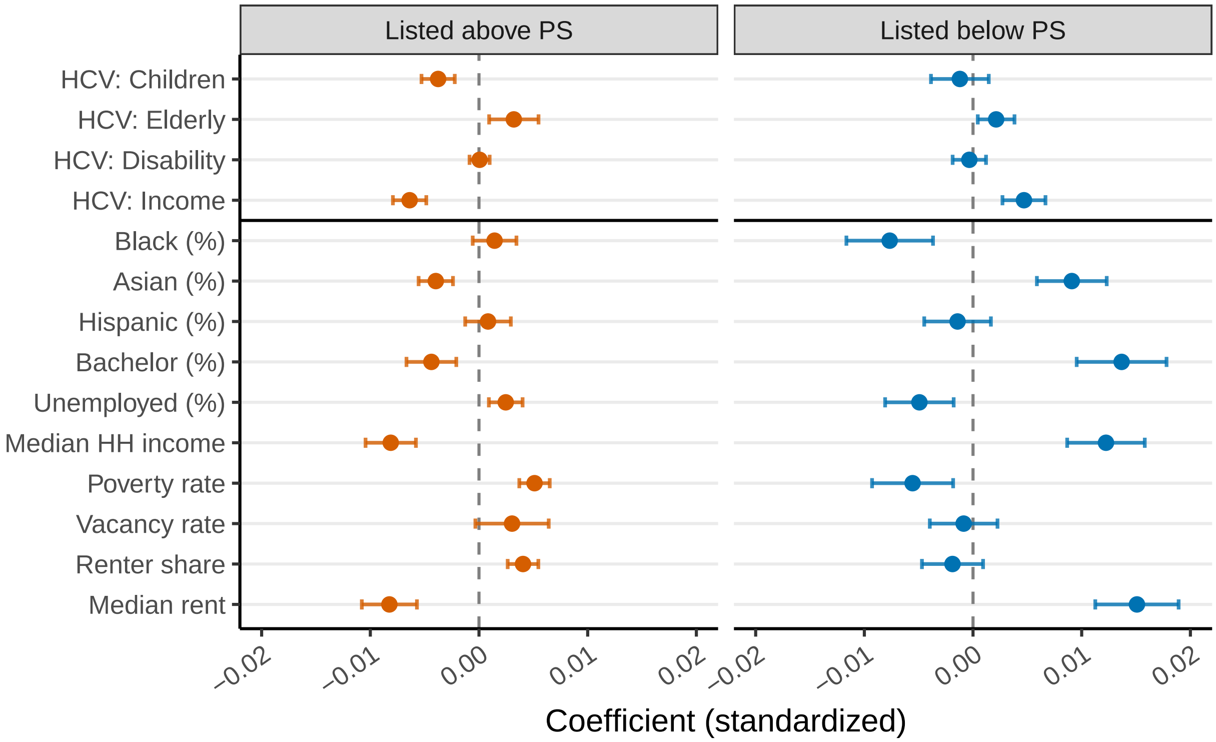

Note. Each panel reports the coefficient from a separate regression of the relative difference between contract and listed rent, defined as (contract rent/listed rent) − 1, on a single standardized characteristic. All regressions control for the listed rent and include PHA-by-year and PHA-by-bedroom fixed effects; standard errors are clustered by PHA. Regressions are estimated separately by whether the listed rent lies above the payment standard. Panel (a) uses the negative of the relative difference; panel (b) uses the relative difference itself. Census characteristics are drawn from the 2016–2020 American Community Survey (5-year estimates).

Figure 8: Event studies: The impacts of payment standard adoption at the ZIP-level, by sign of the change of payment standard

(a) Outcome: Payment standard, log

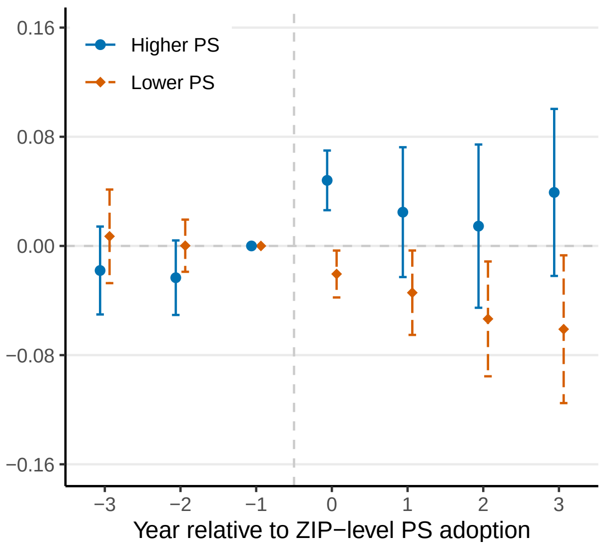

(b) Outcome: Targeted listing (hard or soft)

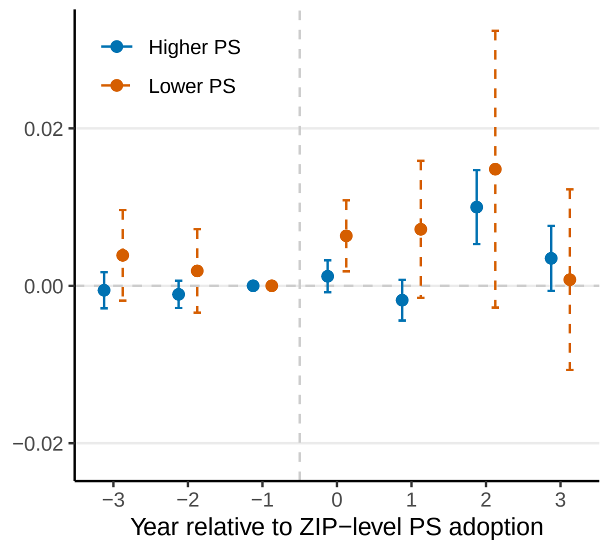

(c) Outcome: $\mathbb{1}$(Contract rent = Payment standard)

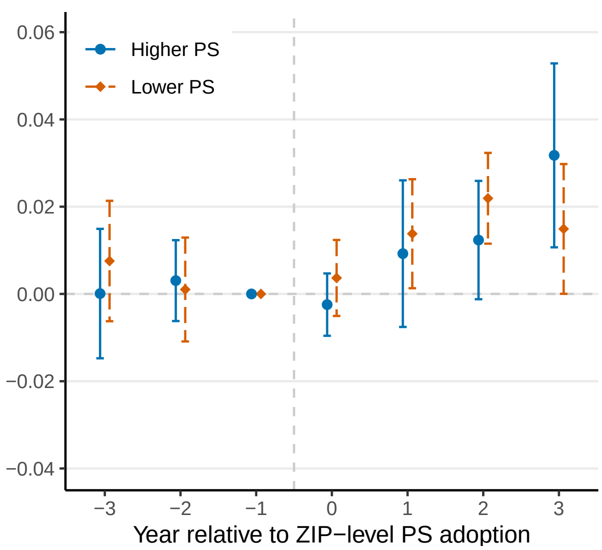

(d) Outcome: Tract poverty rate

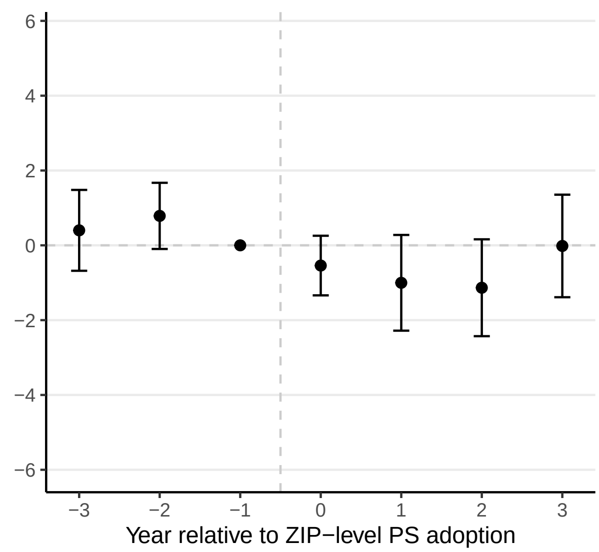

Note. The figure plots the event study coefficients β_e from Equation (2), estimated with the interaction-weighted estimator of Sun and Abraham (2021), together with 95% confidence intervals. Each panel uses a different outcome: panel (a), the log payment standard reported by HUD; panel (b), an indicator for a listing being hard- or soft-targeted; panel (c), an indicator for the contract rent equaling the payment standard; and panel (d), the tract poverty rate. Panels (a), (c), and (d) use HCV new admissions over 2015–2023; panel (b) uses RentHub listing spells over the same period. In panels (a)–(c), each panel reports two separate regressions: one for ZIP codes where ZIP-level PS adoption is predicted to raise the payment standard (black) and one where it is predicted to lower it (gray). For this purpose, we classify ZIP codes by the difference between a ZCTA's 40th-percentile rent and its CBSA's 40th-percentile rent, both measured in the 2011–2015 American Community Survey (5-year estimates).

Table 2: Stage 3 and Stage 2 parameter estimates

Stage 3, upcharging ($R \leq S$)		**Stage 3, downcharging** ($R > S$)		**Stage 2, voucher allocation**	
Constant (μ_0^+)	3.410	Constant (μ_0^-)	2.408	$(S-R)_+$ per \$ (θ_1)	0.00038
	(0.180)		(0.336)		(0.00009)
Headroom d^+ per \$ (μ_1^+)	0.00248	Gap d^- per \$ (μ_1^-)	0.00017	$(R-S)_+$ per \$ (θ_2)	−0.00222
	(0.00003)		(0.00003)		(0.00022)
Poverty	0.142	Poverty	0.002	Poverty	2.003
	(0.044)		(0.080)		(0.103)
Vacancy	0.268	Vacancy	−0.017	Vacancy	−0.514
	(0.067)		(0.111)		(0.193)
Log median rent	0.074	Log median rent	0.238	Log median rent	−1.152
	(0.022)		(0.033)		(0.040)
Soft targeting	0.006	Soft targeting	−0.049	Soft targeting	1.657
	(0.019)		(0.037)		(0.042)
Hard targeting	0.155	Hard targeting	−0.121	Hard targeting	2.399
	(0.020)		(0.032)		(0.054)
No markup (π^{0+})	0.559	No reduction below S (π^{0-})	0.367	Beds = 1	0.630
	(0.002)		(0.003)		(0.080)
Shape (α^+)	1.357	Shape (α^-)	1.244	Beds = 2	0.159
	(0.006)		(0.009)		(0.075)
				Beds = 3	−0.393
					(0.081)
				Beds = 4	−0.687
					(0.089)
Contracts	82,276	Contracts	19,155	Observations	15,134,940

Notes. This table reports Stage 3 and Stage 2 parameter estimates. Stage 3 is maximum likelihood, estimated on the linked listing–contract sample, with bedroom coefficients suppressed. Stage 2 is the fixed-effects logit, estimated on the 2015–2023 listings panel restricted to the Stage 3 estimation universe, with CBSA × year fixed effects and bedroom indicators. Standard errors are clustered at the CBSA × year level.

Table 3: Stage 1 estimates

Parameter	Estimate	SE
Program participation cost, c^{HCV} ($)	$1{,}245.2$	64.5
Soft-targeting cost, c^V_{soft} ($)	5.45	1.97
Hard-targeting cost, c^V_{hard} ($)	26.01	6.34
Leasing prospect intercept, γ_0	0.159	0.001
Leasing prospect rent premium, γ_R	-2.206	0.003
tract poverty rate, γ_{pov}	-1.056	0.464
tract vacancy rate, γ_{vac}	0.074	0.522
tract log median rent, γ_{rent}	-0.094	0.094
1 bedroom, γ_{B1}	0.357	0.539
2 bedrooms, γ_{B2}	-0.326	0.488
3 bedrooms, γ_{B3}	-0.712	0.390
4 bedrooms, γ_{B4}	0.295	0.875
Shock scale, σ_q (calibrated)	0.19	–

Notes. This table reports Stage 1 parameter estimates, obtained from the minimum-distance estimates of Equation (10) at the best converged multi-start solution (objective 8.54×10^{-4}) on 181 CBSAs, 2015–2023, and 1.64 million listing cells. Costs in dollars per listing episode are estimated on the model's monthly-rent revenue scale. Covariates are centered at listing-weighted means, so γ_0 provides the prospect at the benchmark rent for the average listing. Standard errors are conditional on the Stage 2–3 estimates and σ_q and thus understate total sampling uncertainty.

Figure 9: Model fit, stages 3 and 2

(a) Stage 3, mean rent adjustment by distance to the payment standard

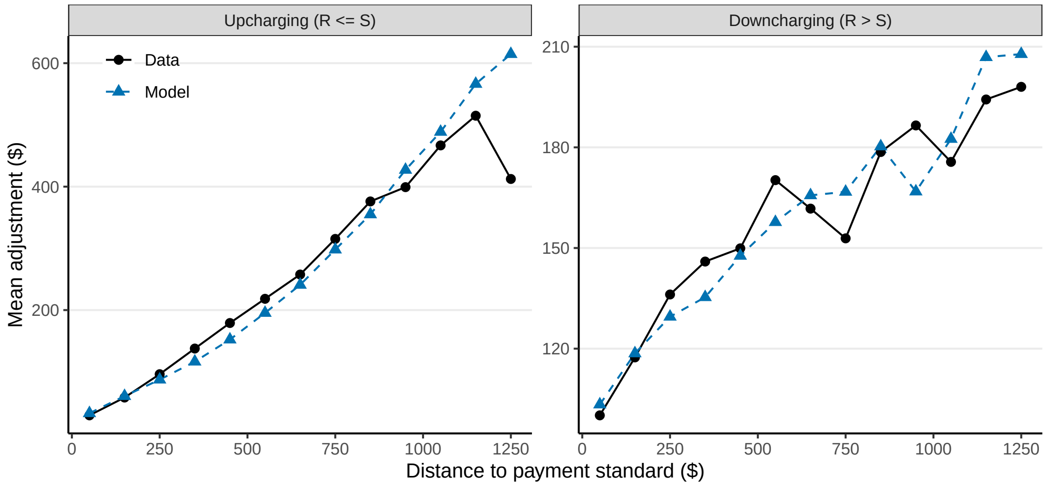

(b) Stage 2, match rate by tract covariate decile

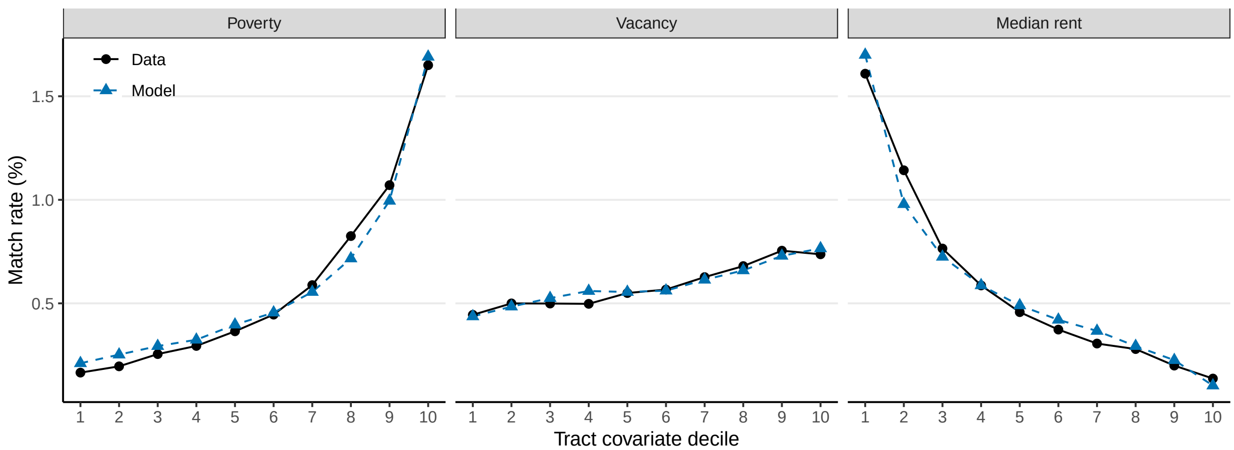

Notes. Panel (a) plots the mean contract rent adjustment among interior adjustments by $100 bins of the distance to the payment standard (headroom and gap), for the upcharging and downcharging regions of the linked contract sample, with bins beyond $1,200 omitted and the point masses discussed in the text. Panel (b) plots the share of listings matched to a new voucher admission by listing-weighted deciles of the tract poverty rate, vacancy rate, and log median rent, against the allocation-model prediction π_{jt}^{HCV} evaluated on the same listings.

Figure 10: Model fit, stage 1

(a) Stage 1, targeting and bunching by within-CBSA poverty quartile

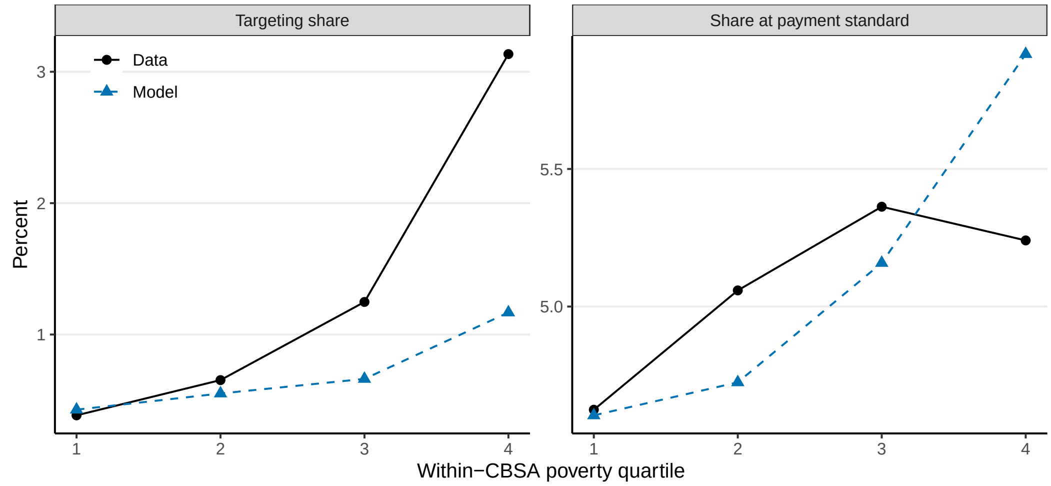

(b) Stage 1, fit across CBSA

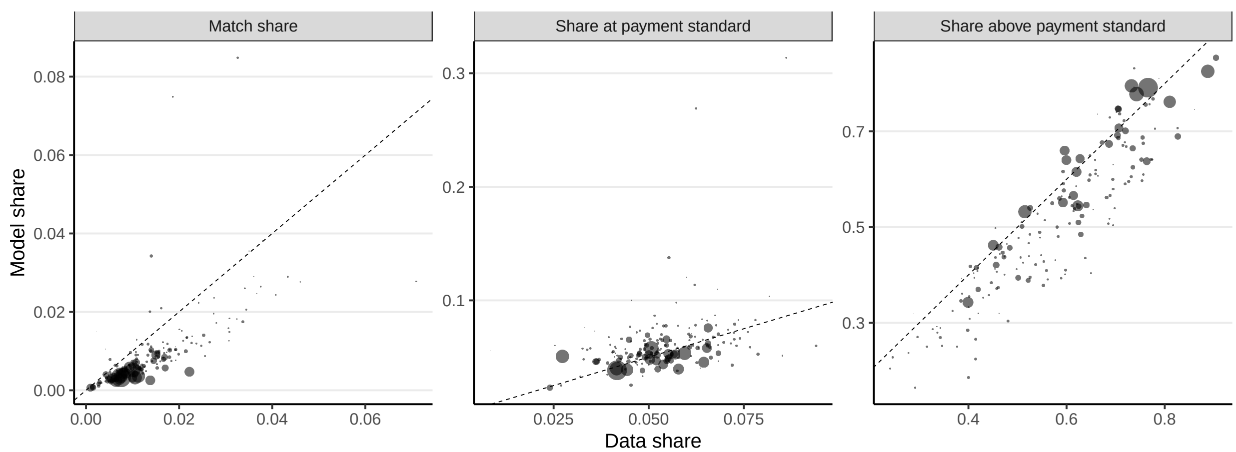

Notes. Panel (a) plots the share of listings that target voucher tenants and the share priced within 2 percent of the payment standard by within-CBSA quartile of the tract poverty rate, with listing-weighted quartile cutoffs computed from the data and held fixed for the model. In panel (b) each point is one of the 181 estimation CBSAs, with marker area proportional to the CBSA's listing count, the dashed line marking equality, and axis ranges set per panel.

Table 4: Counterfactual outcomes

	(1) Baseline	(2) Standards (+20%), all tracts	(3) Standards (+20%), Q1 tracts	(4) Rent-indexed	(5) Rent-indexed, hold-harmless	(6) Cost offset (25%), all	(7) Cost offset (25%), Q1 tracts	(8) Enforcement (half markup)
Matches								
Matches (thousands)	74.0	84.9	75.7	69.9	76.7	78.3	74.4	73.7
change (%)	–	+14.6	+2.2	−5.5	+3.5	+5.7	+0.5	−0.5
Change in matches by poverty quartile (%)								
Q1, lowest (base 7.1k)	–	+18.1	+57.8	+5.2	+8.2	−0.2	+16.0	−0.8
Q2 (base 11.7k)	–	+16.3	−3.8	−0.9	+5.0	+0.5	−1.1	−0.7
Q3 (base 18.3k)	–	+13.8	−3.6	−6.8	+2.2	+2.4	−1.0	−0.5
Q4, highest (base 36.9k)	–	+13.9	−3.7	−8.4	+2.8	+10.2	−1.1	−0.4
Rent manipulation, matched listings ($)								
Mean markup	63	137	83	39	63	63	64	26
Mean concession	414	370	409	436	399	404	413	414
Budget and incidence, per matched voucher ($/month)								
Δ rent payments	–	+418	+83	−94	+126	+128	+16	−29
Transfers to landlords	–	–	–	–	–	+311	+34	–
Δ landlord surplus	–	+90.7	+17.2	−13.4	+16.3	+298.8	+32.0	−17.9
Cost effectiveness ($)								
Cost per marginal match	–	3,270	3,853	–	3,689	8,103	9,286	–
Cost per match moved into Q1	–	–	1,540	–	–	–	3,306	–

Note. The table reports the model-baseline equilibrium and seven counterfactual equilibria under $\mu = 0.5$, in which half of the unobserved outside pool responds to policy symmetrically with observed listings. Column (1) is the status-quo baseline. Column (2) raises payment standards by 20 percent in all tracts and column (3) only in the lowest-poverty within-metro quartile of tracts. Column (4) replaces payment standards with rent-indexed standards holding metro-level generosity fixed, and column (5) adds a hold-harmless floor at the baseline standard. Column (6) offsets 25 percent of landlords' voucher participation cost on all matches and column (7) only in the lowest-poverty quartile. Column (8) halves the expected upcharging markup through rent-reasonableness enforcement. Poverty quartiles are baseline within-metro quartiles of tract poverty, held fixed across scenarios. Mean markup and mean concession are the mean expected upward and downward rent adjustments among matched upcharging and downcharging listings, respectively. Budget and incidence entries are per matched voucher per month, and transfers appear only for the offset scenarios. Cost per marginal match divides incremental spending, rent payments plus transfers, by the change in matches. Cost per match moved into Q1 divides the same spending by matches gained in that quartile. Both are reported only where spending rises. Details and definitions are in Appendix D.

Appendix

A Payment Standards: Dataset Construction

A.1 Overview

Payment standards are the maximum gross rent a Public Housing Authority (PHA) will subsidize under the Housing Choice Voucher (HCV) program. Each PHA sets its own payment standards, expressed as a dollar amount per bedroom size, within an area defined by HUD's Fair Market Rents (FMR) or, for designated metropolitan areas, Small Area Fair Market Rents (SAFMR). Because payment standards are not directly reported in any centralized HUD administrative file, we construct them from two complementary sources: (i) systematic inference from HUD's Form 50058, and (ii) manual collection from PHA websites and Freedom of Information Act requests. We reconcile the two sources into a single dataset. The HUD-inferred data cover February 2015 through September 2023 (the span of available Form 50058 microdata); the manually collected data extend through the end of 2025. The remainder of this section describes the sample of PHAs covered by the dataset and the decisions we made to assign a payment standard to each PHA-period combination.

A.2 Sample of PHAs

We start from the sample of PHAs present in the HUD Public and Indian Housing (PIH) 50058 administrative records. We impose the following sequential restrictions:

1. *Size*: We retain PHAs that administered at least 1,000 voucher units in 2015, yielding 481 PHAs (after dropping the Virgin Islands authority, which is out of our scope). This restriction is made to guarantee a minimum of voucher lease-up activity to reduce noise in our inference methodology.

2. *Moving to Work (MTW) exclusion*: When inferring payment standards from HUD data, we exclude MTW PHAs because these PHAs have missing information in the payment standard field. This exclusion leaves a sample of 472 PHAs for inferred payment standards.

We classify each PHA as *regional* or *local*. Regional and state PHAs are identified by the third character of the PHA code being "9," or by the keywords "state" or "region" appearing in the formal PHA name or contact email. Regional PHAs receive additional restrictions described below.

A.3 HUD-Inferred Payment Standards

We infer payment standards from HUD administrative data for each PHA-period combination, where a period is the interval during which a given set of FMR or SAFMR values is in effect according to HUD. These periods usually begin in October of each year, when HUD issues updated FMRs/SAFMRs, and end either when mid-year updates take effect or at the following October.

A.3.1 Modal Method

For each PHA, period, and bedroom size, we identify the *most frequent* (modal) payment standard from the observed distribution of contracts among leased-up units, dropping any payment standards outside the \$100–\$8,000 range. Formally, let $\mathcal{S}_{j,t,b}$ denote the multiset of payment standards reported in PHA j's lease-up records during FMR period t for bedroom size b, and let $\hat{s}_{j,t,b}$ denote the modal value. A PHA-period cell is declared "valid" if, for either 1-bedroom *or* 2-bedroom units, the modal share clears a minimum threshold τ, which we set at 55 percent, and the cell has at least 10 observations. Formally, the restriction can be expressed as

$$\frac{\#\{r \in \mathcal{S}_{j,t,b} : r = \hat{s}_{j,t,b}\}}{|\mathcal{S}_{j,t,b}|} \geq \tau, \qquad \tau \geq 0.55, \tag{12}$$

together with $|\mathcal{S}_{j,t,b}| \geq 10$. When a cell qualifies, we assign the modal value to all bedroom sizes 0–4 in that PHA-period. When inference fails for isolated periods, we fill gaps as follows: a single missing period takes the prior period's value; a two-period gap is filled only if the values on both sides agree. We use newly admitted voucher holders (action type code = 1 with admission year $\geq$ 2015) as the primary sample, and keep results using all voucher admissions as a fallback.

Overall, this method yields valid payment standards for approximately 454 PHAs across at least some portion of the study period.

A.3.2 FMR– and SAFMR–Multiples Method

As a complementary approach, and given that most PHAs set payment standards as a fixed percentage of FMR, we infer the "FMR multiple" that PHAs use to set payment standards. For each PHA, FMR period, and bedroom size, we test whether a given share of observed payment standards matches any of the canonical multiples:

$$m \in \{0.90,\ 0.95,\ 1.00,\ 1.05,\ 1.10,\ 1.15,\ 1.20,\ 1.25\}$$

applied to the county-level FMR. We assign a multiple m if:

$$\frac{\#\{s \in \mathcal{S}_{j,t,b} : s = m \times \text{FMR}_{c,t,b}\}}{|\mathcal{S}_{j,t,b}|} \geq \tfrac{1}{3}, \tag{13}$$

subject to $|\mathcal{S}_{j,t,b}| \geq 10$, retaining the highest-share multiple when several qualify. FMR values come from HUD's publicly available FMR schedule. We fill 1- and 2-period gaps using the same logic as in the modal method, and likewise run the inference separately on new admissions (primary) and all admissions (fallback). For regional PHAs operating across multiple counties, we restrict inference to PHA–by–county cells with at least 100 active vouchers in 2015 to ensure sufficient data density for reliable inference. This method yields valid inferences for 401 PHAs.

For PHAs in SAFMR-designated metropolitan areas, we follow an analogous procedure that infers the SAFMR multiple using HUD's ZIP-code-level Small Area Fair Market Rents. SAFMR-multiple inference is run as a single pass (not split into new vs. all admissions). SAFMR inference begins in October 2018, the month in which the SAFMR mandate took effect for most designated PHAs.[26] Inferred SAFMR multiples are then mapped from each PHA to the ZIP codes within the PHA's HUD metro area using HUD's published ZIP-to-metro crosswalk.

Overall, this method yields valid payment standards for approximately 401 and 63 PHAs for FMR–multiples and SAFMR–multiples, respectively, across at least some portion of the study period. Figure A.1 shows the frequency of several multiples of FMR and SAFMR in our final dataset using 2-bedroom payment standards for which this variable is non-missing.

A.4 Manual Collection of Payment Standards

For a subset of large or SAFMR-covered PHAs, we supplement inference with payment standards collected directly from PHA documents, including administrative plans, FOIA responses, and archived web pages. We target the largest 110 PHAs considered for the SAFMR implementation in 2018, of which we obtain data for 76. The collected data record the number of bedrooms, the payment standard, the geographic unit at which it applies, and the effective date.

Most of these PHAs set payment standards at the metropolitan area or ZIP code level, though several publish them at other geographies: county, submarket/tier, or Census tract. We account for this variation by constructing the manual panel separately for two sets of geographies.

Metro, county, and submarkets. We merge the metro, county, and submarket datasets into a single PHA-geography-bedroom panel. This is feasible because submarket definitions generally

[26] Eight PHAs with documented early adoption are assigned SAFMR-inferred standards starting earlier, based on visual inspection of their lease-up records: IL101, TN004, TX009, TX392, TX434, TX436, TX559, and TX901.

map to county boundaries.[27]

ZIP codes. We separately process PHAs that publish ZIP code-specific standards, the majority of which are PHAs in SAFMR-mandated metropolitan areas. For agencies that list payment standards for explicitly named ZIP codes alongside an *"All Other ZIP codes"* category (`CA006`, `CA072`, `NY409`), we expand the *"All Other ZIP codes"* category to all remaining ZIP codes in the PHA's county jurisdiction using HUD's USPS 2015 ZIP-to-county crosswalk.

A.5 Final Payment Standard Dataset

We combine the manual and inferred datasets into a unified monthly panel in three steps. First, we construct a PHA-level timeline by pooling all effective dates from the manual schedules and the inferred FMR/SAFMR period boundaries. Within each PHA, we assign the geographic level (metro or ZIP) in each time segment to the most granular source available, giving precedence to manual over inferred data. A PHA may therefore transition from metro-level to ZIP-level coverage mid-panel as manual ZIP-specific schedules become available. We also apply a small number of PHA-specific corrections at this stage, dropping inferred metro-level values when we can verify from the schedule that the PHA was using ZIP-level standards, and dropping legacy metro rows from the manual data once ZIP-level information takes over.

Second, we expand the timeline to monthly resolution and populate each row with every available source-specific payment standard value. After this step, each row carries all inferred and manual estimates alongside one another, regardless of which source we ultimately treat as canonical.

Third, we apply the priority hierarchy in Table A.1 to select a canonical payment standard from all three methods, assigning each observation the lowest-priority non-missing value. Figure A.2 reports the number of PHAs covered by each source. Figure A.3 assesses agreement across sources by plotting the distribution of percent deviations between modal-inferred and manually collected payment standards, restricted to PHA-month cells where both are available. The two sources match exactly in more than 60 percent of these observations.

The resulting dataset is our final payment standards dataset, which contains payment standard information across a total of 474 PHAs.

[27] A small number of PHAs whose submarket definitions are based on custom maps without correspondence to any administrative boundary, such as Atlanta, cannot be included in this step.

A.6 ZIP-Level Version of the Payment Standard Dataset

For the analyses with rental listings, where it is not clear which PHA serves the listing's unit, we build a curated dataset of payment standards at the ZIP code level from the dataset in the previous subsection.

We first expand metro, county, and submarket rows to the ZIP-code level using a ZCTA-to-county crosswalk (or, for New England states where FMR values vary at the county-subdivision level, a ZCTA-to-county-subdivision crosswalk). For metro-level payment standards, we only capture ZIP codes of the county where the PHA is located.[28]

For ZIP codes covered by more than one PHA, we retain the PHA with the most voucher holders in that ZIP code; if no PHA has voucher holders there, we retain the largest PHA by total Section 8 unit count. Each (ZIP code, bedroom size, month) cell appears exactly once in this curated panel.

Figure A.1: Frequency of FMR/SAFMR multiples in inferred payment standards

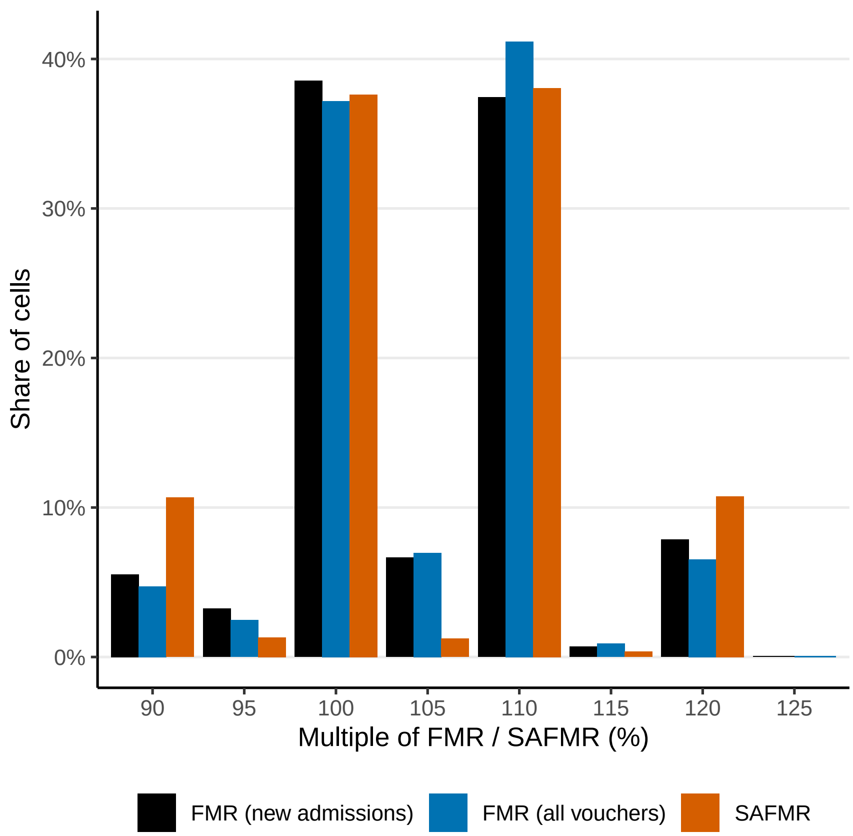

Note. The figure shows the share of PHA-month combinations assigned to each canonical multiple of FMR or SAFMR, restricted to the sample of 2-bedrooms and PHA-months for which the multiples method yields a non-missing payment standard.

[28]For multi-county metro PHAs such as the New York City Housing Authority, which spans five boroughs, we replicate metro-level rows across all counties served.

Table A.1: Payment standard priority hierarchy

Priority	Source	Column
1	Manual (directly collected)	`ps_manual`
2	HUD-inferred, modal PS, new admissions	`ps_infer_new`
3	HUD-inferred, FMR multiple, new admissions	`ps_infer_fmr_new`
4	HUD-inferred, modal PS, all vouchers	`ps_infer_all`
5	HUD-inferred, FMR multiple, all vouchers	`ps_infer_fmr_all`
6	HUD-inferred, SAFMR multiple (ZIP-level only)	`ps_infer_safmr`

Figure A.2: Number of PHAs by source of payment standard data

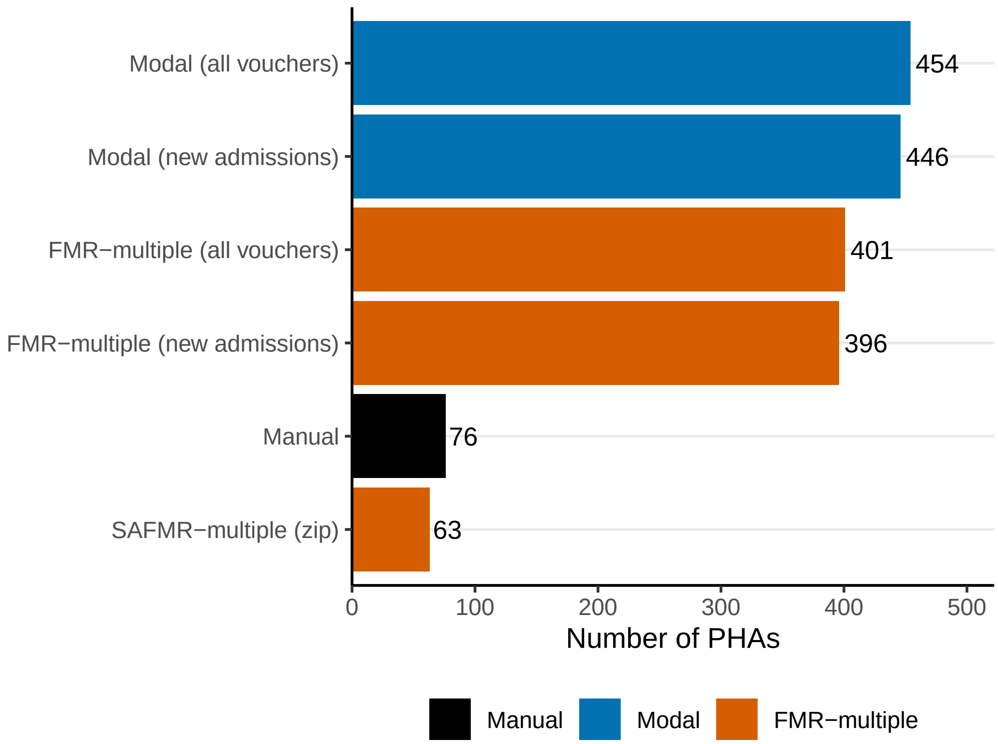

Note. The figure shows the number of unique PHAs with non-missing payment standard information from each source used to construct the final dataset.

Figure A.3: Agreement between modal-inferred and manually collected payment standards

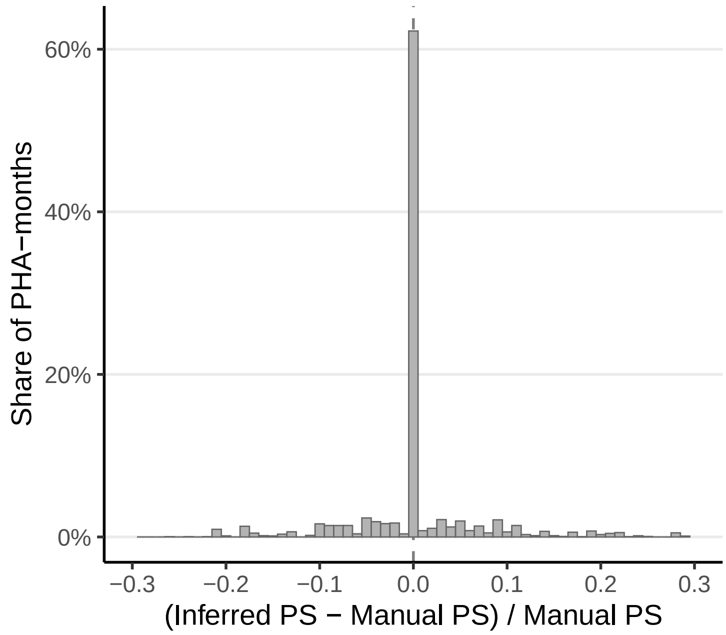

Note. The figure plots the histogram of the relative deviation between modal-inferred and manually collected payment standards, using 1 percentage point bins. The sample includes all PHA-month observations for 2-bedroom units with non-missing values for both sources; each PHA-month receives equal weight.

B RentHub Listing Data Details

This appendix documents the RentHub listing panel used in Section 3.2. We describe how coverage varies across states and years (Appendix B.1), how we chain scrapes into listing spells (Appendix B.2), and how we classify voucher targeting (Appendix B.3).

B.1 Coverage statistics

Figure A.4 shows the number of RentHub scrapes by state and year. Coverage is broad but uneven. It concentrates in the most populous states, with Texas, Florida, California, and New York contributing the most scrapes, and it is thinner through the middle of the sample before recovering in 2023. Several state-year cells are sparse or empty, shown in gray. The District of Columbia is absent until 2023, and Hawaii is missing in 2019 and thin in the adjacent years. Sixteen states have at least one stretch of statewide scraping silence longer than 45 days.

B.2 Constructing listing spells

As described in Section 3.2, we chain a unit's consecutive scrapes into a listing spell and start a new spell whenever the gap between scrapes exceeds 60 days. Figure A.5 plots the distribution of these within-unit gaps, measured across covered periods so that statewide coverage holes do not inflate the tail. The mass sits at short gaps, with a median of 6 days and about 91 percent of gaps at 21 days or fewer, which reflects the weekly to bi-weekly scrape cadence of an active listing. Beyond this cadence the distribution has a pronounced valley. Gaps in the 46-to-90-day range are rare, under about 1 percent, before a long thin tail of genuine re-listings. The 60-day break sits inside this valley, so the exact threshold is not consequential, and it leaves roughly 97 percent of consecutive gaps within a spell. A 60-day cut partitions a unit's scrapes into spells that average about eight scrapes spanning roughly seven weeks on the market.

We anchor each spell on its first scrape date rather than on the listing's posting date. The posting date is almost always present, but only about 39 percent of units carry a single posting date across their scrapes, since the field is refreshed over the life of a listing, so it is an unreliable marker of when a unit first appeared.

B.3 Voucher targeting classification

Voucher Hard-Targeting

RentHub records a management or listing company for many scrapes. We map each distinct company string to one of six groups using a hand-built crosswalk. The groups are affordable-housing platforms (e.g., AffordableHousing.com), affordable-housing management organizations, mainstream rental platforms, institutional owners or managers, brokerages, and a residual other category. We aggregate to the spell by counting the distinct companies in each group. We classify a spell as hard-targeted if it lists on an affordable-housing platform or with an affordable-housing management organization, since these channels exist specifically to serve voucher holders.

The company and platform metadata are sparse in the early years and fill in sharply from 2019 onward. Figure A.6 plots the share of spells with any company or platform metadata by anchor year. Coverage is near 4 percent in 2015 and stays below 12 percent through 2018, then jumps to 69 percent in 2019 and holds between roughly 70 and 88 percent thereafter. Because a hard-targeted listing can be identified only through these fields, hard targeting is effectively unmeasured before 2019, at below 0.02 percent of spells, and those early-year shares should be read as uninformative rather than as lower bounds. From 2019 on, the hard-targeted share runs from about 0.2 to 1.0 percent of spells. It is not monotonic across years and dips in 2022 even though coverage peaks that year, so single-year comparisons of the hard share should be made cautiously.

Figure A.7 shows the composition of company groups among the spells that carry company information. Affordable-housing platforms and organizations together account for under 1 percent of these spells. The bulk are mainstream platforms at about 36 percent, the residual other category at 27 percent, brokerages at 22 percent, and institutional managers at 14 percent. Across all spells, including the roughly 49 percent with no company information, affordable platforms and organizations make up 0.35 percent.

Voucher Soft-Targeting

We detect voucher language in the listing description with a dictionary of regular expressions developed through manual review of a random sample of listings. Each listing spell inherits a voucher flag if any of its scrapes match. The dictionary is organized as a general voucher-mention flag, `i_hcv`, which fires on any voucher or subsidy reference such as "section 8," "housing choice voucher," or "HAP contract," and excludes non-housing senses of the word such as gift, promo, parking, meal, and travel vouchers. Nested under it are four mutually informative children. `i_hcv_yes` captures acceptance or welcome, as in "section 8 welcome," "vouchers accepted," or "voucher holders welcome." `i_hcv_no` captures explicit refusal, as in "no section 8," "private pay only," or "section 8 need not apply.” `i_hcv_sec8_only` captures project-based units that require a voucher, as in “section 8 only” or “must have a voucher.” `i_hcv_boilerplate` captures fair-housing and compliance template text, which carries no acceptance or refusal signal. Our soft-targeted category is the set of listings that explicitly welcome vouchers (`i_hcv_yes`), with acceptance keywords screened against a refusal list so that phrases such as “not authorized to accept section 8” are not read as acceptance. The full pattern set is in the replication files.

Figure A.4: RentHub scrape volume by state and year

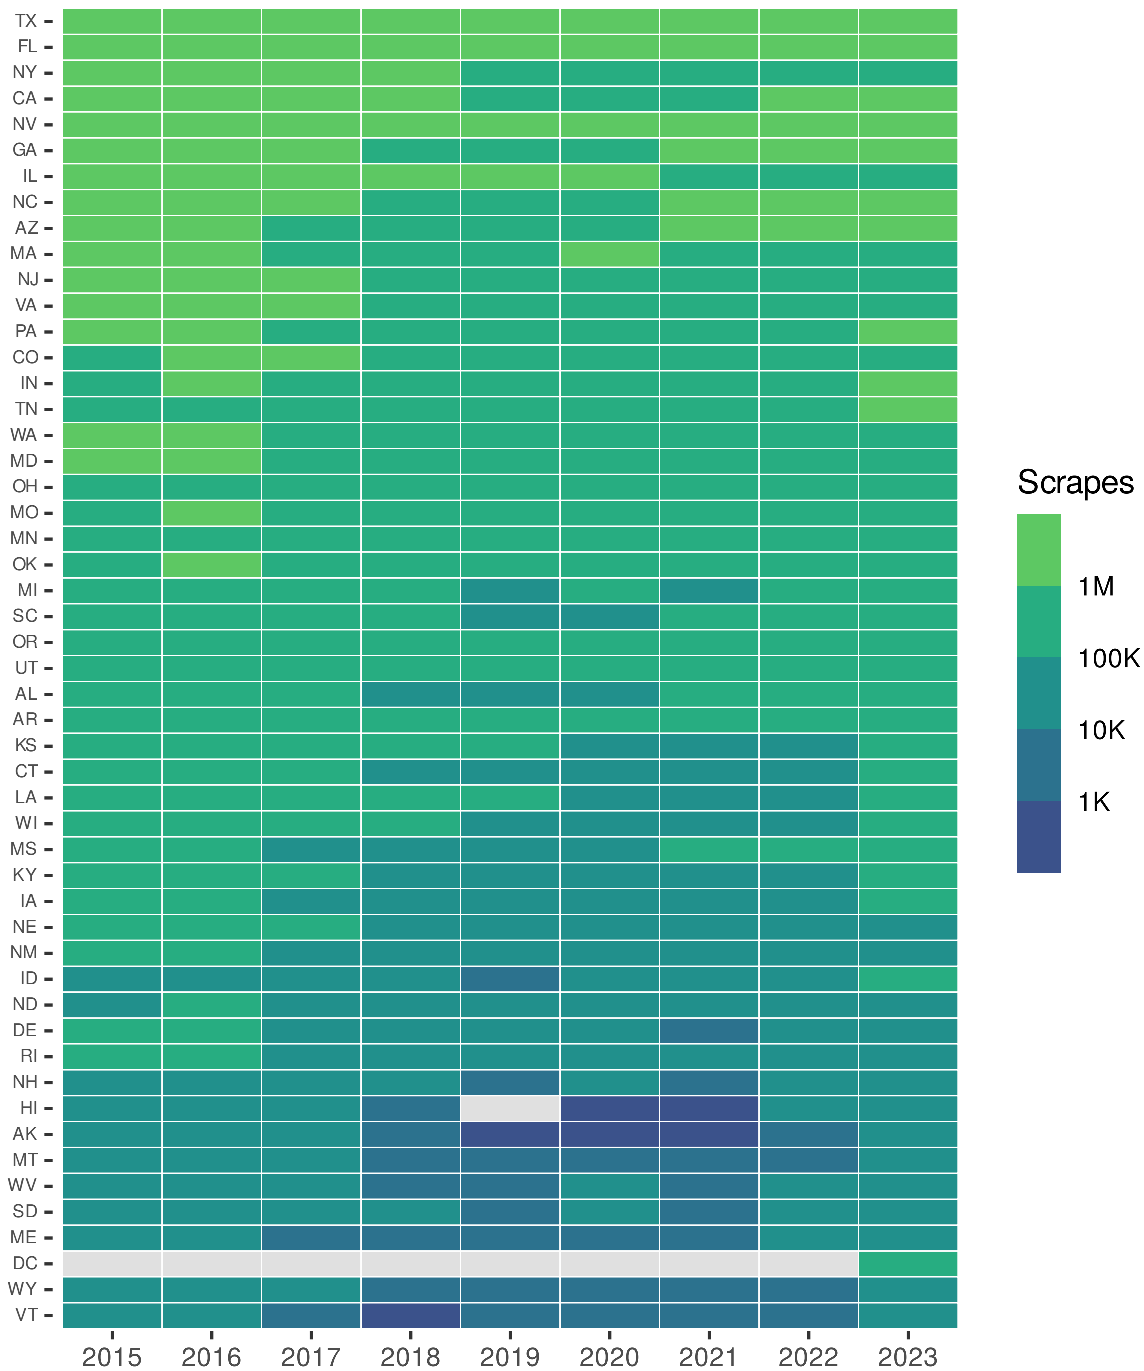

Note. The figure depicts the number of RentHub scrapes per state-year, on a log color scale. Grey cells have no coverage. States are ordered by total scrape volume.

Figure A.5: Distribution of within-unit gaps between consecutive scrapes

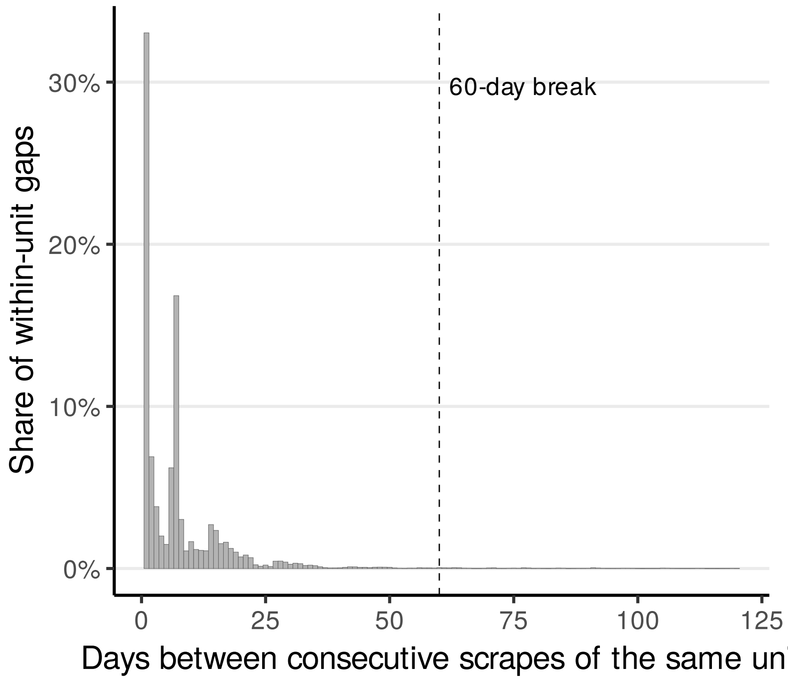

Note. The figure depicts the share of within-unit gaps by number of days between consecutive scrapes of the same unit, pooled across units and truncated at 120 days. Gaps spanning a statewide coverage hole are excluded. The dashed line marks the 60-day spell break.

Figure A.6: Company metadata coverage and hard-targeting share by year

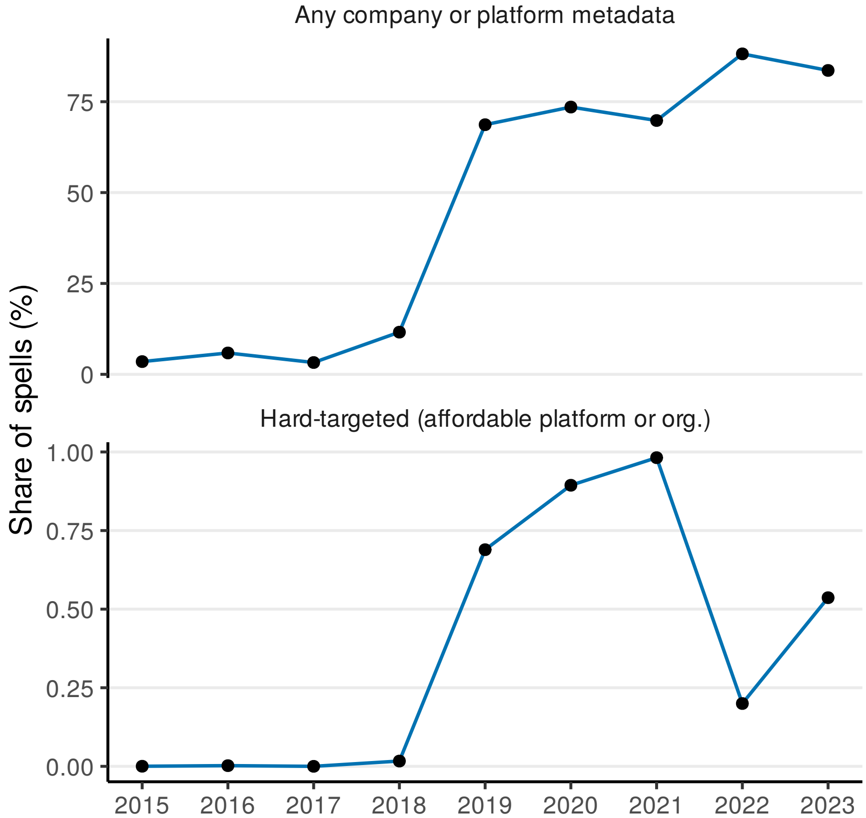

Note. The top panel depicts the share of spells with any non-missing company or platform metadata. The bottom panel depicts share of spells flagged hard-targeted. Axes use separate scales.

Figure A.7: Company-group composition among spells with company information

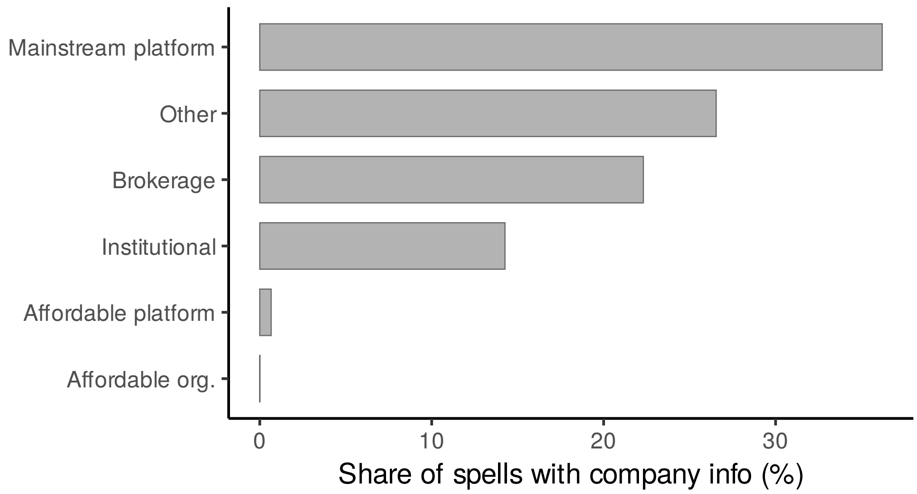

Note. The figure depicts the shares among spells with any non-missing company information. The 49 percent of spells with no company information are excluded from the figure and are dominated by pre-2019 years.

C Estimation Details

This appendix presents the estimation details underlying Section 7. Appendix C.1 states the Stage 3 likelihood and its optimization, Appendix C.2 the Stage 2 scarcity calibration, Appendix C.3 the Stage 1 computational implementation and the exact moment conditions.

C.1 Stage 3 likelihood and optimization

The per-contract log-likelihood ℓ^s_{jt} entering the Stage 3 estimator is a mixture over the three regimes of Section 6.4. In the upcharging region, writing the headroom $d^+_{jt} = S_{jt} - R_{jt}$ and the markup $m^+_{jt} = r_{jt} - R_{jt}$,

$$\ell^+_{jt} = \begin{cases} \log \pi^{0+} & m^+_{jt} = 0 \quad (r = R), \\ \log(1 - \pi^{0+}) + \log f^+(m^+_{jt}) & 0 < m^+_{jt} < d^+_{jt} \quad (R < r < S), \\ \log(1 - \pi^{0+}) + \log \bar{F}^+(d^+_{jt}) & m^+_{jt} = d^+_{jt} \quad (r = S), \end{cases} \tag{14}$$

whose three branches are the no-adjustment mass, the truncated interior density, and the censoring mass at the ceiling. In the downcharging region, writing the reduction $m^-_{jt} = S_{jt} - r_{jt}$,

$$\ell^-_{jt} = \begin{cases} \log \pi^{0-} & m^-_{jt} = 0 \quad (r = S), \\ \log(1 - \pi^{0-}) + \log f^-(m^-_{jt}) & m^-_{jt} > 0 \quad (r < S), \end{cases} \tag{15}$$

which has no ceiling mass because the reduction is not censored from above. Here f^s is the density of the latent adjustment and $\bar{F}^s(x)$ is its survival function, $\mathbb{P}(\tilde{m}^s > x)$. The density evaluates the exactly observed interior adjustments and the survival function the censored ones. With region-specific scale λ^s_{jt} from Equation (9) and shape α^s, the Weibull provides closed-form density and survival functions,

$$f^s(x) = \frac{\alpha^s}{\lambda^s_{jt}} \Big(\frac{x}{\lambda^s_{jt}}\Big)^{\alpha^s - 1} \exp\Big\{ - \big(x/\lambda^s_{jt}\big)^{\alpha^s} \Big\}, \qquad \bar{F}^s(x) = \exp\Big\{ - \big(x/\lambda^s_{jt}\big)^{\alpha^s} \Big\}. \tag{16}$$

We assign a contract to a mass point when its rent lies within one dollar of the listing rent or of the payment standard, and to the interior otherwise.

We maximize the resulting log-likelihood by BFGS with closed-form scores, parameterizing $\pi^{0s} = \Lambda(\cdot)$ and $\alpha^s = \exp(\cdot)$ to enforce their bounds, and we finish the optimization with damped Newton steps so that the score vanishes at the reported estimates. The two regions share no parameters, so each is maximized separately.

C.2 Stage 2 two-step estimator

In the first step, we estimate the attractiveness slopes θ as a weighted logit of the match indicator D^{HCV}_{jt} on the listing's position relative to the payment standard, its neighborhood attributes, and its targeting status, with interacted CBSA-by-year intercepts absorbing $\log A_{ct}$,

$$(\hat{\theta},\, \hat{a}) = \arg\max_{\theta,\, a} \sum_{jt} w_{jt} \Big[D^{\mathrm{HCV}}_{jt} \log \Lambda\big(a_{ct} + x'_{jt}\theta\big) + \big(1 - D^{\mathrm{HCV}}_{jt}\big) \log\big(1 - \Lambda(a_{ct} + x'_{jt}\theta)\big) \Big], \quad (17)$$

where $\Lambda(z) = e^z/(1+e^z)$ and w_{jt} is the linkage weight of Section 3.4, so that each contract counts once.

In the second step, we discipline the level of match probabilities with administrative totals. For each CBSA-year we obtain the flow M_{ct} of new voucher admissions and the linkage rate L_{ct}, the share of those admissions linked to a listing in the estimation panel.[29] Under Equation (5), the outside mass s^0_{ct} determines the share of the voucher flow landing on observed listings, and equating that share to L_{ct} gives the closed form

$$\tilde{s}^{\,0}_{ct} = \frac{1 - L_{ct}}{L_{ct}} \sum_{k \in ct} \exp\big(x'_{kt}\hat{\theta}\big), \quad (18)$$

whose log we smooth with additive metro and year components on the interior metro-years $\mathcal{I}$, those with $0 < L_{ct} < 1$, weighting by the flow,

$$(\hat{\kappa}_c,\, \hat{\kappa}_t) = \arg\min_{\kappa_c,\, \kappa_t} \sum_{ct \in \mathcal{I}} M_{ct} \left(\log \tilde{s}^{\,0}_{ct} - \kappa_c - \kappa_t\right)^2, \qquad s^0_{ct} = \exp\big(\hat{\kappa}_c + \hat{\kappa}_t\big). \quad (19)$$

The smoothing allows metro-years with few or zero links to receive finite low-linkage levels rather than dropping out of the calibration.

C.3 Stage 1 implementation

Cells and the landlord problem. A cell is a tract-by-bedroom-by-payment standard-by-year combination. Its listings share the payment standard S_{jt}, the market-rent anchor R^{mkt}_{jt}, and the tract covariates Y_{jt}, and the cell enters all model aggregates weighted by its listing count. Within a cell, listings differ only through the market shock q, so the cell's solution describes all of its listings. For

[29] The linkage rate is rebuilt against the final estimation panel after sample restrictions and outlier removals, counting a contract as linked when at least one of its linked spells survives the sample screens, so the calibration target and the platform mass refer to the same listing set.

a candidate ψ, every cell solves the listing problem of Equation (3) on a discrete choice set,

$$\big(R^*_{jt}(q),\, V^*_{jt}(q)\big) \in \arg\max_{R\in\mathcal{G}_{jt},\, V} \; \pi^{\mathrm{HCV}} \cdot \Pi^{\mathrm{HCV}} + \big[1-\pi^{\mathrm{HCV}}\big] \cdot \Pi^{\mathrm{MKT}}(R, Y_{jt}, q) - c^V(V), \tag{20}$$

with arguments suppressed and the rent grid $\mathcal{G}_{jt}$ of Section 7.1. We integrate the resulting choices over the market shock $q \sim N(0, \sigma_q^2)$ by seven-node Gauss–Hermite quadrature, and calibrate the shock scale to $\sigma_q = 0.19$, the listing-weighted standard deviation of log listing rents within cells. Because listings in a cell differ only through q, this calibration makes the shock account for exactly the rent dispersion that the cell structure leaves unexplained.

Data-side leasing prospect. The leasing-prospect moments in Table A.2 compare the model's $\rho(R^*, V^*, q)$ to a data-side counterpart implied by tract vacancy. Consider a unit that alternates between occupied spells of mean length T months and vacant spells that end each month with probability ρ, so a vacant spell lasts $1/\rho$ months on average. The long-run share of time the unit sits vacant is then $v = (1/\rho)/(T + 1/\rho)$, and inverting this relation expresses the monthly leasing probability in terms of the observed vacancy rate, $\rho_{jt} = \min\{(1 - v_{jt})/(v_{jt}\, T),\ 0.99\}$. We set the mean tenancy duration to $T = 24$ months, clamp v_{jt} to $[0.01, 0.95]$ and cap the implied probability at 0.99, a bound that binds whenever v_{jt} falls below roughly four percent.

SAFMR responses. The model counterpart of each SAFMR response applies the estimated first-stage shift in the payment standard,

$$g_o(\psi) = \bar{y}_o\big(\psi,\ S\, e^{\widehat{\Delta \log S}}\big) - \bar{y}_o\big(\psi,\ S\big), \tag{21}$$

where $\bar{y}_o(\psi, \cdot)$ is the listing-weighted model mean of outcome y_o at the solution under the indicated payment standards, $\widehat{\Delta \log S}$ is the event study first stage from Section 5, and the outcomes are the log listing rent and the shares of listings at and above the standard. The empirical targets are Sun and Abraham (2021) post-period effects re-estimated at the listing level with CBSA, year, and bedroom fixed effects, on the same CBSA universe as the model solve.

Weighting. The weighting matrix W in Equation (10) is diagonal, with entries proportional to the inverse sampling variances of the corresponding data moments (influence-function variances for the conditional moments and a delta-method variance for the dispersion moment). Inverse-variance weighting makes a given relative miss on a small share, such as the hard-targeting share, as costly as one on a large share, which is essential for the cost parameters. Under an identity weight, the optimizer can sacrifice exactly the small targeting moments that pin $(c^V_{\mathrm{soft}}, c^V_{\mathrm{hard}}, c^{\mathrm{HCV}})$ at negligible objective cost. Within the cross-sectional block, we cap the ratio of the largest to smallest diagonal entry at 10^4 around the median so that no near-degenerate moment dominates. We then normalize the cross-sectional and SAFMR blocks each to a unit mean diagonal, because the raw inverse variances are not comparable across the two blocks. The cross-sectional variances reflect listing-level sampling variation across millions of listings, while the SAFMR weights are inverse squared event study standard errors, so without the balancing the quasi-experimental moments would carry essentially no weight. The normalization preserves relative weights within each block while placing the two blocks on the same per-moment scale.

Optimization and reparameterization. We minimize Equation (10) over $(\tilde{c}, \tilde{\gamma}_R, \gamma_0, \gamma_Y)$ while reparameterizing

$$c = e^{\tilde{c}} \text{ for } c \in \{c^{\text{HCV}}, c^V_{\text{soft}}, c^V_{\text{hard}}\}, \qquad \gamma_R = -\, e^{\tilde{\gamma}_R}, \tag{22}$$

imposing positive costs and a leasing prospect that declines in the rent premium. The objective is minimized by parameter-scaled Nelder–Mead with simplex restarts, run from eight parallel starts, six seeded at prior interior optima ranked by objective and two dispersed fresh starts. The search proceeds in two phases, first on a coarse rent grid (2.5 percent steps with five quadrature nodes) and then re-optimizing the leading fits at the production grid of Equation (20) (1 percent steps, seven nodes). Each evaluation of the objective solves the 1.64 million cells twice, at the observed and shifted payment standards.

Stage 1 moment conditions Tables A.2 and A.3 list the 36 cross-sectional moments and the 3 quasi-experimental SAFMR moments.

Table A.2: Stage 1 moment conditions, Panels A–E

#	Moment	Definition	Primarily disciplines
Panel A. Levels of matching, targeting, and pricing			
1	HCV match share	$\mathbb{E}[D]$	fit of π^{HCV} level
2	Soft-targeting share	$\mathbb{E}[V_s]$	c^V_{soft}
3	Hard-targeting share	$\mathbb{E}[V_h]$	c^V_{hard}
4	Share at S	$\mathbb{E}[\mathbf{1}\{\lvert R-S\rvert/S < 0.02\}]$	bunching incentive
5	Share above S	$\mathbb{E}[\mathbf{1}\{R > S\}]$	γ_R, ceiling penalty
Panel B. Match share by listing and neighborhood type			
6	...high-rent listings	$\mathbb{E}[D \mid R > \text{med}(R)]$	$\gamma_R \times$ matching
7	...high-poverty tracts	$\mathbb{E}[D \mid \text{pov hi}]$	γ_Y (poverty)
8	...high-vacancy tracts	$\mathbb{E}[D \mid \text{vac hi}]$	γ_Y (vacancy)
9	...high-median rent tracts	$\mathbb{E}[D \mid \text{rent hi}]$	γ_Y (median rent)
Panel C. Joint pricing–targeting margins			
10	At S and hard-targeted	$\mathbb{E}[\mathbf{1}\{\text{at } S\} \cdot V_h]$	$c^V_{\text{hard}} \times$ pricing
11	At S and soft-targeted	$\mathbb{E}[\mathbf{1}\{\text{at } S\} \cdot V_s]$	$c^V_{\text{soft}} \times$ pricing
Panel D. Shape of the $R - S$ distribution, $\zeta \equiv (R - S)/S$			
12	Far below S	$\mathbb{E}[\mathbf{1}\{\zeta \leq -0.15\}]$	pricing dispersion vs. S
13	Just below S	$\mathbb{E}[\mathbf{1}\{-0.15 < \zeta \leq -0.05\}]$	pricing dispersion vs. S
14	Just above S	$\mathbb{E}[\mathbf{1}\{0.05 \leq \zeta < 0.15\}]$	pricing dispersion vs. S
15	Far above S	$\mathbb{E}[\mathbf{1}\{\zeta \geq 0.15\}]$	pricing dispersion vs. S
Panel E. Targeting and bunching by neighborhood			
16	Soft share, high poverty	$\mathbb{E}[V_s \mid \text{pov hi}]$	targeting gradients
17	Hard share, high poverty	$\mathbb{E}[V_h \mid \text{pov hi}]$	targeting gradients
18	Soft share, high vacancy	$\mathbb{E}[V_s \mid \text{vac hi}]$	targeting gradients
19	Hard share, high vacancy	$\mathbb{E}[V_h \mid \text{vac hi}]$	targeting gradients
20	At-S share, high poverty	$\mathbb{E}[\mathbf{1}\{\text{at } S\} \mid \text{pov hi}]$	bunching gradient
21	Soft share, high median rent	$\mathbb{E}[V_s \mid \text{rent hi}]$	targeting gradients
22	Hard share, high median rent	$\mathbb{E}[V_h \mid \text{rent hi}]$	targeting gradients

Notes. Continued in Table A.3, below which the notation and construction notes for both parts appear.

Table A.3: Stage 1 moment conditions, Panels F–I

#	Moment	Definition	Primarily disciplines
Panel F. Location, scale, and gradients of the rent premium $\ell \equiv \log(R/R^{\mathrm{mkt}})$			
23	Mean premium	$\mathbb{E}[\ell]$	γ_0
24	SD of premium	$\mathrm{SD}[\ell]$	γ_R
25	Mean premium, hard-targeted	$\mathbb{E}[\ell \mid V_h = 1]$	pricing–targeting interaction
26	Mean premium, soft-targeted	$\mathbb{E}[\ell \mid V_s = 1]$	pricing–targeting interaction
27	Mean premium, high poverty	$\mathbb{E}[\ell \mid \text{pov hi}]$	γ_Y (poverty)
28	Mean premium, high vacancy	$\mathbb{E}[\ell \mid \text{vac hi}]$	γ_Y (vacancy)
29	Mean premium, high median rent	$\mathbb{E}[\ell \mid \text{rent hi}]$	γ_Y (median rent)
30	Above-S share, high poverty	$\mathbb{E}[\mathbf{1}\{R > S\} \mid \text{pov hi}]$	pricing gradient vs. S
Panel G. Matching and bunching conditional on targeting			
31	Match share, hard-targeted	$\mathbb{E}[D \mid V_h = 1]$	c^{HCV}
32	Match share, soft-targeted	$\mathbb{E}[D \mid V_s = 1]$	c^{HCV}
33	At-S share, hard-targeted	$\mathbb{E}[\mathbf{1}\{\text{at } S\} \mid V_h = 1]$	pricing–targeting interaction
Panel H. Vacancy-implied leasing prospect ρ			
34	Leasing-prospect level	$\mathbb{E}[\rho]$	γ_0
35	. . . high poverty	$\mathbb{E}[\rho \mid \text{pov hi}]$	γ_Y (poverty)
36	. . . high median rent	$\mathbb{E}[\rho \mid \text{rent hi}]$	γ_Y (median rent)
Panel I. Quasi-experimental SAFMR adoption responses			
37	Log listing rent	post-adoption ATT on $\log R$	pass-through of S to rents
38	Share above S	post-adoption ATT on $\mathbf{1}\{R > S\}$	repricing across the ceiling
39	Share at S	post-adoption ATT on $\mathbf{1}\{\text{at } S\}$	bunching response to S

Notes. All expectations in Panels A–H are listing-weighted means over the estimation panel, and the identical function computes the model counterpart on the cell solutions, with D equal to the observed match indicator in the data and the model probability π^{HCV} on the model side. Here R is the listing rent, S the payment standard, V_s and V_h the soft- and hard-targeting indicators, "at S" the fractional band $|R - S|/S < 0.02$, $\zeta = (R - S)/S$, and $\ell = \log(R/R^{\mathrm{mkt}})$ the rent premium over the market anchor of Appendix C.3. Conditioning events "pov hi," "vac hi," and "rent hi" denote tracts above the listing-weighted median of the tract poverty rate, vacancy rate, and log median rent, and $R > \mathrm{med}(R)$ listings above the median listing rent. All cuts are computed once from the data and held fixed on the model side. The leasing prospect ρ is implied by tract vacancy in the data ($\rho = \min\{(1 - v)/(24v), 0.99\}$) and equals $\rho(R^*, V^*, q)$ at the model solution. Panel I targets are staggered-robust post-period ATTs of SAFMR adoption (Section 5), and the model counterpart shifts every cell's payment standard by the event study first stage $\widehat{\Delta \log S}$ and differences the listing-weighted outcome, Equation (21). The "primarily disciplines" column is heuristic, as all parameters are estimated jointly from all moments. Weights are the inverse sampling variances of the data moments (Appendix C.3).

D Counterfactual Details

This appendix documents the counterfactual experiments of Section 8. Appendix D.1 presents additional details, Appendix D.2 describes the equilibrium computation, Appendix D.3 details the mobility costing and the budget-neutral recycling experiments, and Appendix D.4 lists the scenario, dose, and outside-response grid and presents the full results.

D.1 Design Details

Poverty quartile definition. Poverty quartiles are computed once as within-CBSA quartiles of the tract poverty rate, and the resulting tract poverty classification is held fixed across all scenarios. Targeted policies treat the tracts according to this classification, and all quartile outcomes aggregate matches over the same classification, so compositional changes reflect reallocation across a fixed geography rather than tracts switching quartiles. We do not weight tracts by listing counts or voucher contract counts.

Rent-indexed standards. Within each CBSA, year, and bedroom count, the indexed schedule sets every tract's standard to ωR^{mkt}, where R^{mkt} is the tract-level market-rent benchmark of Section 7.1 and ω solves for an unchanged listing-weighted mean standard at the CBSA-year-bedroom level. The hold-harmless variant replaces the indexed standard with the baseline standard wherever the indexed value is lower, which raises average generosity by construction. Composite variants first index the schedule and then apply the uniform or targeted uplift factors to it.

Participation-cost offsets. An offset reduces the estimated per-contract participation cost c^{HCV} by 12.5, 18, or 25 percent for matches in treated tracts, entering the listing problem exactly as a reduction in c^{HCV}. Because the estimated cost absorbs administrative burden and non-pecuniary aversion alike, and a cash transfer offsets the net cost dollar for dollar regardless of its composition, the experiment does not require taking a stand on what the cost consists of. Doses are stated as shares of the estimated cost because dollar amounts would inherit the revenue-horizon normalization of the Stage 1 payoffs. The rungs are dose-matched to the standards family per contract at the mean standard, with the 12.5 percent offset worth about \$156 against roughly \$155 under the 10 percent uplift and the 25 percent offset worth \$311 against \$310 under the 20 percent uplift. The 18 percent rung, worth about \$224, fills the middle of the ladder.

Rent-reasonableness enforcement. Enforcement scales the expected upward adjustment of the contracting stage to 75, 50, 25, or 0 percent of its estimated size, implemented as a proportional rescaling of the Stage 3 latent markup. The rescaling lowers expected voucher revenue for upcharging listings and feeds back into listing rents, targeting, and matching through the Stage 1 problem. The downcharging adjustment is unchanged.

D.2 Implementation

Every scenario re-solves the model on the estimation universe of Section 7, the 1.64 million tract-by-bedroom-by-payment standard-by-year cells, at the production settings of the Stage 1 estimator, a 1 percent rent grid augmented with the payment standard and seven Gauss–Hermite nodes for the market shock, as in Appendix C.3.

Policies feed back into the market through the Stage 2 scarcity terms A_{ct}, defined in Section C.2. Given candidate metro-year scarcity levels, each cell solves the listing problem, the implied attractiveness sums update the scarcity levels, and the equilibrium is the fixed point of this map, computed by damped iteration in logs with an adaptive step. A few small metro-years flip discrete listing choices as the scarcity level moves and settle into limit cycles rather than exact fixed points. We accept flow-settled solutions in these cases and track the share of the voucher flow carried by cycling metro-years, which stays below 0.1 percent in every scenario.

The outside-market responses are governed by μ. A fraction μ of the outside mass re-prices in proportion to the attractiveness response of observed listings in the same metro-year,

$$s^0_{ct}(\mu) \;=\; s^0_{ct}\left[(1-\mu) \;+\; \mu\,\frac{\sum_{k\in ct}\exp\!\left(x'_{kt}\theta\right)}{\sum_{k\in ct}\exp\!\left(x^{\mathrm{b}\prime}_{kt}\theta\right)}\right], \tag{23}$$

where the numerator sums the attractiveness index over observed listings at their counterfactual choices and the denominator evaluates the same sum at the baseline choices x^{b}_{kt}. Under $\mu = 0$ the outside mass is fixed at s^0_{ct}, and under $\mu = 1$ the mapping becomes scale-invariant, so the aggregate match level is additionally normalized to its baseline value each iteration and those results are informative about composition only. All three values of μ are solved for every scenario, with each solution initializing the next, and the three cases coincide at the baseline.

D.3 Mobility Costing and Budget-Neutral Recycling

Cost per match moved into the target set. For each targeted scenario, this measure divides incremental spending, the change in rent payments plus any transfer outlays, by the gross gain in matches located in the targeted tracts relative to baseline, with the target set Q1 or Q1–Q2 according to the scenario. The numerator is the same as in the cost per marginal match, but the denominators differ conceptually. The aggregate denominator is a net change, in which matches displaced away from non-targeted quartiles subtract, while the target denominator is a gross gain, to which displaced matches that relocate into the target set contribute. For targeted instruments the gross gain therefore exceeds the net change, and the cost per match moved into the target set lies below the cost per marginal match. Both ratios are reported only where incremental spending is positive.

Recycling calibration. The recycling experiments pair enforcement with a payment standard uplift whose cost exhausts the enforcement saving. Cutting the expected markup to 50 and 0 percent of its estimated size saves approximately \$3.2 and \$5.2 million per month in the listed market, and the uplift factor for each variant is interpolated on the cost ladders of the standards experiments so that the payment standard increase cost matches the saving. The calibrated factors are 1.011 economy-wide, 1.083 in Q1, and 1.039 in Q1–Q2 for the 50 percent markup cut, and 1.017, 1.119, and 1.058 for full elimination. Because the factors are calibrated under $\mu = 0$ by interpolation, budget neutrality holds only approximately, and we verify that the residual budget change stays within \$20 per matched voucher in absolute value at $\mu = 0$ and within \$15 at $\mu = 0.5$. The recycling scenarios report the target-set match gains but no cost ratio, since incremental spending is zero by design.

D.4 Scenario, Dose, and Outside-Response Grid

Table A.4 lists the full grid of 34 scenarios, every one solved under all three outside-market responses $\mu \in \{0, 0.5, 1\}$. Tables A.5, A.6, and A.7 report the complete results for the payment standard family, the participation-cost offsets, and the enforcement and recycling experiments, respectively.

Table A.4: Counterfactual scenario grid

Family	Doses	Targeting	Runs
Baseline	–	–	1
Uniform standards	$+10$, $+20$ percent	all, Q1, Q1–Q2	6
Rent-indexed standards	indexed, hold-harmless	economy-wide	2
Indexed plus uplift	$+10$, $+20$ percent on indexed	all, Q1, Q1–Q2	6
Participation-cost offsets	12.5, 18, 25 percent	all, Q1, Q1–Q2	9
Enforcement	markup at 75, 50, 25, 0 percent	economy-wide	4
Enforcement recycling	markup at 50, 0 percent, calibrated uplift	all, Q1, Q1–Q2	6
Total			34

Notes. Each run is solved under all three outside-market responses $\mu \in \{0, 0.5, 1\}$, which coincide at the baseline. Targeting labels refer to the baseline within-metro poverty-quartile masks of Appendix D.1. Recycling uplift factors are calibrated in Appendix D.3.

Table A.5: Full results, payment standard experiments

	Matches (%)		Q1 matches (%)			Fiscal, $\mu = 0.5$ ($)		
	$\mu = 0$	$\mu = 0.5$	$\mu = 0$	$\mu = 0.5$	$\mu = 1$	Δ bill per match	Cost per marg. match	Cost per Q1(–Q2) match
Uniform $+10$, all	$+19.7$	$+8.1$	$+19.9$	$+9.6$	$+1.8$	$+231$	3,093	–
Uniform $+20$, all	$+40.4$	$+14.6$	$+41.9$	$+18.1$	$+3.6$	$+418$	3,270	–
Uniform $+10$, Q1	$+2.0$	$+1.0$	$+27.3$	$+25.5$	$+23.9$	$+35$	3,458	1,444
Uniform $+20$, Q1	$+4.8$	$+2.2$	$+63.8$	$+57.8$	$+53.0$	$+83$	3,853	1,540
Uniform $+10$, Q1–Q2	$+5.3$	$+2.5$	$+25.5$	$+21.5$	$+18.0$	$+81$	3,347	1,519
Uniform $+20$, Q1–Q2	$+11.7$	$+5.1$	$+57.9$	$+45.9$	$+37.1$	$+178$	3,666	1,637
Rent-indexed	-10.1	-5.5	$+0.2$	$+5.2$	$+11.1$	-94	–	–
Rent-indexed, hold-harmless	$+7.7$	$+3.5$	$+12.2$	$+8.2$	$+4.5$	$+126$	3,689	–
Indexed $+10$, all	$+8.5$	$+3.6$	$+20.9$	$+15.7$	$+11.6$	$+193$	5,504	–
Indexed $+20$, all	$+28.6$	$+11.2$	$+43.1$	$+24.6$	$+12.2$	$+422$	4,191	–
Indexed $+10$, Q1	-7.9	-4.3	$+27.6$	$+33.3$	$+39.8$	-44	–	–
Indexed $+20$, Q1	-5.1	-2.8	$+63.3$	$+67.8$	$+73.5$	$+25$	–	377
Indexed $+10$, Q1–Q2	-4.7	-2.6	$+26.0$	$+28.8$	$+32.1$	$+18$	–	288
Indexed $+20$, Q1–Q2	$+1.9$	$+0.7$	$+57.6$	$+54.4$	$+52.4$	$+152$	23,117†	1,234

Notes. Matches columns report the percent change in listed-market matches relative to baseline. Aggregate changes under $\mu = 1$ are held at zero by construction and are omitted. Q1 columns report the percent change in matches located in the lowest-poverty within-metro quartile. Fiscal columns are evaluated at the baseline outside-market response $\mu = 0.5$. The bill change is per matched voucher per month. Cost per marginal match divides incremental spending by the net change in matches and is suppressed where spending or the net change is nonpositive. Cost per Q1(–Q2) match divides the same spending by the gross gain in the scenario's target set, Q1 or Q1–Q2 as labeled, and applies to targeted scenarios only. †Net-denominator entries close to zero are unstable, as here where net matches rise only 0.7 percent.

Table A.6: Full results, participation-cost offsets

	Matches (%)		Q1 matches (%)			Fiscal, $\mu = 0.5$ ($)			
	$\mu = 0$	$\mu = 0.5$	$\mu = 0$	$\mu = 0.5$	$\mu = 1$	Δ bill per match	Transfer per match	Cost per marg. match	Cost per Q1(–Q2) match
Offset 12.5%, all	$+6.9$	$+2.6$	$+3.1$	$+0.1$	-2.1	$+64$	$+156$	8,541	–
Offset 18%, all	$+11.0$	$+3.9$	$+5.0$	$+0.0$	-2.9	$+92$	$+224$	8,360	–
Offset 25%, all	$+17.4$	$+5.7$	$+7.6$	-0.2	-4.6	$+128$	$+311$	8,103	–
Offset 12.5%, Q1	$+0.4$	$+0.3$	$+6.3$	$+6.0$	$+5.6$	$+8$	$+16$	8,960	4,064
Offset 18%, Q1	$+0.7$	$+0.4$	$+10.6$	$+10.0$	$+9.3$	$+11$	$+23$	9,549	3,592
Offset 25%, Q1	$+1.2$	$+0.5$	$+17.4$	$+16.0$	$+14.8$	$+16$	$+34$	9,286	3,306
Offset 12.5%, Q1–Q2	$+1.1$	$+0.5$	$+5.7$	$+4.7$	$+4.0$	$+15$	$+41$	10,951	4,185
Offset 18%, Q1–Q2	$+1.9$	$+0.8$	$+9.5$	$+7.9$	$+6.5$	$+24$	$+61$	10,199	3,854
Offset 25%, Q1–Q2	$+3.2$	$+1.3$	$+15.4$	$+12.2$	$+10.0$	$+37$	$+89$	9,630	3,605

Notes. Layout as in Table A.5. Transfers are averaged over all matched vouchers program-wide, so the per-match transfer in targeted scenarios is small even though treated matches receive the full offset. Incremental spending in the cost columns includes both the bill change and the transfer outlays.

Table A.7: Full results, rent-reasonableness enforcement and budget-neutral recycling

	Matches (%)		Q1 matches (%)			Fiscal, $\mu = 0.5$ (\$)
	$\mu = 0$	$\mu = 0.5$	$\mu = 0$	$\mu = 0.5$	$\mu = 1$	Δ bill per match
Markup at 75%	-0.6	-0.3	-0.6	-0.4	-0.2	-17
Markup at 50%	-1.0	-0.5	-1.1	-0.8	-0.4	-29
Markup at 25%	-1.3	-0.6	-1.6	-1.1	-0.6	-39
Markup at 0%	-1.5	-0.7	-1.8	-1.2	-0.7	-48
Recycle, markup 50%, all	$+0.9$	$+0.5$	$+0.8$	$+0.4$	-0.2	-2
Recycle, markup 50%, Q1	$+0.4$	$+0.3$	$+18.1$	$+17.6$	$+17.0$	-4
Recycle, markup 50%, Q1–Q2	$+0.7$	$+0.5$	$+7.3$	$+6.8$	$+6.2$	-2
Recycle, markup 0%, all	$+1.4$	$+0.7$	$+1.1$	$+0.4$	-0.4	-9
Recycle, markup 0%, Q1	$+0.6$	$+0.4$	$+24.9$	$+24.3$	$+23.6$	-15
Recycle, markup 0%, Q1–Q2	$+1.0$	$+0.6$	$+10.1$	$+9.4$	$+8.5$	-11

Notes. Layout as in Table A.5. Enforcement rows scale the expected contracting-stage markup to the indicated share of its estimated size. Recycling rows pair the markup cut with the calibrated payment standard uplift of Appendix D.3, so incremental spending is approximately zero by design and the cost ratios are not defined. The bill column reports the residual budget change per matched voucher.

E Supplementary Figures

Figure A.8: Government and tenant payments under the HCV Program

(a) Government payment, full sample

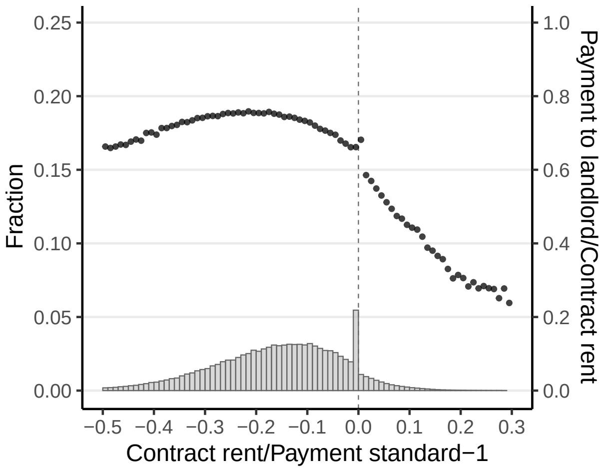

(b) Government payment, no utility allowance

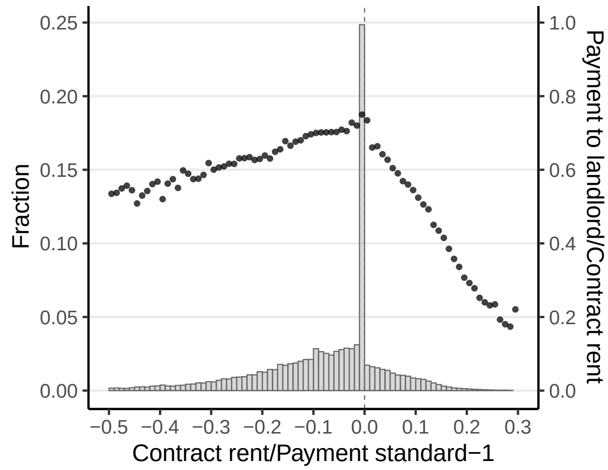

(c) Tenant payment, full sample

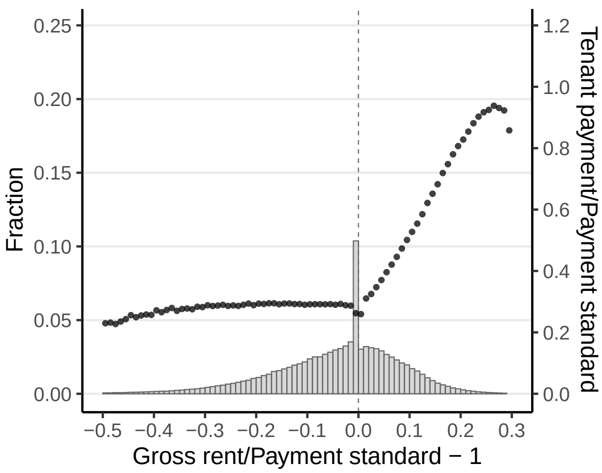

(d) Tenant payment, no utility allowance

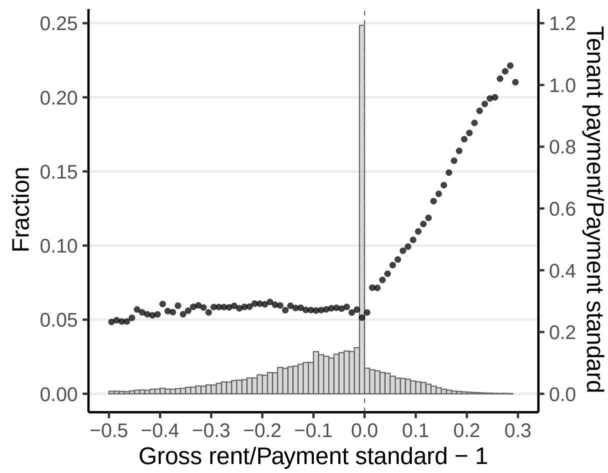

Note. The gray bars show the distribution of contract rents (panels a and b) or gross rents (panels c and d) for newly admitted voucher holders in HUD administrative data, expressed as the percentage deviation from the corresponding payment standard. The sample includes voucher holders admitted between 2015 and 2023 in the year when they were admitted. In panels (b) and (d), we restrict the sample to contracts with no utility allowance (i.e., where contract rent equals gross rent). The histogram uses 1 percentage-point-wide bins. The gray dashed line indicates a zero deviation from the payment standard. The black dots in panel (a) show the average share of the contract rent paid by the government in each bin; in panel (b), the average share of the gross rent paid by voucher holders.

Figure A.9: Regression of linkage indicator on voucher holder and neighborhood characteristics

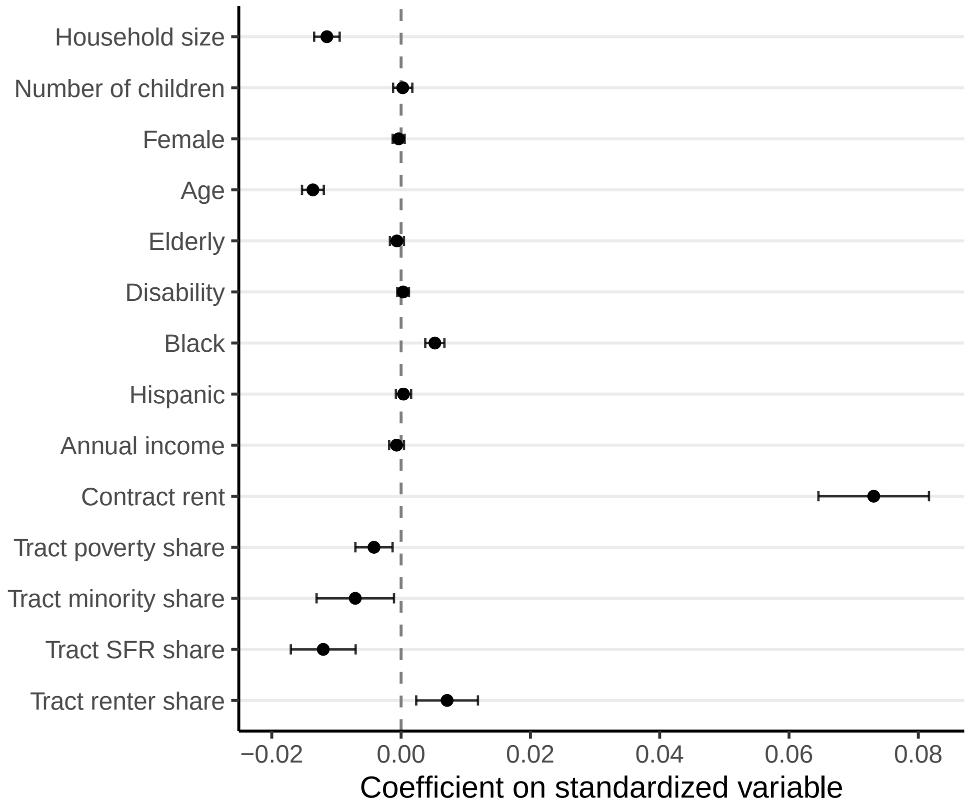

Note. This figure plots coefficients and 95% confidence intervals from pairwise regressions of an indicator for whether an HCV contract is linked to a RentHub listing spell on voucher-holder and neighborhood characteristics, the latter being standardized within PHA. The sample includes HCV new admissions from 2015 to 2023. All regressions include PHA-by-year and PHA-by-bedroom fixed effects, with standard errors clustered at the PHA level. All independent variables are obtained from HUD administrative data.

Figure A.10: Bunching at payment standards in HCV contracts by neighborhood characteristics

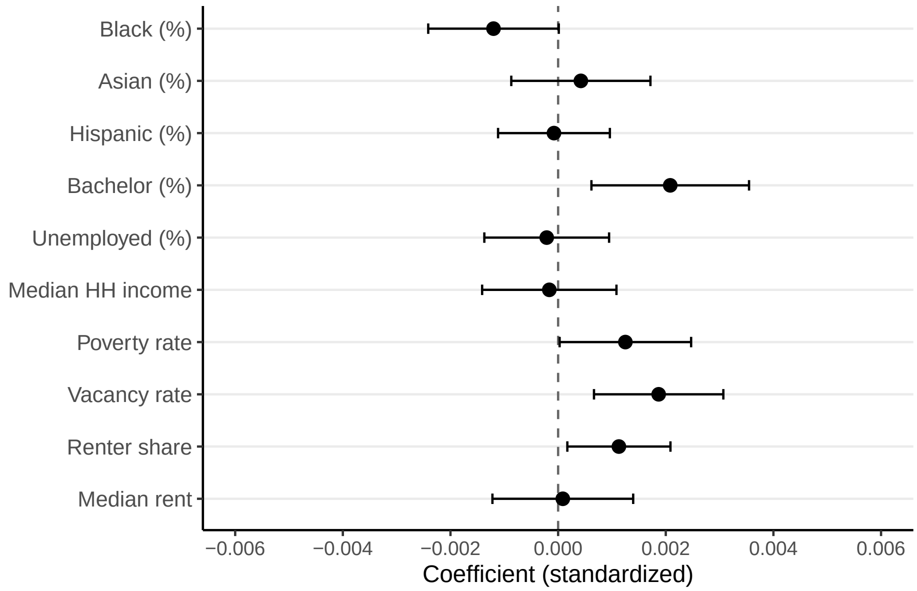

Note. The figure plots coefficients from separate HCV contract-level regressions in which the outcome is an indicator for whether the contract rent equals the payment standard and the regressor is a tract-level characteristic standardized within the PHA. All regressions include PHA-by-year and PHA-by-bedrooms fixed effects; standard errors are clustered at the PHA level. The sample includes all new admissions in HCV administrative data, 2015–2023. Census characteristics are obtained from the 2016–2020 American Community Survey (5-year estimates).

Figure A.11: Histogram of contract and listed rents for linked units (single-unit addresses only)

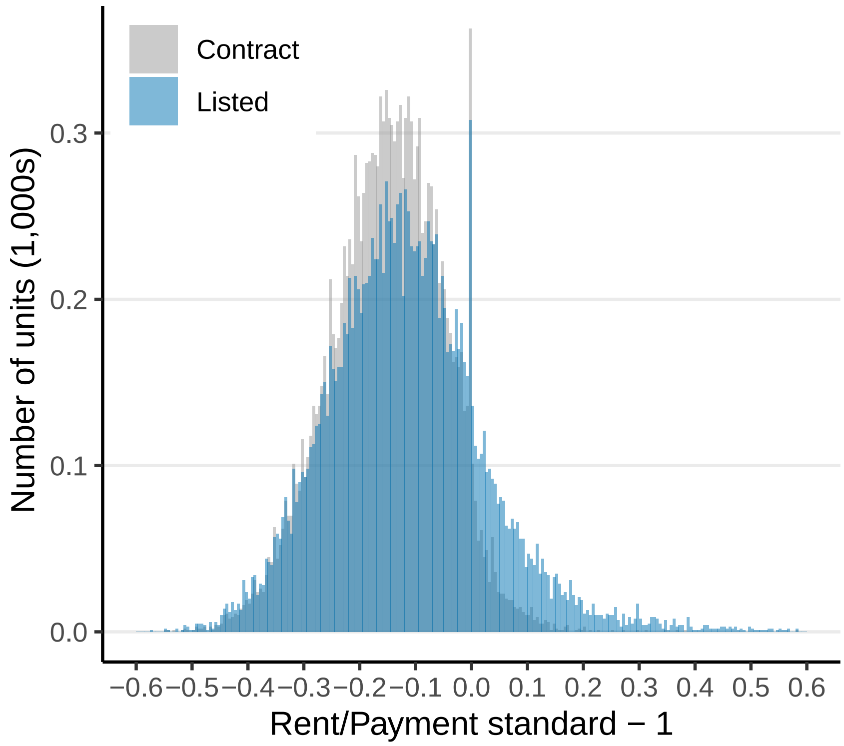

Note. The plot shows the histogram of contract (gray) and listed (blue) rents relative to the payment standard, measured as (rent/payment standard) $-$ 1. The sample includes new admissions in HCV administrative data from 2015 to 2023 that are linked to a RentHub listing spell, as described in Section 3.4, where the address is classified as being single-unit according to both the HCV administrative and RentHub datasets.

Figure A.12: Differences-in-bunching: Listed vs counterfactual listed rents for linked units

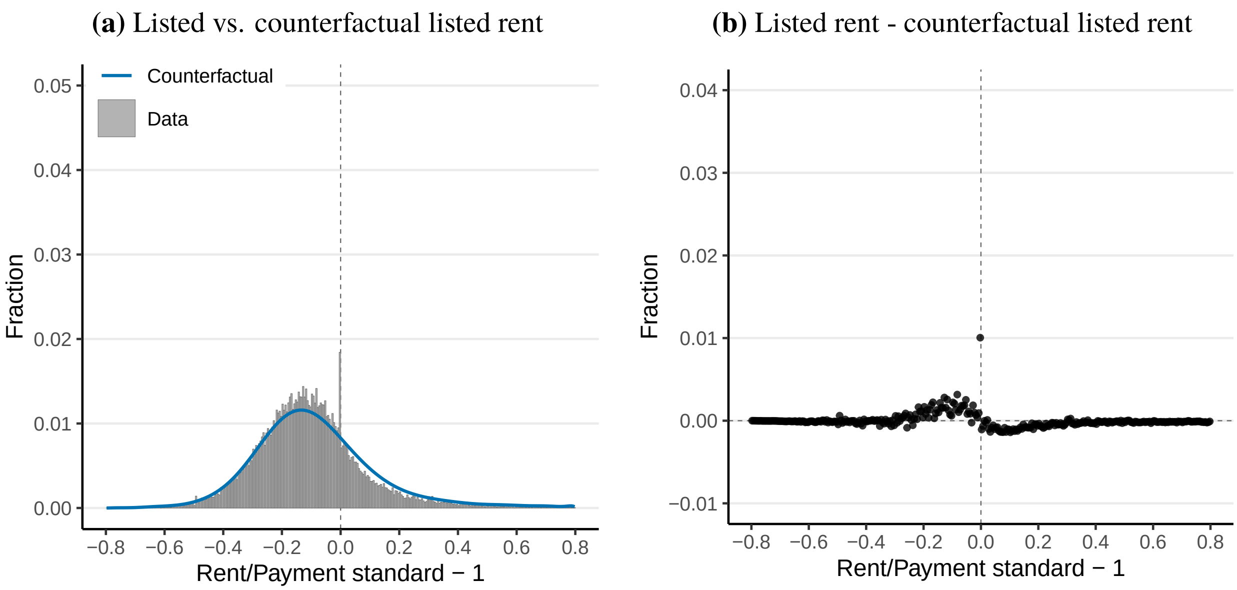

Note. Panel (a) plots the histogram of listed rents for new admissions in HCV administrative data, 2015–2023, that are linked to a RentHub listing spell (gray), together with the counterfactual distribution of listed rents (black), computed from linked listed rents using the differences-in-bunching estimator described in Section 4. Panel (b) plots the difference between the two, with 95% confidence bands (gray) from the iterative bootstrap procedure described there. Bins are 0.5 percentage points wide.

Figure A.13: Differences-in-bunching (single-unit addresses): Contract vs counterfactual rents

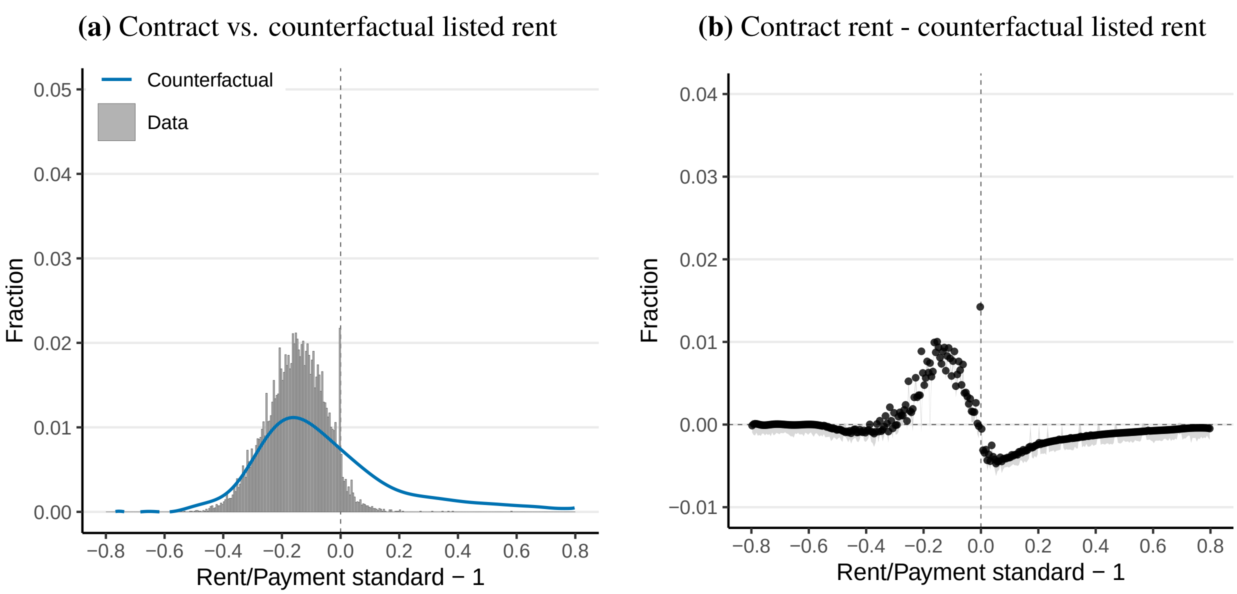

Note. Panel (a) plots the histogram of contract rents for new admissions in HCV administrative data, 2015–2023, that are linked to a RentHub listing spell (gray), together with the counterfactual distribution of listed rents (black), computed from linked listed rents using the differences-in-bunching estimator described in Section 4. The sample is restricted to HCV units that we identify as single-unit addresses using voucher administrative data and RentHub rental listing information. Panel (b) plots the difference between the two, with 95% confidence bands (gray) from the iterative bootstrap procedure described there. Bins are 0.5 percentage points wide.

Figure A.14: Histogram of cumulative number of PHAs adopting ZIP-level payment standards

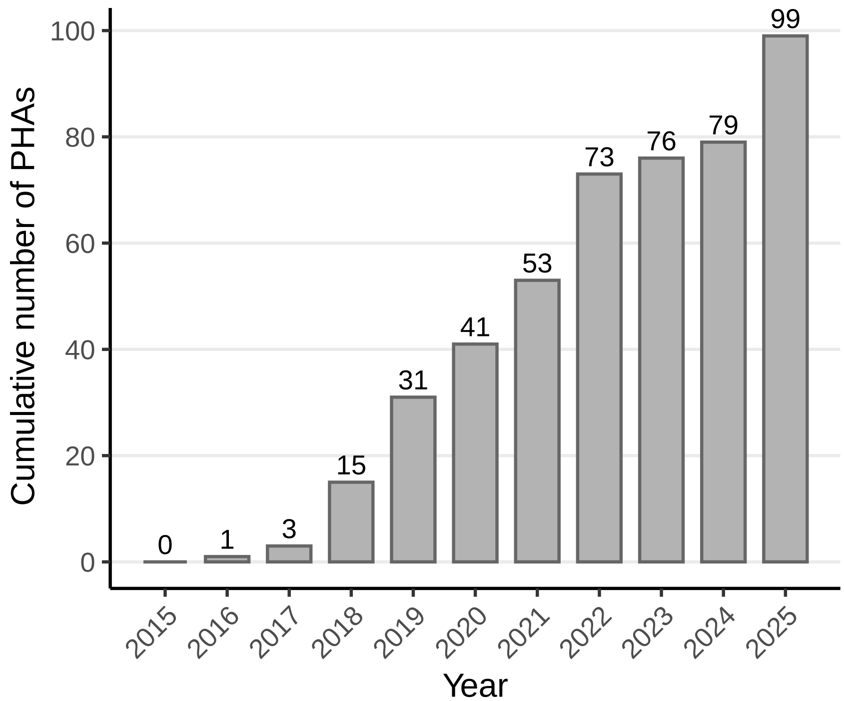

Note. This figure shows a histogram of the cumulative number of PHAs that adopt ZIP-level payment standards between 2015 and 2025 in our payment standard dataset described in Section 3.3 and Appendix A.

Table A.8: Summary Statistics: Payment Standard Subsamples

	HCV New Admissions		**RentHub Listings**	
	All	Linked	Scrape level	Spell level
	(1)	(2)	(3)	(4)
# Observations	792,330	105,002	243,141,388	18,870,624
# CBSAs	510	420	618	618
Rent (USD)				
25th Percentile	669	729	1,235	1,250
Median	860	975	1,725	1,795
75th Percentile	1,200	1,450	2,295	2,495
Mean	1,013.5	1,174.5	1,912.6	2,015.9
Standard Deviation	513.5	618.4	1,056.5	1,177.5
Payment Standard (USD)				
25th Percentile	783	830	1,017	1,139
Median	1,012	1,107	1,342	1,497
75th Percentile	1,380	1,635	1,721	1,950
Mean	1,166.8	1,320.0	1,434.6	1,620.6
Standard Deviation	563.3	683.7	561.9	669.6
Number of Bedrooms				
25th Percentile	1	1	1	1
Median	2	2	2	2
75th Percentile	3	2	3	3
Mean	1.87	1.83	2.07	2.24
Standard Deviation	0.95	0.95	1.05	1.05
Square Footage				
25th Percentile	—	665	793	850
Median	—	852	1,066	1,158
75th Percentile	—	1,084	1,412	1,592
Mean	—	928.1	1,182.8	1,297.2
Standard Deviation	—	463.3	569.7	699.8
Targeting composition (%)				
Hard-targeted	—	7.19	0.13	0.39
Soft-targeted	—	6.97	0.50	0.82

Note. The table reports summary statistics for the subsamples with non-missing payment standards. The sample and variable definitions follow Table 1: Columns (1) and (2) use HUD voucher contract data restricted to new admissions, with Column (2) further restricted to admissions linked to a RentHub listing; Columns (3) and (4) use RentHub listings (spells anchored in 2015–2023) at the level of the underlying raw scrapes and the spell, respectively. Each column is further restricted to observations with a non-missing payment standard.

Table A.9: Robustness: Differences-in-bunching estimates in linked HUD-rental listings analyses

	Main	Polynomials			1 pp	Pre-2020	Absolute	Single
		$k = 12$	$k = 16$	$k = 24$	bins		deviation	unit
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Total (Contract - Counterfactual listed)								
Excess mass (EM)	0.037	0.036	0.037	0.037	0.040	0.034	0.034	0.014
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
Missing mass above (MMA)	-0.189	-0.189	-0.188	-0.190	-0.193	-0.114	-0.200	-0.270
	(0.004)	(0.005)	(0.004)	(0.004)	(0.003)	(0.004)	(0.001)	(0.009)
Missing mass below (MMB)	0.168	0.168	0.167	0.168	0.165	0.095	0.165	0.261
	(0.004)	(0.005)	(0.004)	(0.004)	(0.003)	(0.004)	(0.001)	(0.009)
Panel B: Stage 1 (Listed - Counterfactual listed)								
Excess mass (EM)	0.010	0.010	0.010	0.010	0.011	0.006	0.009	0.010
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Missing mass above (MMA)	-0.060	-0.062	-0.060	-0.061	-0.059	-0.026	-0.067	-0.123
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.005)
Missing mass below (MMB)	0.053	0.055	0.053	0.054	0.052	0.022	0.058	0.122
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)	(0.005)
Panel C: Stage 2 (Contract - Listed)								
Excess mass (EM)	0.027	0.027	0.027	0.027	0.029	0.028	0.026	0.004
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
Missing mass above (MMA)	-0.129	-0.128	-0.128	-0.129	-0.134	-0.088	-0.133	-0.147
	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)	(0.000)	(0.009)
Missing mass below (MMB)	0.114	0.113	0.114	0.115	0.113	0.073	0.107	0.139
	(0.004)	(0.005)	(0.005)	(0.004)	(0.003)	(0.003)	(0.000)	(0.010)

Note. Each column reports estimates of Equation (1) for an alternative sample or specification: (1) the main sample; (2-4) alternative polynomial orders for the counterfactual; (5) 1-percentage-point rent bins; (6) the pre-COVID period only; (7) rent deviations from the payment standard measured in absolute rather than relative terms; and (8) voucher units in single-unit addresses. Panel A reports bunching estimates from the difference between the contract rent and counterfactual listed rent distributions; Panel B, between the listed rent and counterfactual listed rent distributions; Panel C, between the contract rent and listed rent distributions. Bootstrap standard errors (1,000 iterations), following the procedure detailed in Section 4, are reported in parentheses.

Table A.10: Summary statistics: Early vs. later ZIP-level payment standard adopters

	(1) Early adopters (2015–2022)	(2) Later adopters (2023–2025)	(3) Never adopters
Panel A: PHA characteristics			
Voucher units	4,329	9,772	4,028
Panel B: County characteristics 2020			
Bachelor degree share	0.36	0.35	0.33
Black share	0.14	0.18	0.14
Hispanic share	0.18	0.21	0.17
Median HH income	79,294	71,587	69,191
Poverty rate	0.11	0.14	0.14
Vacancy rate	0.09	0.09	0.10
Renter share	0.35	0.42	0.38
Median bdg age	47	49	46
Median rent	1,198	1,110	1,002
Renter turnover	0.24	0.25	0.25
Observations	73	26	366

Note. The table reports sample means for three groups of PHAs in the national payment standards dataset described in Section 3.3: PHAs that adopt ZIP-level payment standards in 2015–2022, those adopting in 2023–2025, and never-adopters. Panel A reports the number of voucher units served by each PHA as of 2024. Panel B reports Census characteristics of the county in which the PHA is located, obtained from the 2016–2020 American Community Survey (5-year estimates).

Figure A.15: Tenant Payment as a Share of Income

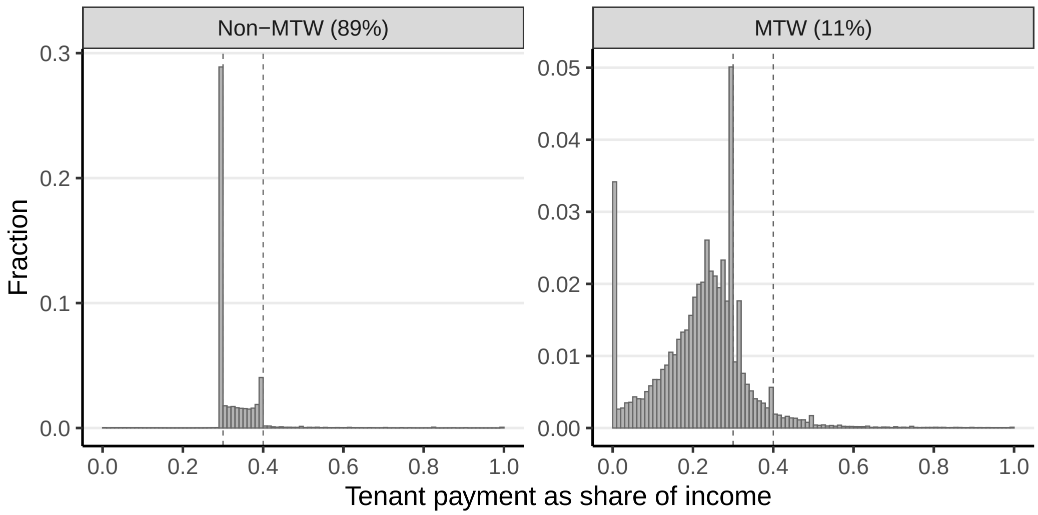

Note. The figure plots the distribution of tenant out-of-pocket rent payment as a share of household income. The left panel shows voucher contracts in non-MTW PHAs, and the right panel shows voucher holders in MTW PHAs. The sample includes all new admissions between 2015 and 2023 in HUD administrative data.

Table A.11: Stage 3 estimation robustness checks: upcharging region ($R \leq S$)

	Baseline	**Extended covariates**	**Binary targeting**	**PCA covariates**	**No fixed effects**	**2023**
Constant (μ_0^+)	3.410	3.930	3.405	3.963	2.032	3.589
	(0.180)	(0.189)	(0.176)	(0.100)	(0.103)	(0.356)
Headroom d^+ per \$ (μ_1^+)	0.00248	0.00247	0.00248	0.00248	0.00248	0.00200
	(0.00003)	(0.00003)	(0.00003)	(0.00003)	(0.00003)	(0.00005)
Poverty	0.142	0.358	0.148	—	0.209	0.098
	(0.044)	(0.055)	(0.044)		(0.039)	(0.105)
Vacancy	0.268	0.324	0.260	—	0.082	0.316
	(0.067)	(0.070)	(0.067)		(0.054)	(0.152)
Log median rent	0.074	−0.018	0.074	—	0.295	0.059
	(0.022)	(0.025)	(0.022)		(0.015)	(0.047)
Soft targeting	0.006	0.020	—	0.011	—	0.008
	(0.019)	(0.020)		(0.020)		(0.036)
Hard targeting	0.155	0.167	—	0.163	—	0.131
	(0.020)	(0.020)		(0.020)		(0.036)
Targeted, any (μ_V^+)	—	—	0.083	—	—	—
			(0.015)			
No markup (π^{0+})	0.559	0.559	0.559	0.559	0.566	0.573
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)
Shape (α^+)	1.357	1.359	1.357	1.357	1.300	1.397
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.015)
Tract covariates	Baseline	Extended	Baseline	3 PCs	Baseline	Baseline
Fixed effects	Bed, yr, CBSA	Bed, yr, CBSA	Bed, yr, CBSA	Bed, yr, CBSA	None	Bed, CBSA
Contracts	82,276	82,276	82,276	82,276	94,350	15,451
Log-likelihood	−232,270	−232,222	−232,286	−232,279	−263,204	−44,622

Notes. This table reports robustness checks for the upcharging branch of the censored Weibull model, equation (14). *Extended covariates* adds tract share Black, share Hispanic, unemployment, renter share, and college share. *Binary targeting* replaces the soft and hard indicators with a single any-targeting indicator. *PCA covariates* replaces the tract controls with their first three principal components. *No fixed effects* drops all fixed effects and targeting controls. *2023* restricts to the 2023 cross-section without year effects. Constants are not comparable across fixed-effect structures. Standard errors in parentheses, from the inverse observed information at the refined optimum.

Table A.12: Stage 3 estimation robustness checks: downcharging region ($R > S$)

	Baseline	**Extended covariates**	**Binary targeting**	**PCA covariates**	**No fixed effects**	**2023**
Constant (μ_0^-)	2.408	3.287	2.413	4.044	1.198	1.110
	(0.336)	(0.355)	(0.336)	(0.258)	(0.162)	(0.535)
Gap d^- per \$ (μ_1^-)	0.00017	0.00012	0.00017	0.00013	0.00044	0.00032
	(0.00003)	(0.00004)	(0.00003)	(0.00003)	(0.00004)	(0.00006)
Poverty	0.002	−0.154	−0.002	—	0.130	0.745
	(0.080)	(0.101)	(0.080)		(0.070)	(0.191)
Vacancy	−0.017	−0.152	−0.017	—	−0.108	−0.537
	(0.111)	(0.115)	(0.111)		(0.103)	(0.220)
Log median rent	0.238	0.074	0.238	—	0.507	0.540
	(0.033)	(0.038)	(0.033)		(0.022)	(0.066)
Soft targeting	−0.049	−0.013	—	−0.017	—	−0.024
	(0.037)	(0.037)		(0.037)		(0.072)
Hard targeting	−0.121	−0.088	—	−0.090	—	−0.177
	(0.032)	(0.032)		(0.032)		(0.059)
Targeted, any (μ_V^-)	—	—	−0.090	—	—	—
			(0.026)			
No reduction below S (π^{0-})	0.367	0.367	0.367	0.367	0.376	0.313
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.008)
Shape (α^-)	1.244	1.252	1.244	1.251	1.145	1.256
	(0.009)	(0.009)	(0.009)	(0.009)	(0.007)	(0.020)
Tract covariates	Baseline	Extended	Baseline	3 PCs	Baseline	Baseline
Fixed effects	Bed, yr, CBSA	Bed, yr, CBSA	Bed, yr, CBSA	Bed, yr, CBSA	None	Bed, CBSA
Contracts	19,155	19,155	19,155	19,155	21,937	3,464
Log-likelihood	−81,072	−81,016	−81,073	−81,024	−92,680	−16,388

Notes. This table reports robustness checks for the downcharging branch, equation (15). Specifications and standard errors as in Table A.11.

Table A.13: Stage 2 estimation robustness checks

	Baseline	**All CBSAs** 2015–2023	**All CBSAs** 2023	**Universe** 2023
$(S-R)_+$ per \$ (θ_1)	0.00030	0.00030	−0.00015	−0.00015
	(0.00015)	(0.00014)	(0.00013)	(0.00013)
$(R-S)_+$ per \$ (θ_2)	−0.00103	−0.00107	−0.00110	−0.00105
	(0.00027)	(0.00027)	(0.00029)	(0.00028)
Poverty	1.743	1.700	1.564	1.658
	(0.183)	(0.171)	(0.171)	(0.169)
Vacancy	0.123	0.087	0.537	0.548
	(0.356)	(0.338)	(0.412)	(0.432)
Log median rent	−0.812	−0.819	−0.577	−0.551
	(0.078)	(0.077)	(0.127)	(0.129)
Soft targeting	1.607	1.607	1.647	1.646
	(0.093)	(0.088)	(0.081)	(0.085)
Hard targeting	2.081	2.077	2.049	2.045
	(0.115)	(0.110)	(0.084)	(0.086)
Beds = 1	0.692	0.699	0.491	0.477
	(0.121)	(0.118)	(0.197)	(0.203)
Beds = 2	0.252	0.261	0.132	0.114
	(0.118)	(0.115)	(0.198)	(0.203)
Beds = 3	−0.399	−0.389	−0.386	−0.405
	(0.109)	(0.105)	(0.181)	(0.185)
Beds = 4	−0.757	−0.743	−0.614	−0.630
	(0.131)	(0.127)	(0.207)	(0.211)
CBSA fixed effects	174	388	326	168
Year fixed effects	9	9	—	—
Observations	15,162,157	16,123,086	4,562,712	4,263,013
Linked (match) rate	0.0101	0.0100	0.0096	0.0097
Log-likelihood	−765,692	−810,106	−224,427	−211,213

Notes. This table reports robustness checks for the Stage 2 logit of equation (17). The baseline column restricts the sample to the Stage 3 estimation universe and is the specification carried into Stage 1, the all-CBSA columns use the unrestricted panel, and the 2023 columns are single cross-sections without year effects. The below-standard slope θ_1 is identified by the pooled panel's payment standard variation and is insignificant in the single-year cross-sections.