# Homeownership as Rental-Rate Risk Insurance*

Hasan Çetin
University of Minnesota

Jakub Pawelczak
Universidad del Pacífico

This version: September 2026. Preliminary and incomplete.

**Abstract**

Homeownership can insure households against rental-rate risk: buying fixes the price of housing services that renters repeatedly face. But rental risk also creates valuable rental positions: a renter with a favorable rent draw gives up that claim by moving or buying. We ask which force dominates: the insurance value of owning or the lock-in value of a favorable rental draw. To measure rental-rate risk, we combine PSID data with long external rent indexes, overcoming a short panel's inability to distinguish persistent rent shocks from permanent heterogeneity. We find three distinct components: a highly persistent market component, persistent unit-specific risk within a tenancy, and substantial re-draw risk when households enter a new tenancy. We estimate a life-cycle model of tenure choice and find that removing rental-rate risk raises homeownership by 3.29 pp. Rental-rate risk raises ownership through insurance but lowers it through lock-in to favorable rent draws; the insurance channel raises ownership by 0.36 pp, but lock-in dominates. Renters with unfavorable rent draws buy at nearly twice the rate of renters with favorable draws, and the welfare effects of rental risk are strongly redistributive, with losses concentrated among poorer households. A policy that protects renters from losing favorable rent draws when they move releases cheap-rent lock-in and raises homeownership, making tenant protection and homeownership complementary forms of insurance against rental-rate risk.

**JEL codes**: D15, E21, R21.

**Keywords**: housing, tenure choice, rental risk, life cycle.

*We thank participants of workshops at the University of Minnesota and of the 2026 SAET Conference in Rio de Janeiro for comments. All errors are our own. Email: cetin019@umn.edu, jm.pawelczak@up.edu.pe.

# 1 Introduction

Tenure choice is partly a risk-management decision: renters remain exposed to future rents, while homeowners hedge that exposure but take on house-price risk and the costs of ownership (Sinai and Souleles, 2005; Cocco, 2005; Rampini and Viswanathan, 2016). But this familiar hedge logic is incomplete when the rent a household pays is persistent. A renter with a favorable rent draw is not only exposed to future rental-rate risk; she also holds a valuable position that she risks losing by moving and gives up by buying. Rental-rate risk can therefore push tenure choice in opposite directions: uncertainty raises the value of ownership as insurance, while favorable rent draws raise the value of remaining a renter through lock-in. Which force dominates is therefore an empirical question about the rent process households face: how large and persistent rent shocks are, and how much of a current rent draw survives a move.

Measuring that process is not as simple as reading rental risk from the dispersion of observed rents. Within a tenancy, residual rent variation combines permanent differences in rent levels across households and units, persistent shocks, and measurement error. Over the short renter spells observed in the PSID, a sufficiently persistent shock looks almost permanent, leaving household data alone with little leverage to separate persistence from heterogeneity. We therefore combine the PSID with external rent series that isolate different parts of the process. Regional CPI rent indexes extend back to 1977 and contain no household-specific component, allowing us to estimate market persistence separately. We absorb region-common movements from PSID rents and use the remaining panel variation to estimate the idiosyncratic process within a tenancy. Finally, new-tenant and all-tenant rent indexes, together with rent changes among movers and stayers in the PSID, discipline what happens at re-entry: whether moving brings a predictable change in rent or instead exposes the household to a new draw. The result is a household rent process with separate market, within-tenancy, and re-entry components—the objects needed to quantify both insurance and lock-in.

The estimated rent process has three main features. First, market rents are highly persistent: we cannot reject a unit root nationally or in three of four census regions (Table 1), although market innovations are small relative to household-specific ones. Second, within a tenancy, idiosyncratic rent shocks are large and persistent; at the model's three-year frequency, their persistence is 0.83 and the innovation standard deviation is 0.285 log points (Table 2). Third, re-entry is primarily a dispersion phenomenon rather than a level effect. The average gap between new-tenant and sitting-tenant rents is small, but rent changes are much more dispersed for movers than for stayers (Table 3), implying that moving exposes households to a substantial re-draw of their rent position. The estimates also show why raw rent changes are a poor measure of risk: measurement error accounts for about 68 percent of the variance of reported within-unit rent changes. Taken together, the estimated dynamics make favorable rent draws durable enough to be valuable, yet vulnerable to moving—the combination that generates lock-in alongside insurance.

We embed the estimated rent process in a life-cycle model in which households choose whether to rent or own, how much housing to consume, and how much to save. Households face income and house-price risk, a down-payment requirement, a mortgage borrowing wedge, moving costs, and a bequest motive. The key departure from a standard tenure model is that the rental-rate process depends on the household's housing decision. A renter who remains in the same unit carries her current rent draw forward according to the estimated within-tenancy process; entering a new rental partially re-draws that position, while buying gives it up altogether (Section 3.2). This single feature generates both forces we want to quantify: ownership insures the household against future rental-rate risk, but a favorable current rent creates an incentive to remain a renter. We estimate the model's preference parameters jointly by simulated minimum distance using moments of aggregate and life-cycle ownership, wealth, rent expenditures, and renter mobility, while taking the income, house-price, and rental-rate processes from the data (Section 3.7).

The main result is that rental-rate risk lowers, rather than raises, aggregate homeownership in the estimated model. Removing all rental-rate risk raises the homeownership rate by 3.29 percentage points (Table 7). The sign reveals which force dominates: if rental risk mattered only because ownership insures it, removing the risk would reduce the incentive to buy. Indeed, when we isolate this pure insurance channel while preserving the cross-sectional dispersion of rent positions, removing uncertainty lowers homeownership by 0.36 percentage points. But eliminating the rent dispersion that creates favorable and unfavorable positions more than offsets this effect; the idiosyncratic component alone accounts for 2.74 percentage points of the total ownership response. The lock-in mechanism is also visible in household behavior: renters with unfavorable rent draws buy at a rate of 0.51, compared with 0.29 among renters with favorable draws (Table 8). Thus the insurance motive is present, but quantitatively it is dominated by the value renters place on keeping favorable rent positions.

The average welfare effect of rental-rate risk masks large distributional consequences. Removing rental-rate risk is worth 1.08 percent of lifetime consumption on average, but the gains are concentrated among households with little ability to absorb unfavorable rent draws: the poorest lifetime-consumption quintile gains 6.8 percent, while the richest would, on average, prefer the risky environment (Figure 9a). The incidence also varies sharply with the rent position a household happens to hold: renters with unfavorable draws would gain 8.05 percent from eliminating the risk, compared with 1.01 percent for those with favorable draws. This distributional pattern motivates a tenant-protection experiment that insures renters against losing favorable rent positions when they move. The policy raises homeownership by 0.82 percentage points while delivering gains equivalent to 0.71 percent of consumption on average and 2.03 percent for the poorest quintile (Table 10). By removing the penalty associated with leaving a favorable rental position, tenant protection releases some renters from lock-in rather than crowding out ownership. In this sense, tenant protection and homeownership can be complementary forms of insurance against rental-rate risk.

Our paper builds most directly on the idea that homeownership provides insurance against future housing costs. Sinai and Souleles (2005) show that owning can hedge rent risk at the cost of exposure to house-price risk, while Amior and Halket (2014) provide complementary evidence on how housing risk shapes homeownership. We take this insurance logic to the level at which rental risk is borne by the household. Rather than infer renters' exposure from market-level housing volatility, we estimate the conditional process of the rent an individual household pays and distinguish what happens when the household remains in its tenancy from what happens when it moves. This distinction reveals a countervailing force that is absent when rental risk is represented only as uncertainty: persistence turns favorable rent realizations into valuable positions that renters give up by buying. Our quantitative contribution is therefore not only to measure the insurance motive, but to separate it from the lock-in generated by the same rental-rate process. More broadly, we add this measured renter-side risk to the life-cycle housing literature, where income risk, house-price risk, borrowing constraints, and adjustment costs shape tenure and portfolio decisions (Cocco, 2005; Bajari et al., 2013; Halket and Vasudev, 2014; Li et al., 2016; Sommer and Sullivan, 2018; Kaplan et al., 2020).

The empirical contribution connects this housing literature to work on measuring idiosyncratic risk from household panels. Income-dynamics papers use panel covariance structures to distinguish persistent shocks, permanent heterogeneity, and measurement error, because these objects have different implications for insurance and welfare (Deaton and Paxson, 1994; Meghir and Pistaferri, 2004; Storesletten et al., 2004; Blundell et al., 2008). We bring that decomposition to housing costs, where short renter spells create a particularly sharp persistence-versus-heterogeneity problem and external rent indexes provide the missing time-series information. Our treatment of re-entry also builds on evidence that rents adjust differently for continuing and new tenants (Genesove, 2003). We use that distinction not primarily to study rent rigidity, but to measure how a household's rental position changes when it moves. Finally, our policy experiment connects to work on tenant protection and rent regulation, including Diamond, McQuade and Qian (2019). Because our counterfactual holds rents, house prices, and housing supply fixed, it isolates a narrower mechanism: protecting renters from losing favorable rent positions can reduce lock-in and increase transitions into homeownership.

The rest of the paper proceeds as follows. Section 2 measures and decomposes household rental-rate risk. Section 3 embeds the estimated rent process in a life-cycle model of tenure choice and describes its estimation. Section 4 uses the model to quantify the insurance and lock-in channels, their welfare consequences, and the policy counterfactuals. Section 5 concludes.

# 2 Measuring Household Rental-Rate Risk

## 2.1 The Object

We measure the rent a household pays for its unit net of what the observed characteristics of the unit and of the occupant predict, and how that rent evolves while the household keeps the unit and when it enters a new one. The end product is the distribution of future rent that a renter faces if it keeps its unit and if it moves (Section 2.7). The rent is measured as the residual of a hedonic regression of log rent,

$$\log Q_{it} = \mu_{r(i,t)t} + \kappa_u' x_{it}^u + \kappa_o' x_{it}^o + \tilde{q}_{it}, \tag{1}$$

where $Q_{it}$ is the annual rent of the unit that household $i$ occupies in PSID wave $t$, deflated as described in Section 2.2; $\mu_{rt}$ is an effect for census region $r$ in the survey year of wave $t$, and $r(i,t)$ is the household's region in that year; $x_{it}^u$ collects characteristics of the unit (the log number of rooms and indicators of dwelling type); $x_{it}^o$ collects characteristics of the occupant (log real family income, the inverse hyperbolic sine of real family net worth, and the head's years of schooling and race); and $\kappa_u$ and $\kappa_o$ are coefficient vectors. We estimate (1) by least squares weighted by the head's PSID core/immigrant individual longitudinal weight. Its residual $\tilde{q}_{it}$ is the log rent residual.

The region-by-year effects absorb every movement common to the renters of a census region in a survey year, so the residual contains no region-level market component; Section 2.4 measures that component from external indexes. The unit characteristics hold the recorded quantity of housing fixed. The occupant controls absorb rent differences that go with the household rather than with the recorded unit: location, condition and amenities are not in the data and plausibly vary with the occupant's resources. The controls have a cost. Within a tenancy the unit is fixed but the occupant's income and net worth change, so the change in the residual of a household that keeps its unit contains $-\kappa_o' \Delta x_{it}^o$ in addition to the repricing of the unit. The term is visible in the data. In a regression of stayers' residual changes on information dated before the change, lagged growth of family income contributes 80 percent of the $R^2$, and it loses its predictive power ($p = 0.99$) when the hedonic regression omits log family income and the inverse hyperbolic sine of net worth (Appendix G). The term adds variation unrelated to the price of the unit, while the controls also remove any genuine co-movement of rent with the occupant's resources, so the net effect on the estimated variance components has no definite sign.

We model the residual as the sum of three independent components,

$$\tilde{q}_{it} = f_i^q + z_{it} + m_{it}, \qquad z_{it} = \rho_w z_{i,t-1} + \eta_{it} \quad \text{within a tenancy}, \tag{2}$$

where consecutive waves are two years apart. The permanent rent type $f_i^q$ has variance $\sigma_f^2 \equiv \mathrm{Var}(f_i^q)$. The within-tenancy component $z_{it}$ is the household's rent position: it has persistence $\rho_w$ per wave, an innovation $\eta_{it}$ with standard deviation $\sigma_{\eta,w}$ per wave, and stationary variance $\sigma_z^2 = \sigma_{\eta,w}^2/(1-\rho_w^2)$. A low $z_{it}$ is a favorable position and a high $z_{it}$ an unfavorable one. The last

component, $m_{it}$, is classical measurement error with variance $\sigma_m^2$. The model of Section 3 runs in three-year periods, over which $z$ has persistence $\rho_q = \rho_w^{3/2}$ and innovation standard deviation $\sigma_\eta$ (Section 2.5).

Rent movements vary along two dimensions: whether they are common to a market or specific to the household's unit, and whether they arrive while the household keeps its unit or when it enters a new one. Within a tenancy, the market cell is the market component $g_t$, which moves the rents of all tenants in a market, and the unit cell is the within-tenancy component $z_{it}$, which reprices the household's own unit. At re-entry, the market cell is a level step, a predictable premium of new leases over the rents of sitting tenants, and the unit cell is the re-draw of the position, with correlation $\rho_{\text{move}}$ between the new position and the old one. Sections 2.4–2.6 take up the cells in turn. For the level step, the data show that it is small next to the dispersion of rent changes at a move but do not identify its sign (Section 2.6). The model receives three risk components, the market component, the within-tenancy component and the re-entry re-draw, and sets the level step to zero. The permanent rent type is heterogeneity known to the household, not risk, and measurement error is not risk either; neither enters the model's shock processes.

One assumption underlies the reading of $z_{it}$ as a position worth holding.

**Assumption 1** (Price wedge)**.** *Conditional on the regressors of* (1)*, $f_i^q$ and $z_{it}$ are differences in the price a household pays for a given flow of housing services, not differences in unmeasured quality.*

Under the assumption, a favorable position is a claim on cheap housing services that the household can lose. Such wedges arise when landlords reprice sitting tenants less often than new ones (Genesove, 2003) and when finding a comparable unit takes search. If a low residual reflected unmeasured quality instead, a cheap rent would buy a worse unit and there would be nothing to lock into. The within-tenancy moments limit the damage of a violation. Quality that stays constant over a tenancy drops out of stayers' rent changes, which identify the innovation $\sigma_{\eta,w}$ (Section 2.5), and in the level autocovariances within a tenancy it shifts every lag by the same amount, so it loads on $\sigma_f^2$ rather than on $z$. Two departures remain. Slow unobserved depreciation or improvement of the unit during a tenancy would appear as a persistent shock to $z$. And at a move the unobserved quality of the new unit differs from that of the old one, so movers' rent changes, which identify $\rho_{\text{move}}$ (Section 2.6), mix the re-draw of the price with a change in quality. The assumption therefore bears more on the persistence of positions and on the re-draw, the objects behind lock-in, than on the size of the innovation, the object behind insurance. Rents set by rule, as in assisted housing, are outside it. We maintain the assumption throughout.

## 2.2 Data

Household rents come from the Panel Study of Income Dynamics (PSID) family files, 2001–2021, 11 biennial waves. The panel contains household heads aged 27–80 who own or rent and carry a positive PSID core/immigrant individual longitudinal weight, 53,432 head-wave observations

from 9,841 households. The renter sample keeps heads who rent in the four census regions and report a positive monthly rent, a positive family income and a positive room count, within the bounds of Appendix A: 21,319 observations from 6,108 households. The extract does not identify public or subsidized housing, so renters whose rent is set by rule remain in the sample. We annualize reported rent and divide it by the consumer price index (CPI) for all items less shelter of the household's census region, averaged over January to June of the survey year to match PSID interview timing, so rent movements are measured against the prices of goods and services other than shelter. Because the region-by-year effects of (1) absorb any deflator that varies only by region and year, the choice of deflator leaves the residual unchanged; it matters only for the region-by-year effects themselves and for the external series. Appendix A lists the filters and variables, and Table A.1 the series.

Two external sources measure what the panel cannot. The first is the CPI for rent of primary residence, which the Bureau of Labor Statistics (BLS) publishes for the nation and for each census region. The regional indexes start in 1977; the annual real series we use cover 1978–2025 (48 observations), and a long national series covers 1941–2025. The second is the pair of BLS research indexes of rents paid by new tenants and by all tenants (Adams et al., 2024), quarterly, from 2005Q1 for new tenants and 1999Q4 for all tenants, both to 2025Q3.

### 2.3 Why a Household Panel Is Not Enough

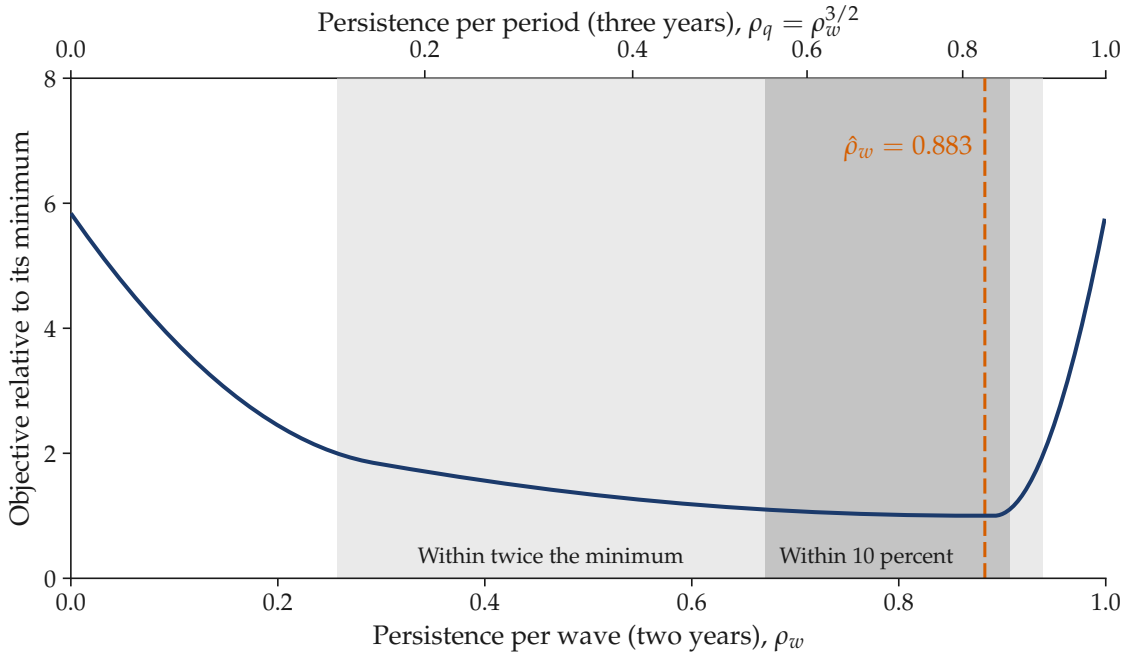
Under (2), the autocovariances of the residual in levels, within a tenancy, are

$$\mathrm{lv}_0 = \sigma_f^2 + \sigma_z^2 + \sigma_m^2, \qquad \mathrm{lv}_k \equiv \mathrm{Cov}(\tilde{q}_{it}, \tilde{q}_{i,t-k}) = \sigma_f^2 + \sigma_z^2 \rho_w^k \quad (k \geq 1). \tag{3}$$

Covariance structures of this kind are the standard tool for household income (Meghir and Pistaferri, 2004; Blundell et al., 2008). A slowly decaying tail of $\mathrm{lv}_k$ fits a large permanent type with moderate persistence, a small permanent type with persistence near one, or any mixture of the two. The distinction is the one the paper turns on. A persistent shock is risk: it creates positions that a household holds and can lose. A permanent type is heterogeneity: it is known when the household enters the rental market, so it affects which tenure the household chooses but carries no uncertainty for ownership to insure. How the dispersion of rents divides between the two fixes how much of it the model treats as positions a household can gain or lose.

Telling the two apart takes long spells in the same unit, and renter spells in the PSID are short: the pairs available fall from 7,306 one wave apart to 1,932 three waves apart. Over the lags the panel covers, the tail barely decays: $\mathrm{lv}_3/\mathrm{lv}_1 = 0.867$, against 0.894 when the hedonic regression has national year effects and an all-items deflator. Netting region-common movements barely changes the ratio, so the slow decay does not come from region-common movements (Appendix C).

**Figure 1:** The equal-weighted minimum-distance objective, profiled over within-tenancy persistence per wave




*Notes.* At each persistence per wave (two years) $\rho_w$ on the grid, the variances $(\sigma_f^2, \sigma_z^2, \sigma_m^2)$ are chosen by nonnegative least squares to fit the seven autocovariance moments of Section 2.5 with equal weights; the curve is the minimized sum of squared deviations divided by its minimum (0.00052). Dark band: within 10 percent of the minimum, $[0.671, 0.907]$. Light band: within twice the minimum, $[0.257, 0.939]$. Dashed line: the estimate. Top axis: persistence per period (three years), $\rho_q = \rho_w^{3/2}$. Residual $\tilde{q}$ of (1); PSID renter heads, 2001–2021.

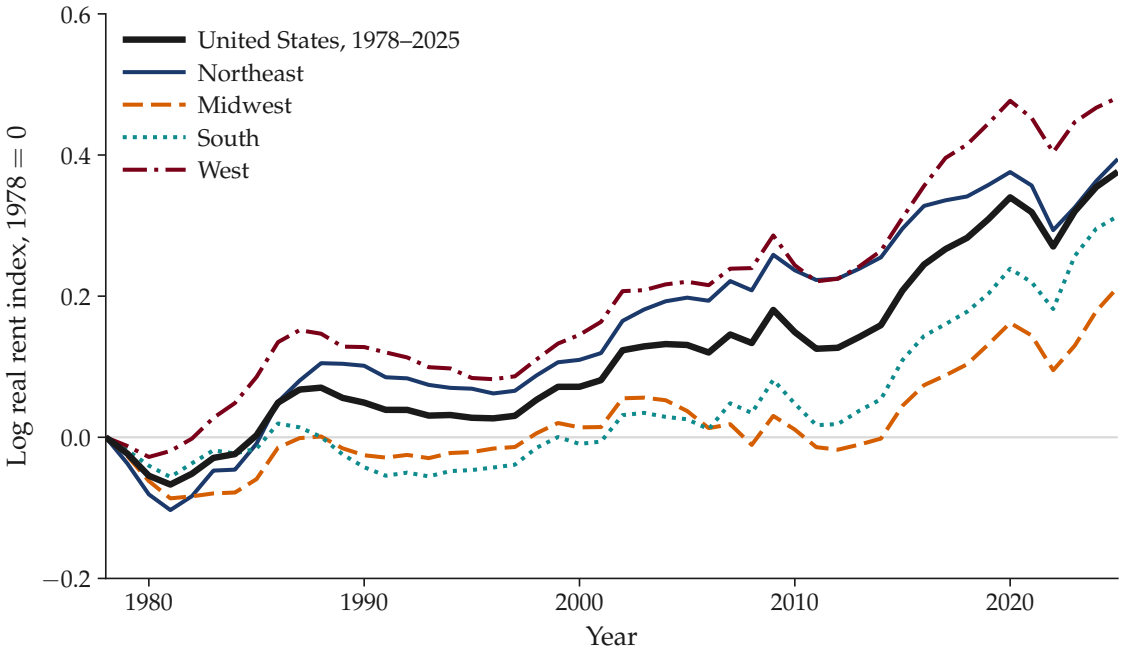
Differencing within a tenancy removes $f_i^q$. The autocovariances of the two-year changes of stayers, $\mathrm{fd}_k \equiv \mathrm{Cov}(\Delta\tilde{q}_{it}, \Delta\tilde{q}_{i,t-k})$, are

$$\mathrm{fd}_0 = 2\sigma_z^2(1-\rho_w) + 2\sigma_m^2, \qquad \mathrm{fd}_1 = -\sigma_z^2(1-\rho_w)^2 - \sigma_m^2, \qquad \mathrm{fd}_2 = -\sigma_z^2\rho_w(1-\rho_w)^2. \tag{4}$$

In population these three moments identify the within-tenancy process with no restriction on the permanent type. With $D \equiv \mathrm{fd}_0/2 + \mathrm{fd}_1 = \sigma_z^2\rho_w(1-\rho_w)$ and $C \equiv \mathrm{fd}_2 = -\sigma_z^2\rho_w(1-\rho_w)^2$, persistence is $\rho_w = 1 + C/D$, and $\sigma_z^2$ and $\sigma_m^2$ follow. For every $\rho_w$ between zero and one the model requires $C$ to be negative. In the data $C = +0.0036$ and $D = 0.0314$, so the closed form returns $\rho_w = 1.114$, above one and outside the range of a stationary process. The lag-two autocovariance rests on 546 pairs from spells of at least four consecutive waves, and its bootstrap standard error, 0.0075, exceeds its value. Differences alone cannot deliver persistence; it has to come from the level autocovariances and from fitting levels and differences jointly.

The levels alone are not decisive either. The first three lagged level autocovariances imply $\rho_w = (\mathrm{lv}_2 - \mathrm{lv}_3)/(\mathrm{lv}_1 - \mathrm{lv}_2) = 0.449$, a ratio of two small differences that lies far below the value at which the joint fit settles. Fitting both blocks together traces a ridge. Figure 1 profiles the estimator of Section 2.5: at each $\rho_w$ the three variances are chosen to fit the seven moments $\mathrm{lv}_0$ to $\mathrm{lv}_3$ and $\mathrm{fd}_0$ to

**Figure 2:** Log real CPI rent of primary residence, United States and census regions, 1978–2025




*Notes.* Log of the January–June average of the monthly ratio of the CPI for rent of primary residence to the CPI for all items less shelter of the same area, set to zero in the first year. The United States line is the national index over the regional window, the series that governs the model's market component. These are the series of Table 1. Source: BLS consumer price indexes, via Federal Reserve Economic Data (FRED).

$\mathrm{fd}_2$, and the curve is the minimized distance relative to its minimum. The objective stays within 10 percent of its minimum for $\rho_w$ from 0.671 to 0.907 per wave, and within twice its minimum from 0.257 to 0.939. It rises slowly toward low persistence and a large permanent type, and steeply beyond $\rho_w = \mathrm{lv}_2/\mathrm{lv}_1 = 0.908$, the persistence at which the permanent variance that fits $\mathrm{lv}_0$ to $\mathrm{lv}_2$ exactly reaches zero (Appendix C). Persistence is point-identified in population and weakly identified in this sample.

The ridge concerns the household component. The market component needs a long time series rather than long household spells, and external indexes supply one (Figure 2, Section 2.4). Within the household component, the size of the innovation changes little along the ridge; the split between persistence and the permanent type is what the ridge leaves open (Section 2.5). The external indexes therefore separate market persistence from the household component but leave that split where the panel leaves it; Sections 2.5 and 4.2 report which results depend on it.

## 2.4 The Market Component

The panel's own measure of market rents, the region-by-year effects $\mu_{rt}$ of (1), has 11 biennial observations per region, too few to tell a unit root from persistence well below one. A market rent index averages over many units, so it carries no household effect, and the regional CPI indexes

supply 48 annual observations each. We use them to estimate the persistence of the market component. For each census region and for the nation, the annual log real index $g_t$, the empirical counterpart of the market component, is the log of the January–June average of the monthly ratio of the CPI for rent of primary residence to the CPI for all items less shelter of the same area. Real rents rose over the period in every region, and the gaps between regions lasted for decades (Figure 2). We estimate a first-order autoregression for $g_t$ with a constant and a linear trend by least squares. Least squares is biased downward near a unit root, so we report the median-unbiased estimate of Andrews (1993), obtained by inverting the simulated median function of the least-squares estimator at the regression sample length, with a 90 percent interval from the simulated quantiles. The unit-root test is the augmented Dickey–Fuller (ADF) test with constant and trend (Dickey and Fuller, 1979). The decision rule, fixed before estimation, classifies a series as trend-stationary if the test rejects at 5 percent and then reports its median-unbiased persistence; otherwise the series keeps a unit root. The verdict for the national index over the regional window, 1978–2025, chosen in advance as the governing series, decides how the market component enters the model. Appendix B gives the construction of the index, the estimator and the test.

Table 1 reports the estimates. The unit root is not rejected for the nation ($p = 0.644$) or in three of four census regions, the Midwest, South and West ($p = 0.443$, 0.948 and 0.439). The long national series, 1941–2025, gives the same verdict ($p = 0.973$). Only the Northeast rejects ($p = 0.028$), with median-unbiased persistence 0.847 per year. The bias correction matters here: least squares gives 0.842 for the nation, while the median-unbiased estimate lies at or above the ceiling of the inversion grid, 1.020. Failing to reject a unit root does not establish one; the national interval runs from 0.799 to the grid ceiling. Under the rule the model's market component is a random walk with innovation standard deviation $\sigma_g = 0.038$ per three-year period: the per-year residual standard deviation of the national trend autoregression, 0.022, aggregated as a random walk. The deterministic drift of real rents is not risk, and the model excludes it (Section 3.2).

The market component is persistent but small. Its innovation standard deviation per period, 0.038, compares with 0.285 for the within-tenancy innovation (Section 2.5), and it accounts for 5 percent of the variance of log rent 30 years ahead for a household that keeps its unit (Section 2.7).

Three caveats apply to the market series. First, the CPI rent index follows the rents of a panel of units, most of them occupied by sitting tenants (Adams et al., 2024). Its annual changes correlate 0.98 with those of the BLS all-tenant index and 0.55 with those of the new-tenant index (Table 3). If the rents of sitting tenants adjust slowly, the index overstates the persistence and understates the volatility of the rents a household meets on a new lease. Second, the PSID region-by-year effects move with the index but by more than one-for-one: regressing the estimated $\hat{\mu}_{rt}$ on the regional index $g_{rt}$, both demeaned within region, gives a slope of 1.215 with standard error 0.095 and $R^2 = 0.81$ over 44 region-years. A slope above one would arise if PSID rents, which include new

**Table 1:** Persistence of real market rents: regional and national CPI rent indexes

| | | | Persistence | | | Innovation s.d. | | Unit-root test | |
|---|---|---|---|---|---|---|---|---|---|
| Series | Years | Obs. | $\hat{\rho}_{\text{OLS}}$ | $\hat{\rho}_{\mu}$[a] | Interval[b] | Per year | Per period | ADF $p$ | Verdict[c] |
| Northeast | 1978–2025 | 48 | 0.729 | 0.847 | [0.642, 1.020] | 0.023 | 0.034 | 0.028 | trend-stationary |
| Midwest | 1978–2025 | 48 | 0.849 | $\geq 1.020$ | [0.809, 1.020] | 0.022 | 0.038 | 0.443 | unit root |
| South | 1978–2025 | 48 | 0.937 | $\geq 1.020$ | [0.961, 1.020] | 0.023 | 0.040 | 0.948 | unit root |
| West | 1978–2025 | 48 | 0.888 | $\geq 1.020$ | [0.870, 1.020] | 0.023 | 0.040 | 0.439 | unit root |
| United States[d] | 1978–2025 | 48 | 0.842 | $\geq 1.020$ | [0.799, 1.020] | 0.022 | 0.038 | 0.644 | unit root |
| United States, long series | 1941–2025 | 83 | 0.931 | $\geq 1.020$ | [0.920, 1.020] | 0.027 | 0.046 | 0.973 | unit root |

*Notes.* Each row is a first-order autoregression with a constant and a linear trend, estimated by least squares on the annual log real rent index $g_t$ of a census region or of the nation: the log of the January–June average of the monthly ratio of the CPI rent of primary residence to the CPI for all items less shelter of the same area (at least three months required). Obs.: annual observations; the long United States series skips 1945–1946, and its ADF test uses the contiguous years 1947–2025. $\hat{\rho}_{\text{OLS}}$: least-squares coefficient on $g_{t-1}$. ADF $p$: augmented Dickey–Fuller $p$-value (Dickey and Fuller, 1979) with constant and trend, lag length chosen by the Akaike criterion. Innovation s.d. per year: standard deviation of the residuals of the least-squares trend autoregression, in every row, including rows in which the unit root is not rejected. Per period (three years, the model's period): the per-year value aggregated as a random walk when the unit root is not rejected (scaled by the square root of the period length) and as a point-sampled autoregression with coefficient $\hat{\rho}_{\mu}$ for the Northeast.

[a] Median-unbiased estimate (Andrews, 1993): the value of $\rho$ at which the median of the least-squares estimator, simulated at the regression sample length (the number of lagged pairs plus one; 20,000 replications per grid point, Gaussian first-order autoregression, same trend regression), equals $\hat{\rho}_{\text{OLS}}$. The grid runs from 0.30 to 1.020 in steps of 0.004. $\geq$: the estimate lies at or above the grid ceiling; the entry is a bound, not a point estimate.

[b] Equal-tailed 90 percent interval from the same inversion applied to the simulated quantiles of the estimator. The upper end of every interval lies at the grid ceiling, 1.020, and is therefore a bound.

[c] Rule fixed before estimation: trend-stationary if the ADF $p$-value is below 0.05; otherwise the unit root is not rejected ("unit root"). The verdict of the governing series decides how the market component enters the model.

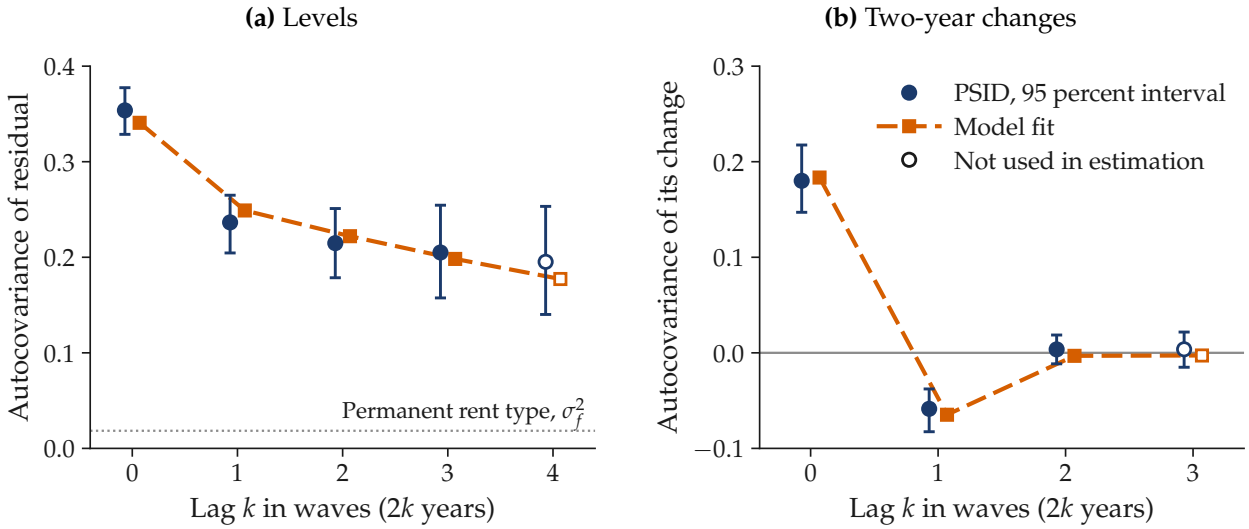
[d] Governing series for the model's market component, fixed before estimation because its window matches the regional series and excludes the postwar rent-control era. The model's market component is a random walk with innovation s.d. per period 0.038. Series identifiers are in Table A.1. Sources: BLS consumer price indexes via FRED; authors' estimates.

leases, move more than an index that mostly follows sitting tenants; in levels it also reflects differences in trend. The regression is unweighted, and its standard error treats $\hat{\mu}_{rt}$ as data. Third, the geography is coarse. The public PSID identifies only four census regions. The region-by-year effects remove all regional movements from the residual, while the model's market component carries the national innovation, so regional deviations from the national index enter neither component, and the model contains no rent risk specific to a region. Movements common to a metropolitan area but not to its region are not absorbed by $\mu_{rt}$ and remain in the residual, where they load on $z_{it}$ or $f_i^q$ according to their persistence.

## 2.5 The Within-Tenancy Component

We estimate $(\rho_w, \sigma_z^2, \sigma_f^2, \sigma_m^2)$ by equal-weighted minimum distance (EWMD) on seven autocovariance moments of the residual, $\text{lv}_0$ to $\text{lv}_3$ and $\text{fd}_0$ to $\text{fd}_2$, at lags measured in waves. The moment set was fixed before estimation; two further moments, $\text{lv}_4$ and $\text{fd}_3$, serve only to test the specification. Equal weights avoid the finite-sample bias of optimally weighted minimum distance in covariance structures (Altonji and Segal, 1996). On raw moments they are not invariant to scale,

**Figure 3:** Autocovariances of the log rent residual: PSID renters and the fitted error-components model




*Notes.* Dots: PSID moments of the residual $\tilde{q}$ of (1), with 95 percent household-block bootstrap intervals. Squares joined by the dashed line: moments implied by the seven-moment estimates of Table 2. Hollow markers: $\text{lv}_4$ and $\text{fd}_3$, not used in estimation. Panel (a): $\text{lv}_k$; the lag-zero model value includes $\sigma_m^2$, and the dotted line marks the variance of the permanent rent type. Panel (b): $\text{fd}_k$ of two-year changes of stayers. Samples differ by moment: $\text{lv}_0$ uses all renter observations (21,319), $\text{lv}_k$ for $k \geq 1$ consecutive waves without a reported move, and $\text{fd}_k$ stayers (3,823 pairs for $\text{fd}_0$).

so the small $\text{fd}_2$ carries little weight, and on a flat profile such choices affect where the estimate lands (Appendix C). The variances are constrained to be nonnegative and $\rho_w$ to lie between zero and 0.999; at each $\rho_w$ the moments are linear in the variances, which gives the profile of Figure 1.

The moments come from three samples: $\text{lv}_0$ from all renter observations, $\text{lv}_k$ for $k \geq 1$ from consecutive waves without a reported move, and $\text{fd}_k$ from stayers, who report no move and an unchanged room count. Of the 7,306 non-mover pairs one wave apart, 3,483 report a changed room count; they enter the level autocovariances but not the stayer sample, and their number suggests that the room count, itself a regressor, is reported with error. Intervals are 95 percent percentile intervals from 999 household-block bootstrap replications, each of which re-estimates the hedonic regression, the moments and the parameters. For the model's three-year period we treat $z$ as an annual autoregression observed every other year, with a stationary variance that does not depend on frequency, so $\rho_q = \rho_w^{3/2}$ and $\sigma_\eta^2 = \sigma_z^2(1 - \rho_q^2)$.

Table 2 reports the estimates. Persistence is 0.883 per wave and 0.830 per three-year period; the innovation standard deviation is 0.240 per wave and 0.285 per period. The permanent rent type has variance 0.018, against a stationary variance of the within-tenancy component of 0.261. Figure 3 compares the fitted moments with the data. The largest misses are on $\text{lv}_0$ and $\text{lv}_1$, which come from different samples; the fitted $\text{fd}_2$ also has the opposite sign to the data.

**Table 2:** The within-tenancy rent component: equal-weighted minimum-distance estimates

| | Estimate | 95 percent interval |
|---|---|---|
| *A. Persistence and innovation size of z* | | |
| Persistence per wave (two years), $\rho_w$ | 0.883 | [0.256, 0.929] |
| Persistence per period (three years), $\rho_q$ | 0.830 | [0.129, 0.896] |
| Innovation s.d. per wave, $\sigma_{\eta,w}$ | 0.240 | [0.179, 0.333] |
| Innovation s.d. per period, $\sigma_\eta$ | 0.285 | [0.209, 0.360] |
| *B. Variance components* | | |
| Within-tenancy component (stationary), $\sigma_z^2$ | 0.261 | [0.074, 0.304] |
| Permanent rent type, $\sigma_f^2$ | 0.018 | [0.000, 0.236] |
| Measurement error, $\sigma_m^2$ | 0.061 | [0.000, 0.079] |
| *C. Measurement error, share of variance (percent)* | | |
| Within-unit rent changes of stayers, $2\sigma_m^2/\mathrm{fd}_0$ | 68.0 | –[a] |
| Rent residual, $\sigma_m^2/\mathrm{lv}_0$ | 17.3 | [0.0, 22.2] |
| *D. Test and sample* | | |
| Overidentification test, bootstrap $p$-value[b] | 0.122 | |
| Renter observations (households) | 21,319 (6,108) | |
| Stayer first-difference pairs | 3,823 | |

*Notes.* Error-components model for the log rent residual, $\tilde{q}_{it} = f_i^q + z_{it} + m_{it}$ with $z_{it} = \rho_w z_{i,t-1} + \eta_{it}$ between consecutive waves: $f_i^q$ is the permanent rent type (variance $\sigma_f^2$), $z_{it}$ the within-tenancy component (innovation $\eta_{it}$, stationary variance $\sigma_z^2$), $m_{it}$ classical measurement error (variance $\sigma_m^2$). $\tilde{q}_{it}$ is the residual of a weighted hedonic regression of log real annual rent on region×year effects, unit characteristics and occupant controls (Appendix A); PSID renter heads aged 27–80, 2001–2021. Estimator: equal-weighted minimum distance on seven autocovariance moments at lags measured in waves, the level autocovariances $\mathrm{lv}_0$–$\mathrm{lv}_3$ and the first-difference autocovariances $\mathrm{fd}_0$–$\mathrm{fd}_2$ (Altonji and Segal, 1996); variances constrained to be nonnegative. $\mathrm{lv}_0$ uses all renter observations, $\mathrm{lv}_k$ for $k \geq 1$ consecutive waves without a reported move, and $\mathrm{fd}_k$ stayers (no reported move and an unchanged room count). Per-period values treat $z$ as an annual autoregression observed every wave: $\rho_q = \rho_w^{3/2}$ and $\sigma_\eta^2 = \sigma_z^2(1-\rho_q^2)$, with $\sigma_z^2$ the same at every frequency. Intervals: 95 percent percentile intervals from 999 household-block bootstrap replications (households resampled from the head panel; hedonic regression, moments and estimator recomputed in each replication). Variances are in squared log points.

[a] Not stored by the bootstrap. The share is sensitive to the specification and imprecise: the interval for $\sigma_m^2$ runs from 0.000 to 0.079.

[b] Distance on nine moments, adding $\mathrm{lv}_4$ and $\mathrm{fd}_3$, with bootstrap critical values recentered as in Hall and Horowitz (1996): a joint test of the whole error-components structure, with low power. Table C.1 reports the fit on all nine moments.

Along the ridge, the size of the within-tenancy innovation changes little. Writing $\Delta\tilde{q}_{it} = -(1-\rho_w)z_{i,t-1} + \eta_{it} + \Delta m_{it}$ and substituting $\sigma_z^2 = \sigma_{\eta,w}^2/(1-\rho_w^2)$ into the variance of stayers' changes gives

$$\sigma_{\eta,w}^2 = \tfrac{1}{2}(1+\rho_w)\left[\mathrm{fd}_0 - 2\sigma_m^2\right]. \tag{5}$$

The permanent type does not enter $\mathrm{fd}_0$; $\sigma_m^2$ is tied mainly to $\mathrm{fd}_1$, of which measurement error makes up 94.5 percent at the estimates; and persistence enters only through $\frac{1}{2}(1+\rho_w)$, which moves from 0.836 to 0.954 across the band in which the objective is within 10 percent of its minimum. The

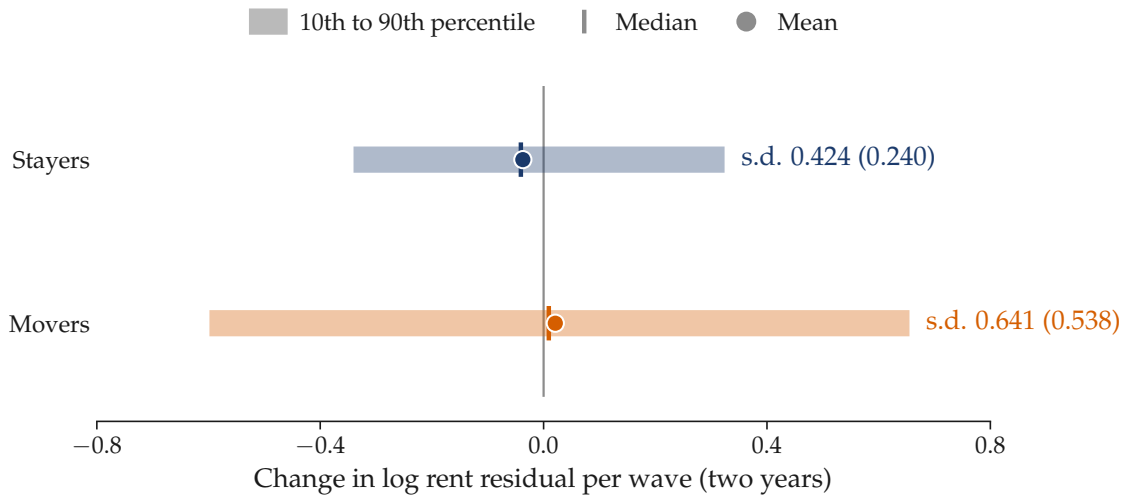
innovation per wave therefore changes little across the band. The per-period innovation that the model uses depends more on persistence, because the conversion to three-year periods involves $\rho_w$ as well: $\sigma_\eta^2 = \frac{1}{2}(1 + \rho_w + \rho_w^2)\left[\mathrm{fd}_0 - 2\sigma_m^2\right]$. Adding $\mathrm{lv}_4$ and $\mathrm{fd}_3$ moves the fit to $\rho_w = 0.765$ per wave and multiplies $\sigma_f^2$ by 7.8, to 0.143, while the innovation variance per wave changes by a factor of 0.99 and $\sigma_\eta$ per period is 0.275 against 0.285 (Appendix C). Robustness to the split is not precision: the bootstrap interval for $\sigma_\eta$ per period is $[0.209, 0.360]$.

The split between persistence and the permanent type is where the panel is weak. The bootstrap interval for $\rho_w$ is $[0.256, 0.929]$ per wave, or $[0.129, 0.896]$ per period, and for $\sigma_f^2$ it is $[0.000, 0.236]$. The design delivers the market component from the external indexes (Section 2.4) and the within-tenancy innovation per wave from within-tenancy changes, which the split barely moves (equation (5)); the split itself is estimated from the panel alone and is weakly identified. The persistence of 0.830 per period that the model uses is the seven-moment estimate, a point on a flat profile. With the innovation nearly fixed, the split moves persistence and the dispersion of positions together, $\sigma_z^2 = \sigma_{\eta,w}^2/(1 - \rho_w^2)$: lower persistence means smaller and shorter-lived positions, and the permanent type absorbs more of the dispersion of rents. The stationary variance $\sigma_z^2$ is 0.261 at the seven-moment estimate and 0.137 at the nine-moment fit, with bootstrap interval $[0.074, 0.304]$. The split therefore reaches the model through $\rho_q$ and the per-period $\sigma_\eta$, through $\sigma_z^2$, through the permanent types that discretize $\sigma_f^2$, and through $\rho_{\mathrm{move}}$, which is computed from $\sigma_z^2$ (Section 2.6).

Measurement error accounts for 68.0 percent of the variance of stayers' within-unit rent changes, $2\sigma_m^2/\mathrm{fd}_0$, and for 17.3 percent of the variance of the residual in levels. The share is imprecise: the interval for $\sigma_m^2$ runs from 0.000 to 0.079. It also counts every transitory component as noise. The variance $\sigma_m^2$ is identified mainly by the negative first-order autocovariance of changes and by the drop from $\mathrm{lv}_0$ to $\mathrm{lv}_1$. A transitory component of rent that households actually pay would produce the same negative autocovariance, as would transitory variation that the occupant controls pass into the residual, so the share is an upper bound on reporting noise. To the extent that $\sigma_m^2$ contains rent that households pay, the model, which excludes measurement error, omits a transitory risk that ownership would insure and that creates no position worth keeping. The drop from $\mathrm{lv}_0$ to $\mathrm{lv}_1$ also compares different samples, all renters against non-mover streaks, which can raise $\sigma_m^2$.

The within-tenancy component is what remains of a stayer's rent change after region-wide movements and measurement error, and not all of it need be repricing of the unit. It can also contain the occupant-control term of Section 2.1, whose net effect has no definite sign; movements common to a metropolitan area but not to its region (Section 2.4); slow change in the unobserved quality of the unit (Assumption 1); and rents set by rule. Rents set as a share of income would move with the tenant's income within a tenancy, yet once the hedonic regression omits the income controls, lagged income growth no longer predicts stayers' rent changes (Section 2.1), which is consistent with a small role for such tenancies in the within-tenancy moments. Whatever its content, the model treats the within-tenancy component as rent risk, and both channels of Section 4 depend on

**Figure 4:** Changes in the log rent residual between consecutive waves, for renters who enter a new unit and for renters who keep theirs




*Notes.* Change in $\tilde{q}$ between consecutive PSID waves, two years apart, for households that rent in both. Movers report a move since January of the previous survey year ($N = 6{,}099$); stayers report none and keep their room count ($N = 3{,}823$). Bars: tenth to ninetieth percentile; circle: mean; tick: median; all weighted. Labels: standard deviation and, in parentheses, the standard deviation net of measurement error, $(s^2 - 2\sigma_m^2)^{1/2}$.

it: the insurance step through the variance of next period's rent, and the dispersion step through $\sigma_z^2 = \sigma_\eta^2/(1-\rho_q^2)$.

The specification test has little power. The overidentification test, the equal-weighted distance on all nine moments with bootstrap critical values recentered as in Hall and Horowitz (1996), does not reject ($p = 0.122$), but it tests the whole error-components structure jointly, and the moments that could reject it are imprecise: the data put positive signs on $\mathrm{fd}_2$ and $\mathrm{fd}_3$ ($+0.0036$ and $+0.0037$), where the model requires negative ones, with standard errors of 0.0075 and 0.0091.

## 2.6 The Re-Entry Margin

Entering a new unit can change a household's rent in two ways: through the level step, a predictable premium of new leases over the rents of sitting tenants (Genesove, 2003), and through a new draw of its position (Figure 4). Because the BLS new-tenant and all-tenant indexes are separately normalized, they show whether that premium changed over time, not its level (Table 3, Panel A). The all-tenant index, like CPI rent, mostly follows sitting tenants (Section 2.4). From 2005Q1 to 2025Q3 the real new-tenant index rose by 14.46 percent in log terms and the all-tenant index by 11.71 percent, and the log gap between them drifted by $+0.13$ percent per year (standard error 0.45, treating quarterly changes as independent); over 2010Q1–2019Q4 the drift is $+0.44$ percent per year (standard error 0.49). A constant premium leaves both indexes growing at the same rate, whatever its size, so the drift measures a change in the premium, not the premium itself. The drift is also an end-point statistic: with the gap set to zero at the start, it is the gap in the last quarter

**Table 3:** Rents at re-entry: new-tenant and all-tenant indexes, and rent changes of movers and stayers

*A. BLS new-tenant and all-tenant rent indexes, 2005Q1–2025Q3*

| | New-tenant | All-tenant |
|---|---|---|
| Cumulative real log change (percent) | 14.46 | 11.71 |
| Correlation of annual changes with CPI rent[a] | 0.55 | 0.98 |
| Log gap, new-tenant minus all-tenant, at 2025Q3 (percent) | 2.75 | |
| Drift of the log gap, percent per year (s.e.)[b] | +0.13 (0.45) | |
| S.d. of the change in the log gap, percent per year | 2.05 | |
| Drift, 2010Q1–2019Q4 (s.e.) | +0.44 (0.49) | |
| S.d., 2010Q1–2019Q4 | 1.53 | |

*B. PSID: change in the rent residual between consecutive waves as a renter*

| | Movers | Stayers |
|---|---|---|
| Mean | +0.021 | −0.037 |
| Standard deviation | 0.641 | 0.424 |
| Net of measurement error | 0.538 | 0.240 |
| Measurement error, percent of variance | 29.8 | 68.0 |
| Pairs of waves | 6,099 | 3,823 |
| Mover minus stayer mean | +0.058 | |
| Bounds on the mover mean[c] | [−0.138, +0.178] | |

*C. Implied correlation of the re-draw of z on entering a new unit*

| | |
|---|---|
| Mover variance net of measurement error, $s^2_{\text{move}} - 2\sigma^2_m$ | 0.289 |
| Variance under an independent re-draw, $2\sigma^2_z$ | 0.522 |
| Re-draw correlation (one-wave estimate, applied per period), $\rho_{\text{move}}$[d] | 0.446 |
| Within-tenancy persistence per wave (two years), $\rho_w$ | 0.883 |
| Within-tenancy persistence per period (three years), $\rho_q$ | 0.830 |
| Difference in the model, $\rho_{\text{move}} - \rho_q$ | −0.384 |

*Notes.* Panel A: BLS research indexes of rents paid by new tenants (R-CPI-NTR) and by all tenants (R-CPI-ATR), quarterly; real changes deflate both by the U.S. CPI for all items less shelter; changes, gaps, drifts and standard deviations are log differences times one hundred (percent); the gap needs no deflator (Figure B.1). Its drift measures relative growth, not the discount a mover gives up (a constant discount of sitting tenants leaves both indexes growing at the same rate). Panel B: change in the rent residual $\tilde{q}$ of Table 2 between consecutive PSID waves as a renter (one wave is two years), weighted. Movers report a move between the waves; stayers report none and keep their room count (3,483 non-mover pairs with a changed room count are excluded). Net of measurement error: variance minus $2\sigma^2_m$. Panel C: $\rho_{\text{move}} = 1 - (s^2_{\text{move}} - 2\sigma^2_m)/(2\sigma^2_z)$, $s_{\text{move}}$ the mover s.d. of Panel B, assuming that the permanent rent type is a household effect kept across the move, that movers' measurement error has the stayers' variance, and that movers' rent dispersion is a draw, not a choice of unit; it holds at the estimates of $\sigma^2_z$ and $\sigma^2_m$, which the data pin down only weakly (Appendix C). In the model, $E[z' \mid z, \text{move}] - E[z' \mid z, \text{stay}] = (\rho_{\text{move}} - \rho_q)\, z$ per period.
[a] Year-to-year changes in the calendar-year mean of the four quarterly log real values of each index, 2005–2024 (19 changes), correlated with those of the log calendar-year mean of the national real CPI rent of primary residence (the index behind the market component, Table 1, here averaged over all twelve months, not January–June).
[b] Mean of the 82 quarterly changes (39 in the shorter window), per year; the s.e. treats quarterly changes as independent. S.d. per year: quarterly s.d. scaled as for a random walk.
[c] Bounds after Lee (2009): the mover distribution trimmed at either end by the share 0.138 of renters who own at the next wave (the exit rate of all renters, not the share of would-be movers whose rent change is unobserved).
[d] Mover changes span one wave; the model applies $\rho_{\text{move}}$ per three-year period by assumption (like for like in the data: $\rho_w$). Sources: BLS; PSID; authors' calculations.

divided by the length of the sample, and the gap moved by about as much in the last quarter as over the whole period (Appendix B). What the indexes show is that, from the start of the sample to its end, rents on new leases did not systematically outgrow the rents of all tenants.

The PSID gives the household-level counterpart (Panel B). Between consecutive waves the residual of renters who move changes on average by $+0.021$ and that of stayers by $-0.037$, a difference of $+0.058$ per wave. A falling stayer mean is what a discount for sitting tenants would produce; it is also what rising occupant income would produce through the controls of (1) if rent rises with income there, and the two cannot be separated here. Both groups are selected. Renters who leave for ownership are not observed as renter movers; trimming the mover distribution at either end by the share of renters who own at the next wave, 13.8 percent, in the spirit of Lee (2009), bounds the mover mean within $[-0.138, +0.178]$. That share is the exit rate of all renters; the share of would-be movers whose rent change is unobserved is plausibly larger, which would widen the bounds. Stayers are households that chose not to move; trimming the stayer distribution by the share of renter pairs with a move, 0.424, bounds $\mathrm{fd}_0$ within $[0.006, 0.310]$, an interval too wide to be informative. The sign of the average level step is therefore not identified, and the model sets the step to zero.

Dispersion is where re-entry shows. The standard deviation of the change in the residual is 0.641 for movers and 0.424 for stayers (Figure 4). The bounds on the mover mean lie well inside one mover standard deviation, so whatever its sign, the level step is small next to the dispersion of rent changes at a move. Subtracting the stayers' measurement-error variance $2\sigma_m^2$ from each variance leaves standard deviations of 0.538 and 0.240. In the model, a household that enters a new unit draws a position correlated with its old one, $z' = \rho_{\text{move}} z + e$, where the innovation $e$ has variance $(1 - \rho_{\text{move}}^2)\sigma_z^2$, so the stationary distribution of $z$ is unchanged. The change $z' - z$ then has variance $2\sigma_z^2(1 - \rho_{\text{move}})$. If the permanent type is kept across the move and movers' measurement error has the stayers' variance, movers' residual changes have variance $s_{\text{move}}^2 = 2\sigma_z^2(1 - \rho_{\text{move}}) + 2\sigma_m^2$, so

$$\rho_{\text{move}} = 1 - \frac{s_{\text{move}}^2 - 2\sigma_m^2}{2\sigma_z^2} = 0.446. \tag{6}$$

The net mover variance, 0.289, is below the 0.522 that an independent re-draw would produce, so the re-draw is partial (Panel C).

Three assumptions deliver (6): (i) the permanent rent type belongs to the household and is kept across a move; (ii) movers' measurement error has the stayers' variance; and (iii) movers' rent dispersion is a draw of the price position, neither a choice nor a change in unrecorded quality, which is Assumption 1 applied at the move. Departures from (ii) and (iii), such as larger measurement error among movers, dispersion that movers choose, or quality that changes at the move, would each mean that less of the mover variance is a re-draw of $z$, so the true $\rho_{\text{move}}$ would be higher. A departure from (i), a permanent type that belongs to the unit and changes at every move, would

also raise the estimate, to 0.517, but the model would then miss that a move replaces the permanent type as well, so the net effect on what a move forfeits has no definite sign (Appendix D). One feature of the data works the other way. The PSID asks whether the family moved since January of the previous survey year, a date that precedes that year's interview, so some reported movers had moved before the earlier interview and paid both observed rents in the same unit; they lower the measured mover dispersion and raise the estimated $\rho_{\text{move}}$. Finally, equation (6) treats movers' positions before the move as distributed like those of all renters. In the model, renters at unfavorable positions move more often (Section 4.3); selection on $z$ biases the formula in a direction that depends on whether movers' positions are more or less dispersed than those of all renters, which the data do not reveal (Appendix D).

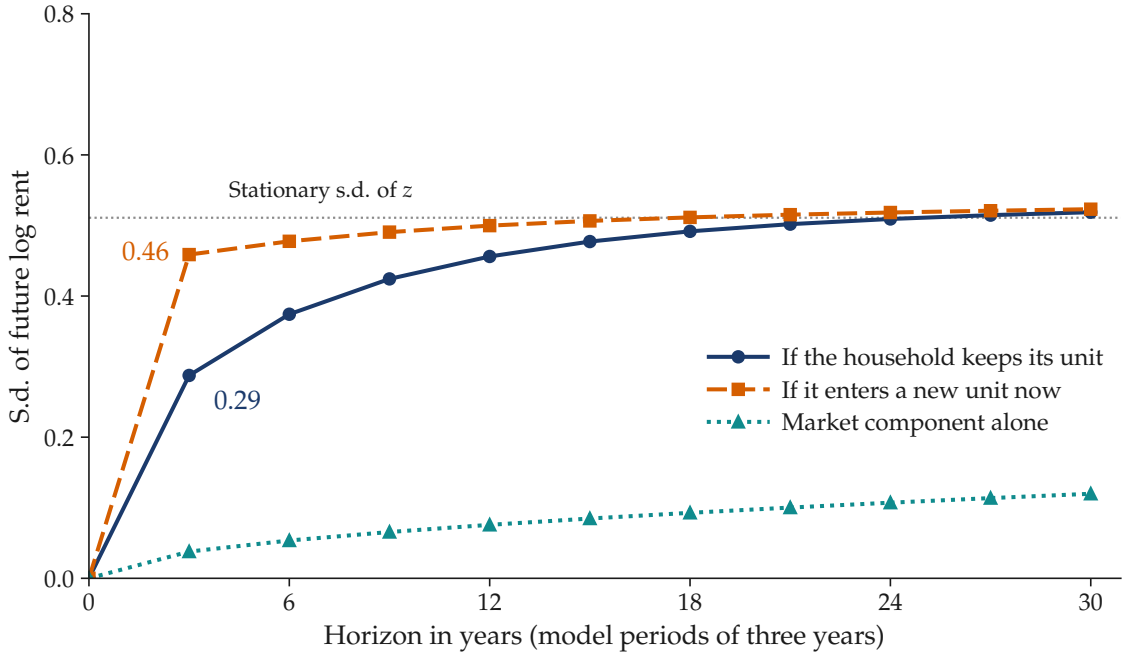
Mover changes span one wave, while the model applies $\rho_{\text{move}}$ per three-year period; the like-for-like comparison in the data is with $\rho_{\tilde{w}} = 0.883$ per wave. The estimate also depends on the point of the ridge, because (6) uses the seven-moment $\sigma_z^2$ and $\sigma_m^2$. At the nine-moment fit, where $\sigma_z^2 = 0.137$, the same formula gives $\rho_{\text{move}} = -0.068$: movers' net rent changes would be more dispersed than even an independent re-draw makes them, which the partial re-draw of the model cannot reproduce. A unit-level permanent type gives 0.975 there. The re-draw correlation is therefore identified only conditional on the point of the ridge, on assumptions (i) to (iii), and on reading a one-wave change as one model period (Appendix D).

Re-entry is a dispersion phenomenon unconditionally, but not conditionally. In the model, a move leaves the expected log rent unchanged on average over the stationary distribution of positions, because the re-draw preserves that distribution and the model sets the level step to zero. Households that choose to move are not drawn at random from that distribution (Section 4.3), so the statement does not describe the average rent change of actual movers. Conditional on the position a household holds, the re-draw pulls its expected log rent toward the mean. Over one model period of three years, with $z$ the current position and $z'$ the next one,

$$E[z' \mid z, \text{move}] - E[z' \mid z, \text{stay}] = (\rho_{\text{move}} - \rho_q)\, z, \tag{7}$$

with $\rho_{\text{move}} - \rho_q = -0.384$. For a household one stationary standard deviation below the mean position, $z = -\sigma_z$ with $\sigma_z = 0.511$, a move raises the expected log rent of the next period by $(\rho_q - \rho_{\text{move}})\sigma_z$ relative to staying. Because $z$ is persistent, a favorable position lowers the rent a household expects to pay for as long as it keeps its unit; the value of holding such a position is lock-in. Because $\rho_{\text{move}} < \rho_q$, entering a new unit gives up part of that value, which is forfeiture at re-entry: a household in a favorable (low) position expects a higher rent if it moves, and one in an unfavorable (high) position a lower rent. Section 3.6 values this wedge in rent levels, as the forfeiture wedge of the renter's one-period problem (equation (23)).

**Figure 5:** Standard deviation of a renter's future log rent by horizon, if it keeps its unit and if it enters a new unit now




*Notes.* Computed analytically from the estimated process, not simulated, conditional on the permanent rent type, the market level and the current position $z$. Parameters per period (three years): $\rho_q = 0.830$, $\sigma_\eta = 0.285$, $\sigma_g = 0.038$ and $\rho_{\text{move}} = 0.446$. Triangles: the market component alone. Dotted line: stationary standard deviation of $z$, 0.511. The two paths also differ in their conditional mean, by $\rho_q^{k-1}(\rho_{\text{move}} - \rho_q)z$; the figure shows the spreads, which do not depend on $z$. Measurement error and the permanent rent type do not enter.

## 2.7 What Households Face

The estimates define the distribution of future rent a renter faces, conditional on what it knows and on whether it keeps its unit. Given its permanent type, the market level and its position $z$, the variance of log rent $k$ periods ahead is

$$k\sigma_g^2 + \sigma_\eta^2 \frac{1-\rho_q^{2k}}{1-\rho_q^2} \qquad \text{if it keeps its unit,}$$

$$k\sigma_g^2 + \rho_q^{2(k-1)}\left(1-\rho_{\text{move}}^2\right)\sigma_z^2 + \sigma_\eta^2 \frac{1-\rho_q^{2(k-1)}}{1-\rho_q^2} \qquad \text{if it enters a new unit now and keeps that one.}$$

The conditional means differ by $\rho_q^{k-1}(\rho_{\text{move}} - \rho_q)z$, the $k$-period version of (7). Figure 5 plots the standard deviations. Three years ahead, the standard deviation is 0.29 for a household that stays and 0.46 for one that moves, which is already close to the stationary standard deviation of positions (the dotted line): a move exposes the household at once to most of the long-run dispersion of positions, which a stayer approaches only gradually. The two spreads rest on different estimates. One period ahead the stayer's variance is $\sigma_g^2 + \sigma_\eta^2$, which rests on the innovation, the object the split between persistence and the permanent type moves least; the mover's is $\sigma_g^2 + (1-\rho_{\text{move}}^2)\sigma_z^2$,

which rests on $\sigma_z^2$ and $\rho_{\text{move}}$, both of which move with the split (Sections 2.5 and 2.6). How much rental-rate risk a renter bears therefore depends on what it does as well as on the rent process.

The model of Section 3 takes these estimates at their point values; Section 3.2 gives the transitions, the events that trigger a re-draw, their timing and what happens to an owner's rent position. Measurement error does not enter, and the permanent rent type enters as three types fixed for life.

# 3 A Life-Cycle Model of Tenure Choice

## 3.1 Environment and Preferences

The economy is a single cohort of households, each solving a finite-horizon problem of consumption, saving and housing. A household enters at age 27 and lives for $T = 18$ periods of three years, to age 81. Households differ in a permanent rent type $f_i^q$ and a permanent income type $f_i^y$, both fixed for life and known to the household, and in the wealth and tenure with which they enter. They do not interact: rents and house prices follow exogenous processes and interest rates are constant (Section 3.5). Every flow, including income, rent, consumption and the costs of owning, is measured per period of three years, and the interest rates, the holding cost of housing and the discount factor are per-period rates. Money is measured in units of mean income per period at age 27. Every ratio of a stock to income, in the data and in the model, divides by annual income, one third of income per period; a ratio of two flows is the same at any frequency.

At the start of period $t$, household $i$ is described by an exogenous state $s_{it} = (g_t, z_{it}, p_{it}, \omega_{it}, \varepsilon_{it})$ and an endogenous state $(a_{it}, h_{it}, o_{it})$. The exogenous state holds the market component $g_t$ and the idiosyncratic component $z_{it}$ of the rent the household would pay, the price $p_{it}$ of owned housing it faces, and the persistent and transitory components $\omega_{it}$ and $\varepsilon_{it}$ of its income (Section 3.2). The endogenous state holds net financial assets $a$, the size $h$ of the dwelling occupied in the previous period, and tenure $o$ in the previous period, with $o = 1$ for an owner and $o = 0$ for a renter. Dwelling size is measured in rooms and chosen from a finite menu $\mathcal{H}$ common to both tenures. We drop the index $i$ where no confusion arises. Within period $t$:

(i) the household observes $(s_t, a_t, h_t, o_t)$ and one taste shock for each tenure;
(ii) it chooses tenure $o'$, a dwelling $h' \in \mathcal{H}$ and next-period assets $a'$, and consumption $c_t$ is what the budget leaves;
(iii) it consumes $c_t$ and the services of $h'$ under tenure $o'$ during period $t$, and $(a', h', o')$ becomes $(a_{t+1}, h_{t+1}, o_{t+1})$;
(iv) the exogenous state of period $t+1$ is drawn, with a law of motion for $z$ that depends on whether $h'$ is a rental unit the household has just entered.

Flow utility has constant relative risk aversion over a constant-elasticity composite of consumption and housing services,

$$u(c,h,o) = \frac{\left[\alpha\, c^{\nu} + \psi_o\,(1-\alpha)\, h^{\nu}\right]^{(1-\gamma)/\nu}}{1-\gamma}, \qquad \psi_1 = 1, \tag{8}$$

where $\gamma$ is relative risk aversion over the composite, $\nu$ sets the elasticity of substitution $1/(1-\nu)$ between consumption and housing services, $\alpha$ is the consumption weight, and the renter housing weight $\psi_0$ scales the services of a rented dwelling relative to an owned dwelling of the same size. Consumption must be positive, and future utility is discounted by $\beta$ per period. After the last period the household values its net worth through a warm-glow bequest,

$$B(a,h,o,p) = \theta_B\, \frac{(a + o\, p\, h)^{1-\gamma}}{1-\gamma}, \tag{9}$$

with bequest weight $\theta_B$, terminal financial assets $a$ and, for an owner, a house of size $h$ valued at the price $p$ of the period after the last. The house counts at its full price, with no holding or moving cost deducted. Unlike the bequest of De Nardi (2004), the function has no shifter that makes bequests a luxury good. With $\gamma$ above one it therefore falls without bound as terminal net worth approaches zero, which gives every household, renters included, a strong motive to leave positive net worth.

### 3.2 Income, Rents, and House Prices

The household faces three exogenous processes: its income, the rent it would pay as a renter, and the price of owned housing. All three are estimated outside the model (Section 2, Appendices E and F, Table E.1), and their innovations are mutually independent.

**Income.** Log income is

$$\log y_{it} = \log \chi_t + f_i^{y} + \omega_{it} + \varepsilon_{it}, \qquad \omega_{it} = \rho_y\, \omega_{i,t-1} + \eta_{it}^{y}, \tag{10}$$

where $\chi_t$ is a deterministic age profile, $\omega_{it}$ a persistent component with persistence $\rho_y = 0.829$ and innovation $\eta_{it}^{y}$ of standard deviation $\sigma_{\eta,y} = 0.326$ per period, and $\varepsilon_{it}$ a transitory component with standard deviation $\sigma_\varepsilon = 0.250$. The permanent income type takes three values that discretize its estimated variance, $\sigma^2_{f^y} = 0.082$, by a three-point quadrature, with factors normalized to mean one. The process is estimated on the family labor income of employed heads aged 27–60 and applied at every age, so the model has no separate retirement income. The profile $\chi_t$ is mean total family income in each three-year age bin relative to its value at 27, pooled over PSID waves; it peaks at 1.62 at age 51 and ends at 0.77 at age 78 (Figure E.2). At 27 the persistent component is drawn with variance 0.390, the residual dispersion of income at the youngest ages of the sample net of the permanent and transitory variances (Appendix E), which exceeds its stationary variance, 0.341.

**Rents.** The rent per room that household $i$ would pay as a renter in period $t$ is

$$q_{it} = \bar{q}\,\frac{e^{g_t}}{E\left[e^{g_t}\right]}\,\frac{e^{f_i^q}}{E\left[e^{f^q}\right]}\,\frac{e^{z_{it}}}{E\left[e^{z}\right]}, \tag{11}$$

where each expectation is over the population distribution of the component, for $g_t$ its distribution at the age of period $t$. Up to these constants, $\log q_{it} = \log \bar{q} + g_t + f_i^q + z_{it}$: a level constant plus the market component, the permanent rent type and the within-tenancy component measured in Section 2. Each factor has mean one, so $\bar{q}$ is the mean rent per room at every age, and it is set so that the mean rent per room at 27, relative to mean income, equals its PSID value. The rent bill for a dwelling of size $h$ is $q_{it}h$, so one rent position prices every size on the menu. Because the hedonic regression (1) conditions on the number of rooms, $\tilde{q}$ measures rent relative to what the size of the unit predicts, which the model reads as a rent per room. The measurement error that Section 2 separates from the rent residual $\tilde{q}_{it} = f_i^q + z_{it} + m_{it}$ does not enter the model.

The market component is a random walk,

$$g_t = g_{t-1} + \eta_t^g, \qquad g_1 = 0, \tag{12}$$

whose innovation $\eta_t^g$ has standard deviation $\sigma_g = 0.038$ per period, estimated from the national CPI rent index (Section 2.4). The random walk has no drift: the model excludes the deterministic trend of real rents. Every household faces this process, and the simulation draws a separate market history for each household (Section 3.5). The process is discretized on a grid that widens with age (Fella et al., 2019), and the normalization in (11) keeps mean rent constant across ages, so the market component spreads rents across market histories without raising the mean rent with age.

The permanent rent type takes three values that discretize $\sigma_f^2 = 0.018$, the variance of $f_i^q$, by a three-point quadrature: rent multipliers 0.78, 0.99 and 1.25, with population shares of 16.7, 66.7 and 16.7 percent. It is heterogeneity, not risk: the household knows it on entry and keeps it for life, including across moves.

**Within-tenancy rent risk.** While a household keeps its rental unit, the idiosyncratic component follows

$$z_{i,t+1} = \rho_q\, z_{it} + \eta_{i,t+1}, \tag{13}$$

with persistence $\rho_q = 0.830$ and innovation $\eta_{i,t+1}$ of standard deviation $\sigma_\eta = 0.285$ per period, converted from the per-wave estimates of Section 2.5 ($\rho_q = \rho_w^{3/2}$, with $\rho_w = 0.883$ per wave of two years). The stationary standard deviation of $z$ is $\sigma_z = \sigma_\eta(1-\rho_q^2)^{-1/2} = 0.511$. The transition matrix $\Pi_{\text{stay}}$ discretizes (13) by the method of Rouwenhorst (1995), which remains accurate for highly persistent processes (Kopecky and Suen, 2010), on five nodes spanning two standard deviations $\sigma_z$ on either side of zero.

**Re-entry.** The transition of $z$ depends on the housing choice. Let $n_t$ indicate that the dwelling chosen in period $t$ is a rental unit the household has just entered. Then

$$z_{t+1} \sim \begin{cases} \Pi_{\text{move}}(\,\cdot \mid z_t) & \text{if } n_t = 1, \\ \Pi_{\text{stay}}(\,\cdot \mid z_t) & \text{otherwise,} \end{cases} \qquad n_t = \mathbf{1}\{o' = 0\}\,\mathbf{1}\{o_t = 1 \text{ or } h' \neq h_t\}. \tag{14}$$

A household enters a new rental unit when it rents after owning or when it rents and changes dwelling size. A renter that keeps its size keeps its unit: the model has no moves between rental units of the same size, so renter mobility in the model is mobility across sizes. The move transition $\Pi_{\text{move}}$ is the Rouwenhorst chain with persistence $\rho_{\text{move}} = 0.446$ on the same grid. A symmetric Rouwenhorst chain has the same binomial stationary distribution whatever its persistence, so the two transitions share the grid and the stationary distribution of $z$, and the move chain has the larger innovation, with standard deviation $\sigma_z(1-\rho_{\text{move}}^2)^{1/2} = 0.457$, or 1.60 times $\sigma_\eta$. Entering a new unit rather than keeping the current one shifts the expected next-period position by $(\rho_{\text{move}} - \rho_q)\,z_t = -0.384\,z_t$, the conditional-mean wedge of equation (7). Averaged over the stationary distribution the re-draw leaves the mean position unchanged; conditional on the current position it does not. The re-draw correlation, estimated from the dispersion of movers' rent changes (equation (6)), carries the qualifications of Section 2.6: it is a one-wave estimate applied per period, and it is conditional on the point on the ridge (Appendix D). Figure 5 shows the implied spread of future rents for a renter who stays and for one who moves.

Two timing conventions complete the law of motion. First, $z_t$ prices whatever rental unit the household occupies in period $t$: a household that enters a new unit in $t$ pays $q_t h'$ there in $t$, and the re-draw governs its position from $t+1$ on, so what a move forfeits takes effect in the period after the move. Second, an owner pays no rent, but its rent state does not disappear. While the household owns, $z$ keeps evolving by $\Pi_{\text{stay}}$; when it returns to renting it enters a new unit, and $z$ is partly re-drawn by $\Pi_{\text{move}}$. A household that buys at position $z$ and rents again after $\ell$ periods of ownership pays rent in its first period back at a position with expected value $\rho_q^\ell z$, which the re-draw then pulls further toward the mean. Buying therefore suspends the use of a position and exposes what remains of it to the re-draw on return. An owner's position matters to it only through this option to return to renting, an option that owners in the model exercise at a rate of 0.104 per three-year period (Section 3.7).

**House prices.** The price of a room of owned housing that household $i$ faces follows

$$\log p_{it} - \log \bar{p} = \rho_p\,(\log p_{i,t-1} - \log \bar{p}) + \eta_{it}^p, \tag{15}$$

with persistence $\rho_p = 0.650$ and innovation $\eta_{it}^p$ of standard deviation $\sigma_{\eta,p} = 0.327$ per period. The process is estimated from owners' self-reported house values net of national year effects (Appendix F), so it has no national component, and the model draws it independently for each household. The permanent component of reported values ($\sigma_{fp}^2 = 0.215$) does not enter the model, and moving does not re-draw the price: a household that changes houses buys at the price it already faces. For a renter the process is the price at which it could buy: a low current price predicts appreciation, so the timing of a purchase has value, and in the estimated economy house-price risk, more than rental-rate risk, keeps households renting (Section 4.4). The constant $\bar{p}$ is set so that the mean house price per room at 27, relative to mean annual income, equals its PSID value. The persistent income component and the house price are also discretized by the Rouwenhorst method (Appendix I).

### 3.3 Assets, Housing, and Frictions

A renter pays the rent bill $q_t h'$ for the dwelling it occupies in period $t$. An owner holds its house at the current price: the house enters the owner's resources at $(1-\delta)\, p_t h_t$, where $\delta$ is the per-period holding cost of owned housing, and keeping the house means buying it back at $p_t$. An owner therefore bears the holding cost, the cost of financing the house and the change in its price. Net financial assets earn

$$r(a) = \begin{cases} r_s & \text{if } a \geq 0, \\ r_m & \text{if } a < 0, \end{cases} \qquad r_m > r_s, \tag{16}$$

where $r_s$ is the real return on saving, $r_m$ the real mortgage rate, and $r_m - r_s$ the mortgage borrowing wedge. The budget constraint in period $t$ is

$$c_t + P_t^{o'} h' + a' + \zeta\, P_t^{o_t} h_t\, \mathbf{1}\{h' \neq h_t \text{ or } o' \neq o_t\} = y_t + \big(1 + r(a_t)\big)\, a_t + o_t\,(1-\delta)\, p_t h_t, \tag{17}$$

where $P_t^1 = p_t$ and $P_t^0 = q_t$ are the prices of a room owned and rented. The last term on the left is the moving cost: a household that changes dwelling size or tenure pays the share $\zeta$ of the price of the unit it leaves, the house value $p_t h_t$ for an owner and the rent bill $q_t h_t$ for a renter. Borrowing is secured by housing:

$$a' \geq -(1-\phi)\, p_t h' \quad \text{if } o' = 1, \qquad a' \geq 0 \quad \text{if } o' = 0. \tag{18}$$

A buyer finances at most the share $1-\phi$ of its house, so it posts a down payment of at least $\phi\, p_t h'$, and a renter cannot borrow. Mortgage debt is one-period debt at $r_m$, rolled over each period subject to (18) at the current price, with no amortization schedule and no cost of repaying early; a fall in the price can force an owner to reduce its debt.

Each friction has a distinct role in the tenure choice. The down payment makes ownership depend on wealth: a young household must save before it can buy, and a renter at an unfavorable position that cannot buy can leave the position only by moving. The mortgage wedge makes leverage

costly and gives owners a reason to repay, and the return on saving governs the buffer stock with which a renter can absorb rent shocks. The holding cost, together with the financing cost and the expected change in the price, is the per-period cost of owning against which the rent bill is compared. The moving cost makes positions sticky. Staying in the unit is free, while a renter pays $\zeta q_t h_t$ both to move and to buy, so the friction separates staying from leaving but not moving from buying. The taste shocks of Section 3.4 smooth the tenure choice. Life-cycle models of housing demand and tenure combine similar credit and transaction frictions with income and house-price risk (Cocco, 2005; Bajari et al., 2013; Halket and Vasudev, 2014; Li et al., 2016; Kaplan et al., 2020); what is specific here is that the rent process depends on the housing choice through (14).

### 3.4 The Household's Problem

Write $x_t = (s_t, a_t, h_t, o_t)$ for the state of a household of given permanent types. For each tenure $o'$ the household chooses its dwelling and saving to solve

$$W_t(x_t; o') = \max_{h' \in \mathcal{H},\, a'} \Big\{ u(c_t, h', o') + \beta\, E\big[ V_{t+1}(x_{t+1}) \mid x_t, o', h' \big] \Big\} \tag{19}$$

subject to (17), (18) and $c_t > 0$, where $x_{t+1} = (s_{t+1}, a', h', o')$. The expectation is over $s_{t+1}$ given $s_t$, with $z_{t+1}$ drawn by (14) and $n_t$ set by $(o_t, h_t)$ and $(o', h')$. The tenure choice carries additive taste shocks: the household chooses the $o'$ with the larger $W_t(x_t; o') + \xi\, v_t(o')$, where $v_t(0)$ and $v_t(1)$ are independent standard type-I extreme-value draws and $\xi$ is their scale. Integrating over the shocks gives the value $V_t$ before the shocks are realized and the probability of owning (Rust, 1987),

$$V_t(x_t) = \xi \Big[ \gamma_E + \log \sum_{o \in \{0,1\}} e^{W_t(x_t; o)/\xi} \Big], \qquad \Pr(o' = 1 \mid x_t) = \frac{e^{W_t(x_t;1)/\xi}}{\sum_{o \in \{0,1\}} e^{W_t(x_t;o)/\xi}}, \tag{20}$$

where $\gamma_E$ is the Euler–Mascheroni constant. The shocks act on the tenure choice only; given tenure, the dwelling and saving choices maximize (19). Staying, moving and buying are all choices of $(o', h')$: a renter stays by choosing $o' = 0$ and $h' = h_t$, moves by choosing $o' = 0$ and $h' \neq h_t$, and buys by choosing $o' = 1$. After the last period the continuation value is the bequest, $V_{T+1}(x_{T+1}) = B(a_{T+1}, h_{T+1}, o_{T+1}, p_{T+1})$, so the expectation in (19) at $t = T$ is over the terminal price.

At age 27 simulated households are assigned to permanent types in proportion to the population shares, and each draws its tenure and net financial assets jointly from the weighted PSID distribution at that age, converted to model units. Every household starts from the same dwelling size, the middle of the menu. Its rent position is drawn from the stationary distribution of $z$, so households enter already holding favorable and unfavorable positions; its house price is drawn from the stationary distribution of $p$, its persistent income from the age-27 distribution, and the market component starts at $g_1 = 0$. The initial position is known on entry, so part of what the counterfactuals of Section 4.1 remove with $z$ is dispersion that a household already faces at 27.

A solution consists of the value functions $\{W_t, V_t\}$, the dwelling and saving policies for each tenure, and the tenure probabilities (20) in every period. We compute it by backward induction on discrete grids for assets and dwelling size; Appendix I describes the grids, the solution method and the treatment of states in which no choice is feasible. Model statistics come from simulating 100,000 lives forward from the initial conditions under these policies.

### 3.5 What the Model Holds Fixed

The model is partial equilibrium. Rents and house prices follow the processes of Section 3.2 and interest rates are constant, whatever households do; dwellings of every size are available in both tenures at those prices, and no market clears. The model describes how households allocate tenure, mobility, saving and consumption given the measured processes; it does not describe how landlords or sellers would reprice a change in those processes, which matters for the policy experiment of Section 4.7. Because no price is determined in equilibrium, a household's problem depends on the market component only through its own rent, and we simulate an independent market history for each household; simulated statistics average over market histories rather than conditioning on one realized path.

Rent, house-price and income risks are mutually independent, so rents and house prices do not co-move. In Sinai and Souleles (2005) owning hedges rent risk because it fixes the cost of housing services, and the price risk an owner takes on is itself partly hedged because house prices capitalize future rents: the value of an owner's house rises with the rents it would otherwise pay. The model keeps the first part and removes the second. An owner's house price is unrelated to its rent position and to the market rent; the price process has no market component, while market rents follow a random walk. The insurance value of ownership measured in this paper is therefore the value of leaving the rental spot market, and it excludes any hedge through the value of the house.

No shock forces a move. A renter enters a new unit only when it chooses to change size or to leave ownership, so every re-draw follows a choice. In the data a move can also follow events outside the model, such as a change of job or of family, and the targeted move rate counts every reported move (Section 3.7). Moves of that kind would shorten the expected tenure over which a favorable position pays off and add re-draws that staying cannot avoid. There is no consumption floor or public transfer: households absorb income and rent shocks through saving, housing and tenure alone.

There is no mortality: every household lives to 81, so model profiles by age follow a cohort of constant composition, while PSID profiles by age also reflect survival and cohort composition. There are no taxes, so the tax treatment of owner-occupied housing (Sommer and Sullivan, 2018) is outside the model.

### 3.6 Two Forces in the Renter's Problem

Rental-rate risk creates two forces in the renter's problem. Uncertainty about future rent makes leaving the rental market valuable as insurance. The persistence of a favorable position makes keeping it valuable, which is lock-in, and the move-contingent transition (14) puts part of that value at stake whenever the renter enters a new unit. We size the two forces over one period of the renter's problem, at the estimated rent process and preferences (Section 3.7), before the life-cycle model combines them with wealth, age and the other risks. The calculation needs only the rent process and the utility function; Appendix H gives the details.

Consider a renter with income $y$ normalized to one per period and a dwelling normalized to one room, so that its rent per room is also its rent bill. It holds position $z$ on the model's grid and pays next period the rent $q(z') = q_0 e^{z'}$, where the scale $q_0$ sets the mean rent under the stationary distribution of $z$ to the median ratio of rent to income among PSID renters, 0.206. The market component and the permanent rent type are held fixed. The renter can keep its unit, so that $z'$ is drawn from $\Pi_{\text{stay}}(\cdot \mid z)$, or enter a new unit of the same size, so that $z'$ is drawn from $\Pi_{\text{move}}(\cdot \mid z)$. In the life-cycle model a renter enters a new unit only by changing size; because one rent position prices every size, the one-period problem sets the size change aside to isolate the rent process. We value each option $d \in \{\text{stay}, \text{move}\}$ by its certainty-equivalent rent $R_d^*(z)$, the certain rent that leaves the renter as well off:

$$u\big(y - R_d^*(z),\, h,\, 0\big) = E\big[u\big(y - q(z'),\, h,\, 0\big) \;\big|\; z, d\big]. \tag{21}$$

Both options are rental, so $\psi_0$ enters both sides of (21) in the same way and creates no wedge between them; the comparison needs no conversion between rented and owned housing. The certainty-equivalent rent of staying splits into its expected value and a risk premium, which define the two forces:

$$\text{Ins}(z) = R_{\text{stay}}^*(z) - E[q' \mid z, \text{stay}], \tag{22}$$

$$\text{Lock}(z) = E[q' \mid z, \text{move}] - E[q' \mid z, \text{stay}]. \tag{23}$$

Insurance is the premium the renter would pay to replace next period's uncertain rent in its unit by its expected value; ownership removes it by taking the household out of the rental spot market. $\text{Lock}(z)$ is the extra rent the renter expects to pay next period if it enters a new unit rather than keeping its own. We call it the forfeiture wedge: over one period it is the part of lock-in that entering a new unit puts at stake, and it leaves out the value of a favorable position against buying. Portable rent positions, $\Pi_{\text{move}} = \Pi_{\text{stay}}$ as in Section 4.7, set it to zero at every $z$. It contains no preference parameter. It is positive at favorable positions and at the mean position, where the larger conditional variance of the move chain raises expected rent in levels, and negative at unfavorable positions, where moving is an escape. At unfavorable positions both forces push the renter out of its unit, and the one-period problem cannot say whether it moves or buys. Freezing each

household's $z$, as one counterfactual of Section 4.1 does, sets both forces to zero while keeping the dispersion of positions, so the problem is silent on the step of Section 4.2 that removes that dispersion.

The two forces differ in order. Fix the renter's quantile in the stationary distribution, $z = k\sigma_z$, and scale the innovation standard deviation $\sigma_\eta$; $\sigma_z$ and the move innovation scale with it. In the continuous-state analog of the two chains,

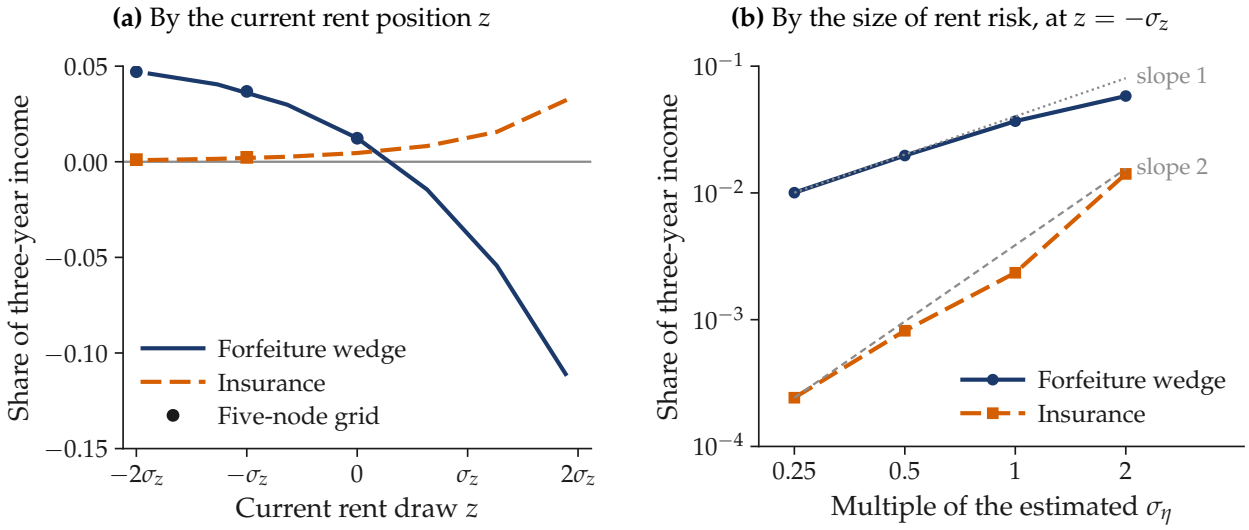
$$\begin{aligned} \text{Lock}(k\sigma_z) &= q_0\,(\rho_{\text{move}} - \rho_q)\,k\,\sigma_z + O(\sigma_z^2), \\ \text{Ins}(z) &\approx \frac{\text{RRA}_c}{2\,(y - E[q' \mid z, \text{stay}])}\,\text{Var}(q' \mid z, \text{stay}) = O(\sigma_z^2), \end{aligned}$$

where $\text{RRA}_c = (\gamma + \nu - 1)\pi + 1 - \nu$ is the relative risk aversion of (8) in consumption and $\pi = \alpha c^\nu / [\alpha c^\nu + \psi_0(1-\alpha)h^\nu]$ is the consumption share of the composite. At any fixed quantile away from the mean position ($k \neq 0$), the forfeiture wedge is first order in the size of rent risk and insurance is second order; at the mean position both are second order. At a fixed $z$ instead, the leading term of the wedge, $q_0(\rho_{\text{move}} - \rho_q)z$, does not shrink with $\sigma_\eta$ at all. The weight $\psi_0$ affects insurance only through $\text{RRA}_c$, which at consumption $y - 0.206$ and a dwelling of one room is 2.12, below $\gamma = 2.33$ because housing dilutes the curvature of utility in consumption.

Figure 6 plots both forces. On the model's grid (the markers of panel (a)), one standard deviation below the mean position the forfeiture wedge is 0.0368 of per-period income and insurance 0.0023, a ratio of 16. At the most favorable grid point, two standard deviations below the mean, the wedge is 0.0471 (23 percent of the mean rent) and insurance 0.0011 (0.5 percent), a ratio of 43. For insurance to match the wedge next period at these two points, $\gamma$ would have to be 14 and 19, against the estimated 2.33 and a consumption risk aversion of 2.12. Panel (b) and Table H.1 confirm the orders approximately at $z = -\sigma_z$; at twice the estimated rent risk, the wedge is still 4 times insurance, although at that multiple most of the premium comes from the highest grid rent, close to income (Appendix H). In logs the conditional-mean wedge is zero at the mean position; in levels a forfeiture wedge of $+0.0123$, 6 percent of the mean rent, remains there, because the move chain has the larger conditional variance and rent is convex in $z$ (Appendix H). Averaged over the stationary distribution, $\text{Lock}(z)$ is zero, since both chains preserve that distribution: what favorable positions forfeit at re-entry, unfavorable positions gain.

The argument has three limits. First, it covers one period. Over a longer horizon the premium grows with the conditional variance of log rent, $\sigma_z^2(1 - \rho_q^{2j})$ at $j$ periods ahead, while the log-rent wedge of a move made today decays as $\rho_q^{j-1}(\rho_{\text{move}} - \rho_q)z$. The orders in $\sigma_\eta$ hold at every horizon, and at leading order the ratio of the wedge $j$ periods ahead to the premium $j$ periods ahead falls with $j$; over a renting spell the calculation does not sign the ratio (Appendix H). A renter who moves also bears a larger premium than one who stays, because its innovation is 1.60 times $\sigma_\eta$; owning removes both. Second, the argument does not settle the tenure margin. Comparing either rental option with buying requires valuing owned housing in rental units, which brings $\psi_0$ back,

**Figure 6:** The forfeiture wedge and the insurance premium in the renter's one-period problem




*Notes.* One-period renter problem of Section 3.6 at the estimated within-tenancy process and preferences, with income per three-year period and the dwelling both normalized to one. The forfeiture wedge is (23); insurance is the premium (22); both are in shares of per-period income. Panel (a): lines are evaluated on a finer Rouwenhorst grid of the same process; markers are the values on the model's five-node grid (the wedge at three nodes, insurance at two), from which the ratios in the text are computed; read at the same $z$, the lines give larger ratios. Panel (b): both forces at $z = -\sigma_z$ with the innovation standard deviation scaled by the multiple on the horizontal axis; log scales; the gray lines have slopes one and two. The multiples stop at two because a slightly larger one puts the highest grid rent above income. Details in Appendix H.

and pricing the costs of owning, the down payment, house-price risk and the household's wealth, none of which a one-period problem contains. The moving cost does not discriminate either, because a renter pays it whether it moves or buys. We therefore compute no static buy threshold, and the comparison leaves out the rest of lock-in: a favorable position is also valuable against buying, since a renter that buys stops paying its low rent, and that part involves the costs of owning. Third, the argument does not sign the aggregate, which depends on how tenure choices respond across positions, wealth and age (Section 4.2).

What the argument establishes is local, and it concerns mobility. Over one period, the rent a renter at a favorable position forfeits by entering a new unit exceeds the premium it would pay to remove next period's rent risk by an order of magnitude, so renters at different positions face sharply different incentives to move; Section 4.3 shows the corresponding pattern in simulated mobility, while buy rates by position barely change when the re-draw is switched off. The ranking is not reproduced on the tenure margin, which the one-period problem leaves out. In the life-cycle model, switching off the re-draw raises ownership by 0.82 points, while freezing positions once the re-draw is off lowers it by more, from 0.766 to 0.754, a step that also makes positions permanent (Section 4.2).

### 3.7 Calibration and Estimation

Table 4 lists the parameters. Four of the seven preference parameters, $\vartheta = (\psi_0, \theta_B, \gamma, \nu)$, are estimated by simulated minimum distance. The other three ($\alpha$, $\beta$, $\xi$) and the five prices and frictions ($r_s$, $r_m$, $\phi$, $\delta$, $\zeta$) are set externally, and the income, rent and house-price processes are estimated outside the model (Section 3.2).

**External parameters.** The discount factor is set to 0.96 per year, $\beta = 0.885$ per period. The interest rates are real 2001–2021 averages: $r_s = 2.92$ and $r_m = 8.39$ percent per period (0.96 and 2.72 percent per year), from the ten-year Treasury yield and the thirty-year mortgage rate, a wedge of 1.76 percentage points per year (Table 4 gives the construction). Saving earns the ten-year yield because the real one-year Treasury rate was negative over the window. The down payment $\phi = 0.20$, the holding cost $\delta = 0.12$ per period (4.2 percent per year, read as the full cost of keeping a house rather than depreciation alone) and the moving cost $\zeta = 0.08$ are set by assumption. The consumption weight $\alpha = 0.80$ is one minus the median renter rent-to-income share, 0.206, rounded. That mapping holds only in the Cobb–Douglas limit with $\psi_0 = \psi_1$; at the estimates the model's rent share depends on all four estimated parameters, and it is itself a targeted moment. The taste-shock scale $\xi = 0.005$ is a smoothing device (Iskhakov et al., 2017) and is not estimated.

**Moments and estimator.** Let $\hat{\mathcal{M}}_j$ be the data value of moment $j$ and $\mathcal{M}_j(\vartheta)$ its simulated counterpart. The criterion is the equal-weighted sum of squared relative gaps,

$$L(\vartheta) = \sum_j \left( \frac{\mathcal{M}_j(\vartheta) - \hat{\mathcal{M}}_j}{\hat{\mathcal{M}}_j} \right)^2, \tag{24}$$

over five moments:

(i) the homeownership rate over all ages;
(ii) the homeownership rate at age 30, the only life-cycle moment in the objective;
(iii) owners' median ratio of net worth, including home equity, to annual family income;
(iv) renters' median ratio of rent to family income;
(v) the renter move rate, the probability that a renter moves to a new rental unit within three years.

The data for moments (i) to (iv) are PSID individuals aged 27–80 in owner- or renter-occupied housing, not only household heads, weighted by the PSID weights (Appendix A). Money values enter only as ratios within an observation, and ownership at 30 is the rate in the three-year age bin that starts there. The move rate comes from the rent-process sample of household heads (Section 2.2): the weighted share of consecutive-wave renter pairs that report a move, 0.424 per two-year wave, converted to three years. The model moments give every simulated household-period the same weight, so the simulated population has a uniform age composition, while the data carry

**Table 4:** Parameters of the life-cycle model

| Parameter | | Value | Annual | Source or target |
|---|---|---|---|---|
| *A. Estimated by simulated minimum distance (Table 5)* | | | | |
| Renter housing weight (owner = 1) | $\psi_0$ | 0.431 | | |
| Bequest weight | $\theta_B$ | 2.15 | | |
| Relative risk aversion over the composite | $\gamma$ | 2.33 | | |
| CES exponent, consumption and housing | $\nu$ | 0.60 | | |
| *B. Externally set preferences* | | | | |
| Consumption weight | $\alpha$ | 0.80 | | One minus the median renter rent-to-income share, 0.206, rounded |
| Discount factor, per period | $\beta$ | 0.885 | 0.96 | Annual value assumed |
| Taste-shock scale, tenure choice | $\xi$ | 0.005 | | Assumed; smooths the tenure choice |
| *C. Prices and frictions* | | | | |
| Real saving rate, percent per period | $r_s$ | 2.92 | 0.96 | Ten-year Treasury yield less CPI inflation, 2001–2021 |
| Real mortgage rate, percent per period | $r_m$ | 8.39 | 2.72 | Thirty-year mortgage rate less CPI inflation, 2001–2021 |
| Down payment, share of house value | $\phi$ | 0.20 | | Assumed |
| Holding cost, share of value per period | $\delta$ | 0.12 | 0.042 | Assumed |
| Moving cost, share of the unit's price | $\zeta$ | 0.08 | | Assumed |

*Notes.* One model period is three years. The annual column converts per-period values: $\beta^{1/3}$, $(1+r)^{1/3}-1$ and $1-(1-\delta)^{1/3}$; each row states its unit. Panel A: the loss is the equal-weighted sum of squared relative gaps between the simulated and data values of the five moments in Table 5. The adopted $\psi_0$, $\theta_B$, $\gamma$ and $\nu$ are the lowest-loss point of a 10-point scrambled Sobol scan (Table I.1), rounded ($\psi_0$ to three decimals, the others to two, as printed) and simulated again with 100,000 lives, at which the loss is 0.186. No standard errors are computed, and $\nu$ lies at the upper edge of the scan box, 0.60. Flow utility is $[\alpha c^{\nu}+\psi_o(1-\alpha)h^{\nu}]^{(1-\gamma)/\nu}/(1-\gamma)$ with tenure $o$ and $\psi_1=1$, so a renter needs $\psi_0^{-1/\nu}=4.07$ units of housing to obtain the services of one owned unit, and the elasticity of substitution between consumption and housing services is $1/(1-\nu)=2.5$. Panel B: $\alpha$ equals one minus the housing expenditure share only in the Cobb–Douglas limit with $\psi_0=1$; the median renter rent share is also a targeted moment. $\xi$ is the scale, in utility units, of the type-I extreme-value taste shocks on the two tenure options. Panel C: real rates are 2001–2021 averages of the nominal ten-year Treasury yield (3.09 percent) and thirty-year mortgage rate (4.85 percent) less average realized CPI inflation (2.13 percent), compounded to three years (series in Table A.1); $r_s$ applies to nonnegative and $r_m$ to negative financial assets. The holding cost $\delta$ is the owner's proportional cost of keeping the house, including depreciation. The moving cost is paid when the household changes dwelling size or tenure, in proportion to the price of the unit it leaves: the house value for an owner, the rent bill for a renter. The income, rent and house-price processes are estimated outside the model (Table E.1).

the PSID age composition. The model's move rate is the share of renter-to-renter transitions into a new unit. Table 5 gives the exact definitions.

We evaluate $L$ at 10 scrambled Sobol points in the box $\psi_0 \in [0.40, 0.90]$, $\theta_B \in [0.20, 3.00]$, $\gamma \in [1.10, 3.00]$ and $\nu \in [0.05, 0.60]$, each with 50,000 simulated lives, and adopt the point with the lowest loss, without local refinement. Its coordinates are rounded to the precision shown in Table 4,

**Table 5:** Targeted moments: data and model at the estimated parameters

| Moment | Data | Model | Gap (percent) |
|---|---|---|---|
| Ownership rate, all ages | 0.731 | 0.758 | +3.6 |
| Ownership rate, age 30 | 0.537 | 0.679 | +26.5 |
| Owners' median net worth to income | 2.707 | 3.186 | +17.7 |
| Renters' median rent to income | 0.206 | 0.242 | +17.6 |
| Renter move rate, per three years | 0.563 | 0.435 | −22.8 |
| Loss | | 0.186 | |

*Notes.* Data: PSID individuals aged 27–80 in owner- or renter-occupied housing, not only household heads, weighted by PSID sample weights; nominal money values enter only as ratios within an observation. Model: 100,000 simulated lives at the estimated parameters (Table 4). Ownership over all ages weights the 18 three-year periods equally; the other moments pool the household-periods of owners (net worth), renters (rent to income) and renters in consecutive periods (move rate). Ownership at 30: data for the three-year age bin starting at 30, model for the period at 30. Net worth includes home equity (model: financial assets plus house value); income is annual total family income (model: annual income). Rent to income: annual rent over income (model: the rent bill over income). Move rate: probability that a renter moves to a new rental unit within three years; in the data, the weighted share of consecutive-wave renter pairs in the rent-process sample (household heads) that report a move, $s = 0.424$ per two-year wave, converted as $1-(1-s)^{3/2}$; in the model, the share of renter-to-renter transitions into a new unit, which requires a change of dwelling size. Gap: model minus data in percent of data, from unrounded values. Loss: sum of the five squared relative gaps.

and the model is simulated again with 100,000 lives, at which the loss is 0.186. Table I.1 reports every point. Three features of the scan bound what the estimates can say. The loss is flat across very different parameter mixes: the second-best point, with $\psi_0 = 0.899$ and $\nu = 0.14$, has a loss of 0.216 against 0.187 at the adopted point, at the same number of lives. That point matches ownership at 30 more closely, 0.530 against 0.679, but puts aggregate ownership at 0.583, against 0.731 in the data. The adopted $\nu$ lies at the upper edge of the box and $\psi_0$ close to its lower edge, so the scan does not show whether the loss falls further beyond them. And we compute no standard errors and make no claim that a particular moment identifies a particular parameter.

**The estimates.** The renter housing weight is $\psi_0 = 0.431$. Because $\psi_0(1-\alpha)h^\nu = (1-\alpha)(\psi_0^{1/\nu}h)^\nu$, a rented dwelling of size $h$ yields the services of an owned dwelling of size $\psi_0^{1/\nu}h$, so a renter needs $\psi_0^{-1/\nu} = 4.07$ rooms of rented housing to obtain the services of one owned room. The weight collects everything that makes owning attractive beyond the costs, constraints and risks the model states: quality differences between the rented and owned stocks, the non-pecuniary value of owning, and the tax treatment the model omits. We therefore do not read its magnitude as a measured taste; Section 4.6 reports how it changes when the model is re-anchored without rental-rate risk.

The CES exponent $\nu = 0.60$ implies an elasticity of substitution between consumption and housing services of $1/(1-\nu) = 2.5$, so the two are strong substitutes; the life-cycle estimate of Li et al. (2016) is below one, the Cobb–Douglas value. The elasticity bears on the insurance margin in two

ways. At the allocation of the one-period problem, a higher $\nu$ lowers $\text{RRA}_c$ directly, because $\text{RRA}_c$ averages $\gamma$ and $1-\nu$, and through a lower consumption share $\pi$. And a high elasticity lets a renter facing a higher rent substitute toward consumption, which in this model means moving to a smaller unit. The adopted $\nu$, at the upper edge of the scan box, therefore works toward a smaller insurance premium; the scan does not show by how much. Relative risk aversion over the composite is $\gamma = 2.33$. Relative risk aversion in consumption varies with the allocation through $\pi$; at the one-period allocation of Section 3.6, consumption one minus the mean rent and a dwelling of one room, it is 2.12. The bequest weight is $\theta_B = 2.15$.

**Fit.** In the targeted fit (Table 5), aggregate ownership is close to the data, 0.758 against 0.731 (+3.6 percent). Ownership at 30 has the largest miss, 0.679 against 0.537 (+26.5 percent): young households own too early. Owners' net worth relative to income and renters' rent shares are both too high, by 17.7 and 17.6 percent. The renter move rate is too low, 0.435 against 0.563 (−22.8 percent); part of that gap is structural, since the data count moves between units of any size and the model has none between units of the same size.

The untargeted balance sheets and transitions in Table 6 miss in one correlated cluster; every ratio to income in it divides by annual income. Owners hold too little debt and too much equity: a share 0.393 of model owners have negative financial assets, against a share 0.693 of PSID owners with housing debt (the two are defined differently; Table 6), median home equity is 2.531 times income against 1.287, and the median model owner holds no positive financial assets. Renters hold far more wealth than in the data, a median net worth of 2.644 times income against 0.106. Model households run their wealth down late in life: median net worth at old age relative to middle age is 0.667 against 2.084 in the data. Tenure changes too often in both directions: per three years, the owner-to-renter rate is 0.104 against 0.064 in the data, and the renter-to-owner rate 0.378 against 0.234. Owners' median house value relative to income is closer, at +16.3 percent.

Several features of the model push in the directions of these misses, and we have not isolated any of them. Mortgage debt is one-period debt that can be repaid at no cost, so nothing ties an owner's debt to an amortization path. The bequest has no shifter, so every household, renters included, has a strong motive to end life with positive net worth, and saving earns a long-term rate. A mean-reverting price that each household faces, and a collateral constraint re-imposed at the current price, make the timing of purchases and sales valuable, in the direction of the excess churn. The renter-wealth miss bears on the paper's question: a renter with a buffer stock can absorb rent shocks by saving, so the insurance margin of Section 4 is measured in a model whose renters are better self-insured than PSID renters.

Figure 7 plots ownership by age at the targeted age and at four untargeted ones. The model starts at 0.606 in its first period, lies above the data at 30 and 39, matches them at 51 (+0.3 percent) and then flattens, falling below the data at 60 and 69 (−5.5 and −10.6 percent). The profile is too flat: ownership at 55 and over minus ownership at 35 and under is 0.123 in the model against 0.311

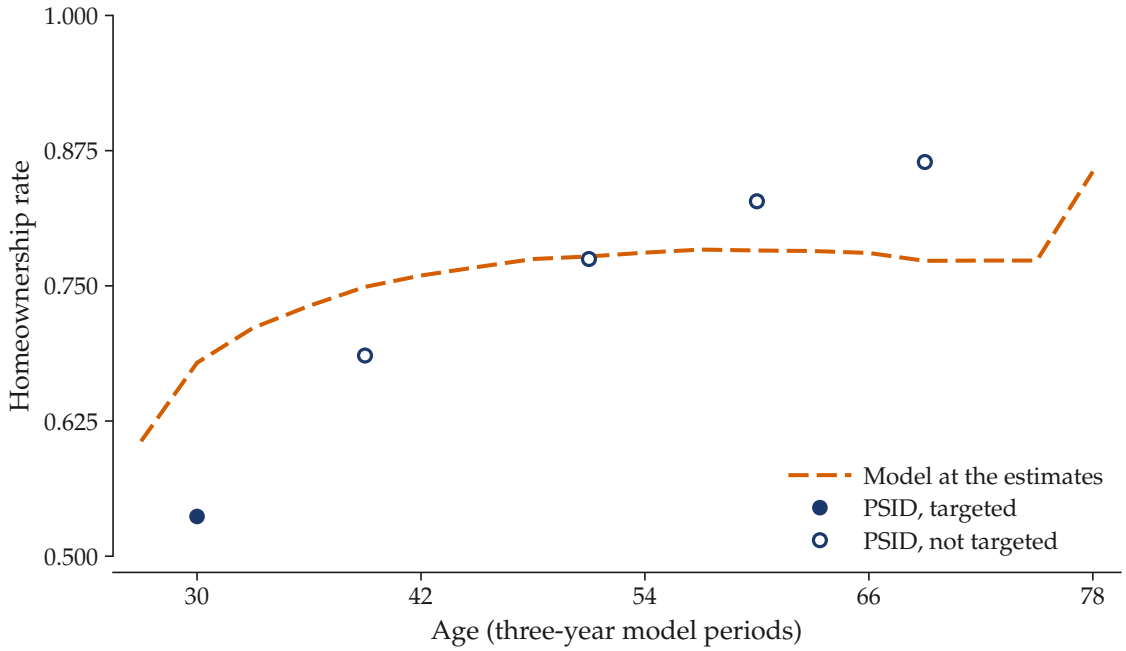
**Table 6:** Untargeted moments: data and model at the estimated parameters

| Moment | Data | Model | Gap (percent) |
|---|---|---|---|
| *A. Owners' balance sheets* | | | |
| Share with a mortgage[a] | 0.693 | 0.393 | −43.3 |
| Median home equity to income | 1.287 | 2.531 | +96.7 |
| Median non-housing wealth to income[a] | 0.867 | 0.000 | —[b] |
| Median house value to income | 2.478 | 2.882 | +16.3 |
| *B. Renters* | | | |
| Median net worth to income | 0.106 | 2.644 | —[b] |
| Standard deviation of log rent[a] | 0.541 | 0.643 | +18.9 |
| *C. Tenure transitions, per three years* | | | |
| Owner to renter | 0.064 | 0.104 | +61.5 |
| Renter to owner | 0.234 | 0.378 | +61.4 |
| *D. Life cycle* | | | |
| Ownership rate, age 39 | 0.686 | 0.749 | +9.3 |
| Ownership rate, age 51 | 0.775 | 0.777 | +0.3 |
| Ownership rate, age 60 | 0.828 | 0.783 | −5.5 |
| Ownership rate, age 69 | 0.864 | 0.773 | −10.6 |
| Ownership, age 55 and over minus age 35 and under | 0.311 | 0.123 | −60.6 |
| Median net worth, old relative to middle age | 2.084 | 0.667 | −68.0 |

*Notes.* No moment in this table enters the loss. Samples, weights and simulation as in Table 5; income is total family income, and model statistics are per three-year period. Mortgage: in the data, owners whose house value exceeds home equity by more than 1,000 (nominal) dollars; in the model, owners with negative financial assets. Home equity: reported equity (model: house value less mortgage debt). Non-housing wealth: net worth less home equity, which can be negative in the data; in the model, financial assets floored at zero. House value: self-reported (model: price times size). Renter net worth: model financial assets. Standard deviation of log rent: in the data, the square root of the variance of the hedonic rent residual in the rent-process sample (household heads) less its measurement-error variance (Table E.1); in the model, the standard deviation of the log rent bill across renters, which also varies with dwelling size, the permanent rent type and the market component, so the two statistics are not like for like. Transitions: share of owners (renters) who rent (own) at the next wave, a two-year rate $s$ converted to three years as $1-(1-s)^{3/2}$; in the model, the same share across consecutive three-year periods. Ownership at an age: the three-year age bin starting at that age (model: the period at that age). The age gradient is the ownership rate at ages 55 and over minus the rate at ages up to 35; it differs from the ownership rate at age 69 minus that at age 30 reported in Table I.1. Old relative to middle age: median net worth at ages 70–78 over median net worth at ages 50–58, each pooled over waves in nominal dollars. For these two rows the model uses the three-year periods that start within the stated ages. Gap: model minus data in percent of data, computed from unrounded values.
[a] Defined differently in the data and the model (definitions above), so the comparison is not like for like.
[b] Not shown: in percent of the data value the gap is uninformative, because the data median is close to zero (renter net worth) or the model value is zero by construction (non-housing wealth). The levels are compared directly.

in the data. In the last period, whose continuation value is the bequest, ownership jumps from 0.773 to 0.856. The bequest counts an owned house at its full price, without the holding or moving cost; no simulation isolates its role in the jump. The aggregate ownership moment weights this last period like any other.

**Figure 7:** Homeownership over the life cycle: model at the estimated parameters and PSID




*Notes.* Line: simulated ownership rate by age at the estimated parameters (100,000 lives), from the first period at 27 to the last at 78. Markers: weighted PSID ownership rate in the three-year age bin that starts at the age shown, individuals aged 27–80 in owner- or renter-occupied housing; filled if targeted (Table 5), hollow if not (Table 6). The last period's continuation value is the terminal bequest, which values an owned house at its full price.

# 4 Insurance, Lock-In, and Policy

## 4.1 How the Counterfactuals Are Built

Each counterfactual removes one or more sources of risk from the estimated economy and simulates the same households again. Preferences stay at their estimates, and households re-solve their problem in each world, so the responses of saving, housing and tenure choices are included. Rents, house prices and interest rates do not respond: the counterfactuals are partial equilibrium (Section 3.5). We write $\text{own}(S)$ for the ownership rate of world $S$, the share of simulated household-periods in which the household owns, and label a world by the set of rent components it removes, with $S = \varnothing$ the baseline, or, for the frozen-draw world defined below, by $F$. The same notation, and the welfare measure $\lambda_i(S)$ below, covers the worlds of Section 4.4. Every household lives from 27 to 81, so every age carries the same weight in $\text{own}(S)$.

The rent a household would pay combines the market component $g_t$, its idiosyncratic position $z_{it}$ and its permanent rent type $f_i^q$ (equation (11)). Four operations on this process define the worlds of Tables 7 and J.1.

*Market component removed (M).* The normalized market factor $e^{g_t} / E[e^{g_t}]$ of equation (11), which has mean one at every age, is set to one, so rents carry no market innovations.

*Within-tenancy component removed (I).* The normalized factor $e^z / E[e^z]$ is set to one, its mean under the stationary distribution of $z$, for every household at every age. The world removes the risk in $z$, within a tenancy and at re-entry, and also the dispersion of positions across households, including the positions households hold at 27. The permanent rent type is heterogeneity, not risk, and every world keeps it.

*Re-draw off (R).* On entering a new unit, $z$ evolves by the within-tenancy transition instead of being partly re-drawn, $\Pi_{\text{move}} := \Pi_{\text{stay}}$ in equation (14), both for renters who change units and for former owners who return to renting. A position then survives a move as it survives a stay: a favorable position is not forfeited at re-entry, and an unfavorable one is not partly reset.

*Each draw frozen (F).* Both transitions become the identity matrix, so each household keeps its age-27 draw of $z$ for life. That draw comes from the stationary distribution, so the dispersion of positions across households is the stationary one at every age; the market component is kept. The world removes all uncertainty about $z$, within a tenancy and at re-entry. It also removes mean reversion: within a tenancy the expected next position becomes $z_t$ instead of $\rho_q z_t$, where $\rho_q$ is the within-tenancy persistence per three-year period, so a position that would revert toward the mean with a half-life of about 11 years is held for life. Because a frozen draw travels with the household into any new unit, the world also switches off the re-draw.

Every removal leaves the unconditional mean rent unchanged. Worlds $M$ and $I$ replace a factor by its mean; worlds $R$ and $F$ keep the stationary distribution of $z$, which $\Pi_{\text{stay}}$ and $\Pi_{\text{move}}$ share and from which the frozen draw is taken. A counterfactual therefore does not change the average price of rented housing, and does not measure a price change as a risk effect. It changes the risk, the dispersion of positions and, through the transitions, the position a household expects given its current one: world $R$ removes the conditional-mean wedge (7) of a move, and world $F$ makes the expected next position equal to the current one. Mean preservation holds at the stationary distribution, not for the positions households actually hold, because households select into positions through their moving and tenure choices; in the baseline, renters are concentrated on favorable positions (Section 4.3).

Each world simulates 100,000 lives, allocated across permanent types in proportion to their population weights, using the same random numbers in every world. The draws behind initial conditions, income, house prices, the market component, the transitions of $z$ and the tenure taste shocks are identical across worlds. A household's rent path differs across worlds only through the components a world removes or the transitions it changes, and through the household's own choices (Appendix I), so every comparison is within a household. All differences are computed at the estimated parameters. They carry no sampling or simulation standard errors, their sensitivity to the solution grids has not been evaluated at these parameters, and they have not been computed at the other low-loss points of the parameter scan (Section 3.7).

We measure welfare by a consumption-equivalent variation (CEV) for each simulated household. Let $c_{it}^S$ be the non-housing consumption of household $i$ in period $t$ of world $S$, and $h_{i,t+1}^S$ and $o_{i,t+1}^S$ the dwelling it occupies and the tenure it holds during period $t$, the choices $h'$ and $o'$ of equation (19). Let $B_i^S$ be the bequest utility (9) of the assets, dwelling and tenure the household chooses in its last period. The CEV $\lambda_i(S)$ of moving from the baseline to world $S$ solves

$$\begin{aligned}\sum_{t=1}^{T} \beta^{t-1}\, u\big((1+\lambda_i(S))\, c_{it}^{\varnothing},\, h_{i,t+1}^{\varnothing},\, o_{i,t+1}^{\varnothing}\big) + \beta^T B_i^{\varnothing} \\ = \sum_{t=1}^{T} \beta^{t-1}\, u\big(c_{it}^{S},\, h_{i,t+1}^{S},\, o_{i,t+1}^{S}\big) + \beta^T B_i^{S},\end{aligned} \tag{25}$$

where $u$ is the flow utility (8), $\beta$ the discount factor per period and $T$ the number of three-year periods. A positive $\lambda_i(S)$ is a gain: baseline non-housing consumption would have to rise by the fraction $\lambda_i(S)$ in every period for the baseline life to be worth as much as the life in world $S$. Housing services and the bequest stay at their baseline values, and the realized taste shocks of the tenure choice (20) are not part of the comparison. We solve for $\lambda_i(S)$ by bisection between $-95$ and 5,000 percent (Appendix I).

The CEV is ex post. It compares two realized lives of the same household, paired by the common random numbers, so its mean over households is not the consumption equivalent of a change in expected utility. Means weight households by the population weights of their types, and percentiles are over simulated households, whose numbers are proportional to those weights. Quintiles sort households by realized mean non-housing consumption over the baseline life. Means by permanent rent type are over all households of the type, owners and renters.

### 4.2 Rental-Rate Risk Lowers Homeownership

Removing all three components of rental-rate risk raises the ownership rate from 0.758 to 0.791, by 3.29 percentage points (Table 7, Panel A). The insurance motive alone points the other way: removing the risk that ownership insures would lower the demand for ownership. The sign shows that other forces dominate; the decompositions below separate them.

The Shapley decomposition averages each component's marginal effect over the six orders in which the three components can be removed (Shorrocks, 2013). The within-tenancy (idiosyncratic) component $z$ accounts for 2.74 percentage points, the re-entry re-draw for 0.38 and the market component for 0.16. Removed one at a time, the three raise ownership by 2.97, 0.82 and 0.05 points (Table J.1). The components interact. Once $z$ is removed every position is the same and the re-draw has nothing to reset, so the re-draw's Shapley value is below its stand-alone effect. The market component matters more once $z$ is absent (0.31 points) than in the baseline (0.05). Each simulated household draws its own market history (Section 3.5), so the market component also spreads rents across households. Once $z$ is removed, that spread is the only dispersion of rents left besides the permanent type, and its larger effect there is consistent with the dispersion step below.

**Table 7:** Removing rental-rate risk: homeownership and welfare

| | Associated channel | Value |
|---|---|---|
| *A. Homeownership* | | |
| Baseline ownership rate | | 0.758 |
| Change when all rental-rate risk is removed (pp) | | +3.29 |
| Shapley contributions (pp) | | |
| Market component | | +0.16 |
| Within-tenancy component | | +2.74 |
| Re-entry re-draw | | +0.38 |
| Sequential decomposition (pp) | | |
| Freeze each draw $z$: uncertainty about $z$ and re-draw off; dispersion of positions and market component kept | Insurance (re-draw also off) | −0.36 |
| Then remove the dispersion of rent positions | Lock-in | +3.33 |
| Then remove the market component | | +0.31 |
| Decomposition through the re-draw (pp) | | |
| Switch off the re-draw only | Forfeiture at re-entry | +0.82 |
| Then remove the remaining rental-rate risk | | +2.47 |
| *B. Welfare: CEV of removing all rental-rate risk (percent)* | | |
| Mean | | +1.08 |
| Median | | −0.61 |
| Tenth percentile | | −9.38 |
| Ninetieth percentile | | +4.70 |
| Households at the lower bisection bound (percent) | | 0.58 |

*Notes.* Simulated model at the estimated parameters, 100,000 lives in each world, common random numbers; the worlds and the CEV are defined in Section 4.1. Ownership rate: share of simulated household-periods owning, ages 27 to 81; changes in percentage points (pp). Shapley contributions average each component's marginal effect over all orders of removal. Freezing each draw keeps each household's age-27 rent position $z$ for life ($\Pi_{\text{stay}} = \Pi_{\text{move}} =$ identity): uncertainty about $z$ and the re-draw are removed, while the dispersion of positions across households and the market component are kept, which is why the market component is removed in a separate, last step. The re-draw has no effect once the within-tenancy component is removed. Switching off the re-draw alone ($\Pi_{\text{move}} := \Pi_{\text{stay}}$) is the portable world of Table 10. Associated channel: the channel of which a step is the closest simulated counterpart (Section 4.2); no step removes one channel alone. Insurance: the value of removing uncertainty about $z$; freezing also removes the re-draw and makes positions permanent. Lock-in: the value of a favorable position. Forfeiture at re-entry: the part of it lost on entering a new unit; the re-draw also partly resets unfavorable positions. Each block sums to the total up to rounding. Panel B: CEV in percent of baseline non-housing consumption, per household along paired simulated paths (ex post); positive values are gains. Figure 9 reports the mean CEV by quintile of baseline lifetime consumption and by permanent rent type. Bisection bounds $[-95, 5{,}000]$ percent; the share at the upper bound is not recorded. Source: model simulations.

Two sequences of worlds separate the channels. The first splits the uncertainty about $z$ from the dispersion of positions:

$$\begin{aligned}\text{own}(\{M,I,R\}) - \text{own}(\emptyset) = &\underbrace{\text{own}(F) - \text{own}(\emptyset)}_{\text{freeze each draw}} + \underbrace{\text{own}(\{I\}) - \text{own}(F)}_{\text{remove the dispersion}} \\ &+ \underbrace{\text{own}(\{M,I,R\}) - \text{own}(\{I\})}_{\text{remove the market}},\end{aligned} \tag{26}$$

that is, $3.29 = -0.36 + 3.33 + 0.31$ percentage points, up to rounding. The first two steps sum to the effect of removing $z$ alone, 2.97 points. The identity holds by construction; its content lies in what each step removes.

The first step freezes each household's draw at its age-27 value. It removes the uncertainty about $z$ that ownership insures against, and ownership falls by 0.36 points. The step is the closest simulated counterpart of the insurance channel, through which rental-rate risk raises ownership. It changes two other things (Section 4.1): it switches off the re-draw, which on its own raises ownership (below), and it makes each position permanent, which raises the value of a favorable position and the cost of an unfavorable one. Starting instead from world $R$, where the re-draw is already off, freezing each draw lowers ownership from 0.766 to 0.754: the frozen-draw step nets the positive effect of switching off the re-draw against a larger negative effect of freezing positions. No simulated world removes the uncertainty and keeps the re-draw. Because the effects of the rent components do not add up, we cannot sign how the effect of such a world differs from $-0.36$ points.

The second step removes the dispersion of positions and raises ownership by 3.33 points. This is the lock-in channel: the value of holding a favorable rent position, which a household uses only while it rents. In world $F$ a position is permanent, known at 27 and carried into every unit, so nothing is lost by moving or buying. In every period in which it rents, a household on a favorable position pays less than one of the same permanent type at the mean position, and a household on an unfavorable position pays more. The step therefore measures how the dispersion of rents across households moves tenure, with positions valued as if they lasted for life. Its sign fits a threshold reading, which is intuition rather than a decomposition we compute. In world $I$, where every household holds the mean position, the ownership rate is 0.788: in most household-periods a household at the mean position owns, so the rent at which it would be indifferent between renting and owning lies below the mean rent. Spreading positions then pulls into renting more households that draw favorable positions than it pushes into owning households that draw unfavorable ones, and in the baseline renter-periods are indeed concentrated on favorable positions (Section 4.3).

Neither world of the dispersion step has transitions in $z$. In world $F$ each household draws its position at 27 and keeps it for life, so the step removes a dispersion of rents known at entry, with variance $\sigma_z^2$, and it would be the same if that dispersion were a permanent type like the one every world keeps. We count it as part of rental-rate risk because in the estimated process a position after 27 combines the entry draw with the innovations since, and the innovations account for a growing share of its dispersion with age; the step values that dispersion as if each position lasted for life. How much of the dispersion of rents the process assigns to such innovations rather than to the permanent type is the split that Section 2.5 shows is weakly identified. Given $\sigma_z^2$, the persistence of $z$ and the re-draw correlation enter the decomposition only through the first step and through world $R$ below. The last step removes the market component, which adds 0.31 points once $z$ is gone; it too removes a dispersion of rents across households (above).

The second sequence isolates the re-draw. It switches off the re-draw first and then removes the rest of rental-rate risk: $3.29 = 0.82 + 2.47$ percentage points up to rounding, where the second term is $\text{own}(\{M, I, R\}) - \text{own}(\{R\})$. Switching off the re-draw alone raises ownership by 0.82 points, 25 percent of the total. The re-draw cuts both ways: it takes part of a favorable position from a household that enters a new unit, which is forfeiture at re-entry, and it partly resets an unfavorable one. World $R$ removes both, and Section 4.3 shows how each appears in simulated mobility. This share is the most the re-draw accounts for in any order of removal: removed after $z$, it has no effect, and its Shapley value is 0.38 points. World $R$ is the counterpart of the forfeiture wedge (23) of Section 3.6; the one-period problem is silent on the dispersion step, and its ranking of the wedge over the insurance premium concerns mobility, not tenure.

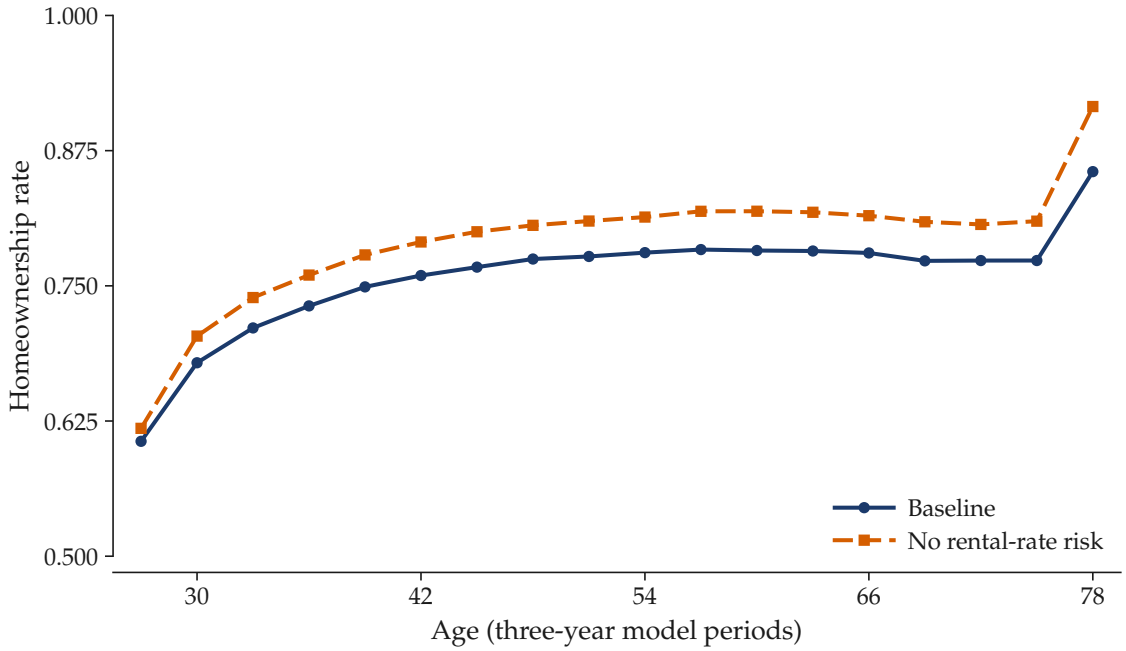
With the re-draw already off, removing the remaining rental-rate risk still raises ownership by 2.47 points, 75 percent of the total. That remainder is the sum of three steps: freezing positions from world $R$, removing their dispersion and removing the market component. At most a quarter of the headline therefore runs through the re-draw, which carries forfeiture at re-entry, a sub-channel of lock-in. The rest runs through the uncertainty, dispersion and market steps, of which the dispersion step is the largest.

Figure 8 shows that the gap forms early and persists. Without rental-rate risk, ownership is higher by 1.20 points at 27, before any innovation after entry; that gap reflects the positions households hold on entry, which world $I$ removes, and the risk they anticipate. The gap is 2.47 points at 30, larger at every later age and 3.64 points in the penultimate period, so the ownership that rental-rate risk removes early in life is not made up later. The gap of 6.02 points in the last period comes with the jump in ownership at that age in both economies (Section 3.7), and that period enters $\text{own}(\cdot)$ with the same weight as every other period.

Four features of the estimated economy bear on these magnitudes. First, the dispersion step scales with the stationary variance of $z$. The adopted point lies near the upper edge of the ridge, toward which the permanent type shrinks and $\sigma_z^2$ grows (Appendix C): $\sigma_z^2$ is 0.261 there, against 0.137 at the nine-moment fit, with bootstrap interval $[0.074, 0.304]$. A point with more permanent dispersion leaves less dispersion in $z$ to remove, and the decomposition has not been computed at such a point. The re-draw step depends on $\rho_{\text{move}}$, which Section 2.6 identifies only conditional on the point on the ridge.

Second, most features of the model whose effect we can sign plausibly make the insurance step smaller than it would be for PSID renters. The median model renter holds a net worth of 2.644 times annual income, against 0.106 in the data, and can absorb rent shocks by saving (Section 3.7). The adopted $\nu$, at the upper edge of the scan box, makes risk aversion in consumption lower than a smaller $\nu$ would. Rent variation that Section 2.5 books as measurement error, part of which may be transitory rent that households pay, does not enter. And rents and house prices do not co-move, so the insurance channel excludes any hedge through the value of the house (Section 3.5). The absence of a consumption floor works the other way. Third, tenure in the estimated economy is

**Figure 8:** Homeownership by age with and without rental-rate risk




*Notes.* Simulated share of households that own, by age, in the estimated economy (baseline) and with all three components of rental-rate risk removed (Table 7). Each economy simulates 100,000 lives with common random numbers, preferences at their estimates and the permanent rent type kept. Ages are the first year of each three-year period, from 27 to 78. The jump at the last age appears in both economies; that period's continuation value is the terminal bequest, which values an owned house at its full price, without the holding or moving cost.

governed by house-price risk (Section 4.4), and the rent effects are measured against it. Fourth, these numbers come from one point of a flat loss: at the second-ranked point of the scan, aggregate ownership is 0.583 rather than 0.758 (Section 3.7), and none of the counterfactuals has been computed there.

### 4.3 Lock-In in Simulated Renter Behavior

Table 8 follows simulated renters over one three-year period, grouped by their current position $z$ on the grid of the within-tenancy process (Section 3.2): favorable positions are the two lowest nodes, the middle position is the central node and unfavorable positions are the two highest nodes. The rates are model output, not data. The favorable and unfavorable groups have equal mass in the stationary distribution of $z$, yet favorable positions account for 40.1 percent of baseline renter-periods and unfavorable positions for 23.4 percent: renters are concentrated on favorable positions.

Renters on unfavorable positions buy at a rate of 0.51 per period, against 0.29 on favorable positions, 1.75 times as often. This gradient is the behavioral counterpart of lock-in as defined in Section 4.2: a renter on a favorable position pays less for its unit, gains less from buying and stays a renter longer. A gradient of this kind requires only that current rents differ across renters; it does

**Table 8:** Simulated renter transitions by current rent position, baseline and portable rent positions

| | Renter-periods (percent) | | Buy rate | | Move rate | |
|---|---|---|---|---|---|---|
| Current rent position | Baseline | Portable | Baseline | Portable | Baseline | Portable |
| Favorable (low $z$) | 40.1 | 42.3 | 0.29 | 0.30 | 0.16 | 0.20 |
| Middle | 36.5 | 37.7 | 0.38 | 0.39 | 0.31 | 0.19 |
| Unfavorable (high $z$) | 23.4 | 20.0 | 0.51 | 0.52 | 0.39 | 0.15 |
| All groups (count) | 479,241 | 463,982 | | | | |

*Notes.* Model output, not data: simulated households that rent at the start of a three-year period, all ages, 100,000 lives at the estimated parameters. Groups are defined by the household's current idiosyncratic rent position $z$ (market component and permanent rent type divided out) on the five-node grid of the within-tenancy process: favorable is the two lowest nodes, middle the central node and unfavorable the two highest nodes. The groups are grid-node groups, not terciles. Buy rate: probability that a household that rented in the previous period chooses to own in the current one. Move rate: probability that it chooses to rent again in the current period and enters a new unit, which in the model requires a change of dwelling size; buyers are in the denominator. Both rates are per three-year period, averaged within each combination of permanent income and rent type and then weighted by population type weights. Renter-periods: each group's share of all simulated renter-periods; the last row gives the total count. Portable: on entering a new unit $z$ evolves by the within-tenancy transition ($\Pi_{\text{move}} := \Pi_{\text{stay}}$, so the correlation 0.446 becomes 0.830 per three-year period), as in Table 10; the move rates under portability are the mobility response to that policy. Source: model simulations.

not show that a position is at stake, which the move rates below do. The gradient is a level effect of the current rent, not a response to forfeiture at re-entry. When positions become portable, so that entering a new unit forfeits nothing, the buy rates barely change: they are 0.30, 0.39 and 0.52, a ratio of 1.74.

Forfeiture shows in mobility. The move rate is the probability that a renter remains a renter and enters a new unit, which in the model requires a change of dwelling size (equation (14)). Buyers stay in the denominator, so these rates are not the targeted move rate of Section 3.7, which conditions on remaining a renter. In the baseline the move rate rises with the position, from 0.16 on favorable positions to 0.31 on the middle position and 0.39 on unfavorable ones. Entering a new unit pulls the expected position further toward the mean than staying does (equation (7)). That costs a household on a favorable position part of its position and gives one on an unfavorable position a partial reset. Under portability the gradient reverses. Renters on favorable positions move more, at 0.20, because a move no longer costs them their position. Renters on unfavorable positions move much less, at 0.15, because a move no longer resets theirs, and the rate on the middle position falls to 0.19. The first change is the release from lock-in on the mobility margin; the second removes the escape that the re-draw offered households on unfavorable positions.

Portability raises ownership (Section 4.7) while buy rates at each position stay almost unchanged. What changes is the composition of renters. The fall in renter-periods on unfavorable positions, from 112,253 to 92,648, exceeds the fall in all renter-periods, from 479,241 to 463,982. These counts do not separate the rise in ownership into more purchases and fewer returns to renting.

### 4.4 Rental-Rate Risk among the Household's Risks

Households in the model face three risks: rental-rate risk (all three of its components), house-price risk and income risk. House-price risk is removed by fixing the price at its stationary mean level. Income risk is removed by fixing the persistent and transitory components at their stationary mean level, which keeps the permanent income types and the age profile. Because the persistent component starts at 27 with more than its stationary dispersion (Section 3.2), this removal also lowers mean income slightly early in life, unlike the removals of rent components. All eight combinations are simulated with common random numbers (Table J.2), and Table 9 reports Shapley decompositions of the ownership rate and of the mean CEV.

The three risks act on different margins. In the Shapley decomposition, house-price risk governs tenure: it accounts for 19.80 of the 23.72 percentage points by which ownership rises when all three risks are removed, and removing it alone raises ownership to 0.943. House-price risk here is the dispersion of owners' self-reported house values net of national year effects (Appendix F), independent of rents, so the ranking is conditional on that process. Income risk acts on tenure in the opposite direction: without it, ownership in the estimated economy changes by $-2.66$ points (Table 9, Panel B). Its Shapley contribution is smaller, $-0.93$ points, because its effect changes sign across starting worlds: with house-price and rental-rate risk already removed, taking away income risk changes ownership by $+1.14$ points (Table J.2). We have not isolated why income risk raises ownership in the estimated economy.

Income risk governs welfare. Its Shapley contribution to the mean CEV, 58.75 percent, exceeds the 43.43 percent of removing all three risks, because house-price risk's contribution is negative, $-15.95$ percent. Removing income risk alone has a mean CEV of 61.63 percent. That world removes all persistent and transitory earnings risk, with the income process of working ages applied at every age to 81 and with no consumption floor or public transfer (Section 3.5), and we report it for scale. Removing house-price risk alone has a mean CEV of $-13.90$ percent. With a stationary price process, dispersion in the price a household faces is also an opportunity to buy when the price is low, which the removal takes away; we have not isolated this channel. For the worlds of this subsection only the mean CEV is recorded, with no median, percentiles or bound shares, and Section 4.5 shows that for rental-rate risk the mean and the median differ in sign.

Rental-rate risk is secondary on both aggregate margins. House-price risk's contribution to ownership is 4.1 times that of rental-rate risk, 4.85 points, and rental-rate risk contributes 0.63 percent to the mean CEV. Its welfare effect is distributional, with gains and losses across households that are large relative to its mean (Section 4.5). The distribution of welfare effects is recorded only for rental-rate risk, so across the three risks we compare means.

The size and form of the three exposures fit this division. Take one-standard-deviation innovations over a three-year period, measured in units of annual income. An innovation to $z$ changes the annual rent of a renter at the median rent-to-income ratio by about 6 percent of income, a flow

**Table 9:** Rental-rate risk among the household's risks: Shapley decomposition of ownership and welfare

| | Ownership (pp) | Mean CEV (percent) |
|---|---|---|
| *A. Shapley contributions* | | |
| House-price risk | +19.80 | −15.95 |
| Rental-rate risk | +4.85 | +0.63 |
| Income risk | −0.93 | +58.75 |
| Total: all three removed (sum of contributions) | +23.72 | +43.43 |
| *B. Removing one risk alone* | | |
| House-price risk | +18.56 | −13.90 |
| Rental-rate risk | +3.29 | +1.08 |
| Income risk | −2.66 | +61.63 |
| *C. Removing rental-rate risk, by starting world (Shapley weight)* | | |
| All risks present ($\frac{1}{3}$) | +3.29 | +1.08 |
| House-price risk already removed ($\frac{1}{6}$) | +4.01 | +0.81 |
| Income risk already removed ($\frac{1}{6}$) | +4.98 | −0.86 |
| Both already removed ($\frac{1}{3}$) | +6.78 | +0.84 |
| Weighted sum: Shapley value | +4.85 | +0.63 |

*Notes.* Simulated model economy at the estimated parameters, 100,000 lives in each world, common random numbers. A world removes some of three risks: house-price risk (price set to its stationary mean level), rental-rate risk (all three components, as in Table 7) and income risk (persistent and transitory innovations removed; permanent income types and the age profile kept). Table J.2 reports all eight worlds. Ownership: change in the ownership rate (share of simulated household-periods owning, all ages) relative to the baseline, in percentage points (pp). Mean CEV: mean over households of the consumption-equivalent variation of moving from the baseline to the world, defined as in Table 7; positive values are gains. Panel A: each risk's marginal effect averaged over all orders of removal; the three contributions sum to the effect of removing all three, up to rounding. Panel B: the world with only that risk removed; removing rental-rate risk alone is the total effect in Table 7. Panel C: the effect of removing rental-rate risk from each starting world, that is, the world without rental-rate risk minus the starting world, for ownership and for the mean CEV against the baseline; weighting the four rows by the Shapley weights in parentheses gives the Shapley value of Panel A. Rent, house prices and income follow independent processes, so the model has no rent–price co-movement. Source: model simulations.

that decays with a half-life of about 11 years within a tenancy. An innovation to the house price revalues the house of an owner at the median value-to-income ratio by about 81 percent of income, and the price deviation has a half-life of about 5 years. That is a change in the value of a stock, which moves the owner's net worth and borrowing capacity, not a recurring cost. An innovation to persistent income changes income by about 33 percent, a flow that reaches consumption in every period. Buying exchanges a recurring flow exposure to rents for a stock exposure to the house price, taken on at the tenure decision.

Rental-rate risk's Shapley contribution to ownership, 4.85 points, exceeds the 3.29 points of Section 4.2. The two numbers answer different questions. The headline removes rental-rate risk from the estimated economy, with the other two risks present. The Shapley value averages the effect of

removing it over four starting worlds (Table 9, Panel C): 3.29 points with the other risks present, 4.01 without house-price risk, 4.98 without income risk and 6.78 without either, with weights $\frac{1}{3}$, $\frac{1}{6}$, $\frac{1}{6}$ and $\frac{1}{3}$. Rental-rate risk moves ownership more when the other risks are absent, and the largest effect is measured where ownership starts at 0.927 and rises to 0.995. The headline is the number that describes the estimated economy.

The same construction for the mean CEV gives $+1.08$, $+0.81$, $-0.86$ and $+0.84$ percent. With income risk already removed, removing rental-rate risk lowers the mean CEV. The effect of rental-rate risk on the mean CEV is no larger in absolute value in any starting world than in the estimated economy, and its sign depends on the other risks the household faces; we have not isolated why it turns negative.

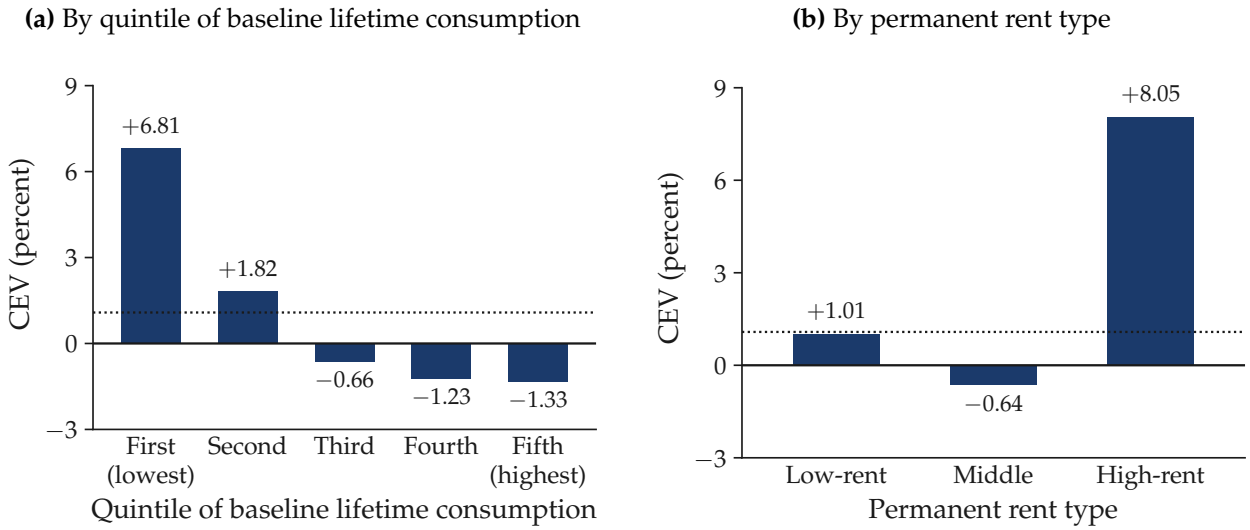
### 4.5 Welfare

Removing rental-rate risk is worth 1.08 percent of baseline non-housing consumption on average, but the median household loses: its CEV is $-0.61$ percent (Table 7, Panel B). The distribution is wide, from $-9.38$ percent at the tenth percentile to $+4.70$ percent at the ninetieth. At least half of households lose from the removal, and the gains of the others are large enough to make the mean positive.

The CEV compares realized lives, and two forces set its sign for a household. Removing the risk sets every household's position to the mean, so a household whose baseline positions were favorable would pay more for the same rented dwelling, and one whose positions were unfavorable would pay less; which of the two a household is depends on its luck in the baseline. And because households choose tenure, dwelling size and moves after observing their positions, the dispersion of positions also has an option value: a household on a favorable position can rent cheaply, and one on an unfavorable position can buy or move. Removing the dispersion takes that option away. The same force may explain why removing rental-rate risk lowers the mean CEV once income risk is gone (Section 4.4). We have isolated neither force.

Figure 9 breaks the CEV down by quintile of baseline lifetime consumption (panel (a)) and by permanent rent type (panel (b)). The first quintile gains 6.81 percent and the second 1.82 percent; the third, fourth and fifth quintiles have CEVs of $-0.66$, $-1.23$ and $-1.33$ percent. The sorting variable is itself an outcome of the baseline rent draws. A household with an unfavorable rent history consumes less in the baseline and so tends to fall into a lower quintile, and it is also a household that removing the risk helps. The quintile gradient therefore mixes households with little ability to absorb rent shocks and households that were unlucky in the baseline, and an ex post CEV does not separate the two. The simulations record no split by a resource measure fixed at entry, such as the permanent income type or wealth at 27, which would separate them; the only split by a characteristic fixed at entry is by permanent rent type, below.

**Figure 9:** Welfare effect of removing rental-rate risk, by quintile of baseline lifetime consumption and by permanent rent type




*Notes.* Consumption-equivalent variation (CEV) of moving from the estimated economy to the economy without rental-rate risk, in percent (equation (25)); positive values are gains. Means over households with population weights. Panel (a): quintiles, first (lowest) to fifth (highest), of each household's mean non-housing consumption over its baseline life. Panel (b): permanent rent type, with rent multipliers 0.78, 0.99 and 1.25 and population shares 16.7, 66.7 and 16.7 percent for the low-rent, middle and high-rent types, over all households, owners and renters. Dotted line: the mean CEV, +1.08 percent; the median is −0.61 percent.

By permanent rent type (panel (b)), which is the same in both economies, households of the high-rent type, whose rent multiplier is 1.25, gain 8.05 percent, and those of the low-rent type (multiplier 0.78) gain 1.01 percent. The middle type (multiplier 0.99), 66.7 percent of households, has a CEV of −0.64 percent. The pattern is not monotone in the rent level. The mean is carried by the high-rent type, 16.7 percent of households: the population-weighted mean over the other two types is negative. Rent is proportional to $e^{f^q} e^z$, so the same innovation to $z$ moves the rent bill of a high-rent household by more, which fits the larger gain of that type; we have not isolated why the middle type loses.

The mean is sensitive to the tails of the distribution. The left side of equation (25) is bounded in $\lambda$: with $\gamma$ above one, flow utility stays below zero however large consumption becomes and stays finite as consumption goes to zero, because housing services are not scaled; the bequest is not scaled either. For some households, therefore, no $\lambda$ solves the equation, and for others the solution lies outside the bisection bracket; both are recorded at an end of the bracket. A share of 0.58 percent of households sits at the lower end; for them any solution lies below it, so the recorded value understates the loss. The share at the upper end of the bracket, a CEV of 5,000 percent, is not recorded. A small mass there would raise the mean, and by more the mean of any group in which it is concentrated; the simulations do not record which groups those are. We therefore report the median and the percentiles alongside the mean.

### 4.6 A Re-Anchored Benchmark

With preferences held at their estimates, the economy without rental-rate risk has more owners than the baseline (Section 4.2). A model that omitted rental-rate risk and matched the same ownership rate would therefore attribute part of the tenure split to different preferences. We compute that attribution for the renter housing weight $\psi_0$: the value at which the economy without rental-rate risk reproduces the baseline ownership rate, with the other parameters at their estimates. The target is the model's baseline rate, 0.758, not the data rate, 0.731, and the exercise re-anchors one parameter on one moment rather than re-estimating the model. The weight rises from 0.431 to 0.455: without rental-rate risk, renting must be made more attractive to keep the same share of households renting. On the aggregate tenure margin, rental-rate risk acts like a stronger taste for renting, and a model that omitted it would read that part of the tenure split as a preference for renting.

### 4.7 Portable Rent Positions

The tenant-protection experiment makes rent positions portable. When a household enters a new rental unit, its position evolves by the within-tenancy transition instead of being partly re-drawn, $\Pi_{\text{move}} := \Pi_{\text{stay}}$, so the correlation of $z$ across a move rises from 0.446 to 0.830 per three-year period. The rule applies to renters who change units and to former owners who return to renting, and risk within a tenancy is unchanged. The experiment is world $R$ of Section 4.1.

The closest real-world analog is an assignable lease, under which a tenant who leaves a unit at a below-market rent can transfer the lease and capture its value instead of forfeiting it. The analogy is loose: an assignable lease passes the terms of the old unit to that unit's next tenant, while in the model the position travels with the household into its next unit, and we know of no instrument that does that. Nor is the experiment security of tenure, which protects a sitting tenant's occupancy: no household in the model is forced out of its unit, so a sitting tenant already keeps its position. The rule is also two-sided. A household on an unfavorable position carries it into its next unit and loses the partial reset that the re-draw gives (Section 4.3). A protection that a tenant could invoke only when it helps, keeping a favorable position and taking the re-draw otherwise, would instead add an option, and at the model's fixed prices it could not lower any household's expected lifetime utility at entry; we have not simulated it. The experiment is a stylized tenant protection whose effects measure what the re-draw costs and gives households.

Portability raises the ownership rate by 0.82 percentage points, from 0.758 to 0.766, and by 1.14 points at age 30, from 0.679 to 0.690 (Table 10). That is 25 percent of the ownership effect of removing all rental-rate risk, obtained with within-tenancy risk left in place. Portability is worth 0.71 percent of baseline non-housing consumption on average and 2.03 percent in the first quintile of baseline lifetime consumption. These two statistics are the only welfare results recorded for this world, with no median, percentiles or bound shares. Behind them (Section 4.3), renters on favorable positions move more, renters on unfavorable positions move much less, and buy rates

**Table 10:** Portable rent positions: homeownership and welfare (partial equilibrium)

| | Baseline | Portable | Change |
|---|---|---|---|
| *A. Ownership rate (change in pp)* | | | |
| All ages | 0.758 | 0.766 | +0.82 |
| At age 30 | 0.679 | 0.690 | +1.14 |
| *B. Welfare: CEV of portability (percent)* | | | |
| Mean | | | +0.71 |
| First (lowest) quintile of baseline lifetime consumption | | | +2.03 |

*Notes.* Portable rent positions: on entering a new rental unit the household's idiosyncratic rent position evolves by the within-tenancy transition instead of being partly re-drawn ($\Pi_{\text{move}} := \Pi_{\text{stay}}$; the correlation 0.446 becomes 0.830 per three-year period). The rule applies to renters who change units and to former owners who return to renting; within-tenancy risk is unchanged. It is the re-draw-off world of Table 7. Partial equilibrium: rents, house prices and interest rates are held fixed, so landlords do not price portability. Panel A: share of simulated household-periods owning, over all ages or at age 30; changes in percentage points (pp). Panel B: consumption-equivalent variation of moving from the baseline to portable positions, defined as in Table 7; the quintile sorts households by realized mean baseline non-housing consumption. Renter mobility by current rent position, baseline and portable, is in Table 8. Estimated parameters, 100,000 lives, common random numbers. Source: model simulations.

at a given position barely change. In the model, portability frees renters on favorable positions to move and raises ownership at the same time. In this sense tenant protection of this form and homeownership are complementary.

The experiment is partial equilibrium: rents, house prices and interest rates are held fixed, so landlords do not respond. A portable position obliges the landlord of the new unit to accept the terms the tenant brings, and under portability the renters who move are more often those on favorable positions (Table 8). Landlords would price that obligation into rents at entry, and evidence from rent control shows that landlords also respond to tenant protection on the supply margin (Diamond et al., 2019). Both responses would tend to raise the rents households face at entry and to lower the welfare gains reported here. How the tenure response would change depends on how landlords set rents at entry, which the model does not determine. The experiment measures the household-side mechanism, not the incidence of a policy at scale.

## 5 Conclusion

Renting leaves a household exposed to future rents; owning hedges that exposure at the cost of house-price risk. This paper shows that when the rent a household pays is persistent, rental-rate risk does more than create demand for that hedge. It also creates rental positions worth keeping: a favorable rent draw is partly lost by moving and entirely given up by buying. Whether rental-rate risk pushes households toward ownership therefore depends on the rent process they face, and we measure that process at the level at which it is borne.

A household panel alone cannot separate persistent rent shocks from permanent differences in rent levels over short renter spells. Combining the PSID with external rent indexes resolves this. Long regional CPI rent series identify a highly persistent market component with small innovations; absorbing region-common movements isolates a large and persistent within-tenancy component; and new-tenant and sitting-tenant indexes, together with rent changes of movers and stayers, show that re-entry re-draws a household's rent position rather than shifting its level predictably. Because measurement error accounts for about 68 percent of the variance of reported within-unit rent changes, risk read directly off observed rent changes overstates its size and is silent on its persistence.

In the estimated life-cycle model, rental-rate risk lowers homeownership: removing it raises the ownership rate by 3.29 percentage points. The insurance motive is present—removing uncertainty while keeping the dispersion of rent positions lowers ownership by 0.36 points—but it is dominated by lock-in, visible in renters on unfavorable draws buying at a rate of 0.51 against 0.29 on favorable ones. The welfare cost of the risk is moderate on average, 1.08 percent of lifetime consumption, and concentrated at the bottom, at 6.8 percent for the poorest quintile. A tenant protection that lets renters keep their rent position when they move raises homeownership by 0.82 points and is worth 0.71 percent of consumption on average and 2.03 percent to the poorest quintile. In this economy, tenant protection and homeownership are complementary forms of insurance against rental-rate risk.

These policy results are partial equilibrium. They hold rents, house prices, and the housing stock fixed, while evidence from rent control shows that landlords respond to tenant protection on the supply margin (Diamond et al., 2019); the experiment isolates the household-side mechanism, not the incidence of a policy at scale. Three extensions follow. The restricted-geocode PSID, for which our application is pending, would replace the four census regions with metropolitan areas and sharpen the separation of market-level from unit-level rent risk. A mortgage amortization schedule would address the model's main systematic misfit, owners who deleverage too fast. And endogenizing rents and house prices would give the policy experiments the general-equilibrium content the present exercise leaves aside. The central point does not depend on these extensions: the effect of rental-rate risk on tenure choice turns on how persistent a household's rent is and how much of it survives a move, and both can be measured.

## References


**Adams, Brian, Lara Loewenstein, Hugh Montag, and Randal Verbrugge**, "Disentangling Rent Index Differences: Data, Methods, and Scope," *American Economic Review: Insights*, 2024, *6* (2), 230–245.

**Altonji, Joseph G. and Lewis M. Segal**, "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business & Economic Statistics*, 1996, *14* (3), 353–366.

**Ambrose, Brent W., N. Edward Coulson, and Jiro Yoshida**, "The Repeat Rent Index," *The Review of Economics and Statistics*, 2015, *97* (5), 939–950.

**Amior, Michael and Jonathan Halket**, "Do households use home-ownership to insure themselves? Evidence across U.S. cities," *Quantitative Economics*, 2014, *5* (3), 631–674.

**Andrews, Donald W. K.**, "Exactly Median-Unbiased Estimation of First Order Autoregressive/Unit Root Models," *Econometrica*, 1993, *61* (1), 139–165.

_ , "Estimation When a Parameter Is on a Boundary," *Econometrica*, 1999, *67* (6), 1341–1383.

_ **and Hong-Yuan Chen**, "Approximately Median-Unbiased Estimation of Autoregressive Models," *Journal of Business & Economic Statistics*, 1994, *12* (2), 187–204.

**Bajari, Patrick, Phoebe Chan, Dirk Krueger, and Daniel Miller**, "A Dynamic Model of Housing Demand: Estimation and Policy Implications," *International Economic Review*, 2013, *54* (2), 409–442.

**Blundell, Richard, Luigi Pistaferri, and Ian Preston**, "Consumption Inequality and Partial Insurance," *American Economic Review*, December 2008, *98* (5), 1887–1921.

**Cocco, João F.**, "Portfolio Choice in the Presence of Housing," *The Review of Financial Studies*, 2005, *18* (2), 535–567.

**De Nardi, Mariacristina**, "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 2004, *71* (3), 743–768.

**Deaton, Angus and Christina Paxson**, "Intertemporal Choice and Inequality," *Journal of Political Economy*, 1994, *102* (3), 437–467.

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

**Dickey, David A. and Wayne A. Fuller**, "Distribution of the Estimators for Autoregressive Time Series with a Unit Root," *Journal of the American Statistical Association*, 1979, *74* (366), 427–431.

**Fella, Giulio, Giovanni Gallipoli, and Jutong Pan**, "Markov-Chain Approximations for Life-Cycle Models," *Review of Economic Dynamics*, 2019, *34*, 183–201.

**Genesove, David**, "The Nominal Rigidity of Apartment Rents," *The Review of Economics and Statistics*, 2003, *85* (4), 844–853.

**Halket, Jonathan and Santhanagopalan Vasudev**, "Saving up or settling down: Home ownership over the life cycle," *Review of Economic Dynamics*, April 2014, *17* (2), 345–366.

**Hall, Peter and Joel L. Horowitz**, "Bootstrap Critical Values for Tests Based on Generalized-Method-of-Moments Estimators," *Econometrica*, 1996, *64* (4), 891–916.

**Iskhakov, Fedor, Thomas H. Jørgensen, John Rust, and Bertel Schjerning**, "The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks," *Quantitative Economics*, 2017, *8* (2), 317–365.

**Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante**, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 2020, *128* (9), 3285–3345.

**Kopecky, Karen A. and Richard M. H. Suen**, "Finite State Markov-Chain Approximations to Highly Persistent Processes," *Review of Economic Dynamics*, 2010, *13* (3), 701–714.

**Lee, David S.**, "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects," *Review of Economic Studies*, 2009, *76* (3), 1071–1102.

**Li, Wenli, Haiyong Liu, Fang Yang, and Rui Yao**, "Housing Over Time and Over the Life Cycle: A Structural Estimation," *International Economic Review*, 2016, *57* (4), 1237–1260.

**Meghir, Costas and Luigi Pistaferri**, "Income Variance Dynamics and Heterogeneity," *Econometrica*, 2004, *72* (1), 1–32.

**Rampini, Adriano A. and S. Viswanathan**, "Household Risk Management," Working Paper 22293, National Bureau of Economic Research May 2016.

**Rouwenhorst, K. Geert**, "Asset Pricing Implications of Equilibrium Business Cycle Models," in Thomas F. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press, 1995, pp. 294–330.

**Rust, John**, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55* (5), 999–1033.

**Shorrocks, Anthony F.**, "Decomposition Procedures for Distributional Analysis: A Unified Framework Based on the Shapley Value," *Journal of Economic Inequality*, 2013, *11* (1), 99–126.

**Sinai, Todd and Nicholas S. Souleles**, "Owner-Occupied Housing as a Hedge Against Rent Risk," *The Quarterly Journal of Economics*, May 2005, *120* (2), 763–789.

**Sommer, Kamila and Paul Sullivan**, "Implications of U.S. Tax Policy for House Prices, Rents, and Homeownership," *American Economic Review*, 2018, *108* (2), 241–274.

**Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron**, "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, 2004, *51* (3), 609–633.

# A Data and Samples

## A.1 Samples

All household data come from the PSID family files, 11 biennial waves from 2001 to 2021 (data-center extract J346293). The rent, income and house-price processes are estimated on one panel of household heads aged 27–80 who own or rent their home and have a positive PSID core/immigrant individual longitudinal weight: 53,432 head-wave observations from 9,841 households. Each observation carries the head's weight, normalized to mean one, in the hedonic regression and in every moment; a pair of observations carries the weight of its later wave. In the empirical work the unit of time is the wave, two years. The income process uses the heads of this panel who work (Appendix E) and the house-price process its owners (Appendix F).

**Renters.** A head-wave enters the renter sample if the family rents and reports a monthly rent from \$1 to \$99,996 in current dollars (higher values are the PSID's top code and its codes for missing answers), a positive room count below a cap that removes extreme counts, and a positive family income below the top code, since log income is a regressor. Rents are not trimmed within these bounds, so very low reported rents are kept. The extract does not identify public or subsidized housing, so tenancies whose rent is set as a share of the tenant's income remain in the sample. The filters leave 21,497 observations from 6,147 households. Dropping the 178 observations that the PSID codes outside its four census regions (Alaska and Hawaii, or abroad), which the region-by-year design does not assign to a region, gives the renter sample: 21,319 observations from 6,108 households in 44 region-wave cells.

**Movers and stayers.** Rent changes are taken between consecutive waves in which the household rents. At the later wave of every such pair, the PSID asks whether the head has lived anywhere else since January of the previous survey year; a pair is a mover pair if the answer is yes. That January precedes the previous interview, so a move made between the two dates is reported at both waves, and the later of the two pairs spans no change of unit. Of the 13,405 consecutive renter pairs with a valid answer, 6,099 report a move (mover pairs), 3,823 report no move and the same room count in both waves (stayer pairs), and 3,483 report no move but a changed room count and belong to neither group. A changed room count without a reported move points to error in the reported count, which is a regressor of (1). The weighted share of pairs with a move is 0.424 per wave; because a move in the overlap counts at both waves, the share exceeds the share of pairs that span a change of unit. Appendix C states which observations enter each moment.

**The sample behind the model's targets.** The moments that discipline the life-cycle model (Section 3.7) come from a second sample: PSID individuals aged 27–80 in families that own or rent, with a positive individual longitudinal weight and a room count in the same range as the renter sample, pooled over the waves. The sample is not restricted to heads; each individual carries the

tenure, income and wealth of his or her family at his or her own age. Every target taken from it is a share, or a ratio of two money values reported in the same wave. The same sample supplies the joint distribution of assets and tenure at age 27 that initializes the simulated households, and the ownership rates by age of Figure 7. The renter move rate is the one target taken from the head panel: the weighted share of consecutive renter pairs with a move, $s = 0.424$ per wave, converted to the model's period as $1 - (1 - s)^{3/2}$.

### A.2 The Rent Residual

The dependent variable of the hedonic regression (1) is $\log Q_{it}$, where $Q_{it}$ is twelve times the monthly rent the family reports, divided by the rent deflator of its census region and survey year (below). This is the rent of the whole unit; the log number of rooms is among the regressors. The other regressors are the region-by-year effects $\mu_{rt}$, one for each region-wave cell; indicators of dwelling type, which complete the unit characteristics $x^u_{it}$; and the occupant characteristics $x^o_{it}$: the log of real family income, the inverse hyperbolic sine of real family net worth, the head's years of schooling, with a missing value coded as zero, and indicators of the head's race. We estimate the regression by least squares with the weights above. Section 2.1 discusses the role of each block, including the cost of the occupant controls. The earlier protocol compared in Appendix C is identical except that national year effects replace $\mu_{rt}$, the annual national all-items index deflates rent, and the 178 observations outside the four census regions are kept.

### A.3 Deflators

In the head panel, PSID family income, net worth and house values are in constant 2010 dollars, deflated by the annual consumer price index for all items. The rent deflator of region $r$ in survey year $t$ is the average of the monthly consumer price index for all items less shelter of the region over January to June of that year, divided by the same average for 2010. Rent is thus measured against the prices of goods and services other than shelter. Any deflator that varies only by region and year is absorbed by $\mu_{rt}$, so the residual and every moment computed from it are the same under the all-items and the less-shelter deflator; the seven moments of Appendix C agree to machine precision under the two. The choice matters only for the region-by-year effects themselves and for the external market index and comparisons with it (Appendix B).

### A.4 External Series

Table A.1 lists every external series with its source, frequency, the years that enter an estimate and its identifier. The consumer price indexes of Panel B supply the market component (Section 2.4, Appendix B) and the rent deflator; the BLS research indexes of Panel C supply the market-level evidence on re-entry (Section 2.6); the series of Panel D supply the model's saving and mortgage rates (Section 3.7), as average real rates over 2001–2021 compounded to the model's three-year period.

**Table A.1:** Data sources

| Series | Source | Frequency | Years used | Identifier |
|---|---|---|---|---|
| *A. Household panel* | | | | |
| PSID family files: rent, tenure, moves, rooms, income, wealth, house value | Panel Study of Income Dynamics, public release | Biennial | 2001–2021 | Extract J346293 |
| CPI for all urban consumers, all items (deflator of PSID incomes, wealth and house values; 2010 = 100) | BLS | Annual | 2001–2021 | Not recorded[a] |
| *B. Consumer price indexes: rent of primary residence; all items less shelter* | | | | |
| Northeast | BLS, via FRED | Monthly | 1978–2025 | CUUR0100SEHA<br>CUUR0100SA0L2 |
| Midwest | BLS, via FRED | Monthly | 1978–2025 | CUUR0200SEHA<br>CUUR0200SA0L2 |
| South | BLS, via FRED | Monthly | 1978–2025 | CUUR0300SEHA<br>CUUR0300SA0L2 |
| West | BLS, via FRED | Monthly | 1978–2025 | CUUR0400SEHA<br>CUUR0400SA0L2 |
| U.S. city average | BLS, via FRED | Monthly | 1941–2025 | CUUR0000SEHA<br>CUUR0000SA0L2 |
| *C. Rents of new and of all tenants* | | | | |
| New-tenant rent index | BLS research series | Quarterly | 2005Q1–2025Q3 | R-CPI-NTR |
| All-tenant rent index | BLS research series | Quarterly | 2005Q1–2025Q3 | R-CPI-ATR |
| *D. Interest rates and inflation* | | | | |
| Thirty-year fixed mortgage rate | Freddie Mac, via FRED | Weekly | 2001–2021 | MORTGAGE30US |
| Ten-year Treasury yield, constant maturity | Federal Reserve Board, via FRED | Monthly | 2001–2021 | GS10 |
| CPI for all urban consumers, all items | BLS, via FRED | Monthly | 2001–2021 | CPIAUCSL |

*Notes.* Years used: the years that enter an estimate. Panel A: PSID waves 2001–2021 (11 waves); the samples drawn from it are defined in Appendix A. PSID incomes, wealth and house values are in constant 2010 dollars, deflated by the annual all-items CPI; PSID rents are deflated by the regional CPI less shelter of Panel B. Panel B: not seasonally adjusted; identifiers are the BLS series codes, rent of primary residence above, all items less shelter below. The regional series are bimonthly in their early years; the annual market index averages the available January–June months of a year and requires at least three. The U.S. rent index enters in two samples, the long series and the regional window 1978–2025; the less-shelter index of each region deflates the PSID rents of that region, averaged over January–June of the survey year. Panel C: quarterly research indexes of the BLS; the all-tenant index starts in 1999Q4, and only the quarters common to both indexes are used. Panel D: real rates are the average nominal rate over the window minus the average realized CPI inflation, compounded to the model's three-year period. FRED: Federal Reserve Economic Data, Federal Reserve Bank of St. Louis. All external series were downloaded once and are stored with checksums.
[a] The annual index is read by the PSID extract pipeline; no result file records its series identifier.

# B Persistence of the Market Component

## B.1 The Annual Market Index

For each census region and for the nation, the monthly real rent index is the ratio of the consumer price index for rent of primary residence to the consumer price index for all items less shelter of the same area (Table A.1, Panel B). The annual index $g_t$ is the log of its average over the months of January to June of year $t$ in which both series are available, and a year enters only if at least three

such months exist. The regional series begin in the December before the first year of the regional window and were published every other month in their early years, which leaves exactly three matched months in each early window; no regional year is lost. The long national series covers 1941–2025; the years 1945–1946 have fewer than three matched months and drop out, which leaves 83 annual observations. Every series ends in the last year for which all six months from January to June exist in both series.

### B.2 Median-Unbiased Persistence

For each series we estimate, by least squares on all pairs of consecutive years,

$$g_t = a + b\,t + \rho\,g_{t-1} + \varepsilon_t,$$

where $\varepsilon_t$ is the residual. Near a unit root the least-squares estimate $\hat{\rho}$ is biased downward, and more so with a trend. We use the exact median-unbiased estimator of Andrews (1993), computed by simulation. For each $\rho$ on a grid from 0.30 to 1.020 in steps of 0.004 (181 points), we simulate 20,000 Gaussian first-order autoregressions as long as the regression sample (the number of consecutive-year pairs plus one), started from the stationary distribution when $\rho < 1$ and from zero otherwise, and apply the same trend regression to each. The same random draws serve at every grid point. The median-unbiased estimate $\hat{\rho}_\mu$ is the $\rho$ at which the simulated median of $\hat{\rho}$ equals the estimate in the data, and the equal-tailed 90 percent interval inverts the simulated fifth and ninety-fifth percentiles in the same way. We invert the running maximum of each simulated curve, because above one the median and quantile functions of the detrended estimator are nearly flat and need not increase strictly.

When the estimate in the data exceeds the simulated median at the top of the grid, the method returns the ceiling, 1.020, as a bound: the data do not distinguish the series from a unit root, and the entry is not an estimate above one. This happens for $\hat{\rho}_\mu$ in every series except the Northeast and for the upper end of every interval (Table 1).

### B.3 The Unit-Root Test and the Decision Rule

The unit-root test is the augmented Dickey–Fuller test (Dickey and Fuller, 1979) with a constant and a trend and the lag length chosen by the Akaike criterion, run on the longest stretch of consecutive years (1947–2025 for the long national series). The governing national series and the South select no lagged difference, so for them the test and the simulated first-order autoregression share the lag structure. The Northeast, the Midwest, the West and the long national series select one; we do not apply the median-unbiased estimator for higher-order autoregressions of Andrews and Chen (1994). The two procedures also use different statistics, the test the $t$-ratio of the lagged level and the interval the estimate itself, so the Northeast interval can reach the ceiling although the test rejects a unit root.

A rule fixed before estimation maps the test into the model. A series whose $p$-value is below 0.05 is classified as trend-stationary, and its median-unbiased persistence is reported with its interval. Any other series keeps a unit root: the rule assigns it a unit root with drift and reads its persistence only as an interval, so its median-unbiased entry in Table 1 is the grid ceiling, a bound. The model's market component follows the classification of one series fixed in advance, the national index over the regional window, 1978–2025, chosen because its window matches the regional series and excludes the era of postwar rent control; the long national series is reported alongside. Only the Northeast rejects; its persistence is 0.847 per year, or 0.609 per period (three years). Against the stationary values in the national interval, a random walk with the same innovation standard deviation has the larger variance at every horizon beyond one year, so the rule does not understate the market component relative to those values.

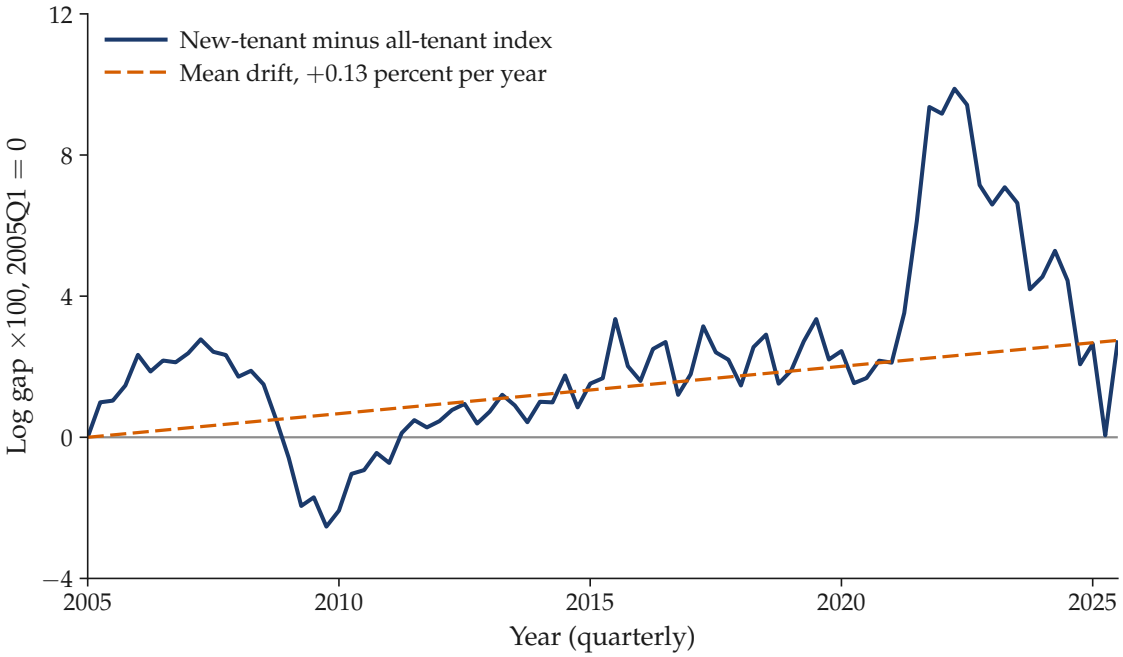
The innovation standard deviation per year, $\sigma$, is the residual standard deviation of the trend regression, with a correction for its three coefficients. Over $k$ years it is $\sigma\sqrt{k}$ for a unit-root series, and $\sigma\left[(1-\hat{\rho}_{\mu}^{2k})/(1-\hat{\rho}_{\mu}^{2})\right]^{1/2}$, with persistence $\hat{\rho}_{\mu}^{k}$, for a trend-stationary one. For the governing series $\sigma$ is 0.022 per year, and $\sigma_g$ is 0.038 per three-year period. Under the imposed unit root this $\sigma$ still comes from the residuals of the unrestricted trend regression, not from the first differences of $g_t$. That regression nests the random walk with drift, so its residual sum of squares is at most that of the demeaned first differences; apart from the correction for degrees of freedom, $\sigma$ is no larger than the random walk's own estimate. The model keeps the unit root and sets the deterministic drift of real rents to zero (Section 3.2).

### B.4 New-Tenant and All-Tenant Rents

The CPI rent index follows the rents of a panel of occupied units, and rent indexes built from new leases move differently from it (Ambrose et al., 2015; Adams et al., 2024); the rents of sitting tenants adjust less often than rents set at new leases (Genesove, 2003). The BLS research indexes of rents for new tenants and for all tenants separate the two at the market level (Table A.1, Panel C). Figure B.1 plots the log difference between them, quarterly over 2005Q1–2025Q3 (the all-tenant index starts in 1999Q4, and only the quarters common to both enter), normalized to zero in the first quarter. A deflator common to both indexes cancels.

The drift of Section 2.6 is the mean of the quarterly changes of the gap, multiplied by four. Because the gap starts at zero, that mean equals the gap in the last quarter, 2.75 percent in log terms, divided by the 82 quarterly changes: the drift, +0.13 percent per year, depends only on the last quarter, and the dashed line in Figure B.1 ends at the last observation by construction. The drift over 2010Q1–2019Q4 likewise depends only on the end points of that window. Around the line the gap swings widely: it turns negative around the financial crisis, rises sharply late in the sample, falls back to about zero in the second-to-last quarter and jumps in the last, so it moved by about as much in the last quarter as over the whole period. The sign and size of the drift therefore depend on the quarter that ends the sample. Because the two indexes are normalized separately, the gap measures the

**Figure B.1:** Log difference of the BLS new-tenant and all-tenant rent indexes, 2005Q1–2025Q3




*Notes.* Solid line: log of the BLS new-tenant rent index minus log of the BLS all-tenant rent index, times one hundred, quarterly, normalized to zero in 2005Q1. Dashed line: the straight line from the first to the last observation, whose slope is the drift (+0.13 percent per year). The gap measures relative growth, not the level difference between rents on new leases and rents of sitting tenants. Source: BLS research series R-CPI-NTR and R-CPI-ATR (Table A.1).

growth of rents on new leases relative to the rents of all tenancies, not the level difference between them (Section 2.6). The rent change at re-entry is measured in the PSID (Appendix D).

# C Identification of the Error-Components Model

## C.1 Moments

In this appendix $t$ indexes PSID waves. Under (2), with $\eta_{it}$ and $m_{it}$ serially uncorrelated and the three components independent, the level autocovariances are those of (3), and the autocovariances of changes within a tenancy are

$$\mathrm{fd}_0 = 2\sigma_z^2(1-\rho_w) + 2\sigma_m^2, \qquad \mathrm{fd}_1 = -\sigma_z^2(1-\rho_w)^2 - \sigma_m^2, \qquad \mathrm{fd}_k = -\sigma_z^2\rho_w^{k-1}(1-\rho_w)^2 \quad (k \geq 2),$$

which extends (4). Differencing removes $f_i^q$, so the first-difference moments carry no information on $\sigma_f^2$. Measurement error enters $\mathrm{lv}_0$, $\mathrm{fd}_0$ and $\mathrm{fd}_1$ and no other moment. The estimator fits seven moments, $\mathrm{lv}_0$–$\mathrm{lv}_3$ and $\mathrm{fd}_0$–$\mathrm{fd}_2$; $\mathrm{lv}_4$ and $\mathrm{fd}_3$ enter only the overidentification test. Table C.1 reports all nine with their bootstrap standard errors and counts.

**Table C.1:** Error-components fits without and with the overidentifying moments

| | Fitted on | | Data | | |
|---|---|---|---|---|---|
| | 7 moments | 9 moments | Value | s.e. | Obs. |
| *A. Parameters* | | | | | |
| Persistence per wave (two years), $\rho_w$ | 0.883 | 0.765 | | | |
| Persistence per period (three years), $\rho_q$ | 0.830 | 0.669 | | | |
| Innovation s.d. per wave, $\sigma_{\eta,w}$ | 0.240 | 0.238 | | | |
| Innovation s.d. per period, $\sigma_\eta$ | 0.285 | 0.275 | | | |
| Within-tenancy component, $\sigma_z^2$ | 0.261 | 0.137 | | | |
| Permanent rent type, $\sigma_f^2$ | 0.018 | 0.143 | | | |
| Measurement error, $\sigma_m^2$ | 0.061 | 0.059 | | | |
| Measurement error, percent of $\mathrm{fd}_0$ | 68.0 | 65.9 | | | |
| *B. Moments per wave* | | | | | |
| $\mathrm{lv}_0$ | 0.3407 | 0.3395 | 0.3537 | 0.0127 | 21,319 |
| $\mathrm{lv}_1$ | 0.2490 | 0.2480 | 0.2365 | 0.0156 | 7,306 |
| $\mathrm{lv}_2$ | 0.2221 | 0.2233 | 0.2148 | 0.0192 | 3,560 |
| $\mathrm{lv}_3$ | 0.1983 | 0.2045 | 0.2050 | 0.0244 | 1,932 |
| $\mathrm{lv}_4$ | 0.1773[a] | 0.1900 | 0.1952 | 0.0282 | 1,128 |
| $\mathrm{fd}_0$ | 0.1834 | 0.1830 | 0.1800 | 0.0181 | 3,823 |
| $\mathrm{fd}_1$ | −0.0648 | −0.0669 | −0.0586 | 0.0118 | 1,282 |
| $\mathrm{fd}_2$ | −0.0031 | −0.0058 | 0.0036 | 0.0075 | 546 |
| $\mathrm{fd}_3$ | −0.0028[a] | −0.0044 | 0.0037 | 0.0091 | 261 |
| *C. Innovation s.d. from the invariance identity* | | | | | |
| Per wave, $\sigma_{\eta,w}$ | 0.233 | 0.233 | | | |
| Per period, $\sigma_\eta$ | 0.277 | 0.269 | | | |
| *D. Implied re-draw correlation* | | | | | |
| Net mover variance, $s^2_{\mathrm{move}} - 2\sigma_m^2$ | 0.289 | 0.293 | | | |
| less $2\sigma_f^2$ | 0.252 | 0.007 | | | |
| Independent re-draw, $2\sigma_z^2$ | 0.522 | 0.274 | | | |
| $\rho_{\mathrm{move}}$, $f^q$ kept at a move (model) | 0.446 | −0.068[b] | | | |
| $\rho_{\mathrm{move}}$, $f^q$ new at each move | 0.517 | 0.975 | | | |
| *E. Fit* | | | | | |
| Equal-weight distance on the fitted moments | 0.00052 | 0.00067 | | | |
| Overidentification test, bootstrap $p$-value | | 0.122 | | | |

*Notes.* Error-components model of Table 2. First column: the adopted equal-weighted minimum-distance estimates, fitted on $\mathrm{lv}_0$–$\mathrm{lv}_3$ and $\mathrm{fd}_0$–$\mathrm{fd}_2$. Second column: the same estimator on all 9 moments, adding $\mathrm{lv}_4$ and $\mathrm{fd}_3$; its distance is the statistic of the overidentification test. Panel B: $\mathrm{lv}_k = \mathrm{Cov}(\tilde{q}_{it}, \tilde{q}_{i,t-k})$ and $\mathrm{fd}_k = \mathrm{Cov}(\Delta\tilde{q}_{it}, \Delta\tilde{q}_{i,t-k})$, $k$ in waves, fitted at each column's estimates; standard errors from 999 household-block bootstrap replications. Obs.: renter observations for $\mathrm{lv}_0$, pairs of observations $k$ waves apart without a reported move for $\mathrm{lv}_k$, stayer first-difference pairs for $\mathrm{fd}_k$. The model implies negative $\mathrm{fd}_2$ and $\mathrm{fd}_3$; both data values are positive, and their bootstrap intervals include zero. Panel C: $\sigma^2_{\eta,w} = \frac{1}{2}(1+\rho_w)(\mathrm{fd}_0 - 2\sigma_m^2)$, valid at any parameter vector of the model, evaluated at the data $\mathrm{fd}_0$ with each column's $\rho_w$ and $\sigma_m^2$; per period, $\sigma_\eta^2 = \sigma^2_{\eta,w}(1-\rho_w^3)/(1-\rho_w^2)$. Panel D: $\rho_{\mathrm{move}} = 1 - (s^2_{\mathrm{move}} - 2\sigma_m^2)/(2\sigma_z^2)$, $s_{\mathrm{move}}$ the mover s.d. of Table 3, when the permanent rent type $f_i^q$ is a household effect kept at a move, as in the model (if $f^q$ is a unit effect, new at each move, $2\sigma_f^2$ is also subtracted), from unrounded inputs. Panel E: sum of squared deviations between fitted and data moments over each column's own moment set; $p$-value from the bootstrap distribution of the second column's distance, recentered as in Hall and Horowitz (1996).

[a] Not used in the fit; implied by the first column's estimates.

[b] Negative: outside the range the model's re-draw construction accepts.

## C.2 The Samples behind the Moments

Each moment uses its own subset of the renter sample of Appendix A.

- $\text{lv}_0$ uses all 21,319 renter observations, movers and households observed once included.
- $\text{lv}_k$, $k \geq 1$, uses pairs of observations $k$ waves apart of a household that rents in every wave in between and reports no move at any of them, so that it occupies the same unit throughout; its room count may change. There are 7,306 such pairs one wave apart and 1,128 four waves apart.
- $\text{fd}_k$ uses the stayer pairs of Appendix A, with no reported move and an unchanged room count: 3,823 changes for $\text{fd}_0$, and pairs of such changes $k$ waves apart within a household for $k \geq 1$, 1,282 at the first lag and 261 at the third.

The level variance, the level autocovariances and the first-difference moments therefore come from three different renter populations, and the estimator imposes one parameter vector on all three. The fit's two largest deviations are on $\text{lv}_0$, fitted at 0.3407 against 0.3537 in the data, and on $\text{lv}_1$, fitted at 0.2490 against 0.2365, the two moments on either side of the change from all renters to renters who do not move. Because a move made between January of a survey year and that year's interview is reported at that wave and at the next (Appendix A), the no-move condition also drops some pairs that span no change of unit. Stayers are selected as well, since they chose not to move. The bound of Section 2.6 trims, in the spirit of Lee (2009), either the largest or the smallest absolute stayer changes by the weighted share of pairs with a move, 0.424; the resulting bounds on $\text{fd}_0$ are too wide to be informative.

## C.3 Estimator and Inference

Let $\theta = (\rho_w, \sigma_f^2, \sigma_z^2, \sigma_m^2)$, let $\hat{\mathcal{M}}_k$ denote a data moment and $\mathcal{M}_k(\theta)$ its value in the model, and let $\mathcal{K}$ be the set of seven moments. The equal-weighted minimum-distance estimator is

$$\hat{\theta} = \arg\min_{\theta} \sum_{k \in \mathcal{K}} \left(\mathcal{M}_k(\theta) - \hat{\mathcal{M}}_k\right)^2 \quad \text{subject to} \quad \sigma_f^2, \sigma_z^2, \sigma_m^2 \geq 0, \quad 0 \leq \rho_w \leq 0.999.$$

At a given $\rho_w$ the moments are linear in the three variances, which we estimate by nonnegative least squares. The outer search runs over a grid for $\rho_w$ in steps of 0.001 (1,000 points), followed by a bounded one-dimensional search within one step of the best grid point; the minimized objective at each grid point is the profile of Figure 1. Section 2.5 gives the reason for equal weights; on raw moments they give the smallest fitted moment, $\text{fd}_2$, little weight, as they do $\text{fd}_3$ in the nine-moment fit below.

Inference uses 999 household-block bootstrap replications. Households are resampled from the head panel, owners included, and a household drawn several times enters with its weight multiplied. Each replication rebuilds the renter sample and re-estimates the hedonic regression, the moments and the parameters. Intervals are 95 percent percentile intervals. Those for $\sigma_f^2$ and $\sigma_m^2$

start at zero because the nonnegativity constraints bind in some replications. At the boundary of the parameter space the estimator's distribution is nonstandard, and a percentile interval need not have its nominal coverage there (Andrews, 1999).

The overidentification statistic is the equal-weighted distance of the same estimator fitted to all nine moments, 0.00067. Its bootstrap distribution refits the nine-moment model to each replication's moments recentered as $\hat{\mathcal{M}}_k^b - \hat{\mathcal{M}}_k + \mathcal{M}_k(\hat{\theta}_9)$, where $\hat{\theta}_9$ is the nine-moment estimate, so that the model holds in the bootstrap population (Hall and Horowitz, 1996). The $p$-value, 0.122, is the share of bootstrap statistics at or above the sample value. The statistic weights deviations equally rather than by the inverse covariance of the moments. It tests the whole structure jointly and has little power against misfit in the smallest moments, $\mathrm{fd}_2$ and $\mathrm{fd}_3$, whose signs contradict the model.

### C.4 Closed Forms from Differences and from Levels

The first-difference block identifies $\rho_w$ in population. With $D = \mathrm{fd}_0/2 + \mathrm{fd}_1$ and $C = \mathrm{fd}_2$, the model implies $D = \sigma_z^2\rho_w(1-\rho_w)$ and $C = -\sigma_z^2\rho_w(1-\rho_w)^2$, so

$$\rho_w = 1 + \frac{C}{D}, \qquad \sigma_z^2 = \frac{D}{\rho_w(1-\rho_w)}, \qquad \sigma_m^2 = \frac{\mathrm{fd}_0}{2} - \sigma_z^2(1-\rho_w).$$

A persistence between zero and one requires $D > 0 > C$. In the data $C$ is positive, so the closed form exceeds one (Section 2.3); $C$ is positive under national year effects as well, $+0.0028$ (Table C.2). The level block identifies $\rho_w$ in population through $\mathrm{lv}_1 - \mathrm{lv}_2 = \sigma_z^2\rho_w(1-\rho_w)$ and $\mathrm{lv}_2 - \mathrm{lv}_3 = \sigma_z^2\rho_w^2(1-\rho_w)$, whose ratio is $\rho_w$. In the data the ratio is 0.449, against the joint estimate 0.883. The bootstrap interval for $\rho_w$, $[0.256, 0.929]$ per wave, and the profile of Figure 1 measure how little the sample pins persistence down. The model is point-identified in population, but in this sample the objective is nearly flat along one direction.

### C.5 The Ridge and Its Edges

Fix $\rho_w \in (0,1)$. The moments $\mathrm{lv}_0$, $\mathrm{lv}_1$ and $\mathrm{lv}_2$ then determine the three variances exactly:

$$\sigma_z^2(\rho_w) = \frac{\mathrm{lv}_1 - \mathrm{lv}_2}{\rho_w(1-\rho_w)}, \qquad \sigma_f^2(\rho_w) = \mathrm{lv}_1 - \frac{\mathrm{lv}_1 - \mathrm{lv}_2}{1-\rho_w}, \qquad \sigma_m^2(\rho_w) = \mathrm{lv}_0 - \mathrm{lv}_1 - \frac{\mathrm{lv}_1 - \mathrm{lv}_2}{\rho_w}.$$

Along this exactly identified ridge a higher $\rho_w$ moves dispersion from the permanent type to the within-tenancy component: $\sigma_f^2(\rho_w)$ falls everywhere, and $\sigma_z^2(\rho_w)$ rises for $\rho_w$ above one half. Given $\mathrm{lv}_0 > \mathrm{lv}_1 > \mathrm{lv}_2$, as in the data, all three variances are nonnegative if and only if

$$\frac{\mathrm{lv}_1 - \mathrm{lv}_2}{\mathrm{lv}_0 - \mathrm{lv}_1} \le \rho_w \le \frac{\mathrm{lv}_2}{\mathrm{lv}_1},$$

and every $\rho_w$ in this range fits the three moments exactly. At the lower edge the measurement-error variance reaches zero, at the upper edge the permanent variance. The lower edge lies below

**Table C.2:** Error-components estimates under national year effects and under region-by-year effects

| | National year effects | | Region×year effects |
|---|---|---|---|
| | Scale-free (first step) | Two-step (second step) | Equal weights (adopted) |
| *A. Estimates* | | | |
| Persistence per wave (two years), $\rho_w$ | 0.959 | 0.905 | 0.883 |
| Persistence per period (three years), $\rho_q$ | 0.939 | 0.861 | 0.830 |
| Innovation s.d. per period, $\sigma_\eta$ | 0.179 | 0.279 | 0.285 |
| Within-tenancy component, $\sigma_z^2$ | 0.268 | 0.302 | 0.261 |
| Permanent rent type, $\sigma_f^2$ | 0.000 | 0.000 | 0.018 |
| Measurement error, $\sigma_m^2$ | 0.066 | 0.064 | 0.061 |
| Measurement error, percent of $\mathrm{fd}_0$ | 71.0 | 69.2 | 68.0 |
| Measurement error, percent of $\mathrm{lv}_0$ | 17.5 | 17.0 | 17.3 |
| *B. Data moments (per wave)* | | | |
| $\mathrm{lv}_0$ | | 0.3758 | 0.3537 |
| $\mathrm{lv}_1$ | | 0.2563 | 0.2365 |
| $\mathrm{lv}_2$ | | 0.2387 | 0.2148 |
| $\mathrm{lv}_3$ | | 0.2292 | 0.2050 |
| $\mathrm{lv}_4$ | | 0.2242 | 0.1952 |
| $\mathrm{lv}_3/\mathrm{lv}_1$ | | 0.894 | 0.867 |
| $\mathrm{fd}_0$ | | 0.1849 | 0.1800 |
| $\mathrm{fd}_1$ | | −0.0587 | −0.0586 |
| $\mathrm{fd}_2$ | | 0.0028 | 0.0036 |
| *C. What differs between the columns* | | | |
| Time effects in the hedonic regression | National year | | Region×year |
| Rent deflator | All items, national | | Less shelter, by region |
| Moment deviations divided by | \|moment\| | Bootstrap s.d. | None |
| Moments fitted | 8 | | 7 |
| Variance parameters | Positive | | Nonnegative |
| Renter observations | 21,497 | | 21,319 |
| Renter households | 6,147 | | 6,108 |
| Stayer first-difference pairs | 3,852 | | 3,823 |

*Notes.* The error-components model of Table 2 under national year effects and under region×year effects (the adopted estimates of that table). No column uses external data; the external rent indexes identify only the market component (Table 1). The national-year columns share the rent residual and the data moments and differ only in how deviations between model and data moments are scaled. The scale-free column, the first step of a two-step estimator, divides each deviation by the absolute value of the data moment (with a floor); the two-step column, its second step, divides by the bootstrap standard deviation of the moment. The region×year column differs from both in the time effects, the deflator, the scaling, the moment set and the sample: renter observations outside the four census regions (178) are dropped. Moments fitted: $\mathrm{lv}_0$–$\mathrm{lv}_4$ and $\mathrm{fd}_0$–$\mathrm{fd}_2$ under national year effects; $\mathrm{lv}_0$–$\mathrm{lv}_3$ and $\mathrm{fd}_0$–$\mathrm{fd}_2$ under region×year effects, where $\mathrm{lv}_4$ and $\mathrm{fd}_3$ are kept for the overidentification test. Variance parameters: exponential parameterization under national year effects (strictly positive, so $\sigma_f^2$ can approach zero but not reach it; it prints as zero), nonnegativity constraints under region×year effects, under which the region-year deflator drops out of every moment. Measurement-error shares: $2\sigma_m^2/\mathrm{fd}_0$ and $\sigma_m^2/\mathrm{lv}_0$, with each column's own data moments. $\mathrm{lv}_3/\mathrm{lv}_1$: ratio of data moments; a slow decay of the level autocovariances is what leaves persistence and permanent heterogeneity weakly separated.

both bands of Figure 1; the upper edge is 0.908 per wave at the data moments. Above it no non-negative permanent variance fits $\mathrm{lv}_0$ to $\mathrm{lv}_2$: with $\sigma_f^2$ at zero the implied ratio $\mathrm{lv}_2/\mathrm{lv}_1$ equals $\rho_w$ and exceeds the ratio in the data. The upper end of the set on which the equal-weighted objective is within 10 percent of its minimum, 0.907 per wave, lies just below this edge. Beyond it the objective rises steeply toward a random walk; below the estimate it rises slowly (Figure 1). The other four moments tilt the ridge without removing its flat direction. The adopted estimate, $\rho_w = 0.883$, lies close to the upper edge, where the permanent type is small, $\sigma_f^2 = 0.018$.

### C.6 The Invariance Identity

Within a tenancy, $\Delta\tilde{q}_{it} = -(1-\rho_w)z_{i,t-1} + \eta_{it} + m_{it} - m_{i,t-1}$, and the three terms are uncorrelated, so

$$\mathrm{fd}_0 = (1-\rho_w)^2\sigma_z^2 + \sigma_{\eta,w}^2 + 2\sigma_m^2.$$

Substituting $\sigma_z^2 = \sigma_{\eta,w}^2/(1-\rho_w^2)$ and using $(1-\rho_w)^2/(1-\rho_w^2) = (1-\rho_w)/(1+\rho_w)$ gives $\mathrm{fd}_0 = 2\sigma_{\eta,w}^2/(1+\rho_w) + 2\sigma_m^2$, which is (5). The identity holds at every parameter vector of the model. The permanent variance does not enter it, and persistence enters only through $\frac{1}{2}(1+\rho_w)$, whose range across the band is given in Section 2.5. On the exactly identified ridge the innovation variance per wave is $\sigma_{\eta,w}^2(\rho_w) = \sigma_z^2(\rho_w)(1-\rho_w^2) = (\mathrm{lv}_1 - \mathrm{lv}_2)(1+\rho_w)/\rho_w$, which changes far less with $\rho_w$ than $\sigma_z^2(\rho_w)$ does.

The identity makes the innovation robust to the split between persistence and the permanent type, not to measurement error. The variance $\sigma_m^2$ is tied to $\mathrm{fd}_1$, of which measurement error makes up 94.5 percent at the estimates, so it moves little along the ridge; but its bootstrap interval runs from 0.000 to 0.079, and it absorbs any transitory component of rent (Section 2.5). Differentiating (5) at a given $\rho_w$ gives

$$\frac{\partial \log \sigma_{\eta,w}^2}{\partial \log \sigma_m^2} = -\frac{s_m}{1-s_m}, \qquad s_m = \frac{2\sigma_m^2}{\mathrm{fd}_0},$$

where $s_m$ is the measurement-error share of $\mathrm{fd}_0$, 68.0 percent at the estimates. Because the share exceeds one half, an error in $\sigma_m^2$ moves the innovation variance by a larger percentage than its own.

Panel C of Table C.1 evaluates (5) at the data $\mathrm{fd}_0$ with the persistence and measurement error of each fit. Per wave it gives 0.233 at the adopted point and 0.233 at the nine-moment point. Per period the values are 0.277 and 0.269, because the conversion $\sigma_\eta^2 = \sigma_{\eta,w}^2(1-\rho_w^3)/(1-\rho_w^2)$ depends on $\rho_w$. The fitted values per wave, 0.240 and 0.238, are slightly higher because each fit overpredicts $\mathrm{fd}_0$; at the adopted point the residual is +0.0034. The identity does not extend to the stationary variance $\sigma_z^2 = \sigma_{\eta,w}^2/(1-\rho_w^2)$, which moves with persistence along the ridge. The dispersion of rent positions and the re-draw correlation of Appendix D, which both depend on $\sigma_z^2$, inherit the ridge.

## C.7 Seven and Nine Moments

Table C.1 sets the adopted fit beside the same estimator on all nine moments. Adding $\mathrm{lv}_4$ and $\mathrm{fd}_3$ moves the estimate along the ridge. Under equal weights the added misfit is mostly on $\mathrm{lv}_4$, which the adopted point puts at 0.1773 against 0.1952 in the data, while it puts $\mathrm{fd}_3$ at $-0.0028$ against 0.0037. The level autocovariances decay more slowly at long lags than the adopted fit allows, and a larger permanent component flattens the implied tail. Persistence falls from 0.883 to 0.765 per wave (0.830 to 0.669 per period), $\sigma_f^2$ rises from 0.018 to 0.143, and $\sigma_z^2$ falls from 0.261 to 0.137; the innovation variance per wave changes by a factor of 0.99, and $\sigma_m^2$ moves from 0.061 to 0.059. The nine-moment point lies inside the band of the previous subsections, and its measurement-error share of $\mathrm{fd}_0$, 65.9 percent, is close to the adopted 68.0 percent.

## C.8 National Year Effects and Region-by-Year Effects

Table C.2 sets the adopted estimates beside two estimates under the earlier protocol, whose hedonic regression has national year effects. No column uses the external rent indexes. Each column deflates rent by a price index, but a deflator that varies only by year, or only by region and year, is absorbed by the time effects of its column, so the change of deflator is inert. The two national-year columns share the residual, the eight moments $\mathrm{lv}_0$–$\mathrm{lv}_4$ and $\mathrm{fd}_0$–$\mathrm{fd}_2$ and a parameterization with strictly positive variances; they differ only in how deviations between model and data moments are scaled (Panel C), and they place persistence at 0.959 and 0.905 per wave. On a flat profile the weighting decides where the estimate lands. The adopted column differs from both at once in the time effects, the scaling, the moment set, the variance constraints and the sample. These changes cannot be separated, so the move of $\sigma_f^2$ from zero to 0.018 cannot be assigned to any one of them; the bootstrap interval for $\sigma_f^2$ in the adopted column, $[0.000, 0.236]$, contains zero.

Panel B shows what region-by-year effects, together with the drop of the 178 observations outside the four regions, do to the data. Every level autocovariance falls, as removing a common persistent component should: $\mathrm{lv}_1$ from 0.2563 to 0.2365 and $\mathrm{lv}_4$ from 0.2242 to 0.1952. The decay of the tail barely changes, from 0.894 to 0.867 for $\mathrm{lv}_3/\mathrm{lv}_1$, so the slow decay does not come from region-common movements. It comes from the residual, which contains household-specific movements and movements common to a metropolitan area (Section 2.4). The first-difference moments change little.

The innovation standard deviation per period is 0.179, 0.279 and 0.285 across the three columns. The scale-free estimator divides each deviation by the size of its moment, so the small $\mathrm{fd}_2$ dominates its objective; the fit drives $\rho_w$ toward one and leaves $\mathrm{fd}_0$ underfitted, and by (5), which holds at the fitted moments of any parameter vector, the innovation follows the fitted $\mathrm{fd}_0$. The two estimators that fit $\mathrm{fd}_0$ closely agree on $\sigma_\eta$. The measurement-error share of $\mathrm{fd}_0$ is 71.0, 69.2 and 68.0 percent.

# D The Re-Entry Re-Draw

## D.1 Derivation

Consider a renter observed in two consecutive waves who reports a move between them, a mover pair of Appendix A. Suppose that the permanent rent type $f_i^q$ belongs to the household and is kept across the move, that the region-by-year effects remove the change in region-common rents, and that across the move the within-tenancy component is partly re-drawn,

$$z_{it} = \rho_{\text{move}}\, z_{i,t-1} + e_{it}, \qquad \text{Var}(e_{it}) = \sigma_z^2\big(1 - \rho_{\text{move}}^2\big),$$

where $z_{i,t-1}$ has mean zero and the stationary variance $\sigma_z^2$, as for a renter drawn at random, and the move innovation $e_{it}$ is independent of it. The variance of $e_{it}$ is the one that keeps $\text{Var}(z_{it}) = \rho_{\text{move}}^2\sigma_z^2 + \text{Var}(e_{it})$ equal to $\sigma_z^2$, so a move leaves the cross-sectional distribution of positions unchanged. The change in the residual is $\Delta\tilde{q}_{it} = -(1-\rho_{\text{move}})z_{i,t-1} + e_{it} + m_{it} - m_{i,t-1}$, and if movers' measurement error has the stayers' variance $\sigma_m^2$,

$$s_{\text{move}}^2 = (1-\rho_{\text{move}})^2\sigma_z^2 + \sigma_z^2\big(1-\rho_{\text{move}}^2\big) + 2\sigma_m^2 = 2\sigma_z^2(1-\rho_{\text{move}}) + 2\sigma_m^2.$$

Solving for $\rho_{\text{move}}$ gives (6). With $s_{\text{move}} = 0.641$, so that $s_{\text{move}}^2 = 0.411$, and the adopted $\sigma_m^2 = 0.061$ and $\sigma_z^2 = 0.261$, the mover variance net of measurement error is 0.289, against 0.522 for an independent re-draw, and $\rho_{\text{move}} = 0.446$.

The same algebra for a stayer, whose position follows the within-tenancy autoregression, gives $2\sigma_z^2(1-\rho_w) + 2\sigma_m^2 = \text{fd}_0$, a moment the estimator of Appendix C already fits, so the stayers provide no separate check. Set side by side, the two expressions show what $\rho_{\text{move}}$ measures: the correlation of a renter's position across a wave in which it changes unit, to be compared with $\rho_w = 0.883$ across a wave in which it does not.

## D.2 Three Assumptions about the Mover's Rent Change

Three assumptions about what a mover's rent change contains deliver (6), beyond the error-components structure of Section 2.1 and the condition, taken up in the last subsection of this appendix, that movers' positions before the move are distributed like those of all renters. None of the three is tested by the data used here.

(i) *The permanent rent type belongs to the household.* If $f_i^q$ belonged to the unit and changed at each move, the variance of movers' changes would also contain $2\sigma_f^2$, and $\rho_{\text{move}} = 1 - (s_{\text{move}}^2 - 2\sigma_m^2 - 2\sigma_f^2)/(2\sigma_z^2) = 0.517$ at the adopted estimates. The model would then overstate how much of $z$ a move replaces but omit that a move also replaces the permanent type, so the net effect on what a move forfeits has no definite sign. Assumption (i) is also what makes $f_i^q$ heterogeneity rather than risk: a type that changed at each move would be part of re-entry risk.

(ii) *Movers report rents with the stayers' measurement error.* Measurement error then accounts for 29.8 percent of the variance of movers' changes, against 68.0 percent for stayers. If movers report with more error, the net mover variance and the re-draw are overstated, and $\rho_{\text{move}}$ is understated.

(iii) *The dispersion of movers' rent changes is a draw of the price position, not a choice.* This is Assumption 1 applied at the move. Movers choose their new unit. Changes in unmeasured quality, and in location within the census region, enter $\Delta\tilde{q}_{it}$ and are read as re-draw risk; to that extent the re-draw is overstated and $\rho_{\text{move}}$ understated. The change also contains the occupant-control term $-\kappa_o'\Delta x_{it}^o$ of Section 2.1: moves that coincide with larger changes in income or net worth than stayers experience add variance that (6) reads as re-draw, again understating $\rho_{\text{move}}$.

The definitions of a move also differ. A mover in the data is any renter who reports having lived elsewhere since January of the previous survey year. That includes a move to a unit of the same size, and a move made before the previous interview, after which the household occupies the same unit at both waves (Appendix A); pairs of the second kind lower the mover variance and raise the estimated $\rho_{\text{move}}$. A stayer in the data also keeps its room count. In the model a household enters a new rental unit only when it changes its dwelling size as a renter or rents after owning (Section 3.2). An owner's position keeps evolving by $\Pi_{\text{stay}}$ while it owns and is partly re-drawn by $\Pi_{\text{move}}$ when the household rents again.

### D.3 From One Wave to One Period

Movers' changes span one wave of two years, and the model applies $\rho_{\text{move}}$ per period of three years. This is an assumption, not a conversion: the derivation treats the whole change between the two waves as the re-draw and ignores the evolution of $z$ within the tenancies before and after the move. If $z$ evolves within those tenancies as it does within any tenancy, the correlation across the move itself exceeds the one-wave estimate, and the correlation over a three-year period that contains the same move falls below it, because a period holds more time within a tenancy than a wave does. The one-wave value applied per period then overstates the part of a position that survives a move and understates forfeiture at re-entry. The like-for-like comparison in the data is with the per-wave $\rho_w$; in the model the gap is $\rho_{\text{move}} - \rho_q = -0.384$ per period. The timing also differs. The data measure the change between the last wave in the old unit and the first wave in the new one; in the model a household that enters a new unit pays the rent of its current position there, and the re-draw governs its position from the next period on (Section 3.2).

### D.4 Dependence on the Point on the Ridge

The re-draw correlation is a function of $\sigma_z^2$ and $\sigma_m^2$, and the rent data pin down $\sigma_z^2$ only up to the ridge of Appendix C. Panel D of Table C.1 evaluates the formula at the nine-moment fit. There the net mover variance, 0.293, exceeds the variance of an independent re-draw, 0.274, and

$\rho_{\text{move}} = -0.068$: a re-draw that places the new position on the other side of the mean from the old one, outside the range the model's construction admits. If the permanent type belongs to the unit instead, twice the larger $\sigma_f^2$ of that fit accounts for nearly all of the net mover variance (0.007 remains), and $\rho_{\text{move}} = 0.975$: a move then leaves $z$ almost unchanged but replaces the permanent component of the unit's rent, which the model holds fixed. The dispersion of movers' rent changes is the same data moment in every case. What the ridge and assumption (i) leave open is how it divides between the re-draw of $z$ and a change in the permanent component; across the ridge and the two readings of the permanent type, the correlation of $z$ across a move ranges from below zero to nearly one. The re-draw correlation is identified only conditional on the point on the ridge and on assumption (i).

The counterfactuals of Section 4 use the adopted point. Of their results, the effect of switching off the re-draw, 0.82 percentage points of ownership, and the portable-position experiment, which is the same world, depend on $\rho_{\text{move}}$ directly (Sections 4.2 and 4.7). The 2.47 points from removing the rest of rental-rate risk with the re-draw already off do not depend on $\rho_{\text{move}}$, but they depend on $\sigma_z^2$, which inherits the ridge.

### D.5 The Move Innovation in the Model

The model applies the same construction per period. The move chain $\Pi_{\text{move}}$ shares the grid and the stationary distribution of $\Pi_{\text{stay}}$, so its innovation has the variance $\sigma_z^2(1 - \rho_{\text{move}}^2)$ that keeps the distribution of positions unchanged (Section 3.2, equation (H.3)). The re-draw therefore leaves the cross-sectional distribution of positions as it is. For a given household it pulls the expected position toward the mean, by (7), and adds dispersion around it; averaged over positions it does not shift the mean. In rent levels the larger move innovation adds a convexity term (Appendix H).

### D.6 Selection into Moving

The mover statistics condition on remaining a renter. A renter who buys between two waves leaves no rent change, and if buyers' rent changes would have differed from those of movers who keep renting, the mover distribution is selected. In the spirit of Lee (2009), we bound the mean change of movers by trimming their distribution at either end by 13.8 percent, the weighted share of renters observed in two consecutive waves who own at the second (16,792 renter observations at risk). The bounds, $[-0.138, +0.178]$ per wave around a mean of $+0.021$, contain zero: the sign of the average rent change at a move is not identified, and the model's zero level step lies inside the bounds. The trimming share is not derived from a selection model; it is the exit rate of all renters. Among renters who change units, by moving to another rental or by buying, exits to ownership are a larger share, and trimming at that share would widen the bounds, which would then also contain zero.

The bounds concern the mean. The dispersion that identifies $\rho_{\text{move}}$ rests on the condition that movers' positions before the move are distributed like those of all renters, with mean zero and

variance $\sigma_z^2$, and selection on $z$ breaks it. The model's own movers are selected on $z$: its renters at unfavorable positions enter a new unit with probability 0.39 per period, and those at favorable positions with probability 0.16 (Table 8). If movers' positions before the move have mean $\bar{z}_{\text{mv}}$ and variance $v_{\text{mv}}$, their changes have mean $-(1-\rho_{\text{move}})\bar{z}_{\text{mv}}$ and variance $(1-\rho_{\text{move}})^2 v_{\text{mv}} + (1-\rho_{\text{move}}^2)\sigma_z^2 + 2\sigma_m^2$, and (6) returns

$$\rho_{\text{move}} - \frac{(1-\rho_{\text{move}})^2\,(v_{\text{mv}} - \sigma_z^2)}{2\sigma_z^2}$$

instead of $\rho_{\text{move}}$. It overstates the correlation when movers' positions are less dispersed than those of all renters and understates it when they are more dispersed. Neither the data used here nor the stored output of the model say which holds.

# E The Income Process

This appendix documents the income process of the model. Table E.1 sets it beside the within-tenancy rent process of Section 2.5 and the house-price process of Appendix F. Each column fits an error-components model to within-household autocovariances of a log residual at the biennial frequency of the PSID and converts the estimates to the model's three-year period; the columns differ in sample, residual and weighting, as the last rows of the table record. In Appendices E–G, $t$ indexes survey waves, two years apart, unless a model period is stated.

## E.1 Sample and Residual

The sample consists of the heads of the panel of Appendix A who are aged 27–60 and employed, that is, working at the interview. The income concept is family labor income, the head's plus the spouse's, in constant 2010 dollars, because the model's budget constraint carries the resources of the household rather than those of one earner. We keep observations in which the head's labor income is positive and below the PSID top code and family labor income is at least 1,000 dollars: 31,950 observations of 7,647 households.

The residual $\tilde{y}_{it}$ is log family labor income net of a least-squares regression, weighted by PSID weights, on age and its square, years of education (a missing value is coded as zero), an indicator for a female head, and indicators for race, census region, family size and survey year. The year effects are national, unlike the region-by-year effects of the rent hedonic regression (1), so regional income movements stay in the residual. The age terms only remove the mean life-cycle path from the residual; the model's age profile of income is a separate object, described at the end of this appendix.

**Table E.1:** Error-components estimates of the rent, income and house-price processes

| | Rent, within tenancy | Income | House price |
|---|---|---|---|
| Persistence per period, $\rho$ | 0.830 | 0.829 | 0.650 |
| | [0.129, 0.896] | [0.718, 0.904] | [0.401, 0.879] |
| Innovation s.d. per period, $\sigma_\eta$ | 0.285 | 0.326 | 0.327 |
| | [0.209, 0.360] | [0.270, 0.368] | [0.287, 0.376] |
| Permanent-component variance, $\sigma_f^2$ | 0.018 | 0.082 | 0.215 |
| | [0.000, 0.236] | [0.000, 0.178] | [0.000, 0.276] |
| Serially uncorrelated component variance, $\sigma_e^2$ | 0.061 | 0.063 | 0.037 |
| | [0.000, 0.079] | [0.051, 0.073] | [0.008, 0.059] |
| Sample | Renter heads | Employed heads | Owner heads |
| Minimum-distance weights | Equal | Relative | Inverse s.d. |
| Bootstrap replications | 999 | 200 | 500 |
| Interval coverage, percent | 95 | 90 | 95 |

*Notes.* Each column fits an error-components model to within-household autocovariances of a log residual in the PSID, biennial waves 2001–2021: a permanent component (variance $\sigma_f^2$), a persistent AR(1) component at wave frequency with persistence $\rho_w$ and stationary variance $\sigma_c^2$, and a serially uncorrelated component (variance $\sigma_e^2$). In the paper's notation the residuals are $\tilde{q} = f^q + z + m$ for rent, $\tilde{y} = f^y + \omega + \varepsilon$ for income and $\tilde{p} = f^p + p + m$ for house prices, so the rows are $(\rho_q, \sigma_\eta, \sigma^2_{f^q}, \sigma^2_m)$ for rent, $(\rho_y, \sigma_{\eta,y}, \sigma^2_{f^y}, \sigma^2_\varepsilon)$ for income and $(\rho_p, \sigma_{\eta,p}, \sigma^2_{f^p}, \sigma^2_m)$ for house prices. Estimates are converted to the three-year model period as $\rho = \rho_w^{3/2}$ and $\sigma_\eta^2 = \sigma_c^2(1-\rho^2)$; per wave, $\rho_w$ is 0.883 for rent, 0.883 for income and 0.750 for house prices. Brackets are percentile intervals from a bootstrap that resamples households, at the coverage and number of replications in the last two rows. Rent: renter heads aged 27–80 (21,319 observations, 6,108 households); log annual rent deflated by the regional CPI less shelter, net of dwelling and occupant characteristics and region-by-year effects; equal-weighted minimum distance (EWMD) on seven moments: the level variance, the level autocovariances at lags one to three within tenancies, and the variance and first two autocovariances of first differences within a unit; the serially uncorrelated component is measurement error. Income: employed heads aged 27–60 (31,950 observations, 7,647 households); log family labor income of head and spouse, at least 1,000 constant 2010 dollars, net of age, education, sex, race, region, family size and year effects; minimum distance with each moment's gap divided by the moment's absolute value, on level autocovariances at lags zero to 8 waves and the variance and first two autocovariances of first differences; the serially uncorrelated component is transitory income, which includes measurement error. House price: owner heads' self-reported house values deflated by the all-items CPI, net of dwelling and occupant characteristics and national year effects; two-step minimum distance with each moment's gap divided by its first-step bootstrap standard deviation, on level autocovariances at lags zero to four within spells without a move and the variance and first two autocovariances of first differences of owners who did not move and kept their room count (9,611 first-difference pairs); the serially uncorrelated component is measurement error. The model uses $\rho$ and $\sigma_\eta$ of all three processes, $\sigma_f^2$ of rent and income as three permanent types each, and the transitory income variance; the house-price $\sigma_f^2$ and the two measurement-error variances do not enter it. In the model, rent also has a market random-walk component with innovation s.d. 0.038 per period, and entering a new rental unit re-draws $z$ with correlation 0.446.

## E.2 Moments and Estimator

We model the residual as

$$\tilde{y}_{it} = f_i^y + \omega_{it} + \varepsilon_{it}, \qquad \omega_{it} = \rho_{y,w}\,\omega_{i,t-1} + \eta_{it}^y, \tag{E.1}$$

where $f_i^y$ is the permanent income type, with variance $\sigma^2_{fy}$; $\omega_{it}$ is a persistent component with persistence $\rho_{y,w}$ per wave, stationary variance $\sigma^2_\omega$ and innovation $\eta_{it}^y$, whose variance per wave is $\sigma^2_\omega(1-\rho^2_{y,w})$; and $\varepsilon_{it}$, with variance $\sigma^2_\varepsilon$, is serially uncorrelated and combines transitory income with measurement error. The components are mutually independent. In the model's income process (10), $\eta_{it}^y$ is the innovation per three-year period. Storesletten et al. (2004) use the same specification. The covariance-structure approach follows Meghir and Pistaferri (2004) and Blundell et al. (2008), who work with income changes; we add level autocovariances because differencing removes the permanent type. The moments are those of the rent model, (3) and (4), with $(\sigma^2_{fy}, \sigma^2_\omega, \sigma^2_\varepsilon, \rho_{y,w})$ in place of $(\sigma^2_f, \sigma^2_z, \sigma^2_m, \rho_w)$; we write them $\mathrm{lv}_k^y$ and $\mathrm{fd}_k^y$.

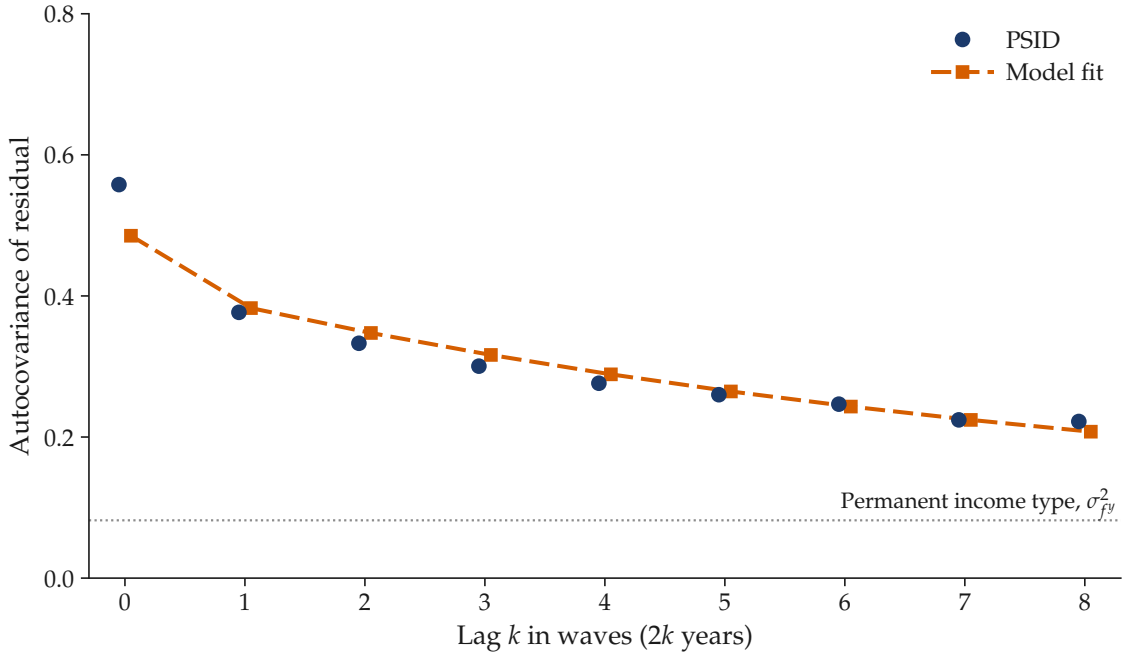
All moments are weighted by PSID weights. The level autocovariances run from lag zero to lag eight waves (16 years). The lag-zero variance uses all observations; the autocovariance at lag $k$ uses pairs $k$ waves apart within runs of consecutive waves in the sample. The change moments $\mathrm{fd}_0^y$, $\mathrm{fd}_1^y$ and $\mathrm{fd}_2^y$ use one-wave changes of households observed in consecutive waves, after dropping the 3 changes with $|\Delta\tilde{y}_{it}| > 4.0$.

We estimate the four parameters $\theta$ of (E.1) by minimizing, over the nine level moments and the three change moments, the sum of squared relative gaps,

$$\sum_j \left[ \frac{\mathcal{M}_j(\theta) - \hat{\mathcal{M}}_j}{\max\{|\hat{\mathcal{M}}_j|,\, c\}} \right]^2,$$

where $\hat{\mathcal{M}}_j$ is a data moment, $\mathcal{M}_j(\theta)$ its model counterpart and $c$ a small floor that binds only for $\mathrm{fd}_2^y$. Apart from $\mathrm{fd}_2^y$, a given relative miss costs the same for every moment, large or small. This is the scale-free first step of Table C.2 and of the house-price estimator (Appendix F), here on a longer set of level moments and with no second step. Compared with the equal weights used for rents (Section 2.5), it puts far more weight on the small change moments than on the level variance, and the most on $\mathrm{fd}_2^y$, whose gap is divided by the floor. The parameterization keeps the variances positive and $\rho_{y,w}$ inside the unit interval, and the problem is solved by nonlinear least squares. The conversion to the model period is the one used for rents: $\sigma^2_\omega$ is kept, $\rho_y = \rho_{y,w}^{3/2}$ and $\sigma^2_{\eta,y} = \sigma^2_\omega(1-\rho^2_y)$. Inference is a household-block bootstrap: each of 200 replications resamples households of the income sample and repeats the regression, the moments and the fit. The intervals are 90 percent percentile intervals, reported for $\rho_y$, $\sigma_{\eta,y}$, $\sigma^2_{fy}$ and $\sigma^2_\varepsilon$.

**Figure E.1:** Autocovariances of the log family labor income residual: PSID and the fitted error-components model




*Notes.* Dots: weighted autocovariances $\mathrm{Cov}(\tilde{y}_{it}, \tilde{y}_{i,t-k})$ of the residual of log family labor income of employed heads aged 27–60, PSID 2001–2021, at lags of $k$ waves; lag zero uses all 31,950 observations and lag $k$ pairs within runs of consecutive waves. Squares joined by the dashed line: the fitted model (E.1), $\sigma^2_{fy} + \sigma^2_{\omega} + \sigma^2_{\varepsilon}$ at lag zero and $\sigma^2_{fy} + \rho^k_{y,w}\sigma^2_{\omega}$ at lag $k$, with $\rho_{y,w} = 0.883$, $\sigma^2_{\omega} = 0.341$ and $\sigma^2_{\varepsilon} = 0.063$. Dotted line: the permanent variance $\sigma^2_{fy} = 0.082$. Bootstrap intervals for these moments are not available.

## E.3 Estimates and Fit

Persistence is 0.883 per wave. Per three-year period, persistence is $\rho_y = 0.829$ $[0.718, 0.904]$ and the innovation standard deviation is $\sigma_{\eta,y} = 0.326$ $[0.270, 0.368]$, around a stationary variance $\sigma^2_{\omega} = 0.341$. The permanent variance is $\sigma^2_{fy} = 0.082$, with an interval that reaches zero, $[0.000, 0.178]$. The transitory variance is $\sigma^2_{\varepsilon} = 0.063$ $[0.051, 0.073]$, a standard deviation of 0.250.

Figure E.1 compares the level autocovariances with the fit. From lag one to lag eight the fitted values lie close to the data, above them at intermediate lags and below them at the longest. At lag zero the fitted variance, 0.485, is 13 percent below the data, 0.558, the largest relative miss of any moment. The miss reflects a conflict among the data moments. For a stationary process observed on one population, $\mathrm{fd}^y_0 = 2(\mathrm{lv}^y_0 - \mathrm{lv}^y_1)$, so the lag-zero variance equals $\mathrm{fd}^y_0/2 + \mathrm{lv}^y_1$. In the data it is well above that value, and the two sides are not computed on one population: the lag-zero variance uses every observation, including heads not observed working in an adjacent wave, while $\mathrm{lv}^y_1$ and the change moments use pairs of consecutive working waves. The model satisfies the identity by construction, and the weights, which favor the change moments over the level variance, resolve the conflict in favor of the pairs: the fitted lag-zero variance sits close to $\mathrm{fd}^y_0/2 + \mathrm{lv}^y_1$. The rent moments are built the same way (Appendix C).

## E.4 Precision of Income and Rent Persistence

Persistence per wave is 0.883 for income and 0.883 for rent, but the two estimates differ in precision. The interval for income persistence per period is $[0.718, 0.904]$ at 90 percent coverage; for rent persistence it is $[0.129, 0.896]$ at 95 percent coverage. The profiled objectives differ in the same way. The income objective is profiled over a coarse grid of per-wave values, and at every point of that grid, including the two on either side of the estimate, it exceeds its minimum by more than 10 percent. The rent objective stays within that tolerance of its minimum for persistence from 0.671 to 0.907 per wave (Figure 1), a band that contains several points of the income grid.

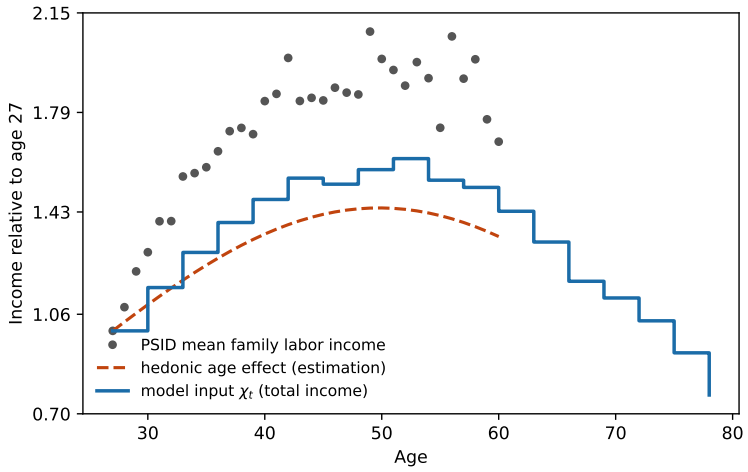
Three features of the data account for the contrast. First, income autocovariances are observed to lag eight over working careers and decline at every lag, which pins down the rate of decline. Within-tenancy rent autocovariances come from short spells, enter the rent estimator only to lag three and decline little over that range, which leaves the ridge of Section 2.3. Second, $\mathrm{fd}_2^y$ has the negative sign the model requires, so for income the change moments alone place persistence inside the unit interval; for rents $\mathrm{fd}_2$ is positive and they do not (Section 2.3). The relative gaps give $\mathrm{fd}_2^y$ the largest weight. Applied to rents under national year effects, the same weighting drives persistence toward one, because no admissible parameter gives $\mathrm{fd}_2$ its positive sign (Appendix C); applied to income, it fits $\mathrm{fd}_2^y$ closely at interior persistence. Third, each income moment that has a rent counterpart, apart from the lag-zero variance, rests on several times as many pairs. The comparison is not controlled, since the two estimators weight their moments differently. The tail pins down how fast the autocovariances decline, not the level to which they decline: at lag eight the autocovariance is still far above the fitted $\sigma^2_{f^y}$, so the permanent variance is an extrapolation beyond the observed tail, and its interval reaches zero.

## E.5 How the Model Uses the Estimates

The model takes $\rho_y$ and $\sigma_{\eta,y}$ as persistent income risk per period, $\sigma_\varepsilon$ as transitory income risk and $\sigma^2_{f^y}$ as three permanent income types known from 27 (Section 3.2; the discretization is in Appendix I). Income is independent of rents and house prices.

Four features of the mapping from estimates to model bear on how much income risk the model carries. The moments cannot separate measurement error from transitory income, so the model's transitory income risk contains measurement error, whereas its rent process excludes it. The transitory variance enters without conversion: the model treats the transitory component of one year's income as the transitory component of a three-year period's income, while transitory shocks that are not perfectly correlated across the years of a period partly average out within it. Both choices can only raise transitory income risk relative to what households face. The other two work toward less risk and dispersion. Heads without work at the interview are outside the sample, so the estimated process omits the income risk of spells without work. And the residual removes the income differences that education, sex, race and family size predict, so the model's income dispersion excludes them.

**Figure E.2:** Income by age: PSID mean family labor income, the age effect of the income regression, and the profile used in the model




*Notes.* Markers (legend: PSID mean family labor income): weighted mean family labor income by single year of age in the income sample (employed heads aged 27–60), relative to age 27. Dashed line (legend: hedonic age effect): the age effect of the income regression (age and its square), relative to age 27. Solid step (legend: model input $\chi_t$): the profile $\chi_t$ used in the model, the weighted mean of total family income of PSID individuals aged 27–80 by three-year age bin, relative to the first bin. The markers are raw means of a working sample; the dashed line holds education, sex, race, region, family size and survey year fixed; the step differs from both in sample and in income concept.

At 27 the persistent component is drawn with variance $\sigma^2_{\omega,27} = 0.390$: the residual variance of heads at the five youngest single-year ages of the sample, 0.535, net of $\sigma^2_{f'y}$ and $\sigma^2_\varepsilon$. It exceeds the stationary variance, $\sigma^2_\omega = 0.341$, because the entry-age variance exceeds the fitted lag-zero variance, 0.485: the difference $\sigma^2_{\omega,27} - \sigma^2_\omega$ equals the difference between the two. The entry-age variance is nonetheless below the data lag-zero variance, 0.558. In the model the dispersion of income therefore falls over the first periods of life, toward the fitted lag-zero variance. In the data, heads at entry are less dispersed than the average observation, consistent with within-cohort dispersion that rises with age (Deaton and Paxson, 1994). Under the stationary specification, dispersion would rise with age only if the entry-age variance were below the fitted lag-zero variance, so the lag-zero miss sets the direction of the model's age pattern.

## E.6 The Deterministic Age Profile

The level of income over the life cycle enters through $\chi_t$ in (10). For the model period that starts at age $a$, $\chi_t$ is the weighted mean of total family income in the three-year age bin that starts at $a$, relative to the bin that starts at 27, computed on PSID individuals aged 27–80 and pooled over the survey waves. It is a ratio of means, not the age effect of the income regression, and because it pools cross-sections it does not separate age from cohort and calendar-time effects. The profile peaks at 1.62 in the bin that starts at 51 and falls to 0.77 in the last bin, which starts at 78; beyond 60, the oldest age of the income sample, its decline carries the fall in income at retirement. Figure E.2 sets it beside the raw PSID means and the age effect of the regression.

The profile and the stochastic components differ in two respects. The profile uses total family income, which includes non-labor income, while the stochastic components come from the labor income of employed heads. And the model applies those components at every age, so income risk after 60 extrapolates working-age risk.

# F The House-Price Process

The model's house-price process takes its persistence and innovation size from the persistent component of the part of an owner's self-reported house value that national year effects and the characteristics of the dwelling and its occupant do not explain. This appendix describes the estimation, the house-price column of Table E.1, and what the model's process contains.

## F.1 Sample and Residual

The sample consists of the owners in the panel of heads of Appendix A who report a positive house value below the PSID top code and pass the room-count and family-income screens of the rent sample. The dependent variable is the log reported value in constant 2010 dollars. A least-squares regression weighted by PSID weights removes log family income, the inverse hyperbolic sine of net worth excluding home equity, log rooms, years of education, indicators for race and dwelling type, and survey-year indicators. The year effects are national: they absorb aggregate house-price movements, and nothing is added back. The regression has no geographic controls, so regional and local price movements stay in the residual and load on its permanent and persistent components. Because the regression conditions on the number of rooms, the residual measures a house's value relative to what its size predicts, which the model reads as a price per room.

## F.2 Moments and Estimator

We model the residual as

$$\tilde{p}_{it} = f_i^p + p_{it} + m_{it}, \qquad p_{it} = \rho_{p,w}\, p_{i,t-1} + \eta_{it}^p, \tag{F.1}$$

where $f_i^p$ is a permanent component of the owner and the house, with variance $\sigma^2_{fp}$; $p_{it}$ is a persistent component with persistence $\rho_{p,w}$ per wave, stationary variance $\sigma_p^2$ and innovation $\eta_{it}^p$, whose variance per wave is $\sigma_p^2(1-\rho_{p,w}^2)$; and $m_{it}$ is serially uncorrelated error in the reported value, with variance $\sigma_m^2$ (in this appendix $m$ and $\sigma_m^2$ refer to house values). The model's price follows the per-period counterpart of $p_{it}$: in (15), $\log p_{it} - \log \bar{p}$ is the log deviation of the price per room from its mean and $\eta_{it}^p$ the innovation per three-year period.

The moments are those of Section 2.3. The level autocovariances run from lag zero to lag four waves: the lag-zero variance uses all owner observations, and the autocovariance at lag $k$ uses pairs $k$ waves apart within spells in which the owner reports no move at any wave of the span.

The variance and first two autocovariances of one-wave changes use owners who report no move and no change in rooms (9,611 changes).

The estimator is two-step diagonal minimum distance. The first step minimizes the sum of squared relative gaps, the scale-free first step used for income (Appendix E) and for rents under national year effects (Table C.2). A household-block bootstrap of 500 replications gives the standard deviation of each moment, and the second step divides each moment's gap by that standard deviation. The same 500 resamples, refitted with the second-step weights held fixed, give 95 percent percentile intervals. Each replication resamples households of the full panel of heads and repeats the sample selection, the regression, the moments and the fit. The per-wave estimates are converted to the model period by the rule used for rents and income.

### F.3 Estimates and Identification

Per wave, persistence is $\rho_{p,w} = 0.750$ $[0.544, 0.917]$ and the stationary variance is $\sigma_p^2 = 0.186$ $[0.144, 0.395]$. Per three-year period, persistence is $\rho_p = 0.650$ $[0.401, 0.879]$ and the innovation standard deviation is $\sigma_{\eta,p} = 0.327$ $[0.287, 0.376]$. The permanent variance is $\sigma_{fp}^2 = 0.215$ $[0.000, 0.276]$ and the reporting-error variance is $\sigma_m^2 = 0.037$ $[0.008, 0.059]$. Figure F.1 compares the level autocovariances with the fit; every fitted value lies inside the bootstrap interval of the corresponding data moment.

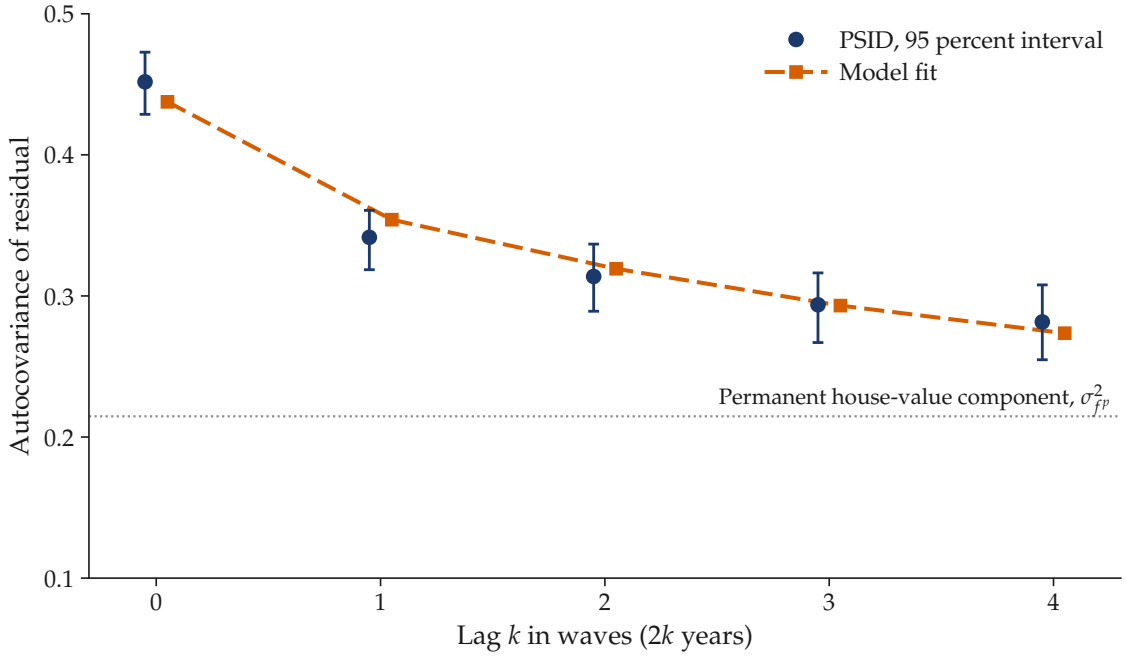
House values do not escape the persistence-versus-heterogeneity problem of rents. Their level autocovariances decline slowly over the four observed lags, about as slowly as the within-tenancy rent autocovariances (Figure 3), and the interval for $\sigma_{fp}^2$ reaches zero, as the interval for the permanent rent variance does, $[0.000, 0.236]$. The change moments do not settle the split either. As for rents, $\mathrm{fd}_2$ is positive at the point estimate, so the change moments alone push persistence to one. The split between permanent and persistent variation therefore depends on the weights. The first step puts the permanent variance at zero and the per-wave persistence well above its second-step value, as the scale-free step does for rents (Table C.2); the precision weights of the second step move both to the interior values reported above.

### F.4 What the Model's Price Process Is, and What It Is Not

The model uses $\rho_p$ and $\sigma_{\eta,p}$, discretized on five nodes around a fixed mean (Section 3.2, Appendix I). Neither $\sigma_{fp}^2$ nor the reporting error enters the model. Four properties of the process bear on the results.

First, the price is stationary. With $\rho_p = 0.650$ per period, a low current price predicts appreciation, which gives a motive to buy while the price is low (Section 4.4). The strength of that motive depends on mean reversion, and so on the split between permanent and persistent variation that the data leave open. At the first-step estimates nearly all variation in reported values other than

**Figure F.1:** Autocovariances of the log house-value residual: PSID owners and the fitted error-components model




*Notes.* Dots: weighted autocovariances $\mathrm{Cov}(\tilde{p}_{it}, \tilde{p}_{i,t-k})$ of the residual of log self-reported house value of owner heads, PSID 2001–2021, at lags of $k$ waves within spells without a move, with 95 percent household-block bootstrap intervals (500 replications). Squares joined by the dashed line: the fitted model (F.1) at the two-step estimates, $\sigma^2_{fp} + \sigma^2_p + \sigma^2_m$ at lag zero and $\sigma^2_{fp} + \rho^k_{p,w}\sigma^2_p$ at lag $k$, with $\rho_{p,w} = 0.750$, $\sigma^2_p = 0.186$ and $\sigma^2_m = 0.037$. Dotted line: the permanent variance $\sigma^2_{fp} = 0.215$, estimated but not used in the model.

reporting error would be persistent, with more persistence per period and less mean reversion than at the second-step point the model uses.

Second, the price has no market component, whereas rents carry a market random walk (Section 2.4). The log ratio of price to rent in the model therefore contains that random walk and is not stationary, although the market innovation is small.

Third, the price is independent of rents and income. The model has no co-movement between rents and house prices, the channel through which the value of an owned house hedges future rent; the insurance the model measures excludes it (Section 3.5).

Fourth, the price is not re-drawn at a move. A household that sells one house and buys another, or changes size, transacts at the price it already faces, so unlike the rent position $z$ the price a household faces does not depend on its housing choices.

The process also leaves out what the permanent component measures. The model has permanent rent types but no permanent price types, so households differ in the price they face only through the persistent component, although in the data the permanent component of reported values is about as large as the persistent one ($\sigma^2_{fp} = 0.215$ against $\sigma^2_p = 0.186$). Nor does the model relate a

household's price to its permanent rent type. Reported values are owners' assessments: persistent misreporting loads on $f_i^p$ or $p_{it}$, not on $m_{it}$, and the model treats what loads on $p_{it}$ as a price at which the household trades.

# G Forecastability of the Rent Innovation

Section 2.5 treats the within-tenancy innovation $\eta_{it}$ as risk. That requires $\eta_{it}$ to be news: it cannot be forecast from what the household knows before it arrives. The error-components model makes $\eta_{it}$ unpredictable from the household's rent history, but the autocovariances cannot distinguish news from changes the household anticipates. This appendix tests the difference with household information that the autocovariances do not use, on the one-wave rent changes of stayers, the renters of Section 2.5 who report no move and an unchanged room count.

## G.1 What Is Tested

Within a tenancy, (2) gives the one-wave change in the rent residual as

$$\Delta\tilde{q}_{it} = -(1-\rho_w)\, z_{i,t-1} + \eta_{it} + (m_{it} - m_{i,t-1}). \tag{G.1}$$

Two parts of the change are predictable without contradicting news. A high residual predicts a fall, $\mathrm{Cov}(\Delta\tilde{q}_{it}, \tilde{q}_{i,t-1}) = \mathrm{lv}_1 - \mathrm{lv}_0 = -(1-\rho_w)\sigma_z^2 - \sigma_m^2$. The first term is mean reversion of the rent position, which the household forecasts itself; the second is measurement error, which also makes consecutive changes negatively correlated: $m_{i,t-1}$ enters the current change with a negative sign and the previous change with a positive one. At the estimates of Section 2.5, measurement error accounts for 94.5 percent of the fitted $\mathrm{fd}_1$. Neither part is information about $\eta_{it}$.

The tests therefore use variables $X_{i,t-1}$ that the household knows at the previous wave and that are not functions of the rent residual. If $\eta_{it}$ is news and $m_{it}$ is classical measurement error, (G.1) implies

$$E\left[\Delta\tilde{q}_{it} \mid X_{i,t-1}\right] = -(1-\rho_w)\, E\left[z_{i,t-1} \mid X_{i,t-1}\right],$$

so $X_{i,t-1}$ predicts the change only through what it reveals about the rent position. Table G.1 reports two tests and two benchmarks. The first test regresses $\Delta\tilde{q}_{it}$ on $X_{i,t-1}$ and asks whether all coefficients $\beta_X$ are zero. This null is stronger than news: it also fails when $X_{i,t-1}$ is correlated with the current position $z_{i,t-1}$. The second test relaxes it by adding the lagged residual $\tilde{q}_{i,t-1}$ to absorb mean reversion, instrumented by the residual at the household's renter observation before $t-1$ (normally wave $t-2$), which is uncorrelated with $m_{i,t-1}$ under classical measurement error. The instrument also carries the permanent rent type, so the lagged level absorbs mean reversion exactly when $\sigma_f^2 = 0$ and approximately when $\sigma_f^2$ is small, as at the estimates of Section 2.5. The benchmarks regress the change on the lagged level alone and on the lagged change alone; they show how much predictability mean reversion and measurement error produce.

**Table G.1:** Forecastability of the within-tenancy rent innovation

| | Test of the null | | Mechanical benchmarks | |
|---|---|---|---|---|
| | Information set only | Adding lagged level (IV) | Lagged level | Lagged change |
| Household information $X_{i,t-1}$ | Yes | Yes | No | No |
| $R^2$ (percent) | 6.64 | — | 12.81 | 20.04 |
| Coefficient on $\tilde{q}_{i,t-1}$ | — | −0.043 | −0.252 | — |
| | | (0.034) | (0.034) | |
| Coefficient on $\Delta\tilde{q}_{i,t-1}$ | — | — | — | −0.397 |
| | | | | (0.035) |
| Wald test of $\beta_X = 0$ | 137.1 | 125.4 | — | — |
| Degrees of freedom | 32 | 32 | — | — |
| $p$-value | <0.001 | <0.001 | — | — |
| First-stage $F$ | — | 521 | — | — |
| First-difference pairs | 2,966 | 2,966 | 2,966 | 2,737 |
| Households | 1,527 | 1,527 | 1,527 | n.r. |

*Notes.* Dependent variable: $\Delta\tilde{q}_{it}$, the change in the log rent residual of a renter head over one wave (two years), for strict stayers (consecutive waves, no reported move, no change in rooms); every regressor is dated $t-1$. Sample: the 2,966 of the 3,823 strict-stayer first-difference pairs with complete lagged information; the lagged-change benchmark also requires the previous change. $X_{i,t-1}$: survey-year, census-region, marital-status, employment-status and college indicators, age and its square, length of the current tenancy, lagged income growth (the change in log real family income between the household's two previous observations), lagged change in family size, family size and number of children (Table G.2). Test columns: regression of $\Delta\tilde{q}_{it}$ on $X_{i,t-1}$ alone, and with the lagged level $\tilde{q}_{i,t-1}$ added and instrumented by the residual at the household's previous renter observation (normally $t-2$); $\beta_X$ denotes the coefficients on $X_{i,t-1}$. In the first column the null $\beta_X = 0$ is stronger than news: it also fails when $X_{i,t-1}$ is correlated with the current position $z_{i,t-1}$. In the second the lagged level absorbs mean reversion, and news with classical measurement error implies $\beta_X = 0$ exactly when $\sigma_f^2 = 0$ and approximately when $\sigma_f^2$ is small, because the instrument also carries the permanent rent type. Benchmark columns: regressions on the lagged level alone and on the lagged change alone, which predict the change through mean reversion and measurement error. Least squares weighted by PSID longitudinal weights; standard errors (in parentheses) and Wald statistics are clustered by household. The lagged-level benchmark uses the pairs and households of the first column; the household count of the lagged-change benchmark is not recorded (n.r.). —: not applicable. $R^2$ in percent, not reported for the instrumental-variables column. Source: PSID, 2001–2021.

$X_{i,t-1}$ contains survey-year, census-region, marital-status, employment-status and college indicators; age and its square; the length of the current tenancy; lagged income growth, the change in log real total family income, the income control of (1), between the household's two previous observations in the renter sample; the lagged change in family size; family size; and the number of children: 32 coefficients in all. Every regressor is dated at the previous wave or earlier. Of the 3,823 stayer changes of Section 2.5, the tests use the 2,966 with complete lagged information, from 1,527 households. Lagged income growth requires a renter observation before the previous wave, so the test sample is a subset of the stayer changes that leans toward longer renter histories. Regressions are weighted by PSID weights, and standard errors are clustered by household.

## G.2 Results

Both tests reject $\beta_X = 0$. In the first test the Wald statistic is 137.1 on 32 degrees of freedom ($p < 0.001$), and the $R^2$ is 6.64 percent. Adding the instrumented lagged level leaves the rejection in place (125.4 on 32 degrees of freedom, $p < 0.001$; first-stage $F$ of 521), so mean reversion does not account for the rejection. The mechanical channels predict more than the household information does: the lagged level alone gives an $R^2$ of 12.81 percent and the lagged change alone 20.04 percent.

The coefficient on the instrumented lagged level is $-0.043$ (standard error 0.034). With the instrument from the same tenancy and $X_{i,t-1}$ left aside, (2) implies that the instrumented regression recovers $-(1-\rho_w)\,\rho_w\sigma_z^2/(\sigma_f^2+\rho_w\sigma_z^2)$, which is close to $-(1-\rho_w)$ at the estimates of Section 2.5, where $\sigma_f^2$ is small next to $\rho_w\sigma_z^2$. The estimate is closer to zero. In the formula, a coefficient closer to zero requires higher persistence, a larger permanent type, or both. The estimate is imprecise, conditions on $X_{i,t-1}$, and comes from the test sample rather than from the samples behind the moments.

Table G.2 locates the rejection. The region-by-year effects of (1) already remove rent movements common to the renters of a region in a year, so the calendar ($p = 0.155$) and region ($p = 0.214$) blocks ask only whether stayers' changes depart from those common movements, and neither is significant. Age ($p = 0.673$) is insignificant, and so is the length of the current tenancy ($p = 0.253$), where scheduled increases, or a sitting-tenant discount that grows with duration, would show up. Tenancy length counts the household's renter observations since its last reported move, or since its first observation in the sample if it reports none, so it is censored and would detect such a pattern only weakly. None of the four blocks adds more than 0.60 percentage points to the $R^2$. The rejection comes from the block of household circumstances ($p < 0.001$) and, within it, from one regressor: lagged income growth alone adds 5.29 of the 6.64 percentage points of $R^2$, or 80 percent. Household composition ($p = 0.156$) and socioeconomic status ($p = 0.160$) are insignificant.

## G.3 The Predictability Loads on the Occupant Controls

Section 2.1 shows that the change in a stayer's residual contains $-\kappa_o'\Delta x_{it}^o$, where $x_{it}^o$ are the occupant controls of (1), among them log family income and the inverse hyperbolic sine of net worth: the residual moves when the occupant's income moves, whether or not the rent does. Changes in income are negatively autocorrelated (for the labor income of Appendix E, $\mathrm{fd}_1^y < 0$), so lagged income growth predicts the current change in the income control and, through it, the change in the residual.

A diagnostic confirms the channel. We recompute the residual from the hedonic regression without the two financial controls, keeping every other regressor, including the region-by-year effects, and rerun the first test on the same pairs. Lagged income growth then has no predictive content ($p = 0.99$): net of region-by-year movements, the rent a stayer pays does not move with its lagged

**Table G.2:** Predictive content of blocks of household information for the next rent change

| Information block | Regressors | Wald | $p$-value | Incremental $R^2$ (percent) |
|---|---|---|---|---|
| Calendar (survey year) | 8 | 11.9 | 0.155 | 0.60 |
| Geography (census region) | 3 | 4.5 | 0.214 | 0.18 |
| Life cycle (age and its square) | 2 | 0.8 | 0.673 | 0.03 |
| Length of the current tenancy | 1 | 1.3 | 0.253 | 0.04 |
| Household circumstances | 18 | 110.1 | <0.001 | 5.84 |
| of which: lagged income growth | 1 | 82.0 | <0.001 | 5.29 |
| of which: household composition | 3 | 5.2 | 0.156 | 0.23 |
| of which: socioeconomic status | 14 | 19.1 | 0.160 | 0.59 |
| All information jointly (total $R^2$) | 32 | 137.1 | <0.001 | 6.64 |

*Notes.* Blocks of the regressors $X_{i,t-1}$ in the test regression of Table G.1 (first column: $\Delta\tilde{q}_{it}$ on $X_{i,t-1}$). Regressors: number of coefficients in the block. Wald: cluster-robust test that the block's coefficients are zero, with degrees of freedom equal to the number of regressors, and its $p$-value. Incremental $R^2$: the regression's $R^2$ minus the $R^2$ with the block dropped, in percent; blocks are correlated, so the increments do not add up to the total. Last row: the joint test of all coefficients and the regression's total $R^2$. Lagged income growth: change in log real total family income between the household's two previous observations in the panel (normally waves $t-2$ and $t-1$; the construction does not require them to be one wave apart). Household composition: lagged change in family size, family size and number of children. Socioeconomic status: marital-status, employment-status and college indicators. Sample, weights and clustering as in Table G.1. Source: PSID, 2001–2021.

income growth. The alternative residual serves only as a diagnostic: the rent process of Section 2.5, and the model built on it, use the residual with the occupant controls.

In the first test on the alternative residual, what remains is small but statistically detectable: $X_{i,t-1}$ explains a share $R^2_{\text{alt}}$, 1.9 percent, of the variance of the change (Wald statistic 55.5 on the same regressors, $p = 0.006$). That share includes any mean reversion that $X_{i,t-1}$ predicts through the current position, and as an in-sample $R^2$ on 32 regressors it overstates the predictable share, so the following is an upper bound. If the share $R^2_{\text{alt}}$ of the stayers' change variance $\text{fd}_0$ came entirely out of the innovation, the innovation variance per wave would fall to $\sigma^2_{\eta,w} - R^2_{\text{alt}}\,\text{fd}_0$, and the innovation standard deviation per wave (two years) from 0.240 to 0.232, or by 3.1 percent. Holding $\rho_w$ fixed, the value per period falls in the same proportion. The model uses the unadjusted $\sigma_\eta = 0.285$ per three-year period (Section 2.5). Apart from the variation that the occupant controls pass into the residual, the household information in $X_{i,t-1}$ predicts little of the within-tenancy rent change.

## G.4 Scope

Four limits apply. The tests condition on staying in the unit, so they concern the within-tenancy innovation and say nothing about the re-draw at re-entry (Section 2.6). Conditioning on staying also selects on the innovation: renters hit by a large increase are more likely to move, so a regressor that predicts moving can predict $\eta_{it}$ among stayers even if $\eta_{it}$ is news. Selection of this kind acts on the rent change and would appear in the alternative residual as well, where lagged income

growth predicts nothing, so it does not explain the income-growth result. $X_{i,t-1}$ is the econometrician's measure of the household's information, and a household may know more: the tests can reject that $\eta_{it}$ is news but cannot establish it, the advance-information problem that Blundell et al. (2008) discuss for income shocks. And the tests inherit the rent sample and residual of Section 2, including the occupant controls whose consequences they document.

# H The Two Forces in Detail

## H.1 The Renter's Options and the Grid

The renter of Section 3.6 holds position $z$, has income $y$ per period of three years and a dwelling of size $h$, both normalized to one, and pays $q(z') = q_0 e^{z'}$ next period. The market component and the permanent rent type are held fixed, so $z'$ is the only source of rent risk. The scale $q_0$, which here plays the role of the level constant $\bar{q}$ of (11), sets mean rent under the stationary distribution of $z$ to 0.206 of income, the median rent-to-income ratio of PSID renters (Table 5).

The two rental options, staying ($d = \text{stay}$) and entering a new unit ($d = \text{move}$), are valued by their certainty-equivalent rents (21). Insurance (22) is the risk premium in the certainty-equivalent rent of staying, and the forfeiture wedge (23) is the gap between the expected rents of moving and staying. The certainty-equivalent rents of the two options therefore differ by

$$R^*_{\text{move}}(z) - R^*_{\text{stay}}(z) = \text{Lock}(z) + \left(R^*_{\text{move}}(z) - E[q' \mid z, \text{move}]\right) - \text{Ins}(z), \tag{H.1}$$

the forfeiture wedge plus the difference between the risk premia of moving and of staying. Both premia are second order in the size of rent risk (below), so as rent risk shrinks the gap between the two options is governed by the wedge. The moving cost of Section 3.3 is paid on leaving the unit, for another rental or for ownership; it depends on the current rent, not on $z'$, and enters neither force.

As in the model, $z$ follows (13) with persistence $\rho_q = 0.830$ and innovation standard deviation $\sigma_\eta = 0.285$ per period, so its stationary standard deviation is $\sigma_z = \sigma_\eta(1-\rho_q^2)^{-1/2} = 0.511$. The Rouwenhorst method (Rouwenhorst, 1995; Kopecky and Suen, 2010) discretizes it on five evenly spaced nodes,

$$z \in \{-2\sigma_z,\ -\sigma_z,\ 0,\ \sigma_z,\ 2\sigma_z\}, \tag{H.2}$$

with transition matrix $\Pi_{\text{stay}}$. Rent at the lowest node is 0.32 times mean rent and at the highest 2.44 times.

The move transition $\Pi_{\text{move}}$ is the Rouwenhorst chain on the same nodes with persistence $\rho_{\text{move}} = 0.446$, the re-draw correlation, estimated from one-wave rent changes and applied per period (Section 2.6, Appendix D). A Rouwenhorst chain with $N$ nodes spans $(N-1)^{1/2}$ stationary standard deviations on either side of zero, and its stationary standard deviation is its innovation standard

deviation divided by $(1-\rho^2)^{1/2}$, where $\rho$ is its persistence. Keeping the nodes of (H.2) at the lower persistence therefore requires the larger innovation

$$\sigma_{\text{move}} = \sigma_z \left(1-\rho_{\text{move}}^2\right)^{1/2} = 0.457, \tag{H.3}$$

1.60 times $\sigma_\eta$. With the innovation held at $\sigma_\eta$, the move chain would sit on a narrower grid than the stay chain. With the rescaled innovation, both chains have the same binomial stationary distribution, which a symmetric Rouwenhorst chain has whatever its persistence (Kopecky and Suen, 2010). Averaged over that distribution, expected rent next period is the same whether households stay or move: the re-draw changes where a given household expects its rent to go, not mean rent across households. The model uses the same construction (Appendix I); in both, the re-drawn position first prices rent in the period after the move.

## H.2 Conditional Means in Logs and in Levels

The conditional mean of a Rouwenhorst chain is linear in the current node, with slope equal to its persistence (Kopecky and Suen, 2010). On the grid, therefore,

$$E[z' \mid z, \text{stay}] = \rho_q\, z, \qquad E[z' \mid z, \text{move}] = \rho_{\text{move}}\, z, \tag{H.4}$$

exactly. The wedge in logs, $(\rho_{\text{move}} - \rho_q)z$, equation (7), is zero at $z = 0$, positive at favorable positions ($z < 0$), where entering a new unit raises the expected rent, and negative at unfavorable ones (Figure H.1, panel (a)). It is the conditional counterpart of the finding of Section 2.6 that re-entry is, unconditionally, a dispersion phenomenon rather than a level one.

Rent is convex in $z$, so levels add a term. In the continuous-state analog of the two chains, $z'$ given $z$ and option $d$ is normal with mean $\rho_d z$ and variance $\sigma_d^2 = \sigma_z^2(1-\rho_d^2)$, where $\rho_{\text{stay}} = \rho_q$, so that $\sigma_{\text{stay}} = \sigma_\eta$ and $\sigma_{\text{move}}$ is (H.3). Then $E[q' \mid z, d] = q_0 \exp(\rho_d z + \sigma_d^2/2)$, and the forfeiture wedge is

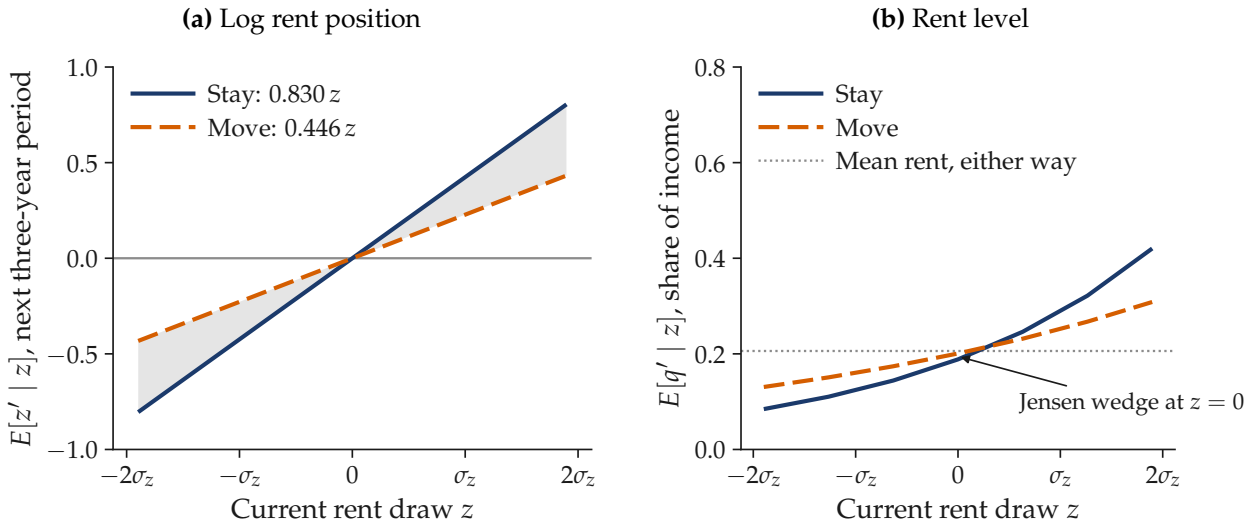
$$\text{Lock}(z) = q_0 \left[ e^{\rho_{\text{move}} z + \sigma_z^2(1-\rho_{\text{move}}^2)/2} - e^{\rho_q z + \sigma_z^2(1-\rho_q^2)/2} \right]. \tag{H.5}$$

At the mean position the terms in $z$ vanish and a Jensen wedge remains,

$$\text{Lock}(0) = q_0 \left[ e^{\sigma_z^2(1-\rho_{\text{move}}^2)/2} - e^{\sigma_z^2(1-\rho_q^2)/2} \right] > 0, \tag{H.6}$$

because $0 < \rho_{\text{move}} < \rho_q$: the move draw has the larger conditional variance. In the analog, the wedge in levels is positive below $z^* = \sigma_z^2(\rho_q + \rho_{\text{move}})/2$, slightly above the mean position, and negative above it. On the five-node grid the expectations are exact finite sums, and the wedge at $z = 0$ is 0.0123 of income, 6 percent of mean rent; Figure H.1, panel (b), shows the same wedge on a finer grid. The wedge does not contradict mean preservation. Both chains have the same stationary distribution, so the average of $\text{Lock}(z)$ over it is zero: the positive values at and below the mean position are offset by negative values at unfavorable positions.

**Figure H.1:** Expected next-period rent given the current draw, when the household keeps its unit and when it enters a new one




*Notes.* One-period renter problem of Section 3.6 at the estimated within-tenancy process. Panel (a): $E[z' \mid z]$, equal to $\rho_q z$ ($\rho_q = 0.830$) if the household keeps its unit and $\rho_{\text{move}} z$ ($\rho_{\text{move}} = 0.446$) if it enters a new one, both as applied per three-year period (the re-draw correlation is estimated from one-wave rent changes, Section 2.6); the shaded area is the gap between them. Panel (b): $E[q' \mid z]$ with $q = q_0 e^z$, in shares of period income. The move innovation is rescaled so that both transitions share the grid and its stationary distribution, so mean rent is 0.206 under either (dotted line). The horizontal axes span the model's five-node grid ($\pm 2\sigma_z$, $\sigma_z = 0.511$); the curves are evaluated at the nodes of a finer Rouwenhorst grid of the same process.

## H.3 Orders at a Fixed Quantile

Fix the renter's quantile, $z = k\sigma_z$ with $k$ constant, and let rent risk shrink by scaling $\sigma_\eta$; $\sigma_z$ and $\sigma_{\text{move}}$ scale with it. Both conditional variances in (H.5) are proportional to $\sigma_z^2$, so expanding the two exponentials gives

$$\text{Lock}(k\sigma_z) = q_0 \, (\rho_{\text{move}} - \rho_q) \, k \, \sigma_z + O(\sigma_z^2). \tag{H.7}$$

The forfeiture wedge is first order in $\sigma_z$ and positive at favorable quantiles ($k < 0$). For insurance, a second-order expansion of (21) around expected consumption $\bar{c} = y - E[q' \mid z, \text{stay}]$ gives the Arrow–Pratt premium

$$\text{Ins}(z) = \frac{\text{RRA}_c(\bar{c})}{2\,\bar{c}} \, \text{Var}(q' \mid z, \text{stay}) + o(\sigma_z^2), \qquad \text{Var}(q' \mid z, \text{stay}) = q_0^2 e^{2\rho_q z} \sigma_\eta^2 + o(\sigma_z^2), \tag{H.8}$$

so insurance is second order. The same expansion applies to the premium of moving in (H.1), with the conditional variance of the move chain. Here $\text{RRA}_c$ is the relative risk aversion in consumption of the renter's utility (8),

$$\text{RRA}_c(c) = \gamma \, \pi(c) + (1 - \nu)\big(1 - \pi(c)\big), \qquad \pi(c) = \frac{\alpha c^\nu}{\alpha c^\nu + \psi_0 (1 - \alpha) h^\nu}, \tag{H.9}$$

the expression of Section 3.6 rearranged as a weighted average of $\gamma$ and $1-\nu$, with the consumption share $\pi$ of the composite as the weight on $\gamma$. At $c = y - 0.206$ it is 2.12, below $\gamma = 2.33$ because $1-\nu < \gamma$. The renter housing weight $\psi_0$ enters insurance only through $\pi$. At a fixed quantile, the ratio of the forfeiture wedge to insurance therefore rises in inverse proportion to $\sigma_z$ as rent risk shrinks. At a fixed $z$ the leading term of the wedge does not shrink at all (Section 3.6), so the fixed-quantile comparison is the one less favorable to the wedge.

The expansions describe limits; the values reported in Section 3.6 use exact certainty equivalents on the five-node grid. Because the forfeiture wedge contains no preference parameter, the grid also gives the curvature at which insurance would match it. Solving $\mathrm{Ins}(z;\gamma) = \mathrm{Lock}(z)$ for $\gamma$, with $\alpha$, $\nu$ and $\psi_0$ at their values in Table 4, gives 14 one standard deviation below the mean position and 19 at the most favorable node, against the estimated 2.33.

### H.4 A Numerical Check of the Orders

Table H.1 evaluates both forces with the innovation standard deviation scaled to $m\sigma_\eta$, for a multiple $m$ from 0.25 to 2.00 (in this subsection $m$ denotes the multiple, not measurement error). For each $m$ the grid and both transitions are rebuilt, so that $\sigma_z$ and $\sigma_{\text{move}}$ scale with $m$; $q_0$ is reset so that mean rent stays at 0.206; and both forces are evaluated one standard deviation below the mean position, $z = -\sigma_z(m)$, a fixed quantile. The check stops at twice the estimated innovation, because at a slightly larger multiple rent at the highest node would exceed income.

Relative to the smallest multiple, both forces rise by less than exact scaling implies at every multiple. Exact scaling holds only in the limit. At finite $m$, rent at the evaluation point falls as $m$ rises, because the point moves further below the mean and $q_0$ falls to keep mean rent fixed as dispersion widens; this lowers both forces by more than the higher-order terms of the expansions raise them. The rise is not even across the range. Each doubling of $m$ multiplies the forfeiture wedge by less than two. Insurance rises by less than a factor of four in each of the first two doublings and by more than four from the estimated innovation to twice it. At twice the innovation, rent at the highest node is close to income. That node is reached from $z = -\sigma_z(m)$ with small probability, but consumption there is close to zero, and most of the premium comes from it rather than from the local curvature that (H.8) uses. The last row of the table therefore lies outside the range of the expansion: its ratio of the wedge to insurance, 4, compares the wedge with a premium driven by the edge of the grid. Across the table the ratio falls monotonically, from 41 at the smallest multiple to 4 at the largest.

The first limit stated in Section 3.6 is that the calculation covers one period. Over a renting spell the two forces accumulate differently. After a move, the gap in expected log rent $j$ periods later, for a household that then keeps its new unit, is $\rho_q^{j-1}(\rho_{\text{move}} - \rho_q)z$, which shrinks by the factor $\rho_q$ each period. The risk premium is paid in every period the household rents, and the spread of future rent widens with the horizon (Figure 5). Saving works in the other direction: a renter with a buffer

**Table H.1:** Insurance and the forfeiture wedge as the rent innovation is scaled

| | Insurance (level $\times 10^3$) | | | Forfeiture wedge (level $\times 10^3$) | | | |
|---|---|---|---|---|---|---|---|
| Multiple $m$ | Level | Relative | $(m/m_0)^2$ | Level | Relative | $m/m_0$ | Wedge / insurance |
| 0.25 | 0.24 | 1.00 | 1 | 10.04 | 1.00 | 1 | 41 |
| 0.50 | 0.82 | 3.38 | 4 | 19.72 | 1.96 | 2 | 24 |
| 1.00 | 2.34 | 9.68 | 16 | 36.80 | 3.67 | 4 | 16 |
| 2.00 | 14.12 | 58.29 | 64 | 58.11 | 5.79 | 8 | 4 |

*Notes.* The certainty-equivalent-rent comparison of Section 3.6, one period ahead, on the five-node grid of the within-tenancy rent position $z$, with rent proportional to $e^z$ and scaled so that mean rent is 0.206 of period income; income and housing are normalized to one; one period is three years. The innovation standard deviation is $m\sigma_\eta$, with $\sigma_\eta = 0.285$ per three-year period, persistence 0.830 within a tenancy and 0.446 on entering a new unit; the move innovation is rescaled so that both transitions share the grid and its stationary distribution. Each row is evaluated one grid step below the center, $z = -\sigma_z(m)$, a fixed quantile of the stationary distribution. Insurance: the risk premium of staying, the certainty-equivalent rent minus the expected rent, $R^*_{\text{stay}}(z) - E[q' \mid z, \text{stay}]$. Forfeiture wedge: the extra rent expected from entering a new unit rather than staying, $E[q' \mid z, \text{move}] - E[q' \mid z, \text{stay}]$. Levels are in thousandths of period income. Relative: the level divided by its value at the smallest multiple $m_0$; $(m/m_0)^2$ and $m/m_0$ are the ratios exact second- and first-order scaling would give. Preferences as in the model ($\gamma = 2.33$, $\nu = 0.60$, $\alpha = 0.80$, $\psi_0 = 0.431$). The renter housing weight $\psi_0$ is common to both options and creates no wedge between them; it affects the insurance premium only through the curvature of utility in consumption (relative risk aversion 2.12 at the mean rent). Larger multiples are not computed because the dearest grid rent would exceed income. Source: calculations from the estimated rent process and preferences.

stock spreads a rent shock over many periods of consumption and bears it at less than the one-period premium, while the forfeiture wedge, a gap in expected rents, does not depend on saving. The one-period ratios therefore do not determine the ratio over a spell, and the calculation does not sign it. Nor does it compare renting with owning, which brings in $\psi_0$, the costs of owning, house-price risk and wealth (Section 3.6). That comparison is made by the life-cycle model.

# I Computation and Welfare Measurement

## I.1 Discretization

The model is solved separately for each of the nine combinations of the three permanent rent types and the three permanent income types (Section 3.2); a type scales the rent grid by its rent factor and the income grid by its income factor. Within a type, the state is $(g, z, p, \omega, \varepsilon; a, h, o)$, as in Section 3.1.

The rent position, the house price and the persistent income component are first-order autoregressions, discretized by the Rouwenhorst method, which matches the conditional and unconditional variances and the first-order autocorrelation of the process for any number of nodes and remains accurate at high persistence (Rouwenhorst, 1995; Kopecky and Suen, 2010). The grid of the rent position has five nodes and that of the house price five. Income combines three nodes of the persistent component with two equally likely nodes of the transitory component, six states in all. The move transition $\Pi_{\text{move}}$ uses the nodes of $\Pi_{\text{stay}}$ with the innovation rescaled as in (H.3), so the two

transitions share the stationary distribution (Appendix H). For each process the levels at the nodes are scaled so that their mean under the discrete stationary distribution equals the normalization of Section 3.2 exactly, with no log-normal correction.

The market component starts at zero at age 27, and its grid widens with age (Fella et al., 2019). In period $t$ its three nodes are zero and $\pm\sqrt{2}\,\sigma_{g,t}$, where $\sigma_{g,t} = \sigma_g(t-1)^{1/2}$ is the standard deviation of the random walk (12) after $t-1$ periods, and the transition from $t$ to $t+1$ is a Rouwenhorst matrix with persistence $\sigma_{g,t}/\sigma_{g,t+1}$. The discrete chain is a martingale with variance $\sigma_{g,t}^2$ in period $t$, as the random walk is. Rent is multiplied by $e^g$ divided by its mean under the period-$t$ distribution of $g$, so the market component spreads rents without changing their mean at any age, and removing it leaves mean rent unchanged.

Net financial assets, in units of mean income at age 27, take 21 values: 20 evenly spaced points and zero, which is added so that a renter can hold exactly zero assets. Dwelling size, in rooms, takes seven geometrically spaced values, so the menu is finer at small sizes; renters and owners choose from the same menu. We have not evaluated how the results depend on the asset and size grids, or on the number of nodes of the rent grid, at the adopted parameters.

### I.2 Solution

We solve the household problem (19) by backward induction from the last period, whose continuation value is the bequest (9). Each period takes three steps.

(i) The value function of period $t+1$ is integrated over next period's exogenous state with the transition matrices of the market component, the rent position, the house price and income, once with $\Pi_{\text{stay}}$ and once with $\Pi_{\text{move}}$ for $z$.

(ii) For each state and each tenure $o'$, the household chooses $(a', h')$ by maximizing flow utility plus discounted expected continuation value over all grid pairs, without interpolation. The continuation value uses $\Pi_{\text{move}}$ when the choice enters a new rental unit ($n_t = 1$ in (14)) and $\Pi_{\text{stay}}$ otherwise, including for every choice to own.

(iii) The value at the start of the period and the probability of owning follow from the logit (20) over the two choice-specific values, with taste-shock scale $\xi = 0.005$.

As in Iskhakov et al. (2017), extreme-value taste shocks smooth the discrete choice; here they act on tenure only, and size and saving are chosen by maximization over the grid. At the adopted parameters, the value function of every type is nondecreasing in net financial assets and in the transitory income state at every pair of adjacent grid points at which both values are feasible.

**Infeasible choices.** A choice that leaves nonpositive consumption or violates the collateral constraint (18) receives a large negative value, as does a terminal position with nonpositive net worth, at which the bequest is not defined. Before the value function of period $t+1$ enters the expectation of step (i), it is bounded below by a finite penalty far below the values of feasible states. The bound

limits, but does not remove, the weight such states carry in the expectation. A next-period state from which no choice is feasible lowers the expected value of every current choice that can reach it, in proportion to the probability of reaching it and to the size of the penalty, which is a numerical choice; because the penalty is far below feasible values, even a small probability of reaching such a state can lower the expected value of a choice substantially. We have not measured how often simulated households reach such states, or how the results depend on the size of the penalty. The model has no consumption floor.

### I.3 Simulation

Each world is simulated with 100,000 lives from age 27 to 81, allocated across the nine types in proportion to their population weights up to integer rounding. The search for the re-anchored benchmark of Section 4.6 instead simulates the same number of lives for every type and weights types by their population weights. At age 27, tenure and net financial assets are drawn as in Section 3.4; initial net financial assets are placed at the nearest grid point, and the persistent income component is drawn from a distribution on its three nodes whose variance is the age-27 variance of Section 3.2. Each period the simulator draws tenure from the logit probabilities, applies the size and saving policies and draws next period's exogenous state. Every household draws its own market history.

Each type has its own random seed, the same in every world and at every point of the parameter scan. Every period the simulator draws the uniform numbers for both rent transitions for every household and uses the one that the household's choice selects, so the stream of random numbers is the same in every world even when choices differ; Section 4.1 describes the resulting pairing of lives. No simulation standard errors are computed. Model statistics weight types by their population weights. Every household lives all 18 periods, so statistics that pool ages weight each age equally.

### I.4 Estimation of the Preference Parameters

Section 3.7 defines the simulated-minimum-distance criterion $L(\vartheta)$, the scan over the box of parameter values and the adopted point, and Table I.1 reports every evaluated point. The estimates of $\vartheta = (\psi_0, \theta_B, \gamma, \nu)$ are the point of the scan with the lowest value of $L$, rounded to the precision of Table 4. Each point of the scan is solved on the grids above and simulated with 50,000 lives, allocated across types in proportion to their population weights except that every type receives a minimum number of lives, which over-samples the four smallest types; the moments weight types by their population weights, as in the worlds. At the rounded point, simulated with 100,000 lives, the loss is 0.186, against 0.187 for the unrounded point in the scan.

The scan adds two facts to those noted in Section 3.7. First, the points ranked second to fourth have losses from 0.216 to 0.220, against 0.187 at the adopted point, at parameters far from it: $\psi_0$ from 0.497 to 0.899, $\theta_B$ from 0.36 to 1.75 and $\nu$ from 0.14 to 0.52. A scan of 10 points cannot trace

**Table I.1:** The quasi-random parameter scan: every evaluated point

| | Parameters | | | | | Targeted moments | | | | | Untargeted |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | Ownership | | Owner | Renter | Renter | Ownership, |
| | $\psi_0$ | $\theta_B$ | $\gamma$ | $\nu$ | Loss | All | Age 30 | NW/Y | Rent/Y | Move | age 69 minus age 30 |
| Adopted | 0.431 | 2.15 | 2.33 | 0.60 | 0.187 | 0.758 | 0.679 | 3.178 | 0.242 | 0.432 | 0.097 |
| | 0.899 | 1.75 | 1.90 | 0.14 | 0.216 | 0.583 | 0.530 | 2.400 | 0.253 | 0.376 | 0.055 |
| | 0.497 | 0.36 | 2.56 | 0.52 | 0.217 | 0.742 | 0.673 | 2.823 | 0.260 | 0.402 | 0.086 |
| | 0.655 | 0.93 | 2.10 | 0.38 | 0.220 | 0.677 | 0.614 | 2.555 | 0.273 | 0.400 | 0.074 |
| | 0.550 | 2.32 | 1.45 | 0.28 | 0.430 | 0.889 | 0.814 | 2.473 | 0.205 | 0.377 | 0.091 |
| | 0.614 | 1.57 | 1.69 | 0.22 | 0.450 | 0.901 | 0.822 | 2.435 | 0.210 | 0.381 | 0.100 |
| | 0.834 | 0.75 | 1.20 | 0.08 | 0.525 | 0.894 | 0.824 | 1.909 | 0.235 | 0.399 | 0.089 |
| | 0.720 | 2.99 | 2.82 | 0.46 | 0.529 | 0.577 | 0.528 | 3.206 | 0.329 | 0.387 | 0.047 |
| | 0.627 | 1.88 | 2.98 | 0.10 | 0.651 | 0.965 | 0.886 | 3.036 | 0.171 | 0.402 | 0.108 |
| | 0.771 | 0.49 | 1.73 | 0.55 | 1.404 | 0.436 | 0.409 | 2.511 | 0.422 | 0.397 | 0.034 |
| PSID data | | | | | | 0.731 | 0.537 | 2.707 | 0.206 | 0.563 | 0.328 |

*Notes.* Each row is one of the 10 points of a scrambled Sobol sequence over the box $\psi_0 \in [0.40, 0.90]$, $\theta_B \in [0.20, 3.00]$, $\gamma \in [1.10, 3.00]$ and $\nu \in [0.05, 0.60]$; rows are sorted by loss, and these are all the points evaluated. Each point is solved on a grid of 21 asset (20 plus zero) and 7 dwelling-size points and simulated with 50,000 lives, with the same random numbers at every point. The loss is the equal-weighted sum of squared relative gaps between the five simulated moments and their data values in the last row (definitions in Table 5): the ownership rate over all ages and at age 30, owners' median net worth to income (NW/Y), renters' median rent to income (Rent/Y) and the renter move rate per three years. The estimated parameters are the first row rounded (Table 4); simulated again with 100,000 lives, their loss is 0.186. No local search follows the scan and no standard errors are computed. The last column, the ownership rate at age 69 minus the rate at age 30, is not targeted; at every evaluated point it is below its data value. It differs from the age gradient of Table 6, which compares age groups.

the shape of a loss over four parameters, and the counterfactuals of Section 4 are computed at the adopted point only. Second, at every evaluated point the rise in ownership from age 30 to age 69, which is not targeted, is below its PSID value of 0.328.

## I.5 Welfare Measurement

The consumption-equivalent variation $\lambda_i(S)$ of (25) is computed household by household from the paired lives. The left side of (25) increases in $\lambda$, so the solution is found by bisection on $[-0.95, 50]$, that is, for a change in baseline consumption between $-95$ and $+5{,}000$ percent. The interval is halved a fixed number of times, enough to make its width negligible at the reported precision. A household whose baseline life, with consumption scaled by the lower end of the interval, is still worth at least as much as its life in world $S$ is recorded at that end. When all rental-rate risk is removed, 0.58 percent of households are recorded there, and for them the recorded value understates the loss from moving to world $S$. The same computation gives the mean welfare effects of the three-risk worlds (Table J.2) and of portable rent positions. As in Section 4.1, means weight

households by the population weights of their types and percentiles are unweighted; the cut-offs of the consumption quintiles are computed with the population weights.

# J Additional Counterfactual Results

This appendix reports every world behind the decompositions of Sections 4.2 and 4.4. All worlds hold the preference parameters at their estimates, remove risk as in Sections 4.1 and 4.4, and simulate 100,000 lives paired by common random numbers (Appendix I).

## J.1 Every Combination of Rental-Rate Risk Components

Table J.1 reports the ownership rate in the eight worlds formed by removing any subset of the three components of rental-rate risk, the market component ($M$), the within-tenancy component ($I$) and the re-entry re-draw ($R$), and in world $F$, in which each household's draw of $z$ is frozen for life. The simulations record welfare only for the world with all three components removed (Section 4.5) and for world $R$, which is the portable-position experiment of Section 4.7 and reproduces its ownership rate exactly, so the table reports ownership alone.

Three features of the table carry the decompositions of Section 4.2. First, world $I$ sets the factor $e^z$ to its stationary mean for every household, keeping the permanent rent types, so all positions are equal and a re-draw changes nothing: $\text{own}(\{I,R\}) = \text{own}(\{I\}) = 0.788$ and $\text{own}(\{M,I,R\}) = \text{own}(\{M,I\}) = 0.791$, exactly. Second, the re-draw's marginal effect depends on the order of removal. It is largest when the re-draw is switched off first, $+0.82$ points or 25 percent of the total effect of $+3.29$ points, smaller once the market component is removed, and zero once the within-tenancy component is removed. In no order of removal does the re-draw account for more than 25 percent of the total. Third, the market component matters little on its own and more once positions are equal. Removing it raises ownership by $+0.05$ points from the baseline and by $+0.31$ points from world $I$. The second number equals the last step of the sequential decomposition (26), because the re-draw has no further effect once $I$ is removed; the two earlier steps are the frozen-draw step ($-0.36$ points) and the removal of the dispersion of positions ($+3.33$ points).

The Shapley contributions of Table 7 average each component's marginal effect over the orders of removal (Shorrocks, 2013),

$$\text{Sh}_j = \sum_{S \subseteq C \setminus \{j\}} \frac{|S|!\,\big(|C| - |S| - 1\big)!}{|C|!} \big[\text{own}(S \cup \{j\}) - \text{own}(S)\big], \tag{J.1}$$

where $C = \{M, I, R\}$ and $\text{own}(\varnothing)$ is baseline ownership. The weight is $\frac{1}{3}$ when $S$ is empty or contains both other components and $\frac{1}{6}$ when it contains one, and the contributions $\text{Sh}_j$ sum to the total. Sums of displayed values may differ from displayed totals in the last digit; the identities hold for the unrounded values.

**Table J.1:** Homeownership in every combination of rental-rate risk components removed

| | Component removed | | | | |
|---|---|---|---|---|---|
| World | Market | Within-tenancy | Re-draw | Ownership rate | Change (pp) |
| Baseline | | | | 0.758 | — |
| Market removed | ✓ | | | 0.758 | +0.05 |
| Within-tenancy removed | | ✓ | | 0.788 | +2.97 |
| Re-draw off (portable positions) | | | ✓ | 0.766 | +0.82 |
| Market and within-tenancy removed | ✓ | ✓ | | 0.791 | +3.29 |
| Market removed, re-draw off | ✓ | | ✓ | 0.765 | +0.72 |
| Within-tenancy removed, re-draw off | | ✓ | ✓ | 0.788 | +2.97 |
| All three removed | ✓ | ✓ | ✓ | 0.791 | +3.29 |
| Each draw $z$ frozen | | $z$ uncertainty off | | 0.754 | −0.36 |

*Notes.* Ownership rate: share of simulated household-periods in which the household owns, over all ages; change relative to the baseline in percentage points (pp). A check mark indicates that the rental-rate risk component is removed, as defined in Table 7. Once the within-tenancy component is removed all rent positions are equal, so switching off the re-draw has no further effect. Each draw frozen: each household keeps its age-27 draw of the idiosyncratic rent position $z$ for life ($\Pi_{\text{stay}} = \Pi_{\text{move}} =$ identity), which removes uncertainty about $z$ and the re-draw but keeps the dispersion of positions across households and the market component (the Market cell is blank). The Shapley contributions in Table 7 are computed from the worlds above the rule. Estimated parameters, 100,000 lives in each world, common random numbers. Source: model simulations.

## J.2 Every Combination of the Household's Three Risks

Table J.2 reports ownership and the mean consumption-equivalent variation in the eight worlds formed by removing any subset of house-price risk, rental-rate risk (all three components) and income risk, each removed as in Section 4.4. The world without rental-rate risk is the world with all three components removed in Table J.1. The table supplies the Shapley decomposition of Table 9 and the four marginal effects of rental-rate risk that Section 4.4 uses to reconcile its Shapley contribution with the headline effect.

The pattern of Section 4.4, in which house-price risk governs tenure and income risk governs welfare, holds world by world. The lowest ownership rate among the four worlds without house-price risk is 0.927, in the world without house-price and income risk; the highest among the four worlds that keep it, the baseline included, is 0.791, in the world without rental-rate risk. The lowest mean CEV among the four worlds without income risk is 42.59 percent, again in the world without house-price and income risk; the highest among the four that keep it is 1.08 percent, in the world without rental-rate risk. Removing house-price risk lowers the mean CEV in each of the four combinations of the other two risks. The pattern is consistent with the buying opportunity that a stationary price process offers (Section 4.4), a channel we have not isolated.

**Table J.2:** Homeownership and welfare in every combination of the three risks removed

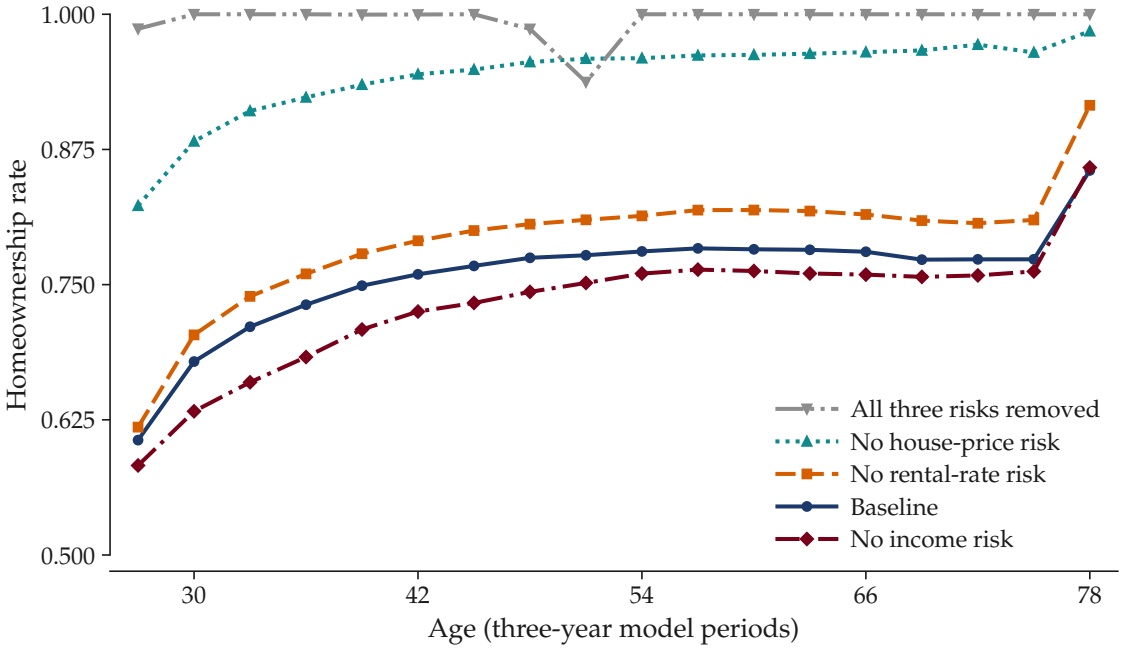
| | Risk removed | | | | | |
|---|---|---|---|---|---|---|
| World | House price | Rental rate | Income | Ownership rate | Change (pp) | Mean CEV (percent) |
| Baseline | | | | 0.758 | — | — |
| House-price risk removed | ✓ | | | 0.943 | +18.56 | −13.90 |
| Rental-rate risk removed | | ✓ | | 0.791 | +3.29 | +1.08 |
| Income risk removed | | | ✓ | 0.731 | −2.66 | +61.63 |
| House-price and rental-rate | ✓ | ✓ | | 0.984 | +22.57 | −13.09 |
| House-price and income | ✓ | | ✓ | 0.927 | +16.94 | +42.59 |
| Rental-rate and income | | ✓ | ✓ | 0.781 | +2.32 | +60.77 |
| All three risks removed | ✓ | ✓ | ✓ | 0.995 | +23.72 | +43.43 |

*Notes.* A check mark indicates that the risk is removed, as defined in Table 9; rental-rate risk comprises all three of its components, so the world without it is the world without all three components in Table 7. Ownership rate: share of simulated household-periods in which the household owns, over all ages; change relative to the baseline in percentage points (pp). Mean CEV: mean over households of the consumption-equivalent variation of moving from the baseline to the world, defined as in Table 7; positive values are gains. The Shapley contributions in Table 9 are computed from these worlds. Estimated parameters, 100,000 lives in each world, common random numbers. Source: model simulations.

## J.3 Ownership by Age

Figure J.1 plots ownership by age in the baseline, in the three worlds that remove one risk and in the world that removes all three. Removing rental-rate risk raises ownership at every age, and removing house-price risk raises it by more at every age. Removing income risk lowers ownership at every age but the last. Without any of the three risks, ownership is close to one at most ages and at least 0.937 at every age, the minimum occurring at age 51. In each of the four worlds that keep some risk, ownership rises at the last age, the period whose continuation value is the bequest (9), which counts an owned house at its full price.

**Figure J.1:** Homeownership by age when each of the household's three risks is removed




*Notes.* Simulated share of households that own, by age, in the estimated economy (baseline), in the three economies that remove one risk and in the economy that removes all three (Table J.2). Rental-rate risk comprises the market, within-tenancy and re-draw components; income risk comprises the persistent and transitory income components, with permanent income types and the age profile of income kept. Ages are the first year of each three-year model period, from 27 to 78. Estimated parameters, 100,000 lives in each economy, common random numbers.