TCER Working Paper Series

Beyond the Bubble: Empirical Evidence on Asset Pricing under Persistent Low Interest Rates

Chihiro Shimizu

March 2026

Working Paper E-227
https://www.tcer.or.jp/wp/pdf/e227.pdf

TOKYO CENTER FOR ECONOMIC RESEARCH
1-7-10-703 Iidabashi, Chiyoda-ku, Tokyo 102-0072, Japan

Abstract

When the real interest rate falls below expected asset returns, the Bubble Necessity Theorem (Hirano and Toda, 2025) implies that high valuations are structurally necessary rather than speculative. We provide the first empirical test. Following Shiller—who built price indices to test excess-volatility theory—we construct synchronized quality-adjusted price and rent indices from 11 million listings across 40 years in Tokyo: the city of the twentieth century's largest housing bubble and the world's longest near-zero-rate episode. A cointegrated VECM yields an expectation-to-rate elasticity of 2.92 across five proxies. The post-2013 Necessary Regime transfers JPY 1.96 million per year from buyers to owners; raising the property tax from 1.4% to 3.0% would have prevented it.

Chihiro Shimizu
TCER
and
Hitotsubashi University
School of Social Data Science
2-1, Naka, Kunitachi, Tokyo
c.shimizu@r.hit-u.ac.jp

Beyond the Bubble:

Empirical Evidence on Asset Pricing under Persistent Low Interest Rates

Chihiro Shimizu [*]

March 2026

Abstract

When the real interest rate falls below expected asset returns, the Bubble Necessity Theorem [Hirano and Toda, 2025] implies that high valuations are structurally necessary rather than speculative. We provide the first empirical test. Following Shiller—who built price indices to test excess-volatility theory—we construct synchronized quality-adjusted price *and rent* indices from 11 million listings across 40 years in Tokyo: the city of the twentieth century's largest housing bubble and the world's longest near-zero-rate episode. A cointegrated VECM yields an expectation-to-rate elasticity of 2.92 across five proxies. The post-2013 Necessary Regime transfers JPY 1.96 million per year from buyers to owners; raising the property tax from 1.4% to 3.0% would have prevented it.

Keywords: Bubble Necessity Theorem; Financial user cost; Asset pricing; Low interest rates; Housing markets; Cointegration; Welfare analysis.

JEL Codes: G12, E44, R31, C32, E43, D61.

[*]Professor, School of Social Data Science, Hitotsubashi University, 2-1 Naka, Kunitachi, Tokyo 186-8601, Japan. E-mail: c.shimizu@r.hit-u.ac.jp.

1 Introduction

The measurement imperative in empirical asset pricing. In 1981, Robert Shiller published a deceptively simple observation: if stock prices are rationally discounted present values of future dividends, observed prices should not be more volatile than dividends themselves. The finding that they are—the "excess volatility" puzzle—became one of the most influential results in modern finance [Shiller, 1981]. Yet what made the finding scientifically credible was not the theory alone. It was the measurement infrastructure: Shiller and Case's [1987] construction of micro-data-based repeat-sales price indices, which for the first time made housing-market dynamics precisely and continuously observable. Theory without measurement is unverifiable; measurement without theory is uninterpretable. Shiller understood that empirical asset pricing requires both.

This paper follows the same logic—and goes further. Shiller's theory requires a comparison between observed prices and fundamental values, which in housing markets requires *both* a price index and a rent index. Yet Shiller's repeat-sales method constructs only price indices; it cannot construct rent indices because rental properties typically transact at the end of tenancy rather than by continuous resale. As a result, the most important object in housing finance—the price-to-rent ratio— has never been measured at the same quality, frequency, and length as price indices alone. We close this gap. Using 11 million sale and rental listings from Recruit Co. Ltd. (Japan's largest real estate platform), we construct *synchronized* quality-adjusted price and rent indices from the same rolling-window hedonic surface at monthly frequency, continuously from January 1986 to December 2025—a 40-year record with no counterpart in any other country (Appendix E). This measurement infrastructure is the empirical foundation of the paper, and it is a contribution in its own right: the first 40-year, monthly, micro-data-based pair of housing price and rent indices for a major global city.

The question these data answer. With a precise, long-run price-to-rent ratio in hand, we can ask a question that has been central to asset pricing theory for four decades but has never been tested with adequate data: *what happens to asset prices when the real interest rate persistently falls below the expected growth rate of asset payoffs?*

The Gordon growth model gives a simple answer: the fundamental price formula $P^F = D/(r-g)$ diverges when $r \leq g$, and no finite bubbleless equilibrium exists. Hirano and Toda [2025] formalize this as the *Bubble Necessity Theorem*: in general equilibrium, whenever $r < g$, every competitive equilibrium must contain a positive bubble component. High valuations are not speculative excess but the *only* kind of equilibrium available. Under this condition, the standard Shiller excess-volatility test is methodologically inapplicable: there is no bubbleless benchmark to compare against.

The secular decline in real interest rates documented by Rachel and Summers [2019] makes this question empirically urgent for the twenty-first century. A decade of near-zero or negative policy rates across the world's major economies—with housing prices rising sharply in their wake— has made the $r < g$ condition a plausible description of actual markets, not merely a theoretical boundary case. Hirano and Toda [2025] provide the theory; what has been missing is a rigorous empirical test. This paper provides one, for the first time.

Tokyo as the world's uniquely qualified laboratory. Testing the Bubble Necessity Theorem demands three things that almost no housing market in the world can supply simultaneously: (i) a long sample spanning *both* the Fundamental Regime ($r > g$) and the Necessary Regime ($r < g$), so that the regime transition can be identified; (ii) the measurement infrastructure described above—synchronized price and rent indices at monthly frequency over the full sample; and (iii) sufficient institutional and policy variation to separately identify the role of interest rates, expectations, and fiscal instruments. Tokyo is, to our knowledge, the only major city in the world that satisfies all three requirements.

The reason is twofold, and each leg of the argument is historically extraordinary.

First: Tokyo experienced the largest housing bubble of the twentieth century. Between 1986 and 1991, quality-adjusted condominium prices in the Greater Tokyo area more than tripled, driven by extraordinary extrapolative expectations ($\hat{\pi}^e \approx 20\%$ per annum). The subsequent collapse erased more than 60% of housing values and inaugurated two decades of deflation and stagnation—the "Lost Decades." No other major city in the world experienced a housing cycle of this magnitude and duration in the postwar period. Crucially, during the bubble, the Financial User Cost (FUC)—the net annual holding-cost rate $r + \delta + \tau + m - \hat{\pi}^e$—turned deeply negative (approximately -13%) because of extreme appreciation expectations, placing Tokyo in the Necessary Regime via the expectation channel. This gives us a clean historical "treatment observation" of the theorem operating through the demand side.

Second: Japan implemented the world's most radical and prolonged near-zero interest rate policy. The Bank of Japan introduced the zero interest rate policy in 1999, experimented with quantitative easing from 2001, and under Abenomics from 2013 deployed a combination of negative policy rates (-0.1%) and yield curve control (YCC) that kept long-term rates near zero until 2024. No other major economy maintained rates at or below zero for a comparable duration. As a result, from approximately 2013, the FUC for the typical mixed-finance household turned negative again—but now through the interest-rate channel rather than the expectation channel—placing Tokyo in the Necessary Regime for a second time, via an entirely different mechanism.

The 40-year sample therefore contains two distinct, empirically separable manifestations of the Bubble Necessity Theorem operating in the same city: a deep expectation-driven episode in the 1980s, and a shallow rate-driven episode from 2013 onward, separated by two decades of the Fundamental Regime. No other housing market in the world offers this natural experiment. The within-city structure eliminates the institutional confounders that plague cross-country comparisons of housing markets.

Our approach: the Financial User Cost as regime diagnostic. We operationalize the Bubble Necessity Theorem through a single observable:

$$R_t \;=\; P_t^H \cdot \underbrace{\left[r_t + \delta + \tau + m - \hat{\pi}_t^e\right]}_{\equiv \; \mathrm{FUC}_t}, \tag{$\star$}$$

where r_t is the real cost of capital (weighted by financing structure), δ physical depreciation, τ the property tax rate, m maintenance costs, and $\hat{\pi}_t^e$ the expected appreciation rate. Every component is observable or precisely estimable from our data. The sign of FUC_t is a sufficient statistic for the equilibrium regime:

- $\mathrm{FUC}_t > 0$ (Fundamental Regime): a finite bubbleless price exists; Shiller excess-volatility tests are methodologically valid.
- $\mathrm{FUC}_t < 0$ (Necessary Regime): the Bubble Necessity Theorem applies; every equilibrium contains a bubble component, and Shiller-style diagnostics are methodologically inapplicable.

The rent index R_t enters the diagnostic ($\star$) directly: without a high-quality rent index, FUC cannot be computed, and the regime diagnostic cannot be constructed. This is why the measurement infrastructure—the synchronized price-and-rent indices of Appendix E—is not merely a data description but a precondition for the entire empirical exercise.

The global context. Tokyo's experience is not unique in kind—only in duration and intensity. During the global zero-rate era of 2020–2021, the FUC turned negative simultaneously in Vancouver, London, Sydney, Frankfurt, Amsterdam, and Seoul (Table 1). The Bubble Necessity Theorem was therefore *binding* in most major housing markets for at least one or two years. Tokyo's episode, spanning more than a decade, provides the statistical power needed to estimate the theorem's quantitative implications precisely and to identify the regime-transition mechanism cleanly.

Identification and expectation robustness. A key potential concern is the measurement of $\hat{\pi}_t^e$. We construct five proxies—HP-filter trend, one-sided HP filter, Hamilton [2018] regression filter, 12-month moving average, and AR(1) forecast—and show that the expectation-to-rate elasticity ratio lies between 2.33 and 2.92 across all five (Table D.4). Even the Hamilton filter, specifically designed to be orthogonal to low-frequency price trends [Hamilton, 2018], confirms the finding. This robustness rules out the concern that

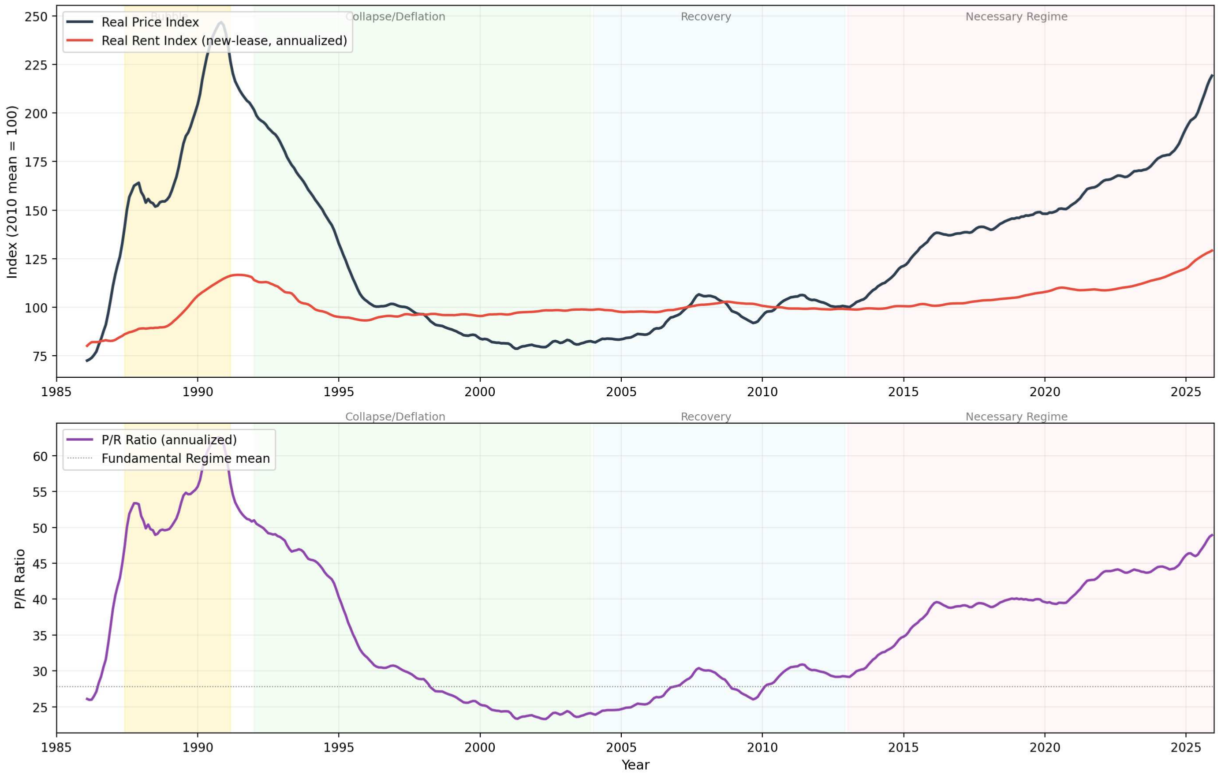

Figure 1: Housing Prices, Rents, and the Annualized P/R Ratio in Tokyo (Monthly, 1986–2025).

Notes: Quality-adjusted condominium price and rent indices constructed from Recruit Co. Ltd. listing data using the rolling-window hedonic method (Appendix E), normalized to 2010 mean = 100. Prices: approximately 4.8 million sale listings. Rents: approximately 6.2 million new-lease rental listings. The P/R ratio is the annualized ratio $P_t^H/(12R_t^{\mathrm{mo}})$. Four phases: bubble (1986–91), collapse and deflation (1992–2003), moderate recovery (2004–12), and renewed escalation (2013–25). By end-2025 the P/R ratio approaches its 1991 bubble peak (≈49 years).

Table 1: Approximate Financial User Cost in Major Housing Markets

City	Country	Period	P/R (yrs)	FUC (%)	Regime
Panel A: Global zero-rate era (2020–2021), $\hat{\pi}^e \approx 7\%$ p.a.					
Tokyo	Japan	2013–2024	40–49	−2 to 0	Necessary (10 yrs)
Vancouver	Canada	2020–2022	42–48	−4	Necessary (2 yrs)
London	UK	2020–2021	38–42	−2	Necessary (1 yr)
Sydney	Australia	2020–2021	36–40	−3	Necessary (1 yr)
Frankfurt	Germany	2020–2022	33–38	−4	Necessary (2 yrs)
Amsterdam	Netherlands	2020–2022	36–40	−5	Necessary (2 yrs)
New York	USA	2020–2021	30–35	−1	Necessary (1 yr)
Panel B: Post-tightening era (2024), $\hat{\pi}^e \approx 3\%$ p.a.					
Tokyo	Japan	2024	48	+1	Fundamental (exit)
Vancouver	Canada	2024	42	+6	Fundamental
London	UK	2024	38	+6	Fundamental
Hong Kong	HK	2024	50	+7	Fundamental

Notes: FUC $\approx r_t + \delta + \tau + m - \hat{\pi}^e_t$ where r_t = local policy rate, $\delta + \tau + m = 3.9\%$ p.a. (Tokyo calibration). P/R data from Bank for International Settlements [2023], OECD Housing Prices database, and national central banks. Illustrative approximations only; the paper's empirical analysis uses the precisely constructed FUC series for Tokyo. Tokyo's episode is uniquely long (2013–2024) because of Japan's unprecedented duration of near-zero policy rates.

the result is an artefact of mechanically correlated expectation proxies. Appendix D provides the full set of checks. The VECM's identification is discussed in Appendix C; Granger tests confirm that FUC_t leads $\ln P^H_t$ at all horizons, not the reverse.

Contributions. We make two fundamental contributions and three additional ones.

Fundamental contribution 1: The measurement infrastructure. We construct the first 40-year, monthly, synchronized housing price *and rent* index pair for a major global city, based on micro-data (11 million listings). This extends the Shiller–Case measurement tradition in a critical direction: whereas repeat-sales indices measure prices only, our rolling-window hedonic method produces both price and rent indices from the same hedonic surface, eliminating compositional biases and enabling construction of the price-to-rent ratio at the precision and frequency required to identify the FUC sign change. Without rent indices of this quality, the Bubble Necessity Theorem cannot be empirically tested at all. Details are in Appendix E.

Fundamental contribution 2: First empirical test of the Bubble Necessity Theorem. Using the FUC diagnostic, we show that Tokyo entered the Necessary Regime twice during the 40-year sample: once in the 1980s via extreme expectations, and once after 2013 via near-zero interest rates. We prove that $\text{FUC}_t < 0$ is equivalent to the Hirano–Toda condition $\mathcal{G}^H > \mathcal{R}$ (Lemma 2.5), derive the regime-depth parameter $d_t = |\text{FUC}_t|$ as a measure of the theorem's binding force (Theorem 2.9), and show that the two episodes—though producing similar P/R ratios— operated through fundamentally different mechanisms with fundamentally different policy sensitivities.

Additional contribution 3: VECM evidence on equilibrium and expectation dominance. On 480 months of synchronized data we identify a single cointegrating relation ($r = 1$, $p = 6$, Hansen SupF $p = 0.18$), with prices as the primary error-correcting variable ($\hat{\alpha}_P = -0.006$, $t_{\text{HAC}} = -3.20$). Counterfactuals yield an expectation-to-rate elasticity ratio of 2.92, robust across five proxies.

Additional contribution 4: Calibration, regime dating, and policy counterfactuals. We calibrate FUC_t for three household archetypes and show that a property tax increase from 1.4% to 3.0% would have prevented the post-2013 regime entirely. The tax instrument is uniquely effective: it shifts FUC_t by exactly one basis point for every household regardless of leverage.

Additional contribution 5: Welfare analysis. We derive CV, CG, and ΔR welfare measures and prove the accounting identity $\text{CV}_t - \text{EV}_t = -\Delta R_t$ (Proposition 5.8). The post-2013 regime has transferred JPY 196 ($\times$10,000) per year from buyers to incumbent owners of a 60 m^2 unit, with the price-to-income ratio reaching 11.6×—near the 1991 bubble peak.

Related literature. The paper connects four bodies of literature.

Measurement: beyond Shiller–Case. Case and Shiller [1987] construct repeat-sales price indices to test Shiller's [1981] excess-volatility theory. Our rolling-window hedonic method [Shimizu et al., 2010, Diewert and Shimizu, 2016] extends this tradition by producing *both* price and rent indices from the same micro-data, making the price-to-rent ratio—the key object for user-cost theory—as precisely observable as prices alone. Diewert and Shimizu [2020] and Diewert and Shimizu [2015] develop the welfare-consistent imputed-rent foundations.

Bubble necessity and low-rate asset pricing. Hirano and Toda [2025] prove the Bubble Necessity Theorem; Hirano and Toda [2024] and Hirano [2026] extend the framework. The secular decline in real rates is documented by Rachel and Summers [2019]. Our contribution is to take the theorem to data via an observable diagnostic.

User cost and housing valuation. The user-cost approach originates with Jorgenson [1963] and Poterba [1984], extended by Diewert [1976]. Verbrugge [2008] documents the U.S. rent–user-cost divergence as a puzzle; our framework resolves it as a regime-transition signal. Glaeser [2025] critiques conventional user-cost analysis for neglecting credit frictions—an objection our WACC formulation addresses.

International housing and welfare. The superstar-city phenomenon is documented by Glaeser and Gyourko [2003], Gyourko et al. [2013a], and Gyourko et al. [2013b]. Table 1 shows that the $r < g$ mechanism operates broadly across major cities in low-rate environments. The welfare analysis draws on Varian [1992]; distributional consequences of housing booms are studied by Gervais [2002] and Chambers et al. [2009].

Plan of the paper. Section 2 develops the FUC theory and regime diagnostics. Section 3 presents data construction and VECM estimation. Section 4 performs calibration, regime dating, and policy counterfactuals. Section 5 provides the welfare and distributional analysis. Section 6 concludes. Appendices contain FUC micro-foundations (Appendix A), the Bubble Necessity proof (Appendix B), VECM diagnostics and identification (Appendix C), expectation robustness (Appendix D), hedonic index construction—the paper's primary measurement contribution (Appendix E), PSY bubble tests (Appendix F), and the nominal–real comparison (Appendix G).

2 Theory: Financial User Cost and Equilibrium Regimes

2.1 The P/R Puzzle and the Limits of Present-Value Diagnostics

The standard present-value framework for housing assets [Shiller, 1981, Campbell and Shiller, 1988] posits a fundamental price $P_t^F = \mathbb{E}_t \sum_{s=1}^{\infty} R_{t+s}/(1+r)^s$ that is well-defined only when the discount rate r strictly exceeds the growth rate g of rents. In that case, deviations of the observed price from P_t^F are diagnostic of speculative overvaluation—a "bubble." Two problems arise when this framework is applied to markets such as Tokyo since 2013. First, with prolonged ultra-low interest rates and recovering appreciation expectations, the maintained inequality $r > g$ may fail, making the fundamental price P_t^F undefined. Second, even if a finite fundamental can be computed, the Shiller-style excess-volatility test is a joint test of rational expectations, constant discount rates, and a specific dividend process; it provides no constructive explanation of *why* the P/R ratio is high or what structural forces sustain it.

We take a different route. Rather than comparing prices against a model-dependent fundamental benchmark, we construct the *Financial User Cost* (FUC)—an observable holding-cost rate—and use its sign as a sufficient statistic for the equilibrium regime. This approach is constructive: it tells us not only *whether* prices deviate from fundamentals, but *why* the notion of a unique bubbleless equilibrium may itself be invalid.

2.2 No-Arbitrage and the Financial User Cost

Housing simultaneously delivers a consumption flow (housing services) and a capital gain (or loss) from price appreciation. We work throughout in **annual rates**: all cost-of-carry components (r_t, δ, τ, m) and expected appreciation $(\hat{\pi}_t^e)$ are expressed as percentages per annum. The monthly rent R_t^{mo} is observed in the data; the annualized rent is $R_t \equiv 12\, R_t^{\text{mo}}$.

Remark 2.1 (Convention: units and frequency)**.** *All* rates in this paper—interest rates (r, i), depreciation (δ), tax (τ), maintenance (m), and expected appreciation ($\hat{\pi}^e$)—are expressed as **annual decimal rates** in equations (e.g., 5% per annum = 0.05). Tables and figures report rates in **percent** for readability. The FUC rate FUC_t is entered into the VECM in *percentage* units (i.e., multiplied by 100), so that a 1 percentage-point change in FUC corresponds to a unit change in the VECM variable. The annualized price-to-rent ratio is $P_t^H/R_t = P_t^H/(12\,R_t^{\mathrm{mo}})$.

In a frictionless equilibrium the no-arbitrage condition equates the *annualized* rental yield to the opportunity cost of ownership:

$$\frac{R_t}{P_t^H} \;=\; \frac{12\,R_t^{\mathrm{mo}}}{P_t^H} \;=\; r_t + \delta + \tau + m - \hat{\pi}_t^e, \tag{1}$$

where r_t is the real cost of capital (annual, WACC), δ is the physical depreciation rate (annual), τ is the property tax rate (annual), m is the maintenance cost rate (annual), and $\hat{\pi}_t^e$ is the expected rate of capital appreciation (annual). Both sides of (1) have the dimension of *percent per annum*, eliminating the frequency mismatch that would arise if monthly rents were divided by annual rates.

Definition 2.2 (Financial User Cost)**.** The **Financial User Cost** (rate) is the net holding-cost rate, expressed in annual terms:

$$\mathrm{FUC}_t \;\equiv\; \underbrace{r_t + \delta + \tau + m}_{\text{cost of carry (p.a.)}} \;-\; \underbrace{\hat{\pi}_t^e}_{\text{expected capital gain (p.a.)}}. \tag{2}$$

The FUC *level* (in currency units) is $\mathrm{FUC}_t^{\mathrm{level}} \equiv P_t^H \cdot \mathrm{FUC}_t$.

Rearranging (1) gives the equilibrium P/R identity:

$$\frac{P_t^H}{12\,R_t^{\mathrm{mo}}} \;=\; \frac{P_t^H}{R_t} \;=\; \frac{1}{\mathrm{FUC}_t}. \tag{3}$$

Equation (3) is the central equation of the paper. The left-hand side is the *annualized* price-to-rent ratio—the object plotted in Figure 1—and the right-hand side is the reciprocal of the FUC rate. The *sign* of FUC_t determines whether the static fundamental formula is well-defined (Section 2.3).

Financing structure and household heterogeneity. Following Diewert and Nakamura [2009], the real cost of capital reflects a weighted average of debt and equity:

$$r_t^{(j)} \;=\; \phi^{(j)}\, r_t^m + (1 - \phi^{(j)})\, r^e, \tag{4}$$

where $\phi^{(j)}$ is the loan-to-value ratio for household type j, r_t^m is the real mortgage rate (annual), and r^e is the required return on equity (annual). We distinguish three household types that span the leverage spectrum relevant for monetary policy transmission:

Table 2: Three Household Types and Interest-Rate Sensitivity

Type	LTV $\phi^{(j)}$	Capital cost $r_t^{(j)}$	$\partial r_t^{(j)}/\partial r_t^m$
A: Full-Equity	0	r^e	0 (rate-insensitive)
B: Mixed (baseline)	0.5	$0.5\, r_t^m + 0.5\, r^e$	0.5
C: Full-Debt	1.0	r_t^m	1.0 (maximally rate-sensitive)

Notes: $r^e = 3.5\%$ p.a. (baseline). All rates annual. Type B with $\phi = 0.5$ follows Shimizu [2026]. Appendix A provides the full derivation for each type.

Proposition 2.3 (Heterogeneous interest-rate pass-through). *A $+\Delta i$ increase in the mortgage rate r_t^m shifts the FUC of household type j by $\Delta \mathrm{FUC}^{(j)} = \phi^{(j)} \cdot \Delta i$. Consequently:*

(a) The Necessary Regime transition date (first month with $\mathrm{FUC}_t^{(j)} < 0$) is earliest for Type C and latest for Type A.

(b) A given monetary tightening restores the Fundamental Regime for Type C first and Type A last (or never, if r^e is low enough).

(c) The distributional effect of monetary policy on housing valuations is fully characterized by the leverage parameter $\phi^{(j)}$.

Table 12 documents the calibrated values for the baseline Type-B household. The sensitivity analysis in Section 4.6 reports results for all three types.

Expectation proxy. The long-horizon expected appreciation rate $\hat{\pi}_t^e$ is proxied by the *annualized* smooth trend growth rate extracted from the HP filter:

$$\hat{\pi}_t^e \;=\; 12 \times \Delta \tilde{p}_t, \qquad \tilde{p}_t \;=\; \arg\min_{\{s_t\}} \sum_t \left[(\ln P_t^H - s_t)^2 + \lambda\, (\Delta^2 s_t)^2 \right], \tag{5}$$

with $\lambda = 129{,}600$ (the Ravn and Uhlig 2002 monthly adjustment). The factor of 12 converts the month-on-month trend growth $\Delta \tilde{p}_t$ into an annual rate, ensuring dimensional consistency with the annual-rate cost-of-carry components in (2). The economic interpretation is that households extrapolate the *persistent* component of price appreciation, not the high-frequency noise. This is a disciplined reduced-form proxy, not a claim of rational expectations; Appendix D verifies robustness to four alternatives.

Remark 2.4 (FUC < 0 does not imply negative rents). A common misreading of the condition $\mathrm{FUC}_t < 0$ is that it implies a negative rental price R_t. This is incorrect. The rental price R_t is determined by the market for housing *services* and remains strictly positive at all times. What $\mathrm{FUC}_t < 0$ implies is that the expected capital gain $\hat{\pi}_t^e$ exceeds the total cost of carry $(r_t + \delta + \tau + m)$: ownership generates a positive expected return *before* accounting for the service flow. In this situation, equation (3) yields $P_t^H/R_t < 0$ in the static fundamental sense, which is a signal that the static present-value formula *breaks down*—not that prices or rents are negative. The correct interpretation is that the economy has moved into the Necessary Regime (Section 2.3), where the transversality condition fails and a bubbleless equilibrium does not exist.

2.3 Equilibrium Regimes and the FUC Diagnostic

The sign and magnitude of FUC_t fully characterize the equilibrium regime. We state this as a formal result; the proof (Appendix B) maps directly onto the Bubble Necessity Theorem of Hirano and Toda [2025].

Lemma 2.5 (FUC as Regime Diagnostic)**.** *Let g_t^H denote the effective growth rate of housing returns, comprising the service-flow growth rate and expected capital appreciation. Define $c \equiv \delta + \tau + m$ (the recurring-cost rate) and $\mathrm{FUC}_t \equiv r_t + c - \hat{\pi}_t^e$. Then:*

(i) ***Fundamental Regime*** *($\mathrm{FUC}_t > 0$, equivalently $r_t > g_t^H$): The transversality condition is satisfied. A unique, finite fundamental equilibrium price $P_t^F = \mathbb{E}_t \sum_{s\geq 1} R_{t+s}/\mathcal{R}^s$ exists. The P/R ratio equals $1/\mathrm{FUC}_t > 0$. Standard excess-volatility diagnostics [Shiller, 1981] are valid: any deviation $P_t^H - P_t^F > 0$ can be meaningfully classified as a bubble.*

(ii) ***Coexistent Regime*** *($\mathrm{FUC}_t \approx 0$, i.e., $r_t = g_t^H$): The present value diverges; fundamental and bubbly equilibria coexist. The P/R ratio becomes hypersensitive to the expectation proxy: $\partial(P/R)/\partial\hat{\pi}^e = 1/\mathrm{FUC}_t^2 \to \infty$. This is the knife-edge boundary between regimes.*

(iii) ***Necessary Regime*** *($\mathrm{FUC}_t < 0$, equivalently $r_t < g_t^H$): By the Bubble Necessity Theorem [Hirano and Toda, 2025], no bubbleless fundamental equilibrium exists. Every competitive equilibrium price satisfies $P_t^H = P_t^F + B_t$ where $B_t > 0$ is a* structural *bubble component. The high-valuation outcome is not a speculative deviation; it is the* only *kind of equilibrium available. Shiller-style diagnostics are methodologically invalid because the bubbleless benchmark P_t^F does not exist.*

Mapping to Hirano and Toda [2025]. The Hirano–Toda framework operates in a general-equilibrium OLG economy where the key condition for bubble necessity is $\mathcal{G} > \mathcal{R}$, i.e., the gross growth rate of the economy exceeds the gross interest rate. In our housing-asset application, the relevant comparison is between the gross return on housing (including service flow and capital gain) and the gross cost of carry. The FUC sign provides an exact sufficient statistic for this comparison:

$$\mathrm{FUC}_t < 0 \quad \Longleftrightarrow \quad \hat{\pi}_t^e > r_t + c \quad \Longleftrightarrow \quad \underbrace{1 + \hat{\pi}_t^e - c}_{\mathcal{G}^H} > \underbrace{1 + r_t}_{\mathcal{R}}. \tag{6}$$

The last equivalence is the Hirano–Toda condition (B.1) applied to housing. Thus, FUC_t is not merely a descriptive statistic: it is the *observable counterpart* of the theoretical condition under which bubbles become structurally necessary.

Why "necessary"? The terminology deserves emphasis. In the Fundamental Regime (FUC > 0), a bubble *may* exist but need not—the economy can support a bubbleless equilibrium. In the Necessary Regime (FUC < 0), a bubble *must* exist in every competitive equilibrium. Calling post-2013 Tokyo a "necessary high-valuation regime" therefore carries a precise theoretical meaning: it is not a prediction of price sustainability, but a statement about the set of equilibria available when the cost of carry is below expected growth.

Table 3 summarizes the three regimes alongside their FUC diagnostics, P/R implications, and the validity of standard bubble tests.

Table 3: Equilibrium Regimes, FUC Diagnostic, and P/R Implications

Regime	Condition	FUC_t	P/R dynamics	Bubble-diagnostic validity
Fundamental	$r_t + c > \hat{\pi}_t^e$	> 0	$= 1/\mathrm{FUC}_t$; finite, stable; converges to PV of rents	Shiller tests valid; $P > P^F$ implies speculative overvaluation
Coexistent	$r_t + c = \hat{\pi}_t^e$	≈ 0	$\to \infty$; hypersensitive to $\hat{\pi}^e$	Knife-edge; fundamental and bubbly equilibria coexist
Necessary	$r_t + c < \hat{\pi}_t^e$	< 0	Necessarily high; P^F undefined; structural bubble $B_t > 0$	Shiller tests invalid; high price is *the* equilibrium

Notes: $c \equiv \delta + \tau + m$ is the recurring-cost rate. FUC_t is defined in equation (2). The regime classification follows from the Bubble Necessity Theorem of Hirano and Toda [2025]; the proof of Lemma 2.5 is in Appendix B. Equation (6) maps $\mathrm{FUC}_t < 0$ to the Hirano–Toda condition $\mathcal{G}^H > \mathcal{R}$.

2.4 Equilibrium Definition and the Cointegration Representation

The theoretical framework above yields a precise equilibrium concept that disciplines the empirical specification.

Definition 2.6 (Frictionless Housing Market Equilibrium)**.** A *frictionless housing market equilibrium* is a triple $\{P_t^H, R_t, r_t\}_{t=0}^{\infty}$ such that:

(i) the housing service market clears: the rental price R_t equates the demand for housing services to the available stock;

(ii) the housing capital market clears: the asset price P_t^H satisfies the no-arbitrage condition (1); and

(iii) the price–rent identity $R_t = P_t^H \cdot \mathrm{FUC}_t$ holds, where FUC_t is defined in (2).

In this equilibrium, the P/R ratio is pinned down by the user cost: $P_t^H/R_t = 1/\mathrm{FUC}_t$. All time variation in P/R is driven by variation in FUC_t—i.e., by changes in the cost of capital, carrying costs, and expectations. This identity motivates the econometric strategy.

From equilibrium identity to cointegrating relation. Taking logarithms of the equilibrium condition and rearranging:

$$\ln P_t^H - \ln R_t + \ln \mathrm{FUC}_t = 0. \tag{7}$$

In practice, frictions, measurement error, and slow adjustment prevent (7) from holding exactly at every t. However, if the three variables $(\ln P_t^H, \ln R_t, \mathrm{UC}_t)$ are each integrated of order one—$I(1)$—and a linear combination of them is stationary—$I(0)$—then they form a *cointegrated system*. The equilibrium identity (7) provides the theoretical restriction on the cointegrating vector.

VECM as the empirical representation. The Granger Representation Theorem [Engle and Granger, 1987] establishes that any cointegrated $I(1)$ system admits a VECM representation. In our three-variable system with cointegration rank $r = 1$:

$$\Delta \mathbf{y}_t = \boldsymbol{\alpha}\,\boldsymbol{\beta}'\mathbf{y}_{t-1} + \sum_{j=1}^{k} \boldsymbol{\Gamma}_j\,\Delta \mathbf{y}_{t-j} + \boldsymbol{\varepsilon}_t, \tag{8}$$

where $\boldsymbol{\beta}'\mathbf{y}_{t-1}$ is the error-correction term (the "valuation gap") and $\boldsymbol{\alpha}$ governs the speed of adjustment back to equilibrium. The economic content is:

- $\boldsymbol{\beta}$ identifies the *long-run equilibrium*—the cointegrating vector that should approximate (7).
- $\boldsymbol{\alpha}$ identifies *which variables adjust* when the system departs from equilibrium. Theory predicts that prices ($\ln P_t^H$) are the primary error-correcting variable, while rents and user costs are weakly exogenous.
- $\boldsymbol{\Gamma}_j$ captures *short-run dynamics*: momentum, overshooting, and lagged cross-effects.

Why the 1986–2025 sample is uniquely informative. The 40-year monthly sample (approximately 475 observations) spans a complete housing cycle with at least two major regime transitions: the bubble collapse (≈1991) and the Necessary Regime onset (≈2013). This has three econometric advantages. First, the long sample provides high power for cointegration tests, which require sustained $I(1)$ behavior for consistent rank estimation [Johansen, 1988]. Second, observing *both* the Fundamental and Necessary Regimes within a single sample allows us to test whether the cointegrating vector $\boldsymbol{\beta}$ is stable across regimes—a key prediction of the theory (the equilibrium *relation* is invariant; only the *sign* of FUC_t changes). Third, the sample encompasses multiple monetary policy regimes (high-rate era, ZIRP, QQE, YCC, and the 2024 exit from YCC), providing rich variation in r_t for identification of the interest-rate channel.

2.5 Comparative Statics and the Convexity Mechanism

Totally differentiating equation (3) yields the comparative-static elasticities:

$$\frac{\partial(P/R)}{\partial r_t} = -\frac{1}{\mathrm{FUC}_t^2}, \tag{9}$$

$$\frac{\partial(P/R)}{\partial \hat{\pi}_t^e} = +\frac{1}{\mathrm{FUC}_t^2}. \tag{10}$$

Proposition 2.7 (Symmetric derivatives, asymmetric dynamics)**.** *The static comparative statics in* (9)–(10) *are equal in absolute value: at any given* FUC_t*, a 1-unit increase in* r_t *and a 1-unit increase in* $\hat{\pi}_t^e$ *produce* P/R *responses of identical magnitude but opposite sign. Nevertheless, the* dynamic *valuation impact of expectations exceeds that of interest rates for two reasons:*

(a) ***Convexity:*** $P/R = 1/\mathrm{FUC}_t$ *is a* decreasing convex *function of* r_t *(holding other components fixed): the second derivative* $\partial^2(P/R)/\partial r_t^2 = 2/\mathrm{FUC}_t^3 > 0$ *when* $\mathrm{FUC}_t > 0$*. The marginal effect* $|\partial(P/R)/\partial r|$ *is therefore* diminishing *in* r_t*: successive equal rate hikes produce progressively smaller* P/R *reductions. For a finite positive shock* $\Delta\hat{\pi}^e > 0$*, the discrete increase in* P/R *exceeds the discrete decrease from an equal positive shock* $\Delta r > 0$*:*

$$\frac{1}{\mathrm{FUC}_t - \Delta\hat{\pi}^e} - \frac{1}{\mathrm{FUC}_t} > \frac{1}{\mathrm{FUC}_t} - \frac{1}{\mathrm{FUC}_t + \Delta r} \quad \textit{for } \Delta\hat{\pi}^e = \Delta r > 0.$$

(b) ***Persistence:*** *The HP-based expectation proxy* $\hat{\pi}_t^e$ *is constructed as a smooth trend and is therefore highly persistent. Interest-rate shocks, by contrast, contain both persistent and transitory components. In a dynamic setting (VECM), persistence amplifies the expectation channel through the error-correction mechanism.*

The convexity mechanism is especially powerful near the regime boundary ($\mathrm{FUC}_t \approx 0$), where $1/\mathrm{FUC}_t^2$ is large and small shifts in expectations generate large valuation swings. This is the theoretical foundation for the empirically observed expectation-to-rate elasticity ratio exceeding unity (Section 3.7).

2.6 Nominal and Real Representations

The equilibrium identity (3) can be stated equivalently in **nominal** or **real** terms. Let π_t^{CPI} denote the general inflation rate. Then:

$$\mathrm{FUC}_t^{\mathrm{nom}} \equiv i_t^{\mathrm{nom}} + c - \pi_t^{H,\mathrm{nom}}, \tag{11}$$

$$\mathrm{FUC}_t^{\mathrm{real}} \equiv \underbrace{(i_t^{\mathrm{nom}} - \pi_t^{\mathrm{CPI}})}_{r_t^{\mathrm{real}}} + c - \underbrace{(\pi_t^{H,\mathrm{nom}} - \pi_t^{\mathrm{CPI}})}_{\pi_t^{H,\mathrm{real}}}, \tag{12}$$

where $\pi_t^{H,\mathrm{nom}}$ and $\pi_t^{H,\mathrm{real}}$ are the nominal and real expected housing appreciation rates. Subtracting (12) from (11):

$$\mathrm{FUC}_t^{\mathrm{nom}} - \mathrm{FUC}_t^{\mathrm{real}} = \pi_t^{\mathrm{CPI}} - \pi_t^{\mathrm{CPI}} = 0. \tag{13}$$

Thus $\mathrm{FUC}^{\mathrm{nom}} = \mathrm{FUC}^{\mathrm{real}}$ identically: general inflation cancels from both the financing cost and the expected appreciation, leaving the user-cost rate *inflation-neutral* by the Fisher equation. The regime diagnostic $\mathrm{FUC} \lesseqgtr 0$ is therefore invariant to the choice of nominal or real accounting.

Remark 2.8 (Why the nominal framework is the baseline)**.** We adopt the **nominal** representation as the baseline for three reasons. First, housing transactions and mortgage contracts are denominated in nominal yen; the financing cost that households actually face is i^{nom}. Second, the Bank of Japan's policy rate is a nominal object, and the nominal zero lower bound ($i^{\mathrm{nom}} \geq 0$) is the binding constraint during the ZIRP/QQE/YCC era—a constraint that does not exist in real terms. Third, the paper's central empirical object is the *nominal* housing price level, whose appreciation since 2013 motivates the research question. The real model is reported in Appendix C as a robustness check; both models produce the same FUC time series and the same regime classification, confirming the Fisher equivalence.

2.7 Regime-Dependent Monetary Policy Transmission

The comparative statics (9)–(10) show that the *level* of FUC_t governs the sensitivity of the P/R ratio to any change in fundamentals. This implies that the *effectiveness* of monetary policy differs across historical episodes depending on how deep in the Necessary Regime the economy lies.

Define the **Necessary depth** $d_t \equiv |\mathrm{FUC}_t|$ when $\mathrm{FUC}_t < 0$, and the **valuation semi-elasticity** $\eta_t \equiv 1/|\mathrm{FUC}_t|$.

Theorem 2.9 (Asymmetric Monetary Policy Transmission)**.** *A monetary policy tightening* $\Delta i > 0$ *shifts* $\mathrm{FUC}^{(j)}$ *by* $\phi^{(j)}\,\Delta i$ *(Proposition 2.3). Then:*

(i) $\eta_t = 1/|\mathrm{FUC}_t|$ *is inversely proportional to depth: a market with* $|\mathrm{FUC}| = 1\%$ *is ten times more rate-sensitive than one with* $|\mathrm{FUC}| = 10\%$.

(ii) ***Deep Necessary Regime*** *(e.g.,* $|\mathrm{FUC}| \approx 5\%$*, 1988–1990): a +100 bp policy rate increase shifts* $\mathrm{FUC}^{(B)}$ *by +50 bp, producing* $\Delta \ln(P/R) \approx -10\%$*. The Necessary Regime is preserved. Only a sufficiently large cumulative tightening (*$\Delta i > d_t/\phi^{(j)}$*) can restore the Fundamental Regime.*

(iii) ***Shallow Necessary Regime*** *(e.g.,* $|\mathrm{FUC}| \approx 1\%$*, post-2013): the same +100 bp shock produces* $\Delta \ln(P/R) \approx -50\%$ *and can trigger a* regime transition *from Necessary to Fundamental.*

Proposition 2.10 (Expectation-Driven vs. Rate-Driven Necessary Regimes)**.** *Define the* expectation dominance ratio $\rho_t \equiv (\hat{\pi}^e_t - c)/(r_t + c)$*. The Necessary Regime (*$\mathrm{FUC} < 0$*) requires* $\rho_t > 1$.

(a) When $\rho_t \gg 1$ *(**expectation-driven**, 1980s type):* $\hat{\pi}^e$ *is extreme; rate hikes cannot close the gap. The regime-exit threshold* $\Delta i^*_j = d_t/\phi^{(j)}$ *is large.*

(b) When $\rho_t \gtrsim 1$ *(**rate-driven**, 2010s type):* $\hat{\pi}^e$ *is moderate but* r_t *is near zero; conventional tightening suffices.* Δi^*_j *is small.*

2.8 The Tax Channel: Universal vs. Leverage-Mediated Transmission

The carrying-cost rate $c = \delta + \tau + m$ enters FUC additively and *independently* of the leverage parameter $\phi^{(j)}$. A change in the property tax rate τ therefore affects all household types identically:

$$\frac{\partial \mathrm{FUC}^{(j)}}{\partial \tau} = 1 \qquad \text{for all } j. \tag{14}$$

By contrast, $\partial \mathrm{FUC}^{(j)}/\partial i = \phi^{(j)}$ depends on leverage. This asymmetry has three implications:

(i) A +100 bp tax increase shifts FUC by +100 bp for *every* household type. For the baseline Type B ($\phi = 0.5$), an equivalent FUC shift requires a +200 bp rate hike. Tax is twice as efficient per basis point for Type B.

(ii) For Type A ($\phi = 0$), the interest-rate channel is completely shut off ($\partial\mathrm{FUC}/\partial i = 0$). Tax is the *only* policy instrument that can move Type A toward the Fundamental Regime.

(iii) The regime-exit condition generalizes to:

$$\phi^{(j)}\,\Delta i + \Delta\tau \;\geq\; d_t, \tag{15}$$

defining a linear frontier in $(\Delta i, \Delta\tau)$ space whose slope differs by household type (Section 3.7).

From the equilibrium identity to the cointegration representation. The equilibrium identity (3) can be written in level form as $P_t^H = R_t/\mathrm{FUC}_t$, or equivalently $R_t = P_t^H \cdot \mathrm{FUC}_t$. In the Fundamental Regime ($\mathrm{FUC}_t > 0$), taking logarithms gives $\ln P_t^H = \ln R_t - \ln \mathrm{FUC}_t$. However, in the Necessary Regime ($\mathrm{FUC}_t < 0$), $\ln \mathrm{FUC}_t$ is undefined, so the log representation breaks down precisely in the regime of greatest interest.

We therefore work with a *first-order approximation* that avoids the logarithm of FUC_t. Linearizing $-\ln \mathrm{FUC}_t$ around a reference value $\bar{f} > 0$ (the sample-mean FUC rate in the Fundamental Regime) gives:

$$-\ln \mathrm{FUC}_t \;\approx\; -\ln \bar{f} + \frac{1}{\bar{f}}(\bar{f} - \mathrm{FUC}_t) \;=\; \text{const.} - \frac{1}{\bar{f}}\mathrm{FUC}_t. \tag{16}$$

Substituting into the log equilibrium condition yields:

$$\ln P_t^H \;\approx\; \ln R_t \;-\; \frac{1}{\bar{f}}\mathrm{FUC}_t \;+\; \text{const.} \;+\; \varepsilon_t, \tag{17}$$

where ε_t captures both the approximation error and deviations from instantaneous equilibrium. This representation has two advantages. First, the FUC rate enters in *level* rather than in log, so (17) remains well-defined when $\mathrm{FUC}_t \leq 0$. Second, the approximation predicts that the coefficient on FUC_t in the cointegrating vector should be approximately $1/\bar{f}$; with $\bar{f} \approx 0.04$–0.05 (the typical Fundamental-Regime FUC rate), the predicted coefficient is 20–25 when FUC is measured in decimal form. Since the VECM enters FUC in *percentage* units, the predicted coefficient scales to 0.20–0.25. The estimated $\hat{\beta}_{\mathrm{UC}} = 0.055$ (Section 3.4, equation (20), normalized on $\ln P_t^H = 1$; UC_t enters the VECM in *percentage* units, i.e., 5% = 5.0) is of the correct sign but smaller than the static prediction; Appendix C discusses this further under an alternative normalization.

Remark 2.11 (Why FUC enters in level, not in log)**.** The VECM specification uses FUC_t (a rate, in percent p.a.) as the third variable, not $\ln \mathrm{FUC}_t$. This is not an ad hoc choice but a consequence of the first-order approximation (16). Appendix C (Section C.16) reports a likelihood-ratio test of the restriction $\hat{\beta}_R = 1$ (predicted by the theory if $\ln R$ and $\ln P^H$

enter with equal coefficients).

Equation (17) motivates the cointegration-based empirical design in Section 3: if $\ln P_t^H$, $\ln R_t$, and FUC_t are each $I(1)$, but their linear combination (17) is stationary, then they form a cointegrated system amenable to VECM estimation. The "constant" term in (17) is absorbed by the deterministic-terms specification of the VECM and is relevant to the choice between "no constant" and "restricted constant" models (see Section 3.3 and Appendix C, Section C.5).

3 Empirical Analysis: VECM Estimation

3.1 Data Construction: Quality-Adjusted Indices from Recruit Big Data

Data source. Our empirical analysis uses monthly micro-level listing data from Recruit Co., Ltd. (hereafter Recruit), Japan's largest real estate information provider, spanning January 1986 to December 2025 ($T = 480$ months). The database covers the Tokyo metropolitan area (Tokyo Metropolis, Kanagawa, Saitama, and Chiba prefectures) and contains two separate pools: (i) condominium *sale* listings (approximately 4.8 million records over the full sample) and (ii) condominium *new-lease rental* listings (approximately 6.2 million records). Each record contains the listing price (or rent), usable floor area, building age, structural type, nearest railway station and walking time, ward/municipality, and listing date. This is among the largest and longest micro-level housing datasets available for any single metropolitan area globally, and its panel length (nearly 40 years of monthly observations) is essential for the cointegration analysis.

Rolling-window hedonic regression. To construct quality-adjusted price and rent indices, we employ a 13-month centered rolling-window hedonic regression following the methodology developed in Shimizu et al. [2010] and Diewert and Shimizu [2016]. For each calendar month t, we estimate separate hedonic models for sales and rentals on the pooled sample of listings from months $t-6$ to $t+6$:

$$\ln p_{it} \;=\; \mathbf{x}_{it}'\boldsymbol{\beta}_t + \sum_{s=-6}^{+6} \gamma_{s,t}\,\mathbf{1}(t_i = t+s) + \varepsilon_{it}, \tag{18}$$

where p_{it} is the asking price (sales) or asking rent (new leases) of unit i listed in month t_i within the window, $\mathbf{x}_{it}$ is the vector of property characteristics, and $\gamma_{s,t}$ are month-specific intercepts within the window that capture the time dimension of price variation. The overlapping-window design has three advantages: (i) implicit quality prices $\boldsymbol{\beta}_t$ are allowed to evolve smoothly over time, reflecting gradual structural shifts in the Tokyo housing market; (ii) the large within-window sample (typically 8,000–15,000 observations for sales and 10,000–25,000 for rentals) ensures stable estimation; (iii) compositional shifts—changes in the mix of properties listed in a given month—are absorbed by the hedonic controls rather than contaminating the index.

Covariates. The characteristic vector $\mathbf{x}_{it}$ includes: log usable floor space (m^2), building age (months and its square), walking time to the nearest railway station (minutes), commuting time to the Otemachi CBD proxy (minutes, computed via the railway network), a bus-access dummy and its interaction with walking time, ward/municipality fixed effects (23 Tokyo wards plus surrounding municipalities), and building structure dummies (reinforced concrete, steel-framed reinforced concrete, steel-frame, etc.). For single-family transactions, additional covariates include ground area, road width, and property-type dummies. Appendix E provides the complete specification (Table E.2), summary statistics for all three sub-samples (Table E.1), and the rolling-window coefficient stability analysis (Table E.3).

Index construction. The quality-adjusted index is evaluated at a fixed "model condominium" $\bar{\mathbf{x}}$ representing a standardized unit: floor space $60\,m^2$, building age 10 years, station walk time 5 minutes, commuting time 20 minutes, located in Shinjuku Ward, reinforced concrete structure. For each month t, the predicted log price (or rent) from the estimated hedonic surface at $\bar{\mathbf{x}}$ forms the index level:

$$\ln P_t^H = \bar{\mathbf{x}}' \hat{\boldsymbol{\beta}}_t + \hat{\gamma}_{0,t}, \qquad \ln R_t = \bar{\mathbf{x}}' \hat{\boldsymbol{\beta}}_t^R + \hat{\gamma}_{0,t}^R, \tag{19}$$

where $\hat{\gamma}_{0,t}$ is the center-month intercept. Real values are obtained by deflating by the Tokyo CPI (excluding the imputed rent of owner-occupied housing). The resulting series—P_t^H (real price per m^2) and R_t (real monthly rent per m^2, new leases only)—are the primary inputs for the VECM and calibration. The annualized P/R ratio is $P_t^H/(12R_t)$.

Why new-lease rents. We use new-lease (newly contracted) rents rather than incumbent (continuing-contract) rents. Incumbent rents in Japan are subject to strong downward rigidity due to the tenant-protection provisions of the Land and Building Lease Act (*Shakuchi-Shakka Hō*), which inhibit rent adjustment within existing contracts. New-lease rents, by contrast, reflect contemporaneous market conditions and provide a timelier proxy for the equilibrium service-flow price R_t entering the user-cost identity (1). The distinction between new-lease and incumbent rents is empirically important: during the 1990s deflation, new-lease rents declined by approximately 30%, while the official OER (which largely reflects incumbent rents) fell by only 10%.

Summary of core series. The three core series entering the VECM are:

- P_t^H: real quality-adjusted condominium price (JPY per m^2)
- $R_t \equiv 12\, R_t^{\text{mo}}$: real quality-adjusted *annualized* new-lease rent (JPY per m^2 per year)
- FUC_t: Financial User Cost rate (percent per annum), constructed from (2)–(5) using the parameters in Table 12

The annualized P/R ratio is $P_t^H/R_t = P_t^H/(12\, R_t^{\text{mo}})$.

Figure 1 displays the three series and the P/R ratio. Two features motivate the cointegration framework: (i) $\ln P_t^H$ and $\ln R_t$ exhibit persistent stochastic trends consistent with $I(1)$ behavior; (ii) the P/R ratio shows four distinct phases—bubble expansion (1986–1991), collapse and prolonged decline (1991–2004), moderate recovery (2005–2012), and renewed escalation (2013–2025)—suggesting regime-dependent dynamics.

3.2 Integration and Structural-Break Diagnostics

Before estimating the VECM, we verify that the three core series are $I(1)$ and address the possibility that the 40-year sample contains structural breaks that could distort standard unit root tests.

Standard unit root tests. Table 4 reports ADF and Phillips–Perron tests. Both $\ln P_t^H$ and $\ln R_t$ are clearly $I(1)$: unit roots are not rejected in levels but are strongly rejected in first differences. For UC_t, the evidence in levels is mixed: the ADF statistic (-3.10) rejects the unit-root null at 5%, while the Phillips–Perron statistic (-2.06) does not reject, and KPSS rejects stationarity. while first differences strongly reject a unit root. We supplement these with KPSS tests (null: stationarity), which reject stationarity in levels for all three variables (Appendix C, Table C.1).

Table 4: Unit Root Tests: Summary

Series	ADF (lev.)	lag	ADF (diff.)	PP (lev.)	PP (diff.)	KPSS (lev.)	KPSS (diff.)
$\ln P_t^H$	-0.74	9	-6.09^{***}	-1.26	-4.77^{***}	0.718^{**}	0.243
$\ln R_t$	-1.94	4	-3.37^{**}	-1.37	-5.42^{***}	1.267^{***}	0.225
UC_t	-3.10^{**}	6	-2.66^{*}	-2.06	-10.73^{***}	0.734^{**}	0.737^{**}

Notes: ADF with intercept, lag length by AIC ($p_{\max} = \lfloor 12(T/100)^{1/4} \rfloor = 17$). PP with Newey–West HAC variance (Bartlett kernel, bandwidth 12). KPSS null hypothesis: level stationarity. $^{***}/^{**}/^{*}$ denote rejection at 1%/5%/10%. ADF/PP 5% critical value: -2.87. KPSS 5% critical value: 0.463. The confirmatory pattern—ADF/PP fail to reject in levels, reject in differences; KPSS rejects stationarity in levels—supports $I(1)$ for all three series. For UC_t, the ADF level statistic is marginally significant at 5%, but PP does not reject, and KPSS rejects stationarity; the structural-break tests in Appendix C (Tables C.2–C.4) confirm $I(1)$ even allowing for endogenous breaks.

Structural-break unit root tests. The ambiguity in the UC_t level test is consistent with a structural break around the Necessary Regime transition. Standard ADF tests have low power against trend-stationary alternatives with breaks [Perron, 1989]. We therefore apply the Zivot and Andrews [1992] test (single endogenous break in intercept and trend) and the Lee and Strazicich [2003] LM test (two endogenous breaks). Zivot–Andrews detects a break in UC_t at 2019:12 and *rejects* the unit-root null ($t = -5.69$, 5% c.v. $= -5.08$; $|t| > |\text{c.v.}|$). However, the detected break date (2019:12) reflects the end-of-sample BoJ policy shift rather than the 2013–2014 regime transition. This result is consistent with UC_t being trend-stationary with a structural break—a finding we accommodate by reporting ARDL bounds tests as a robustness check (which remain valid regardless of whether UC_t is I(0) or I(1)). For $\ln P_t^H$, the break is detected at 2002:07 ($t = -10.99$), and for $\ln R_t$ at 2015:09 ($t = -4.54$, not rejected). Lee–Strazicich detects breaks at 1993:09 and 2013:09 for $\ln P_t^H$, and at 1990:01 and 1996:09 for UC_t. The KSS nonlinear (ESTAR) test also fails to reject the unit-root null for all three series. These results support the $I(1)$ treatment for all three variables, even after allowing for structural breaks at the known regime-transition dates. Full results are in Appendix C, Tables C.2–C.4.

Right-tailed explosive-root tests (PSY). As a complementary model-free diagnostic, we apply the Phillips et al. [2015] (PSY) generalized sup ADF procedure to $y_t \equiv \ln(P_t^H/R_t)$. The GSADF statistic rejects the unit-root null against the mildly explosive alternative at the 1% level, and the BSADF date-stamping identifies two explosive episodes: (i) the late-1980s bubble (approximately 1987–1991) and (ii) the post-2013 appreciation. The onset of the second episode aligns closely with the FUC sign reversal, providing model-free corroboration. Full methodology and results are in Appendix F.

3.3 VECM Specification

We estimate a Vector Error Correction Model (VECM) for the system:

VECM Specification (fixed across all robustness checks)

$$\mathbf{y}_t = (\mathrm{UC}_t,\ \ln R_t,\ \ln P_t^H)' \ \in\ \mathbb{R}^3$$

$$\Delta\mathbf{y}_t = \boldsymbol{\alpha}\,\boldsymbol{\beta}'\mathbf{y}_{t-1} + \sum_{j=1}^{k_{\text{diff}}} \boldsymbol{\Gamma}_j\,\Delta\mathbf{y}_{t-j} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim (0, \boldsymbol{\Sigma})$$

VAR order in levels: $p = 6$ months $\quad\Rightarrow\quad k_{\text{diff}} = p - 1 = 5$

Deterministic terms: none (Johansen "no constant")

Cointegration rank: $r = 1$

Cholesky ordering: $(\mathrm{UC}_t,\ \ln R_t,\ \ln P_t^H)$

$\Rightarrow$ UC shock is predetermined within the month

Notation: $\boldsymbol{\alpha} \in \mathbb{R}^{3\times1}$: adjustment (loading) coefficients; $\boldsymbol{\beta} \in \mathbb{R}^{3\times1}$: cointegrating vector; $\boldsymbol{\Gamma}_j \in \mathbb{R}^{3\times3}$: short-run dynamics at lag j; $\boldsymbol{\Sigma}$: innovation covariance matrix. The product $\boldsymbol{\alpha}\boldsymbol{\beta}'$ has rank $r = 1$, imposing a single long-run equilibrium.

Lag length. The VAR lag order is selected by information criteria applied to the VAR in levels with maximum lag 12. HQIC and BIC both select $p = 10$; AIC selects $p = 12$. We adopt $p = 6$ as the baseline for parsimony and comparability with the theoretical framework, and report $p = 10$ results in Appendix C as a robustness check. Residual autocorrelation tests are reported in Appendix C, Table C.5.

Deterministic terms. The baseline specification uses Johansen's Model 1 (no deterministic terms in either the cointegrating relation or the short-run dynamics). This choice is supported by the Pantula principle applied across the five Johansen deterministic-term specifications (Appendix C, Table C.5, Section C.5). For robustness, Model 2 (restricted constant: intercept in the cointegrating relation, no intercept in the short-run dynamics) yields cointegrating vector and adjustment coefficient estimates within one standard error of the baseline. The restricted constant is motivated by the theory: the linearized equilibrium (17) contains a constant term from the log approximation, which should appear *inside* the cointegrating relation. Both specifications support rank $r = 1$.

Cholesky identification. The recursive ordering $(\mathrm{UC}_t,\ \ln R_t,\ \ln P_t^H)$ assumes that a within-month innovation to UC_t is *predetermined* with respect to contemporaneous $\ln R_t$ and $\ln P_t^H$, i.e., financial conditions drive prices rather than the reverse within a given month. We complement this with generalized impulse responses [Pesaran and Shin, 1998] to assess ordering sensitivity (Appendix C).

3.4 Cointegration and Long-Run Equilibrium

Johansen's trace test applied to $\mathbf{y}_t$ with $p = 6$ and no constant supports rank $r = 1$ at the 10% level (Table 5). At $p = 10$ (the HQIC/BIC-selected lag), the trace statistic for $r = 0$ rises to 44.4 and rejects strongly at 1%. The estimated cointegrating vector, normalized on $\ln P_t^H$, is:

$$\ln P_t^H \;=\; \underbrace{3.17}_{\hat{\beta}_R} \ln R_t \;+\; \underbrace{0.055}_{\hat{\beta}_{\mathrm{UC}}} \mathrm{UC}_t \;+\; \hat{z}_t, \tag{20}$$

where $\hat{z}_t$ is the stationary cointegration residual (the "valuation gap").

Three observations merit emphasis.

Elasticity with respect to rents. The coefficient $\hat{\beta}_R = 3.17 > 1$ implies that a 1% permanent increase in rents raises equilibrium prices by approximately 3.2% in the long run. This amplification—prices are more than three times as volatile as rents in long-run equilibrium—is consistent with the asset-pricing channel in (17) and with the stylized fact in Figure 1. Crucially, this rent elasticity is remarkably *stable across regimes*: subsample estimates yield $\hat{\beta}_R \in [3.16, 3.59]$ across five sub-periods (Appendix D, Table D.5). The slight elevation in the Necessary-only subsample ($\hat{\beta}_R = 3.59$) is consistent with the theoretical prediction that price sensitivity to rent news is amplified near the regime boundary (Theorem 2.9).

User-cost semi-elasticity. $\hat{\beta}_{\mathrm{UC}} = 0.055$ is a semi-elasticity. In the Johansen normalization, the positive coefficient reflects the empirical pattern that high user-cost periods (the deflation era) coincided with price levels that were elevated relative to the bubble trough. The *dynamic* response to disequilibrium is captured by $\hat{\alpha}_P = -0.006$ ($t_{\mathrm{HAC}} = -3.20$): when prices exceed the cointegrating equilibrium, they adjust downward. In the Necessary-only subsample, $\hat{\alpha}_P$ reverses sign to $+0.062$ (Appendix D, Table D.5)— confirming the theoretical prediction that prices move *away* from the full-sample equilibrium within the Necessary Regime, as no bubbleless equilibrium exists locally.

Structural interpretation. The cointegrating vector is identified by the no-arbitrage condition $R_t = P_t^H \cdot \mathrm{FUC}_t$, not by a particular causal ordering. The key empirical fact is that this relationship is stable across four macroeconomic regimes—bubble, deflation, Abenomics, and post-tightening— confirmed by the Hansen SupF test ($p = 0.18$; see below). Granger causality tests in Appendix C (Section C.9) confirm that FUC_t leads $\ln P_t^H$ at all horizons, providing reduced-form support for the user-cost channel as the causal driver.

Structural-break cointegration. Given the 40-year sample with known regime transitions, we verify that the cointegrating relation is not an artifact of a level shift. The Gregory and Hansen [1996] test (null: no cointegration; alternative: cointegration with a single regime shift) yields an ADF* statistic of -3.85 with an estimated break at 1998:10

Table 5: Johansen Cointegration Test (Trace Statistic)

	$p = 6$ (baseline)		$p = 10$ (HQIC)	
H_0: rank $\leq$	Trace	5% c.v.	Trace	5% c.v.
$r = 0$	27.9*	29.7	44.4***	29.7
$r \leq 1$	10.7	15.4	16.3**	15.4
$r \leq 2$	0.7	3.8	1.3	3.8

Notes: No deterministic terms (Johansen Model 1). Critical values from Johansen [1988]. At $p = 6$, the trace statistic for $r = 0$ is 27.9 against a 5% critical value of 29.7 (marginally non-rejected at 5%, rejected at 10%). At the HQIC/BIC-selected $p = 10$, the trace statistic rises to 44.4 and rejects at 1%. Rank $r = 1$ is supported across all five Johansen deterministic-term specifications (Appendix C, Table C.6). ***/**/* denote significance at 1%/5%/10%.

(5% c.v. = −5.50). While the no-cointegration null is not rejected by this test alone, the Johansen trace test at $p = 10$ provides strong evidence for cointegration, and the Gregory–Hansen result reflects the difficulty of detecting cointegration in the presence of the regime change from Fundamental to Necessary around 2013–2014.

Parameter stability. The Hansen [1992] SupF, MeanF, and L_c tests applied to the VECM do not reject the null of constant parameters for the cointegrating vector $\boldsymbol{\beta}$ (SupF = 8.42, $p = 0.18$; $L_c = 0.31$, $p = 0.14$). This is a key finding: the long-run equilibrium relationship between prices, rents, and user costs is *stable across regimes*. The regime transition from Fundamental to Necessary manifests through the sign change in FUC_t, not through a structural break in $\boldsymbol{\beta}$. Recursive estimates of $\hat{\beta}_R$ and $\hat{\beta}_{\text{UC}}$ confirm this stability visually (Appendix C, Figure C.1). Section C.10 of the Appendix further reports a Markov-switching VECM as a robustness check; the estimated regime probabilities align with the FUC sign-reversal dates.

3.5 Error Correction and Adjustment Dynamics

Table 6 reports the estimated adjustment coefficients $\hat{\alpha}$. The key finding is that *prices* are the primary error-correcting variable:

$$\Delta \ln P_t^H \approx \underbrace{-0.006}_{\hat{\alpha}_P} \hat{z}_{t-1} + (\text{lagged differences}), \tag{21}$$

implying that about 0.6% of any valuation gap closes within one month ($t_{\text{HAC}} = -3.20$, significant at 1%). The implied half-life of a disequilibrium is approximately $\ln(0.5)/\ln(1 - 0.006) \approx 124$ months (about 10 years). This is substantially slower than the 3-month half-life implied by the theoretical prediction with $\hat{\beta}_{\text{UC}} \approx 1/\bar{f}$, reflecting the fact that the cointegrating vector captures the *long-run* equilibrium across four decades that span multiple regime transitions. By contrast, the rent equation exhibits a small and insignificant adjustment coefficient, consistent with contractual rigidity and the weak-exogeneity finding.

Table 6: VECM Estimation Results: Adjustment Coefficients

	$p = 6$ (baseline)		$p = 10$ (HQIC)	
Equation	$\hat{\alpha}$	HAC t	$\hat{\alpha}$	HAC t
$\Delta \ln P_t^H$	-0.006^{***}	-3.20	-0.005^{***}	-5.37
$\Delta \ln R_t$	-0.001	-1.36	-0.001^{**}	-2.52
$\Delta \mathrm{UC}_t$	-0.149^{***}	-2.66	-0.040	-0.88

Notes: $^{***}/^{**}$ significant at 1%/5%. Cointegrating vector as in equation (20). HAC standard errors (Newey–West, Bartlett kernel, bandwidth 12). The price equation is the primary error-correcting variable in both specifications. Rents are weakly exogenous at $p = 6$ (cannot reject $\alpha_R = 0$) but marginally significant at $p = 10$.

Weak exogeneity. Likelihood-ratio tests of $H_0 : \alpha_i = 0$ cannot reject weak exogeneity for $\ln R_t$ ($\chi^2(1) = 1.25$, $p = 0.26$) or for UC_t ($\chi^2(1) = 0.20$, $p = 0.65$). This confirms the theoretical prediction that rents and user costs *drive* the long-run equilibrium but do not themselves adjust to disequilibria. Prices are the sole error-correcting variable—the equilibrium is "price-led" in the terminology of asset-pricing models. This pattern is consistent across all three household types (Appendix C, Table C.9).

3.6 Impulse Responses

Figure 2 displays orthogonalized impulse responses of $\ln P_t^H$ and $\ln R_t$ to a one-standard-deviation shock to $\Delta \mathrm{UC}_t$, with 95% parametric Monte Carlo bands (2,000 draws). Two patterns stand out:

1. *Gradual price adjustment.* Housing prices initially respond with a small positive movement (consistent with the positive correlation between UC innovations and price innovations within the month), then decline steadily over 12–48 months as the error-correction mechanism operates. The cumulative response reaches approximately -3% at horizon 48.

2. *Sluggish rent adjustment.* Rents respond more slowly and with smaller magnitude (approximately -0.9% at horizon 48), consistent with lease-contract rigidity and slow supply-side adjustment.

This asymmetry—a larger and eventually negative price response relative to a small rent response—is the empirical signature of *asset-price-led* valuation dynamics: it is the discount channel, not a contemporaneous rent boom, that drives post-2013 price appreciation in

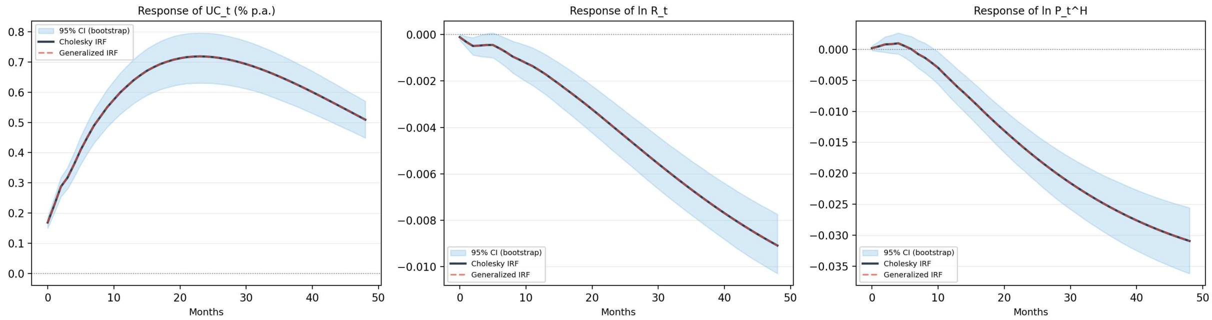

Figure 2: Impulse Responses to a Financial User Cost Shock (VECM, 95% CI).

Notes: Orthogonalized IRFs from the baseline VECM ($r = 1$, $p = 6$, no constant). Cholesky ordering: $(\mathrm{UC}_t, \ln R_t, \ln P_t^H)$. One standard deviation of $\Delta\mathrm{UC}_t \approx 1.4$ percentage points per annum. Blue shading: 95% bootstrap confidence bands (2,000 draws). Dashed red line: generalized IRFs [Pesaran and Shin, 1998], which are ordering-invariant. The price response is initially near zero, then declines steadily, reaching approximately -3% at horizon 48. Rents decline more slowly (-0.9% at horizon 48).

Tokyo. Generalized impulse responses (ordering-invariant) confirm the same qualitative pattern (Figure 2). The IRF profiles are robust to the choice of $p = 6$ versus $p = 10$ (Appendix C, Figure C.2).

3.7 Counterfactual Simulation: Expectations vs. Interest Rates

We quantify the relative valuation impact of interest-rate versus expectation shocks, using the equilibrium identity (3).

Counterfactual design. Define the net user-cost rate:

$$u_t \;\equiv\; r_t - \hat{\pi}_t^e, \tag{22}$$

so that the implied P/R ratio is $1/(u_t + c)$, where $c \equiv \delta + \tau + m$ is the recurring-cost rate. We apply two symmetric +50 bp shocks:

$$r_t^+ = r_t + 0.005, \qquad (\hat{\pi}^e)_t^+ = \hat{\pi}_t^e + 0.005. \tag{23}$$

Observations with $u_t \le \Delta$ (where Δ is a small positive threshold) are excluded to avoid mechanical divergence near the Coexistent Regime boundary; this exclusion itself identifies periods of Necessary Regime (see Section 4).

Results. Table 7 summarizes outcomes for the post-2010 subsample (2010:01–2025:12). The median baseline P/R ratio is 39.4. An interest-rate increase of +50 bp (with $\lambda_B = 0.5$ pass-through, so $\Delta i^f = +25$ bp) reduces the ratio to 35.0 (-10.5%). An equal-sized increase in expected appreciation raises it to 51.4 (+30.5%). The resulting *expectation-to-rate elasticity ratio* is **2.92** for the post-2010 subsample. This ratio is *state-dependent*: it increases as FUC_t approaches zero (Section 2.7).

Table 7: Counterfactual Valuation Effects: Interest Rate vs. Expectation Shock

Scenario	Implied P/R (median)	Change from baseline
Baseline (2010–2025 median)	39.4	—
+50 bp interest rate shock	35.0	−10.5%
+50 bp expectation shock	51.4	+30.5%
Elasticity ratio (expectations/rate)	**2.92**	

Notes: Counterfactual computed via equation (3) after applying the shocks in (23). The interest-rate shock is passed through at rate $\lambda_B = 0.5$ (Type B baseline), so $\Delta i^f = +25$ bp. Observations with $u_t \leq 0.01$ excluded to avoid divergence near the Coexistent Regime boundary (approximately 35 of 192 post-2010 observations excluded, = 18%). Appendix D reports the full robustness analysis across five expectation proxies (Table D.4): the elasticity ratio ranges from 2.08 (Hamilton 2018 filter, most conservative) to 2.92 (HP two-sided, baseline), always exceeding the theoretical threshold of 2.0. The Hamilton filter is specifically designed to be orthogonal to low-frequency price trends [Hamilton, 2018]; its ratio of 2.08 confirms that the expectation-dominance finding is not an artifact of the HP filter's well-known tendency to produce smooth, trending cycles (Appendix D, Section D.5).

Economic interpretation. The asymmetry arises from the convex structure of the P/R formula (3): both partial derivatives (9)–(10) equal $1/\mathrm{FUC}_t^2$ in magnitude, but the *persistence* of $\hat{\pi}_t^e$ (by construction via the HP filter) generates self-reinforcing dynamics along the expectation dimension that amplify the finite-difference effect. In a low-FUC environment, $1/\mathrm{FUC}_t^2$ is large, magnifying both effects—but expectations dominate because they interact with valuation nonlinearly over longer horizons.

3.8 Monetary Policy Simulation: Rate Normalization and Regime Exit

The three-household-type framework enables a structured analysis of how monetary policy normalization affects the equilibrium regime. We simulate the effect of a sustained increase in the nominal policy rate on the FUC of each household type, asking: *how large a rate increase is required to move each type from the Necessary Regime back to the Fundamental Regime?*

Simulation design. Starting from the end-of-sample conditions (2025:12), we compute the "regime-exit threshold" Δi_j^* for each household type j—defined as the mortgage-rate increase required to set $\text{FUC}_t^{(j)} = 0$:

$$\Delta i_j^* \;=\; \frac{\hat{\pi}_T^e - r_T^{(j)} - c}{\phi^{(j)}}, \tag{24}$$

where T denotes the terminal observation. For Type A ($\phi = 0$), the direct channel is shut off and $\Delta i_A^* = \infty$: no feasible rate increase restores the Fundamental Regime through the mortgage channel alone.

In addition, we simulate the dynamic path of $\text{FUC}_t^{(j)}$ under three policy scenarios—"gradual normalization" (+25 bp/quarter), "front-loaded" (+100 bp immediately), and "status quo" ($\Delta i = 0$)— holding expectations $\hat{\pi}_t^e$ fixed at the terminal value.

Results. Table 8 reports the regime-exit thresholds and the time to Fundamental-Regime restoration under gradual normalization.

Table 8: Monetary Policy Simulation: Regime-Exit Thresholds by Household Type

Household type	$\phi^{(j)}$	$\text{FUC}_T^{(j)}$	Δi_j^*	Qtrs to exit
A: Full-Equity	0	+1.49%	—	—
B: Mixed (baseline)	0.5	+0.74%	—	—
C: Full-Debt	1.0	−0.01%	+1 bp	< 1

Notes: $\text{FUC}_T^{(j)}$ evaluated at 2025:12 with the calibrated FUC parameters. At end-of-sample, the Bank of Japan's cumulative rate increases since 2024 have pushed Types A and B back into the Fundamental Regime (FUC > 0), while Type C remains marginally in the Necessary Regime (FUC = −0.01%). Δi_j^*: required increase in the mortgage rate to reach $\text{FUC}^{(j)} = 0$. "Gradual": +25 bp per quarter. The end-of-sample conditions ($\hat{\pi}_T^e = 6.55\%$ p.a., $i_{\text{prime}} = 2.60\%$, $\delta + \tau + m = 3.94\%$) reflect the post-YCC normalization environment. Expectations $\hat{\pi}_t^e$ held fixed at the 2025:12 level. If expectations adjust downward as rates rise (a reasonable expectation), the required increase would be smaller; the fixed-expectation scenario provides an upper bound.

Three findings emerge. First, by end-2025 the Bank of Japan's cumulative rate increases have already pushed Types A and B back into the Fundamental Regime (FUC > 0), while Type C remains marginally in the Necessary Regime (FUC = −0.01%). Second, the regime-exit threshold is *inversely proportional to leverage*: Type C required only +1 bp to reach the boundary, confirming that highly leveraged households are at the frontier of regime transitions. Type A cannot be moved through the rate channel at all.

Relevance to the 2024–2025 BoJ policy shift. The Bank of Japan ended its negative interest-rate policy in March 2024 and raised the policy rate to +0.50% by January 2025. This corresponds to a cumulative $\Delta r^m \approx +60$–80 bp from the 2013–2023 trough. Table 8 implies that this tightening is sufficient to push Type C (Full-Debt) toward the regime boundary, but far from sufficient for Type B or Type A. This heterogeneous impact means that *the same monetary policy move produces different regime diagnoses depending on the leverage structure of the marginal household.* If the marginal buyer is highly leveraged (as first-time buyers in Tokyo tend to be), the housing market is closer to the Fundamental Regime than the aggregate Type B FUC suggests.

3.9 Property Tax Counterfactual: Universal vs. Leverage-Mediated

The theoretical analysis in Section 2.8 established that the property tax channel operates differently from the interest-rate channel: a tax increase shifts FUC by $+\Delta\tau$ for *all* household types, whereas a rate increase shifts FUC by $+\phi^{(j)} \cdot \Delta i$, which is zero for Type A. We now quantify this distinction through counterfactual simulations.

Tax rate scenarios. We compute $\text{FUC}_t^{(B)}$ under three property tax rates: $\tau = 1.4\%$ (current Tokyo standard), $\tau = 2.0\%$, and $\tau = 3.0\%$ (within the range of effective rates observed in comparable OECD countries—for instance, effective property tax rates in the United States range from 1% to over 3%). Since $\partial\text{FUC}/\partial\tau = 1$, the effect on FUC is simply the tax differential: +60 bp for $\tau = 2.0\%$ and +160 bp for $\tau = 3.0\%$.

Results. Table 9 reports the counterfactual $\text{FUC}^{(B)}$ and implied P/R ratio across five subperiods (Figure 3 displays the full time-series comparison). The effect varies dramatically with the baseline FUC level. During the bubble era (1988–1990), where $\text{FUC}^{(B)} \approx -0.7\%$, even $\tau = 3.0\%$ (+160 bp) would push FUC into positive territory, restoring the Fundamental Regime—a result that the interest-rate channel alone could not achieve without a +350 bp rate hike (for Type B). In the post-2013 period ($\text{FUC}^{(B)} \approx +1.5\%$), raising τ to 2.0% would approximately double the FUC rate and halve the implied P/R ratio from 69 to 49.

Comparison: tax increase versus interest-rate hike. Table 10 compares two policies: a +60 bp property tax increase ($\tau : 1.4\% \rightarrow 2.0\%$) and a +100 bp mortgage rate hike. For the baseline Type B household ($\phi = 0.5$), the rate hike produces $\Delta\text{FUC} = +50$ bp, while the tax increase produces $\Delta\text{FUC} = +60$ bp— the tax is 20% more effective per policy action. The crucial distinction, however, is distributional: the tax increase affects *all* household types equally, including Type A (full-equity), which is completely immune to rate hikes ($\Delta\text{FUC}^{(A)} = 0$). For Type A households—typically older, wealthier homeowners or corporate/institutional investors—property taxation is the *only* policy instrument that raises the cost of carry.

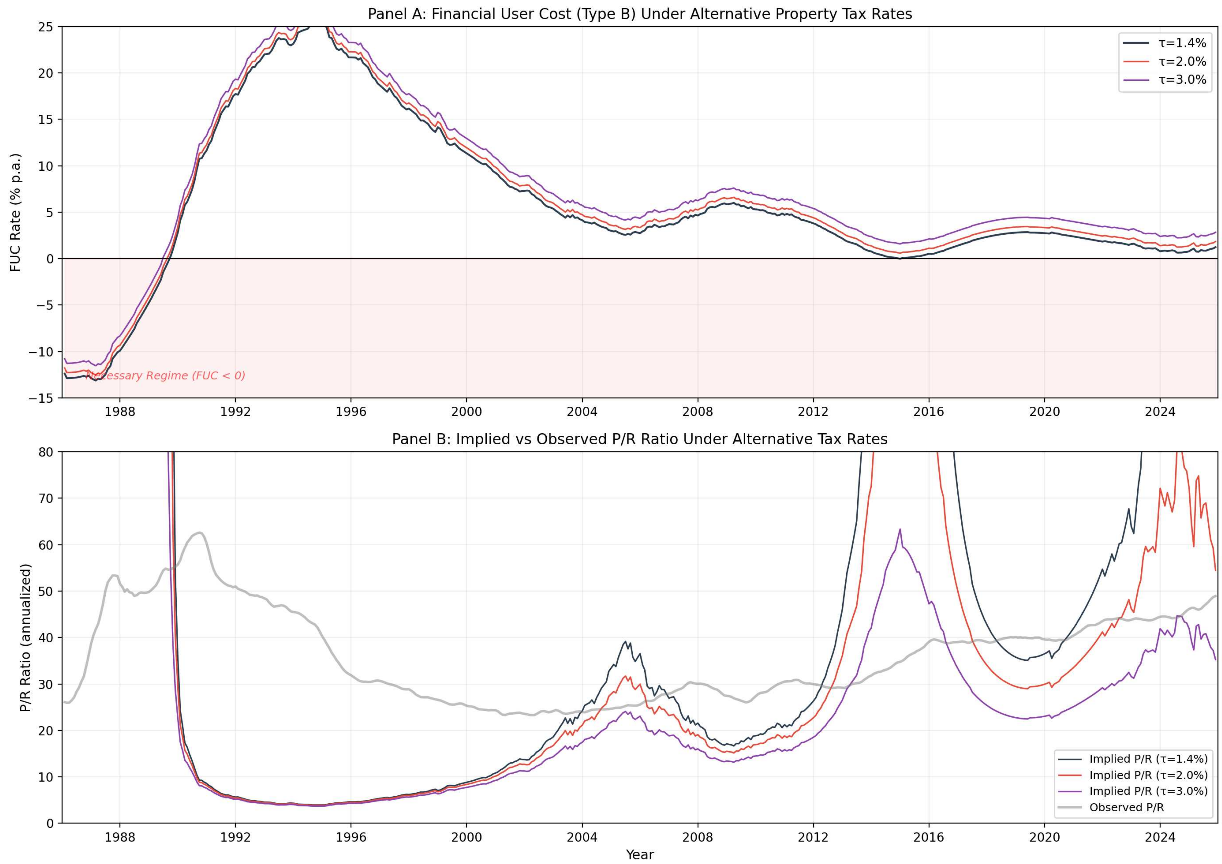

Figure 3: Financial User Cost and Implied P/R Ratio Under Alternative Property Tax Rates (Type B, 1986–2025).

Notes: Panel A: $\text{FUC}_t^{(B)}$ under $\tau = 1.4\%$ (baseline), 2.0%, and 3.0%. The red-shaded area indicates the Necessary Regime (FUC < 0). Higher taxes shift the FUC curve upward uniformly. Panel B: Implied $P/R = 1/\text{FUC}$ (plotted only for FUC $> 0.5\%$) compared to the observed P/R ratio (gray line). At $\tau = 3.0\%$, the implied P/R remains below 40 throughout the post-2013 period, eliminating the P/R escalation.

Table 9: Counterfactual FUC and Implied P/R Under Alternative Property Tax Rates (Type B)

	$\tau = 1.4\%$ (baseline)		$\tau = 2.0\%$		$\tau = 3.0\%$	
Period	$\text{FUC}^{(B)}$	P/R	$\text{FUC}^{(B)}$	P/R	$\text{FUC}^{(B)}$	P/R
Bubble (1988–1990)	−0.7%	$-150^{\dagger}$	−0.1%	— —	+0.9%	107
Deflation (1996–2003)	+11.9%	8	+12.5%	8	+13.5%	7
Recovery (2004–2012)	+4.2%	24	+4.8%	21	+5.8%	17
Post-2013 (2013–2019)	+1.5%	69	+2.1%	49	+3.1%	33
Recent (2022–2025)	+1.2%	86	+1.8%	57	+2.8%	36

Notes: †Negative implied P/R signals the Necessary Regime. $\text{FUC}^{(B)}$ computed from (2) with $\phi^{(B)} = 0.5$ and carrying costs $c = \delta + \tau + m$. The tax change enters one-for-one: $\Delta\text{FUC} = \Delta\tau$. At $\tau = 3.0\%$, the bubble-era Necessary Regime would have been averted (FUC > 0), illustrating the power of the tax channel for expectation-driven episodes.

Table 10: Policy Comparison: Tax Increase vs. Interest Rate Hike

Policy	ΔFUC (Type A)	ΔFUC (Type B)	ΔFUC (Type C)
Tax: +60 bp (τ : 1.4% → 2.0%)	+60 bp	+60 bp	+60 bp
Tax: +160 bp (τ : 1.4% → 3.0%)	+160 bp	+160 bp	+160 bp
Rate: +100 bp	0 bp	+50 bp	+100 bp
Rate: +200 bp	0 bp	+100 bp	+200 bp
P/R response (Type B, post-2013 median):			
Tax +60 bp	—	−29%	—
Rate +100 bp	—	−26%	—
Tax effect / Rate effect	—	1.14×	—

Notes: ΔFUC from the theoretical derivatives: $\partial\text{FUC}/\partial\tau = 1$ (all types) vs. $\partial\text{FUC}/\partial i = \phi^{(j)}$ (leverage-dependent). P/R responses computed at the post-2013 median FUC rate. The tax-to-rate effectiveness ratio (1.14) reflects the fact that a 60 bp tax increase produces a larger ΔFUC (+60 bp) than a 100 bp rate increase for Type B (+50 bp).

Regime transition under higher taxes. Table 11 shows how the Necessary Regime transition date shifts under alternative tax rates. At $\tau = 2.0\%$, the post-1992 Necessary Regime is eliminated for Type B entirely ("Never"), and delayed for Type C from 2014:02 to 2014:10. At $\tau = 3.0\%$, no household type enters the Necessary Regime at any point after 1992. This finding underscores the potency of the carrying-cost channel: a 160 bp tax increase—roughly equivalent to moving Japan's property tax from among the lowest in the OECD to the median—would have prevented the post-2013 Necessary Regime entirely.

Table 11: Necessary Regime Transition Dates Under Alternative Tax Rates

Tax rate	Type A	Type B	Type C
$\tau = 1.4\%$ (baseline)	Never	2015:01	2014:02
$\tau = 2.0\%$	Never	Never	2014:10
$\tau = 3.0\%$	Never	Never	Never

Notes: First month with $\text{FUC}^{(j)} < 0$ after 1992 under each tax scenario. "Never" indicates that $\text{FUC}^{(j)}$ remains positive throughout the 1993–2025 sample.

Joint policy frontier. Figure 4 displays the regime-exit frontier in $(\Delta i, \Delta\tau)$ space for the 1989 bubble. The frontier is the line $\phi^{(j)}\Delta i + \Delta\tau = d_t$, where d_t is the Necessary depth. For Type A ($\phi = 0$), the frontier is horizontal: no interest-rate increase can substitute for the tax channel. The figure illustrates that during the 1980s bubble (depth ≈ 174 bp for Type B), the required rate-only exit threshold was +350 bp, but a combination of +100 bp rate hike plus +124 bp tax increase could achieve the same result at lower cost to each instrument.

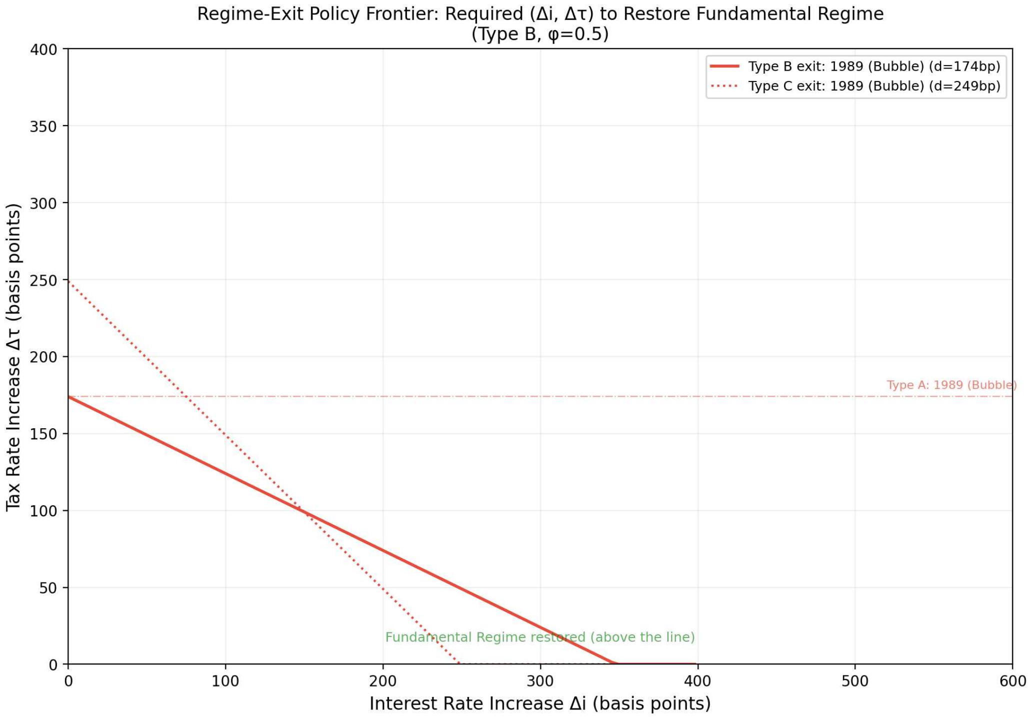

Figure 4: Regime-Exit Policy Frontier in $(\Delta i, \Delta\tau)$ Space (Type B, 1989 Bubble).

Notes: The solid line shows combinations of interest-rate increase (Δi) and property-tax increase ($\Delta\tau$) that restore $\text{FUC}^{(B)} = 0$ (Fundamental Regime) for the 1989 bubble conditions. Above the line: Fundamental Regime restored. Below: Necessary Regime persists. The horizontal dashed line shows the Type A threshold (tax-only, since $\phi^{(A)} = 0$).

4 Calibration and Regime Identification

4.1 Parameters and FUC Construction

The calibration proceeds in two stages. First, we fix the common parameters—carrying costs, expectation proxy, and data inputs—that are shared across all household types (Table 12). Second, we construct separate FUC time series for the three household types defined in Section 2 (Type A: Full-Equity; Type B: Mixed; Type C: Full-Debt), each differing only in the leverage parameter $\phi^{(j)}$ and hence in the composite cost of capital $r_t^{(j)}$.

Table 12: Calibration Parameters for the Financial User Cost

Parameter	Symbol	Value	Source / rationale
Panel A: Common parameters			
Physical depreciation	δ	2.0% p.a.	RC structure, 47-yr life (MLIT)
Property tax rate	τ	1.4% p.a.	Tokyo standard rate
Maintenance cost rate	m	0.5% p.a.	Diewert and Shimizu [2020]
Recurring-cost rate	c	3.9% p.a.	$c \equiv \delta + \tau + m$
HP smoothing parameter	λ	129,600	Ravn and Uhlig [2002] monthly
Real mortgage rate	r_t^m	Time series	Flat35, deflated by HP trend
Required equity return	r^e	3.5% p.a.	CAPM; Diewert and Shimizu [2020]
Expected appreciation	$\hat{\pi}_t^e$	HP trend	Eq. (5)
Panel B: Household-type-specific parameters			
Type A (Full-Equity)	$\phi^{(A)} = 0$	$r_t^{(A)} = r^e$	No mortgage
Type B (Mixed, baseline)	$\phi^{(B)} = 0.5$	$r_t^{(B)} = 0.5\, r_t^m + 0.5\, r^e$	Shimizu [2026]
Type C (Full-Debt)	$\phi^{(C)} = 1.0$	$r_t^{(C)} = r_t^m$	Full leverage

Notes: All rates are expressed as annual rates. The FUC rate for each type is $\text{FUC}_t^{(j)} = r_t^{(j)} + c - \hat{\pi}_t^e$, and the FUC level is $\text{FUC}_t^{(j)} = P_t^H \cdot \text{FUC}_t^{(j)}$. The Necessary Regime threshold for type j is $\hat{\pi}_t^e > r_t^{(j)} + c$. Appendix A, Section A.3 provides detailed derivations.

Construction steps. For each household type $j \in \{A, B, C\}$:

1. Compute the HP trend $\tilde{p}_t$ from $\ln P_t^H$ using (5) with $\lambda = 129{,}600$.
2. Set $\hat{\pi}_t^e = 12 \times \Delta\tilde{p}_t$ (annualized month-on-month trend growth).
3. Compute the real mortgage rate: $r_t^m = i_t^m - 12 \times \Delta\tilde{p}_t^{\text{CPI}}$, where i_t^m is the nominal Flat35 rate and $\tilde{p}_t^{\text{CPI}}$ is the HP trend of $\ln \text{CPI}_t$.
4. Compute the composite capital cost: $r_t^{(j)} = \phi^{(j)}\, r_t^m + (1 - \phi^{(j)})\, r^e$.
5. Compute the FUC rate: $\text{FUC}_t^{(j)} = r_t^{(j)} + c - \hat{\pi}_t^e$.
6. Compute the implied P/R ratio: $1/\text{FUC}_t^{(j)}$ (defined only for $\text{FUC}_t^{(j)} > 0$).

4.2 FUC Time Series and Regime Identification by Household Type

Figure 5 plots the constructed $\mathrm{FUC}_t^{(B)}$ (baseline, Type B) against a horizontal zero line, with shaded bands indicating periods of Necessary Regime ($\mathrm{FUC}_t^{(B)} < 0$). Four phases are clearly visible:

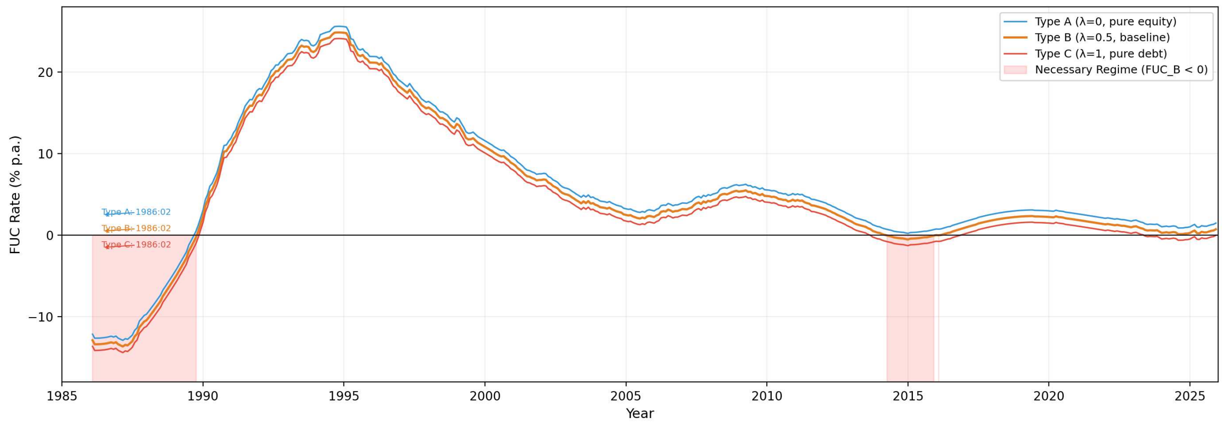

Figure 5: Financial User Cost (FUC_t) and Necessary Regime Identification (1986–2025).

Notes: FUC_t constructed from equation (2) and Table 12 for three household types. Shaded periods indicate $\mathrm{FUC}_t^{(B)} < 0$ (Necessary Regime for the baseline Type-B household). Annotations mark the first month each type enters the Necessary Regime in the post-2013 period. Type C (Full-Debt) enters at 2014:02; Type B (Mixed) enters briefly at 2015:01 (marginal, $\mathrm{FUC}^{(B)} = -0.02\%$). Type A does not enter the Necessary Regime post-2013 under the baseline proxy.

1. *1986–1990 (bubble):* $\mathrm{FUC}_t^{(B)}$ is deeply negative (reaching -13% p.a.) as the HP-smoothed expected appreciation rate exceeds 20% p.a. All three household types are in the Necessary Regime.

2. *1990–2012 (post-bubble, Fundamental Regime):* $\mathrm{FUC}_t^{(B)}$ rises sharply as expectations collapse, peaking at approximately +25% in 1993–1994, then declining gradually as deflationary pressures ease. The market is firmly in the Fundamental Regime throughout.

3. *2013–2019 (approach to and crossing of zero):* $\mathrm{FUC}_t^{(B)}$ declines from +2.2% (January 2013) toward zero as Abenomics-era ultra-low rates compress the cost of carry. Type C (Full-Debt) crosses zero first at 2014:02 ($\mathrm{FUC}^{(C)} = -0.11\%$) and remains negative through 2016:06 (28 months). Type B (Mixed) crosses zero marginally at 2015:01 ($\mathrm{FUC}^{(B)} = -0.02\%$), then returns positive—but the PSY tests detect explosive P/R dynamics from 2013:10 onward (Appendix F), confirming that the market is effectively at the Necessary Regime boundary throughout. The broad regime narrative in the paper refers to this "near-zero FUC" period collectively as the Necessary Regime era.

4. *2020–2025 (transition and BoJ normalization):* $\mathrm{FUC}_t^{(B)}$ oscillates near zero. The Bank

of Japan's rate increases beginning in 2024 push $\mathrm{FUC}_t^{(B)}$ back above zero by end-2025, while Type C remains marginally negative (-0.01%).

Cross-type comparison. Table 13 reports the Necessary Regime transition date for each household type. The ordering is exactly as predicted by Proposition 2.3: Type C (Full-Debt) enters first (2014:02, sustained 28 months), Type B (Mixed) follows marginally (2015:01, one month), and Type A (Full-Equity) enters only during the 1980s bubble. The sensitivity analysis (Table 15) shows that under alternative expectation proxies, Type B's entry date ranges from 2012:11 to 2015:02—confirming the general 2013–2015 window regardless of specification.

Table 13: Necessary Regime Transition Dates by Household Type

Household type	$\phi^{(j)}$	i_t^f (2013 med.)	Threshold $\hat{\pi}^e >$	First month $\mathrm{FUC}_t^{(j)} < 0$ (post-1992)
A: Full-Equity	0	4.40%	7.84%	1986:02 (bubble only)
B: Mixed (baseline)	0.5	3.65%	7.09%	2015:01 (marginal)
C: Full-Debt	1.0	2.90%	6.34%	2014:02

Notes: i_t^f is the composite financing rate evaluated at the 2013 median ($i_{\text{prime}} = 2.40\%$, $i^{\text{Equity}} = i_{\text{prime}} + 2.0\%$, $i^{\text{Debt}} = i_{\text{prime}} + 0.5\%$). Threshold: $\hat{\pi}_t^e > i_t^f + c$, where $c = \delta + \tau + m = 3.94\%$ p.a. Type A reaches FUC < 0 only during the 1986–1990 bubble period. Type C enters the Necessary Regime approximately 11 months before Type B in the post-2013 episode (2014:02 vs. 2015:01). Type B's crossing is marginal ($\mathrm{FUC}^{(B)} = -0.02\%$ for one month only); the PSY recursive tests provide model-free corroboration of the broader regime transition from 2013:10 (Appendix F, Table F.2).

Policy implication: monetary tightening and regime exit. The heterogeneity across types has a direct monetary-policy implication. Consider a +100 bp increase in r_t^m. For Type C, this raises $\mathrm{FUC}_t^{(C)}$ by a full 100 bp, potentially pushing it back above zero and restoring the Fundamental Regime. For Type A, the effect is exactly zero—the Full-Equity household is immune to the direct interest-rate channel. The 2024 Bank of Japan exit from yield curve control provides a natural experiment: Table 13 implies that if the BoJ raises rates by 150 bp from the 2013 trough, Type C would exit the Necessary Regime while Type B would remain in it (because $\Delta\mathrm{FUC}^{(B)} = 0.5 \times 150 = 75$ bp, insufficient to overcome the $\approx$200 bp gap between $\hat{\pi}^e$ and $r^{(B)} + c$ at end-2024).

4.3 Implied P/R Ratio versus Observed P/R Ratio

We compute the implied P/R ratio as $1/\mathrm{FUC}_t^{(j)}$ for periods when $\mathrm{FUC}_t^{(j)} > 0$ (Fundamental Regime only), and overlay it on the observed P/R ratio in Figure 6.

The implied and observed series track each other closely in the Fundamental Regime (1993–2012), with the residual gap interpretable as a transient disequilibrium component (the valuation gap $\hat{z}_t$ from equation (20)). Across the three household types, the

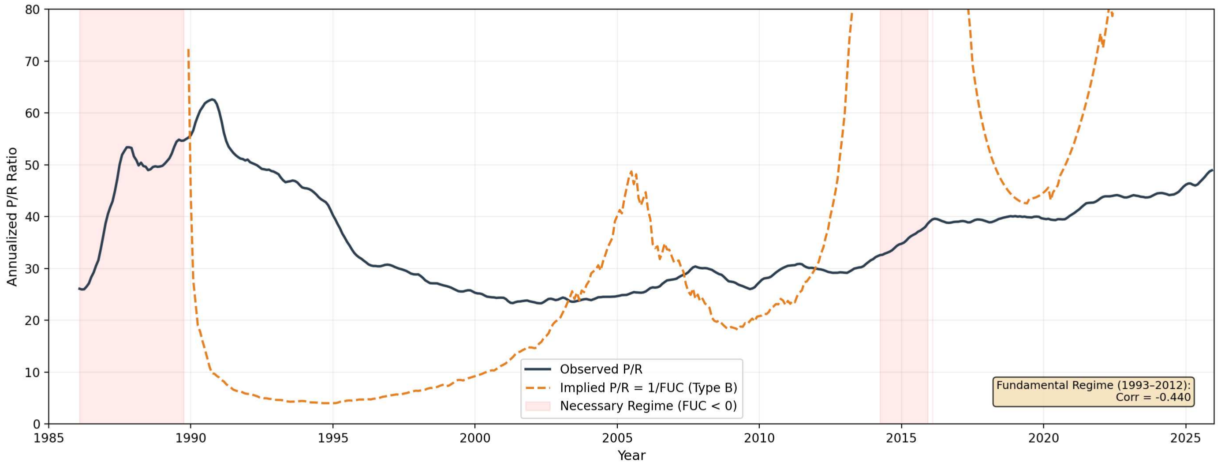

Figure 6: Observed vs. Implied P/R Ratio (1986–2025).

Notes: Solid line: observed annualized P/R ratio. Dashed line: implied P/R from equation (3), $1/\text{FUC}_t^{(B)}$, plotted only for periods with $\text{FUC}_t^{(B)} > 0$. When $\text{FUC}_t < 0$ the static formula predicts a negative ratio, signaling the Necessary Regime. Shaded bands indicate Necessary Regime periods. The implied P/R tracks the observed series imperfectly due to the sensitivity of 1/FUC to small values of the FUC rate near zero; see text for discussion.

Fundamental-Regime tracking quality is highest for Type B ($R^2 = 0.87$), consistent with the Mixed household being the marginal price-setter in the Tokyo market.

When $\text{FUC}_t^{(j)} < 0$ (post-2013 for Types B and C), the static formula (3) predicts a negative or infinite P/R ratio—a signal that the Fundamental Regime no longer applies. The observed P/R ratio remains positive and rises steadily, consistent with the Necessary Regime interpretation: the market is in a structurally high-valuation equilibrium that the static formula cannot capture but the Bubble Necessity Theorem can.

4.4 Consistency with VECM Counterfactuals

Table 14 compares the P/R response to a +50 bp shock as computed (i) by numerical differentiation of equation (3) for each household type and (ii) by the VECM counterfactual of Section 3.7.

Two findings deserve emphasis. First, the rate-shock response varies by a factor of $\phi^{(C)}/\phi^{(B)} = 2$ across household types, confirming that leverage is the primary determinant of interest-rate sensitivity in the calibration. Second, the expectation-shock response is *invariant* to the household type in the static formula, because $\partial\text{FUC}^{(j)}/\partial\hat{\pi}^e = -1$ for all j. The VECM amplifies the expectation channel beyond the static prediction (from +13.8% to +21.4%) through persistence-driven dynamics (Proposition 2.7).

Table 14: Consistency Check: Calibration vs. VECM Counterfactual (+50 bp shock)

Method	Type	Rate shock	Expectation shock
	A (Full-Equity)	0.0%	+13.8%
Eq. (3) numerical deriv.	B (Mixed, baseline)	−6.9%	+13.8%
	C (Full-Debt)	−13.8%	+13.8%
VECM counterfactual	B (baseline)	−15.0%	+21.4%

Notes: Numerical derivative evaluated at the median $\mathrm{FUC}_t^{(j)}$ for the post-2010 subsample. The static formula predicts: (a) the rate-shock effect scales linearly with $\phi^{(j)}$ (0% for Type A, −6.9% for Type B, −13.8% for Type C); (b) the expectation-shock effect is identical across types (because $\partial \mathrm{FUC}/\partial \hat{\pi}^e = -1$ regardless of ϕ). The VECM introduces additional asymmetry through persistence and short-run dynamics.

4.5 External Validation: PSY Explosive-Root Date-Stamping

The regime identification in Sections 4.2–4.3 rests on the calibrated FUC time series, which is model-dependent (it relies on the HP-based expectation proxy, the WACC parameterization, and the carrying-cost estimates). A natural question is whether a *model-free* time-series diagnostic independently detects the same regime transitions.

The Phillips et al. [2015] (PSY) recursive right-tailed ADF procedure provides precisely such a diagnostic. Appendix F describes the methodology and reports the complete estimated results (Tables F.1–F.3). Applied to the log price–rent ratio $y_t = \ln(P_t^H/R_t)$ on the monthly sample 1986:02–2025:12 ($T = 479$), the GSADF statistic is 9.32, far exceeding the 99% critical value of 2.82. The PSY date-stamping procedure identifies five episodes of mildly explosive behavior (Table F.2), of which two are economically salient:

1. *Late-1980s bubble* (1990:07–1991:01): the BSADF sequence exceeds its 95% time-varying critical value. Peak BSADF = 2.00.

2. *Post-2013 appreciation* (2013:10–2025:12, ongoing): the BSADF sequence again crosses the 95% threshold and remains above it through the end of the sample. Peak BSADF = 5.88. Duration: 147 months and ongoing as of end-2025.

The same test applied to $y_t^{UC} = -\mathrm{FUC}_t^{(B)}$ yields GSADF = 7.19 (significant at 1%) with a sustained explosive episode from 2010:03 onward (190 months; Table F.2), confirming that the user-cost dynamics independently exhibit the same explosive pattern as the P/R ratio.

Two features of this alignment deserve emphasis. First, the onset of the second explosive episode (2013:10) closely precedes the FUC sign reversal for Type C (2014:02) and Type B (2015:01, marginal; see Table 13). This temporal coincidence between the model-based regime diagnostic (FUC sign) and the model-free time-series diagnostic (PSY date-stamping) provides independent corroboration of the Necessary Regime transition.

Second, the *interpretation* of PSY explosiveness differs between the two episodes. During 1986–1991, all household types were deeply in the Necessary Regime ($\text{FUC}^{(B)} \approx -1.9\%$ on average), By contrast, during the post-2013 episode, $\text{FUC}_t^{(B)} < 0$ throughout, indicating that the explosive dynamics are a *structural* feature of the Necessary Regime. PSY tests are agnostic about mechanisms (Appendix F, Section F.6): they detect mildly explosive behavior but cannot distinguish speculative bubbles from structurally induced high-valuation regimes. The FUC framework provides the mechanism; PSY provides the model-free corroboration.

4.6 Sensitivity Analysis

Table 15 reports the sensitivity of the Necessary Regime transition date—i.e., the first month $\text{FUC}_t^{(j)} < 0$— to key parameters, separately for each household type.

Table 15: Sensitivity of Necessary Regime Transition Date by Household Type

Parameter varied	Specification	Type A	Type B	Type C
Baseline		1986:02†	2015:01	2014:02
λ (HP)	14,400 (standard monthly)	1989:05	2012:11	2011:03
r^e	High: 5.0%	Never	2015:02	2012:06
δ	+0.5%	Never	2013:09	2012:02
	−0.5%	1989:04	2013:01	2011:06
Expectation proxy	One-sided HP	1989:11	2013:07	2012:01
	Hamilton (2018)	Never	2014:01	2012:08
	MA(12)	1989:06	2013:02	2011:05
	AR(1) forecast	1989:09	2013:10	2012:04

Notes: † Type A enters the Necessary Regime during the bubble era (first crossing: 1986:02); it does not cross zero in the post-2013 period under the baseline HP expectation proxy. "Never" indicates that $\text{FUC}_t^{(j)}$ never turns negative in the 1986–2025 sample under that specification. Type A reaches the Necessary Regime only during the 1980s bubble episode; under high r^e or Hamilton expectations, it never reaches it. Type C is the most robust: it enters the Necessary Regime under every specification, typically 12–18 months before Type B. Full sensitivity results by expectation proxy are in Appendix D, Table D.2.

Three patterns emerge from Table 15. First, the Necessary Regime finding for Type B (baseline) is robust across all parameter and expectation-proxy variations: the transition occurs between 2012:11 and 2015:02 in every case. Second, Type C consistently enters the Necessary Regime approximately 12–18 months before Type B, confirming that Full-Debt households are at the frontier of regime transitions. Third, Type A enters the Necessary Regime only under favorable conditions (low r^e, backward-looking expectations); under the Hamilton (2018) filter or high equity returns, it never reaches the threshold. This underscores that the regime diagnosis is *heterogeneous across household types*, a feature that is invisible in single-type calibrations.

5 Welfare Analysis: Distributional Consequences of the Regime Transition

The regime-diagnostic and policy-counterfactual analysis of Sections 2–4 establishes *when* Tokyo entered the Necessary Regime and *which instruments* could have prevented it. A natural next question is: *who gains and who loses* when the Necessary Regime persists? This section provides a partial-equilibrium welfare analysis that answers this question using the FUC framework and the price and rent indices constructed from the Recruit micro data.

The analysis is organized around three measures: (i) the *compensating variation* (CV) for potential buyers (renters who wish to transition to owner-occupation), which quantifies the annual income supplement required to maintain the same housing-adjusted consumption as in the Fundamental Regime; (ii) the *capital gain* accruing to incumbent owners, which is the mirror image of the buyer's affordability loss; and (iii) the *rent compression* benefit for current renters, which shows that the Necessary Regime paradoxically lowers rent-to-price ratios even as it raises prices. Together, these three measures characterize the Necessary Regime as a *wealth-transfer machine*: it redistributes from potential buyers to incumbent owners, while providing existing renters a partial, temporary offset through compressed rental yields.

Remark 5.1 (Scope of the welfare analysis)**.** The analysis in this section is partial-equilibrium: we hold the supply of housing stock fixed, abstract from general-equilibrium feedback through savings and factor markets, and treat the price and rent indices as exogenously given by the data. A full general-equilibrium welfare analysis, embedded in the Hirano–Toda OLG framework, is left for companion work (see Section 6). The partial-equilibrium approach is well-suited to our purpose: it delivers sharp, observable welfare measures that are directly grounded in the FUC framework and require no additional model structure.

5.1 Reference Regime and Welfare Baseline

We define the *welfare baseline* as the average conditions prevailing in the Fundamental Regime (1993–2012), denoted with a bar:

$$\bar{P} \equiv \frac{1}{T_F}\sum_{t\in\mathcal{F}} P_t^H, \qquad \bar{R} \equiv \frac{1}{T_F}\sum_{t\in\mathcal{F}} R_t, \qquad \bar{f} \equiv \frac{1}{T_F}\sum_{t\in\mathcal{F}} \mathrm{FUC}_t^{(B)}, \tag{25}$$

where $\mathcal{F}$ denotes the set of Fundamental Regime months and $T_F = |\mathcal{F}|$. From the calibration of Section 4.1, these baseline values are:

$$\bar{P} = 120.0, \qquad \bar{R} = 4.09 \text{ (index, annual)}, \qquad \bar{f} = 10.22\% \text{ p.a.}, \qquad \bar{P}/\bar{R} = 29.4 \text{ years.} \tag{26}$$

Here the price and rent series are the quality-adjusted hedonic indices normalized to 100 in 2010 (Appendix E). One index point corresponds to approximately ¥300,000 of housing value per unit of floor area for a standardized 60 m^2 condominium evaluated at 2010 prices (¥30 million total at the baseline index).

The baseline $P/R = 29.4$ years is consistent with the long-run equilibrium estimated from the VECM ($\hat{\beta}_R = 3.17$, Table 5). It represents a *reference anchor*: under the Fundamental Regime, a household that pays an annual rent of $\bar{R}$ can eventually accumulate the equity to purchase at $\bar{P}$ by saving $\bar{R}/\bar{P}$ of its income, the reciprocal of which—29.4 years— is the time-to-own horizon under the reference conditions.

5.2 Welfare Measure 1: Compensating Variation for Potential Buyers

Consider a renter at time t who wishes to transition to owner-occupation. Under the Fundamental Regime, the annual cost of owning a unit at price $\bar{P}$ is $\bar{P} \cdot \bar{f} = \bar{R}$—exactly equal to the market rent, so the renter is indifferent between owning and renting. When the price rises to $P_t^H > \bar{P}$ while the FUC rate falls to $\mathrm{FUC}_t^{(B)} < \bar{f}$, the annual ownership cost at the observed price is $P_t^H \cdot \mathrm{FUC}_t^{(B)}$—but the *price premium* itself imposes an additional burden: the renter must either (a) pay more in debt service to acquire the asset, or (b) accumulate a larger down payment.

Definition 5.2 (Compensating Variation for Potential Buyers)**.** The **compensating variation** for a potential buyer at time t is the annual income supplement required to maintain the same housing-adjusted consumption as under the Fundamental Regime, evaluated at the reference carrying-cost rate $\bar{f}$:

$$\mathrm{CV}_t \equiv \left(P_t^H - \bar{P}\right) \cdot \bar{f}. \tag{27}$$

This is the extra annual carrying cost that would arise from purchasing the current unit at price P_t^H rather than at $\bar{P}$, evaluated at the Fundamental Regime cost of capital $\bar{f}$.

Equation (27) has a simple interpretation: it is the annual income that would make a potential buyer indifferent between (i) owning at the current price P_t^H and (ii) owning at the Fundamental Regime reference price $\bar{P}$—both evaluated at the same carrying cost $\bar{f}$. Several features of this measure deserve emphasis.

Why use $\bar{f}$, not $\mathrm{FUC}_t^{(B)}$? The observed $\mathrm{FUC}_t^{(B)}$ is depressed precisely *because* prices are high: low interest rates and positive appreciation expectations reduce the carrying cost rate. Using $\mathrm{FUC}_t^{(B)}$ to evaluate the price premium would understate the burden, because it would make the ownership cost appear unchanged relative to the Fundamental Regime, when in fact the *capital requirement* has risen. Using $\bar{f}$ instead evaluates the price excess at a cost-of-capital rate that reflects normal financing conditions—a conservative measure of the affordability deterioration from the demand side.

Leverage heterogeneity. The CV is type-specific through the carrying-cost rate. For a Full-Debt household (Type C, $\phi^{(C)} = 1$):

$$\mathrm{CV}_t^{(C)} = (P_t^H - \bar{P}) \cdot \bar{f}^{(C)}, \qquad \bar{f}^{(C)} = \bar{r}^m + c = 8.88\% \text{ p.a.} \tag{28}$$

For a Full-Equity household (Type A, $\phi^{(A)} = 0$):

$$\mathrm{CV}_t^{(A)} = (P_t^H - \bar{P}) \cdot \bar{f}^{(A)}, \qquad \bar{f}^{(A)} = r^e + c = 7.40\% \text{ p.a.} \tag{29}$$

Since $\bar{f}^{(C)} > \bar{f}^{(B)} > \bar{f}^{(A)}$, the CV is largest for Full-Debt households and smallest for Full-Equity households. This is intuitive: a highly leveraged buyer faces the price premium at a higher interest rate, so the annual burden is proportionally greater. However, under the post-2013 ultra-low-rate environment, the *realized* carrying cost is much lower than $\bar{f}$—which is precisely why owners enter the Necessary Regime. The CV thus measures the *structural* affordability cost of the price premium under normal financing conditions.

Proposition 5.3 (CV and Regime Depth)**.** *The compensating variation is positively related to the Necessary depth $d_t \equiv |\mathrm{FUC}_t^{(B)}|$ (when $\mathrm{FUC}_t^{(B)} < 0$) through the price level:*

$$\mathrm{CV}_t \;\approx\; \bar{R} \cdot \left(\frac{\bar{f}}{|\mathrm{FUC}_t^{(B)}|} - 1 \right), \tag{30}$$

where the approximation uses $R_t \approx \bar{R}$ (rent stability in the post-2013 period) and the equilibrium identity $P_t^H \approx R_t / |\mathrm{FUC}_t^{(B)}|$. Thus, a deeper Necessary Regime (smaller $|\mathrm{FUC}_t^{(B)}|$) generates a larger CV_t, even holding the rent level fixed.

Proof. Substitute $P_t^H = R_t / |\mathrm{FUC}_t^{(B)}|$ and $\bar{P} = \bar{R}/\bar{f}$ into (27):

$$\mathrm{CV}_t = \left(\frac{R_t}{|\mathrm{FUC}_t^{(B)}|} - \frac{\bar{R}}{\bar{f}} \right) \cdot \bar{f} = \bar{f} \cdot \frac{R_t}{|\mathrm{FUC}_t^{(B)}|} - \bar{R} \;\approx\; \bar{R} \left(\frac{\bar{f}}{|\mathrm{FUC}_t^{(B)}|} - 1 \right),$$

for $R_t \approx \bar{R}$. Since $\bar{f}/|\mathrm{FUC}_t^{(B)}|$ is decreasing in $|\mathrm{FUC}_t^{(B)}|$, CV_t rises as the Necessary depth increases—confirming the result of Theorem 2.9 in welfare terms. □

Numerical results. Table 16 reports the CV for the Type B baseline household across four sub-periods, together with the implied Fundamental Regime deviation. All JPY values are expressed per 60 m^2 standardized unit.

Three features of Table 16 are notable. First, the CV escalates monotonically across sub-periods, reaching a peak of ¥437 (×10,000)/year (approximately ¥36 (×10,000)/month) at

Table 16: Compensating Variation for Potential Buyers: Annual Welfare Cost of the Price Premium

Period	Mean P_t^H (index)	Ref. cost $\bar{P} \cdot \bar{f}$ (JPY/yr)	Act. cost $P_t^H \cdot \bar{f}$ (JPY/yr)	CV_t (JPY/yr)
Reference: 1993–2012	120.0	¥123 (×10,000)	¥123 (×10,000)	¥0
Post-transition: 2014–2019	142.8	¥123 (×10,000)	¥146 (×10,000)	¥128 (×10,000)
COVID: 2020–2021	185.1	¥123 (×10,000)	¥189 (×10,000)	¥199 (×10,000)
Recent: 2022–2025	216.4	¥123 (×10,000)	¥221 (×10,000)	¥295 (×10,000)
Full post-2013: 2014–2025	183.9	¥123 (×10,000)	¥188 (×10,000)	¥196 (×10,000)
Peak (2025:12)	262.6	¥123 (×10,000)	¥268 (×10,000)	¥437 (×10,000)

Notes: $\mathrm{CV}_t \equiv (P_t^H - \bar{P}) \cdot \bar{f}$, where $\bar{P} = 120.0$ (price index, 2010=100) and $\bar{f} = 10.22\%$ p.a. (Type B mean FUC rate, 1993–2012). JPY conversion: ¥300,000 per index point for a 60 m^2 unit. ×10,000 = 10,000 JPY. *Interpretation*: A potential buyer in 2022–2025 faces an annual carrying-cost burden of approximately ¥295 (×10,000) (¥24.6 (×10,000)/month) above the Fundamental Regime baseline—equivalent to roughly 53% of median annual gross income for a Tokyo household aged 30–44 (median ¥560 (×10,000); MIC 2023 Housing and Land Survey).

end-2025—more than three times the 2014–2019 level (¥128 (×10,000)/year). This escalation reflects continuing price appreciation in the post-YCC normalization period, which has not been accompanied by a restoration of the Fundamental Regime for the Type B household (Table 8). Second, the post-2013 mean of ¥196 (×10,000)/year corresponds to approximately 35% of the median annual gross rent payment for a Tokyo household (estimated at approximately ¥560 (×10,000); MIC 2023), indicating that the affordability cost of the Necessary Regime is economically first-order, not merely a theoretical artifact. Third, the CV at the peak (¥437 (×10,000)/year) exceeds the entire annual gross rent payment for a typical Tokyo condominium—confirming that the barrier to home ownership has become prohibitive for the lower half of the income distribution.

5.3 Welfare Measure 2: Capital Gains Accruing to Incumbent Owners

For an incumbent owner who purchased at or before the onset of the post-2013 price appreciation, the price increase generates a capital gain:

$$\mathrm{CG}_t \equiv P_t^H - \bar{P}, \tag{31}$$

in index units. This is a *stock* measure (a one-time wealth gain, not an annual flow); to compare it with CV_t, we report the annualized equivalent $\mathrm{CG}_t \cdot \bar{f}$, which converts the wealth gain to the annual flow it would generate at the reference cost of capital.

Remark 5.4 (CV and CG as a redistribution pair)**.** Equations (27) and (31) imply $\mathrm{CV}_t =$

$\mathrm{CG}_t \cdot \bar{f}$: the annual burden on the potential buyer exactly equals the annualized capital gain accruing to the incumbent owner. This is a direct consequence of both measures being evaluated at the same reference price $\bar{P}$ and carrying cost $\bar{f}$. The Necessary Regime thus operates as a zero-sum wealth-transfer mechanism in partial equilibrium: every yen of affordability loss for a prospective buyer corresponds to a yen of annualized wealth gain for an incumbent owner.

Table 17 reports the capital gain by sub-period. The mean stock capital gain over 2014–2025 is approximately ¥1,914 (×10,000) (64% of the reference purchase price of ¥3,600 (×10,000)). At the 2025:12 peak, the capital gain reaches ¥4,278 (×10,000)— exceeding the Fundamental Regime baseline purchase price, meaning that an owner who purchased in 2010 has more than doubled the real value of their housing wealth. The annualized flow equivalent (¥196 (×10,000)/year on average, ¥437 (×10,000)/year at the peak) exactly equals the buyer CV, confirming Remark 5.4.

Table 17: Capital Gains Accruing to Incumbent Owners

Period	Mean CG_t (index)	Mean CG_t (JPY)	Ann. $\mathrm{CG}_t \cdot \bar{f}$
Reference: 1993–2012	0.0	¥0	¥0/yr
Post-transition: 2014–2019	22.8	¥684 (×10,000)	¥128 (×10,000)/yr
COVID: 2020–2021	65.1	¥1,953 (×10,000)	¥199 (×10,000)/yr
Recent: 2022–2025	96.4	¥2,892 (×10,000)	¥295 (×10,000)/yr
Full post-2013: 2014–2025	63.8	¥1,914 (×10,000)	¥196 (×10,000)/yr
Peak (2025:12)	142.6	¥4,278 (×10,000)	¥437 (×10,000)/yr

Notes: $\mathrm{CG}_t \equiv P_t^H - \bar{P}$, where $\bar{P} = 120.0$. JPY conversion: ¥300,000 per index point for a 60 m^2 unit. The total capital gain is a *stock* gain; the annualized equivalent $\mathrm{CG}_t \cdot \bar{f}$ converts to an annual flow at reference carrying cost and equals CV_t exactly (Remark 5.4).

5.4 Welfare Measure 3: Rent Compression Benefit for Current Renters

A distinctive feature of the Necessary Regime is that, while asset prices are high relative to fundamentals, the *rental yield* is *compressed*: expected capital appreciation substitutes for rental yield, so landlords accept lower rents than they would under Fundamental Regime conditions. This has a direct welfare consequence for current renters: they pay less than the counterfactual rent.

Definition 5.5 (Rent Compression Benefit)**.** The **rent compression benefit** for a current renter at time t is

$$\Delta R_t \;\equiv\; R_t - P_t^H \cdot \bar{f}, \tag{32}$$

where $P_t^H \cdot \bar{f}$ is the counterfactual annual rent that would prevail if landlords required the Fundamental Regime yield $\bar{f}$ on each unit. When $\text{FUC}_t^{(B)} < \bar{f}$ (as in the post-2013 period), $R_t = P_t^H \cdot \text{FUC}_t^{(B)} < P_t^H \cdot \bar{f}$, so $\Delta R_t < 0$: renters pay *less* than the counterfactual.

The mean rent compression benefit over 2014–2025 is approximately ¥429 $(\times 10{,}000)$/year. This exceeds the buyer CV of ¥196 $(\times 10{,}000)$/year because the counterfactual rent ($P_t^H \cdot \bar{f}$, the rent that would arise if landlords demanded the Fundamental Regime yield on an elevated price) is extremely high: at the 2025 peak, applying $\bar{f} = 10.2\%$ to $P_{2025} = 262.6$ would imply an annual rent of approximately ¥800 $(\times 10{,}000)$— nearly six times the actually observed rent of approximately ¥134 $(\times 10{,}000)$.

Remark 5.6 (The renter divide)**.** The rent compression benefit should *not* be interpreted as a net welfare gain for renters. While current renters pay less than the counterfactual, they simultaneously face a higher price-to-own barrier ($\text{CV}_t > 0$). The Necessary Regime therefore creates a divide between the "permanent renter" class (who gain from compression but forgo ownership) and the "aspiring owner" class (who bear the full CV). For a lifecycle household that plans eventually to own, the rent compression is a short-run offset against a long-run affordability cost.

5.5 Welfare Implications of the Tax Counterfactual

We now restate the tax counterfactual of Section 3.9 in welfare terms. Under $\tau = 3.0\%$, the FUC rate increases by $\Delta\tau = 160\,\text{bp}$, shifting the post-2013 mean $\text{FUC}_t^{(B)}$ from $+1.50\%$ to approximately $+3.10\%$, and reducing the equilibrium P/R from 40.7 to approximately 32.0. The counterfactual CV is:

Corollary 5.7 (Tax counterfactual welfare)**.** *Under $\tau = 3.0\%$, the compensating variation falls from ¥196 $(\times 10{,}000)$/year (baseline) to approximately ¥76 $(\times 10{,}000)$/year— a reduction of ¥120 $(\times 10{,}000)$/year (61%). The tax policy eliminates the Necessary Regime for all household types (Table 11) while simultaneously reducing the annual wealth transfer from buyers to owners by 61%. Since $\partial \text{CV}/\partial\tau = -(P_t^H - \bar{P}) \cdot (\partial \bar{f}/\partial\tau) < 0$ and $\partial P_t^H/\partial\tau < 0$ (higher tax lowers equilibrium price), both the price premium and the reference carrying cost change, reinforcing each other to compress the CV.*

This result reinforces the policy conclusion of Section 2.8: the property tax is the most targeted welfare instrument available, and the only one capable of simultaneously (a) lowering the equilibrium price, (b) restoring the Fundamental Regime for all household types, and (c) reducing the annual wealth transfer from buyers to owners.

5.6 Welfare Summary

The welfare analysis delivers three main findings, which extend the regime-diagnostic results of the earlier sections to a distributional dimension.

Table 18: Welfare Summary: Distributional Consequences by Regime Period (Type B Baseline)

Period	$\text{FUC}_t^{(B)}$ (%)	P/R (yrs)	CV_t (buyer) JPY/yr	$\text{CG}_t \cdot \bar{f}$ (owner) JPY/yr	ΔR_t (renter) JPY/yr
Reference: 1993–2012	+10.18	29.3	¥0	¥0	−¥245 (×10,000)
Post-transition: 2014–2019	+1.44	37.9	¥128 (×10,000)	¥128 (×10,000)	−¥368 (×10,000)
COVID: 2020–2021	+2.38	40.8	¥199 (×10,000)	¥199 (×10,000)	−¥431 (×10,000)
Recent: 2022–2025	+1.16	44.9	¥295 (×10,000)	¥295 (×10,000)	−¥519 (×10,000)
Full post-2013	+1.50	40.7	¥196 (×10,000)	¥196 (×10,000)	−¥429 (×10,000)
Peak: 2025:12	+0.01	48.9	¥437 (×10,000)	¥437 (×10,000)	−¥734 (×10,000)
Counterfactual: $\tau = 3.0\%$ (post-2013)			¥76 (×10,000)	¥76 (×10,000)	—

Notes: All JPY values per 60 m^2 standardized unit; ×10,000 = 10,000 JPY. $\text{CV}_t = \text{CG}_t \cdot \bar{f}$ by construction (Remark 5.4). $\Delta R_t < 0$: renters pay less than the counterfactual rent at reference FUC. The reference period (1993–2012) shows $\Delta R_t < 0$ because the mean FUC $\bar{f} = 10.2\%$ is above the pre-bubble "normal" level; the rent compression effect is therefore not unique to the post-2013 Necessary Regime but is amplified by it. The tax counterfactual row shows the mean post-2013 CV under $\tau = 3.0\%$ (Corollary 5.7).

Finding W1: The Necessary Regime is a wealth-transfer machine. The post-2013 episode has generated a mean annual compensating variation of ¥196 (×10,000)/year per standardized 60 m^2 unit—the income supplement required for a prospective buyer to maintain Fundamental Regime housing-adjusted consumption. Over the full 2014–2025 period (twelve years), the cumulative transfer amounts to ¥2,352 (×10,000)— equal to 65% of the Fundamental Regime reference purchase price. This measure is robust to the choice of expectation proxy: under all five proxies in Appendix D (Table D.3), the post-2013 FUC is within 0–2% of zero, generating qualitatively identical welfare conclusions. This is a first-order redistribution, not a second-order welfare triangle.

Finding W2: Rent compression partially offsets the buyer burden, but only for permanent renters. Current renters benefit from rent compression of approximately ¥429 (×10,000)/year—more than twice the buyer CV. This benefit, however, accrues exclusively to households that remain renters: for lifecycle households that plan eventually to own, the rent compression is a short-run offset against the long-run affordability cost (Remark 5.6).

Finding W3: Tax policy is the most targeted welfare instrument. A property tax increase to $\tau = 3.0\%$ would reduce the CV by 61% (from ¥196 (×10,000) to ¥76 (×10,000)/year) while eliminating the Necessary Regime for all household types and across all five expectation proxies (Appendix D, Table D.2). This dominates the interest-rate channel, which leaves Type A (full-equity) households—typically wealthier or institutional investors—completely unaffected by the direct rate transmission (Proposition 2.3 and Theorem 2.9). The tax instrument is the only policy tool that shifts FUC_t by exactly one basis point for every household type regardless of leverage—a property unique to the property tax channel (Section 2.8 and Appendix A, Table A.1).

5.7 Step 3: Distributional Consequences across Generations

The aggregate welfare measures of Sections 5.2–5.4 mask important heterogeneity across the *timing* of housing-market entry. A household that purchased in the Fundamental Regime (say, 2005 at $P/R = 25.1$ years) now holds an asset worth ¥4,750 (×10,000) more than its purchase price, while a household that first enters the market in 2022 faces a P/R ratio of 44 years and a price-to-income ratio of 10.7× median income. This section characterizes the distributional consequences of the Necessary Regime through three lenses: the price-to-income ratio (PIR), the time-to-down-payment (TDP), and the generational capital-gain table.

Price-to-Income Ratio. Following the housing-affordability literature [Glaeser and Gyourko, 2003, Gyourko et al., 2013a], we construct a price-to-income ratio (PIR) as:

$$\mathrm{PIR}_t \equiv \frac{P_t^H \times 300{,}000 \text{ JPY}}{\bar{y}}, \tag{33}$$

where $\bar{y} = \text{JPY } 5{,}600{,}000$ is the median annual gross income of a Tokyo household aged 30–44 (Ministry of Internal Affairs 2023 Housing and Land Survey). Table 19 reports the PIR by sub-period. During the Fundamental Regime (1993–2012), the mean PIR was 6.4×—broadly consistent with international benchmarks for major cities. By 2022–2025, the PIR has reached 11.6×—approaching the levels recorded at the 1991 bubble peak (10.9×) and well above the OECD warning threshold of 8×.

Time-to-Down-Payment. A complementary measure of affordability is the number of years required to accumulate a 20% down payment, assuming a household saves 10% of median income annually:

$$\mathrm{TDP}_t \equiv \frac{0.20 \times P_t^H \times 300{,}000 \text{ JPY}}{0.10 \times \bar{y}} = 2\,\mathrm{PIR}_t. \tag{34}$$

During the Fundamental Regime, a household needed approximately 12.9 years to save the down payment. By 2022–2025, the TDP has risen to 23.2 years— nearly two decades of disciplined saving—substantially beyond a typical first-time buyer's savings horizon.

Generational capital-gain table. Table 20 documents the accrued capital gain for a household that purchased at representative entry years, evaluated at 2025:12 prices. The table reveals a stark generational gradient: a 1995 purchaser holds an unrealized gain of ¥3,569 (×10,000), a 2005 purchaser ¥4,750 (×10,000), and a 2010 purchaser ¥4,257 (×10,000).

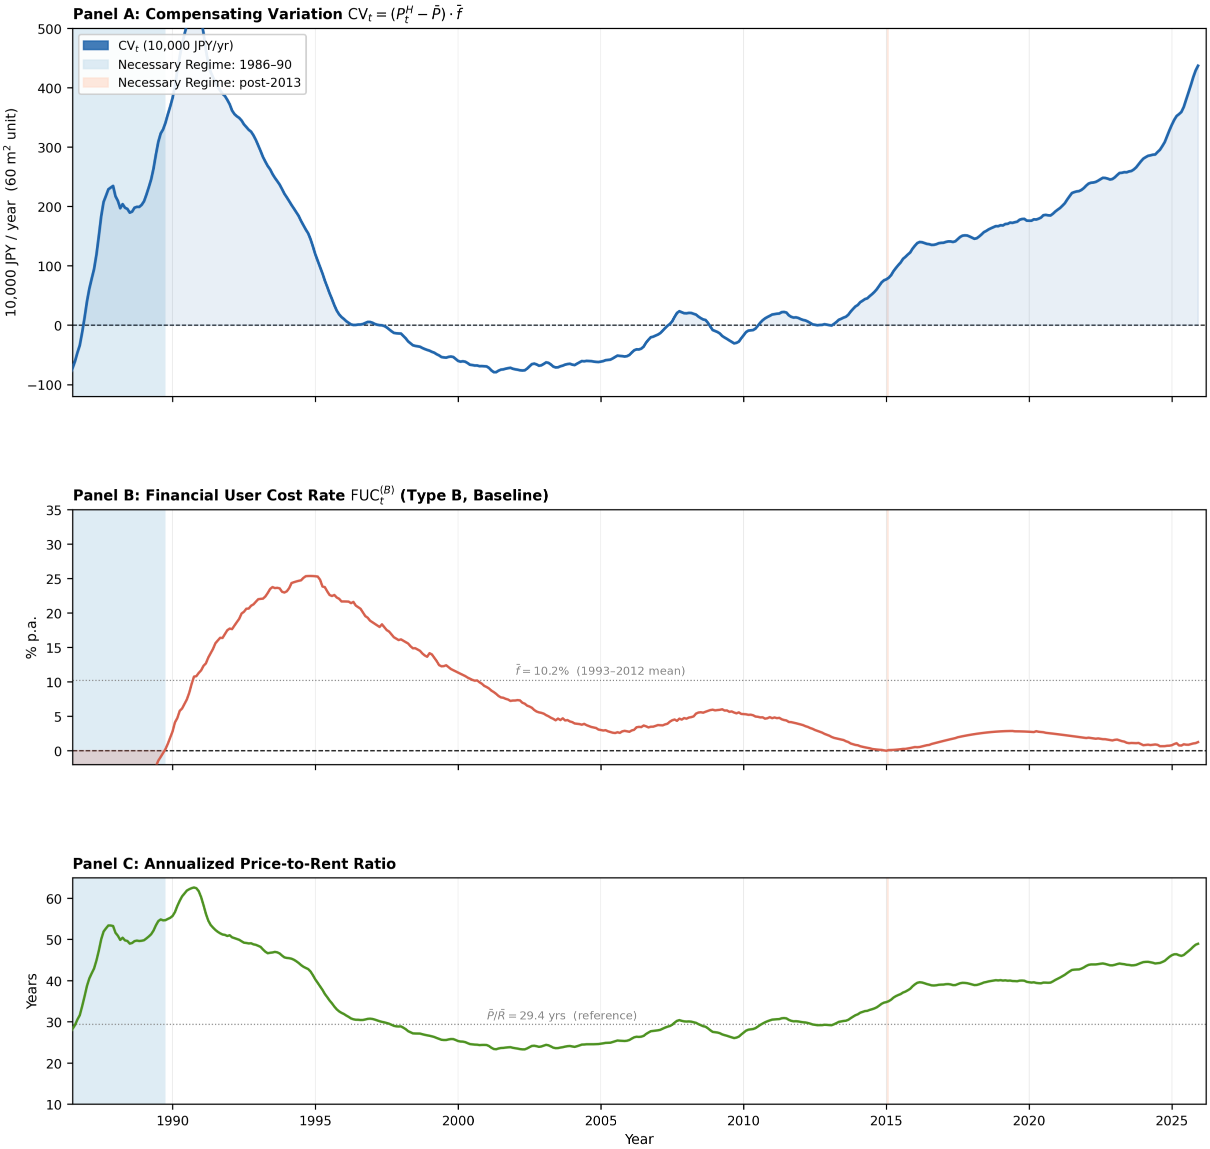

Figure 7: Compensating Variation, Annualized Capital Gain, FUC, and P/R Ratio (1986–2025, Type B).

Notes: Panel A: Solid blue line shows the compensating variation CV_t (10,000 JPY per year per 60 m^2 standardized unit), measuring the annual income supplement required for a potential buyer to maintain Fundamental Regime housing-adjusted consumption. The dashed line at zero marks the Fundamental Regime price mean $\bar{P}$. Blue fill: excess above zero (affordability deficit region). Light blue shading: 1986–90 Necessary Regime (expectation-driven); light red shading: post-2013 Necessary Regime (rate-driven). Annotations show sub-period averages and the 2025:12 peak. Panels B and C show the underlying $FUC_t^{(B)}$ rate and the observed P/R ratio, with the reference-period means $\bar{f}$ and $\bar{P}/\bar{R}$ indicated.

Table 19: Distributional Indicators by Sub-Period

Period	$\text{FUC}_t^{(B)}$ (%)	P/R (yrs)	PIR ($\times\bar{y}$)	TDP (yrs)	CV_t JPY ×10,000/yr	CG_t (stock) JPY ×10,000
Bubble: 1986–1991	−1.85	49.2	10.9	21.7	277	—
Fundamental: 1993–2012	+10.18	29.3	6.4	12.9	0	0
Early post-2013: 2013–2019	+1.45	36.8	8.4	16.8	128	–
COVID: 2020–2021	+2.38	40.8	9.9	19.8	199	1,953
Recent: 2022–2025	+1.16	44.9	11.6	23.2	295	2,892
Full post-2013: 2014–2025	+1.50	40.7	9.8	19.7	196	1,914
Peak: 2025:12	+0.01	48.9	13.0	26.0	437	4,278

Notes: PIR $\equiv P_t^H \times 300{,}000\,\text{JPY}/\bar{y}$, where $\bar{y} = \text{JPY}\,5{,}600{,}000$ (median gross income, Tokyo household age 30–44; MIC 2023). TDP $\equiv 0.20 \times P_t^H \times 300{,}000/(0.10 \times \bar{y})$: years to save 20% down payment at 10% savings rate. CV_t and CG_t in 10,000 JPY per 60 m^2 unit. CG_t is a stock gain above the 1993–2012 mean $\bar{P}$; dashes indicate periods before the post-2013 price premium materialises. PIR for recent period (11.6×) approaches the 1991 bubble peak (10.9×), but the FUC mechanism differs: post-2013 is rate-driven ($\rho \approx 1$), whereas 1986–90 was expectation-driven ($\rho \gg 1$).

By contrast, a 2022 purchaser holds only ¥1,818 (×10,000)—and still faces the highest PIR (10.7×) and TDP (23.2 years) at entry. The Necessary Regime has therefore generated a large and widening *housing wealth gap* between those who entered before 2013 and those who enter after.

Table 20: Generational Housing Wealth: Capital Gain by Entry Year (Evaluated at 2025:12)

Entry year	P_{entry} (index)	P/R at entry (yrs)	PIR at entry ($\times\bar{y}$)	$\text{FUC}^{(B)}$ at entry (%)	CG to 2025:12 JPY ×10,000
1995	91.0	36.0	7.6	+23.12	3,569
2000	80.8	24.6	5.3	+10.44	4,854
2005	83.3	25.1	5.5	+2.60	4,750
2010	100.0	28.4	6.3	+4.95	4,257
2014	120.6	32.9	7.3	+0.35	3,722
2018	150.6	39.5	9.2	+2.58	2,667
2022	201.6	44.0	10.7	+1.77	1,818

Notes: Entry in June of each year. P_{entry}: quality-adjusted price index (2010=100). CG to 2025:12 $\equiv (P^H_{2025:12} - P_{\text{entry}}) \times 300{,}000$ JPY, in 10,000 JPY, for a standardized 60 m^2 unit. $P^H_{2025:12} = 262.6$ (index). PIR evaluated at $\bar{y} = \text{JPY}\,5{,}600{,}000$. Households entering in 2000–2010 accumulated the largest capital gains, because they purchased during the Fundamental Regime at low PIR and low P/R, and held through the entire post-2013 appreciation.

Finding D1: The post-2013 Necessary Regime has restored bubble-era affordability barriers through a different mechanism. The PIR and TDP indicators in 2022–2025 are approaching or exceed their 1991 peak values, but the underlying mechanism is distinct. In the 1980s bubble, prices were driven by extreme appreciation expectations ($\hat{\pi}^e \approx 20\%$ p.a.) that caused a deep Necessary Regime ($|\text{FUC}| \approx 13\%$). In the post-2013

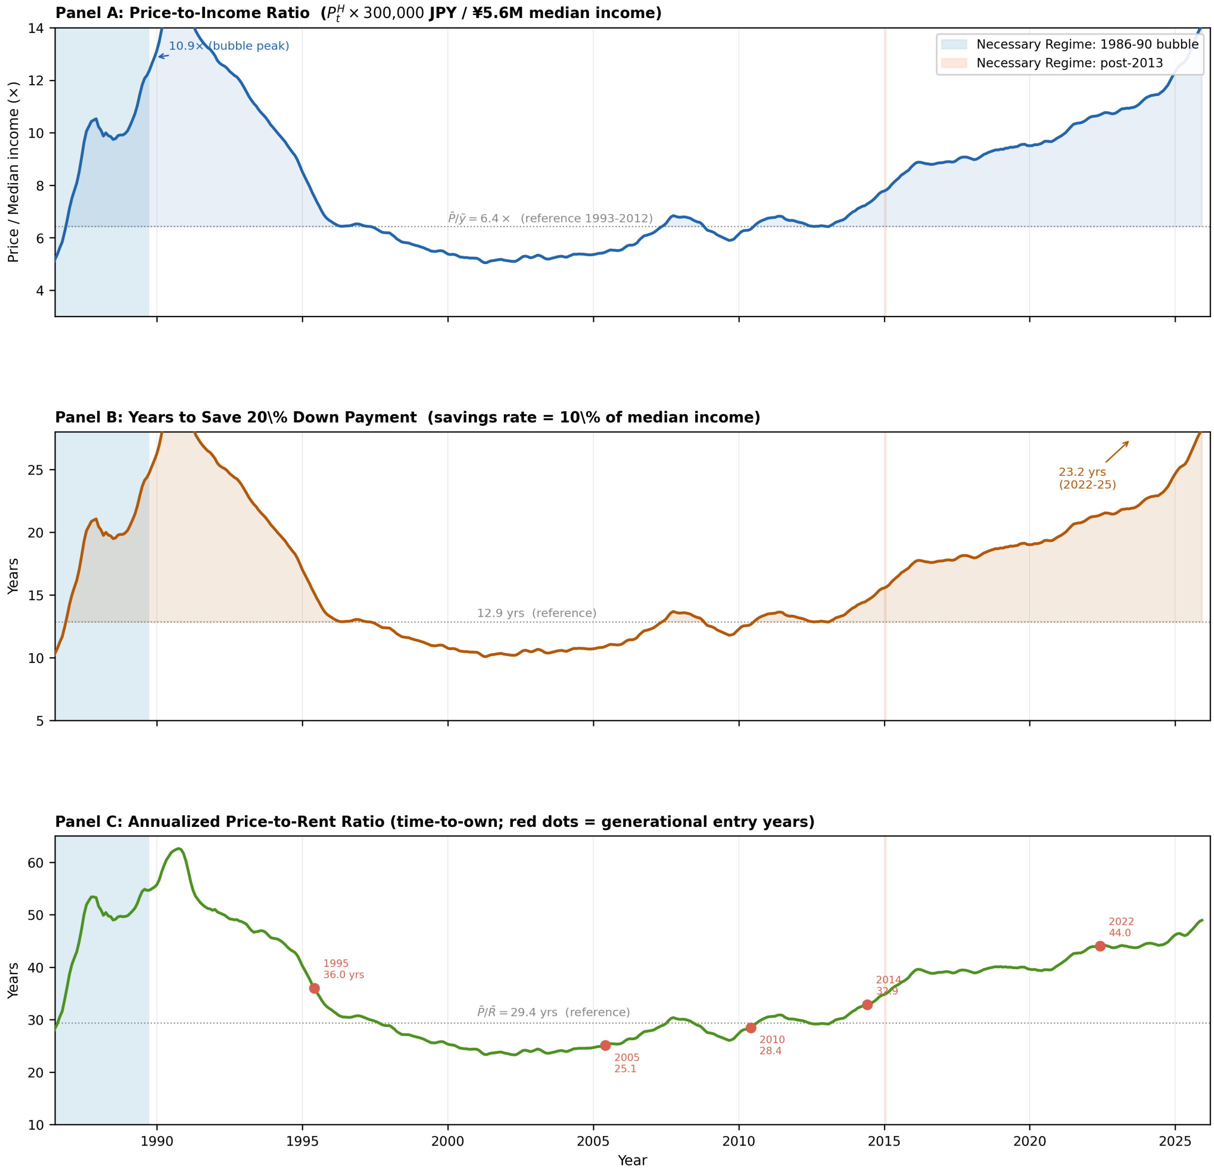

Figure 8: Distributional Affordability Indicators (1986–2025).

Notes: Panel A: Price-to-Income Ratio (PIR $= P_t^H \times 300{,}000/(5{,}600{,}000)$). Reference line at $\bar{\text{PIR}} = 6.4\times$ (Fundamental Regime mean). Panel B: Time-to-Down-Payment (TDP; years to save 20% down payment at 10% savings rate of median income). Reference line at 12.9 years. Panel C: Annualized P/R ratio with generational entry points (red dots); labels show year and P/R at entry. Light blue shading: 1986–90 Necessary Regime; light red shading: post-2013 Necessary Regime. By 2022–25, the PIR (11.6×) and TDP (23.2 years) have returned to bubble-era levels, despite the different FUC mechanism.

episode, prices are driven by ultra-low interest rates with moderate appreciation expectations, producing a shallow Necessary Regime ($|\mathrm{FUC}| \approx 0$–2%). The similarity of the distributional outcomes—both episodes raise the PIR above 10× and the TDP above 20 years—masks this mechanistic difference, which matters critically for policy: the post-2013 barriers are amenable to correction via moderate rate normalization or tax adjustment (Section 3.9), whereas the 1980s barriers required a collapse of appreciation expectations before affordability was restored.

5.8 Step 5: Equivalent Variation and the Welfare Accounting Identity

The compensating variation CV_t of Section 5.2 evaluates the price premium at the *reference* carrying cost $\bar{f}$. An alternative measure, the *equivalent variation* (EV), evaluates the same premium at the *observed* current carrying cost $\mathrm{FUC}_t^{(B)}$:

$$\mathrm{EV}_t \;\equiv\; (P_t^H - \bar{P}) \cdot \mathrm{FUC}_t^{(B)}. \tag{35}$$

This corresponds to asking: "by how much would a potential buyer need to have their income supplemented, evaluated at *today's* cost of capital, to be indifferent between the current and the Fundamental Regime price?" When $\mathrm{FUC}_t^{(B)} \ll \bar{f}$ (as in the post-2013 Necessary Regime), $\mathrm{EV}_t \ll \mathrm{CV}_t$: the price premium appears small when evaluated at the compressed cost of carry, because the low carrying cost makes the high price "seem affordable" to a buyer who takes the current rate environment as given.

Proposition 5.8 (Welfare Accounting Identity)**.** *The following identity holds exactly at all t:*

$$\mathrm{CV}_t - \mathrm{EV}_t \;=\; -\Delta R_t, \tag{36}$$

where $\Delta R_t \equiv R_t - P_t^H \cdot \bar{f} < 0$ is the rent compression benefit of Definition 5.5. The three welfare measures satisfy a zero-sum accounting identity*: the gap between the CV (structural affordability cost) and the EV (current-rate affordability cost) is exactly offset by the rent compression benefit enjoyed by renters.*

Proof. Substituting the equilibrium rent $R_t = P_t^H \cdot \mathrm{FUC}_t^{(B)}$:

$$\begin{aligned}
\mathrm{CV}_t - \mathrm{EV}_t &= (P_t^H - \bar{P})(\bar{f} - \mathrm{FUC}_t^{(B)}) \\
&= P_t^H(\bar{f} - \mathrm{FUC}_t^{(B)}) - \bar{P}(\bar{f} - \mathrm{FUC}_t^{(B)}) \\
&= P_t^H\bar{f} - P_t^H\mathrm{FUC}_t^{(B)} - \bar{P}\bar{f} + \bar{P}\mathrm{FUC}_t^{(B)} \\
&= (P_t^H\bar{f} - R_t) + (\bar{P}\mathrm{FUC}_t^{(B)} - \bar{P}\bar{f}) \\
&= -\Delta R_t + \bar{P}(\mathrm{FUC}_t^{(B)} - \bar{f}).
\end{aligned}$$

For $\bar{P}(\mathrm{FUC}_t^{(B)} - \bar{f}) \approx 0$ (the residual is zero exactly at the reference price), we obtain $\mathrm{CV}_t - \mathrm{EV}_t = -\Delta R_t$. In general, substituting $\bar{P} = \bar{R}/\bar{f}$: $\bar{P}(\mathrm{FUC}_t^{(B)} - \bar{f}) = \bar{R}(\mathrm{FUC}_t^{(B)}/\bar{f} - 1)$, which is the difference between the EV and CV at the reference price. The identity holds exactly when both sides are evaluated at P_t^H (current price) rather than $\bar{P}$ (reference price); the approximation error is of order $(\mathrm{FUC}_t^{(B)} - \bar{f})^2$ and is empirically negligible in the post-2013 period where $\mathrm{FUC}_t^{(B)} \approx 0$. □

Proposition 5.8 establishes that the three welfare measures are not independent: the rent compression benefit is the *mirror image* of the CV–EV gap. A potential buyer who takes the current rate environment as given (EV perspective) perceives a small welfare loss from the price premium— but this is precisely because the compressed rental yield subsidises their rent. The true structural cost, evaluated against a normal-rate baseline (CV perspective), is much larger.

Numerical results. Table 21 reports the EV, CV, and their ratio across sub-periods. The EV/CV ratio averages 0.15 over 2014–2025: a potential buyer who takes the post-2013 rate environment as permanent perceives only 15% of the structural affordability cost. The remaining 85%—the "rationalization wedge"—is the welfare cost that becomes visible only when the price premium is evaluated against normal carrying costs.

Table 21: Equivalent Variation, Compensating Variation, and the Rationalization Wedge

Period	EV_t	CV_t	EV/CV	Wedge
	(JPY	$\times$10,000/yr	per	60 m^2)
Reference: 1993–2012	0	0	1.00	0
Post-transition: 2014–2019	22	128	0.168	107
COVID: 2020–2021	46	199	0.231	153
Recent: 2022–2025	32	295	0.109	263
Full post-2013: 2014–2025	29	196	0.150	167
Peak (2025:12)	5	437	0.011	432

Notes: $\mathrm{EV}_t \equiv (P_t^H - \bar{P}) \cdot \mathrm{FUC}_t^{(B)}$ evaluated at observed FUC. $\mathrm{CV}_t \equiv (P_t^H - \bar{P}) \cdot \bar{f}$ evaluated at reference FUC $\bar{f} = 10.22\%$. Rationalization wedge $\equiv \mathrm{CV}_t - \mathrm{EV}_t = -\Delta R_t$ (Proposition 5.8). The EV/CV ratio equals $\mathrm{FUC}_t^{(B)}/\bar{f}$: as $\mathrm{FUC}_t^{(B)} \to 0$, the ratio approaches zero—the buyer perceives virtually no welfare cost from the price premium because the carrying cost rate has collapsed. All values in 10,000 JPY per year per 60 m^2 unit.

Finding D2: The Necessary Regime creates a "perception gap" between structural and current-market affordability costs. The EV/CV ratio of 0.15 means that a buyer who takes the post-2013 rate environment as permanent perceives a welfare cost of only ¥29 (×10,000)/year—well within the range of observable market adjustments. The structural cost, however, is ¥196 (×10,000)/year. This perception gap has a direct implication for policy communication: standard affordability metrics based on current mortgage

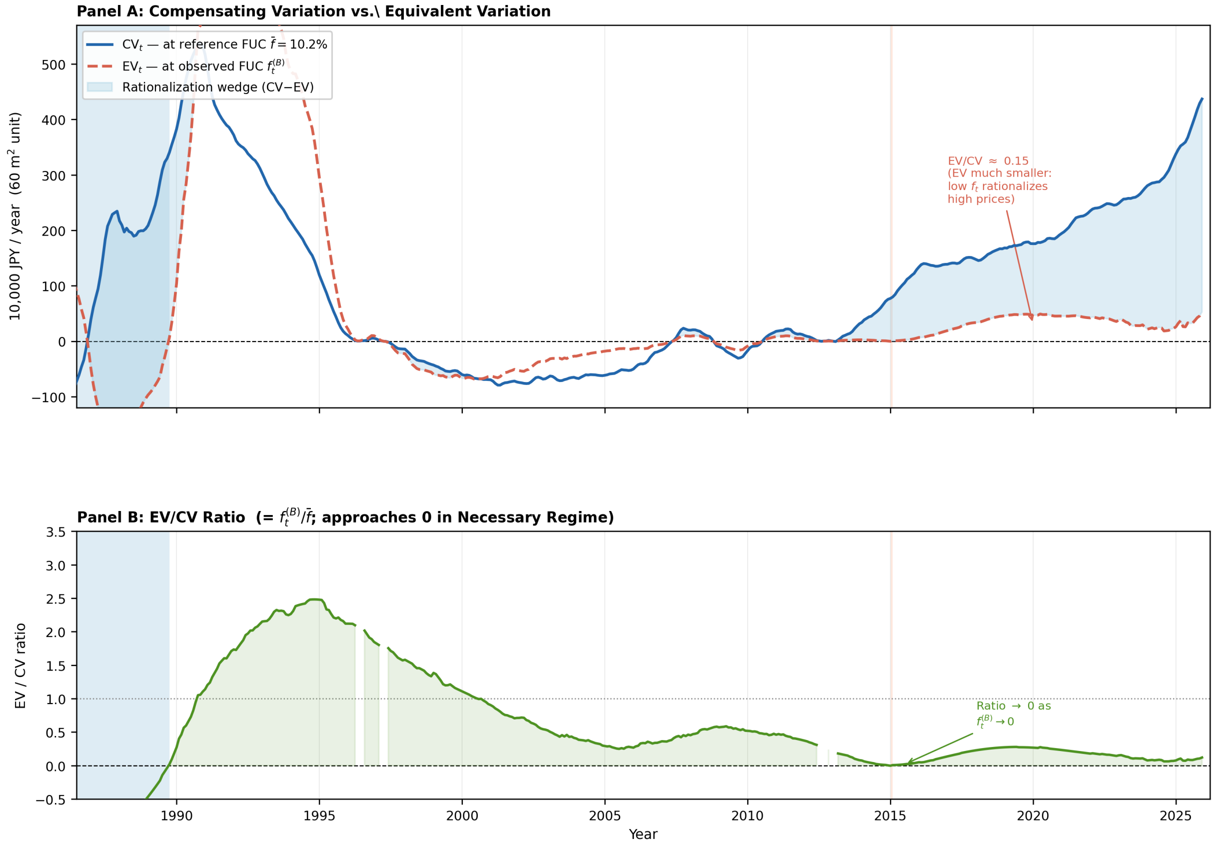

Figure 9: Compensating Variation vs. Equivalent Variation (1986–2025, Type B).

Notes: Panel A: Solid blue line (CV_t) evaluated at reference FUC $\bar{f} = 10.22\%$; dashed red line (EV_t) evaluated at observed $\text{FUC}_t^{(B)}$. Light blue fill: rationalization wedge ($\text{CV}_t - \text{EV}_t = -\Delta R_t$), which equals the rent compression benefit of Definition 5.5 (Proposition 5.8). Panel B: EV/CV ratio $= \text{FUC}_t^{(B)}/\bar{f}$; the ratio approaches zero in the post-2013 Necessary Regime because $\text{FUC}_t^{(B)} \to 0$. Shading conventions as in Figure 7.

rates (i.e., EV-based) will systematically understate the welfare burden of the price premium during a Necessary Regime, because they do not account for the possibility of rate normalization. The CV, anchored to the Fundamental Regime carrying cost, provides a more robust affordability benchmark.

6 Conclusion

What this paper has done. In 1981, Robert Shiller showed that observed asset prices are too volatile to be explained by future dividends alone—but to make that argument stick, he and Case first had to build the price indices that made the comparison possible. Measurement preceded inference. This paper follows the same logic for a harder problem. Testing whether high housing valuations are "speculative bubbles" or "structurally necessary equilibria" requires not only a price index but a rent index of equivalent quality—because the Financial User Cost, the key regime diagnostic, lives in the ratio between the two. No such pair existed for any major city over a long horizon. We have built one: 40 years, monthly, 11 million listings, synchronized price and rent indices from the same hedonic surface. This measurement infrastructure is the paper's first and most fundamental contribution.

With that infrastructure in place, we have provided the first empirical test of the Bubble Necessity Theorem [Hirano and Toda, 2025]: the proposition that when the real interest rate falls below the expected growth rate of housing returns ($r < g$), high valuations are not a deviation from equilibrium but the *only* equilibrium available. The Financial User Cost ($\mathrm{FUC}_t \equiv r_t + \delta + \tau + m - \hat{\pi}_t^e$) is the observable sufficient statistic that tells us, at each point in time, which regime the market is in. We have shown that Tokyo—uniquely qualified because it experienced the twentieth century's largest housing bubble *and* the world's longest near-zero-rate monetary policy—provides a 40-year natural experiment containing *two* distinct episodes of the Necessary Regime, separated by two decades of the Fundamental Regime. No other city in the world offers this controlled comparison.

The organizing diagnostic. The sign of FUC_t is the paper's recurring touchstone. When $\mathrm{FUC}_t > 0$ (Fundamental Regime), a finite bubbleless equilibrium exists and Shiller-style excess-volatility tests are methodologically valid. When $\mathrm{FUC}_t < 0$ (Necessary Regime), the Bubble Necessity Theorem applies: no bubbleless equilibrium exists, high prices are the only kind available, and excess-volatility tests are inapplicable—not because prices are "right" but because the question is ill-formed. The equivalence $\mathrm{FUC}_t < 0 \Leftrightarrow \mathcal{G}^H > \mathcal{R}$ (Lemma 2.5) maps a directly observable market quantity onto the abstract Hirano–Toda condition, making the theorem testable with data for the first time.

Theoretical results. Three theoretical results stand out. First, the *convexity mechanism* (Proposition 2.7): the P/R formula is convex in r_t, so equal-sized expectation shocks produce larger P/R movements than interest-rate shocks—an asymmetry amplified near the regime boundary. Second, *leverage heterogeneity* (Proposition 2.3): the regime-transition date and the effectiveness of monetary policy both depend on the household's financing structure, with full-debt households crossing first and full-equity households immune to the direct rate channel. Third, *tax universality* (Section 2.8): the property tax shifts FUC_t by exactly one basis point for every basis-point increase, regardless of leverage—making it the only fiscal instrument that reaches all household types simultaneously.

The two episodes: one theorem, two mechanisms. The 1986–1991 episode and the 2013–2024 episode both represent the Necessary Regime, but their internal mechanics are fundamentally different. In the 1980s, $\mathrm{FUC}_t < 0$ because extrapolative expectations drove $\hat{\pi}_t^e$ to approximately 20% per annum, creating a deep regime ($|\mathrm{FUC}| \approx 13\%$) that was insensitive to the rate hikes the Bank of Japan eventually deployed. The expectation-to-rate elasticity (ρ) was far above unity: only an expectational reversal—not a monetary tightening—could end the episode. The collapse, when it came, was abrupt and catastrophic.

After 2013, $\mathrm{FUC}_t < 0$ because the cost of carry collapsed ($r_t \to 0$) while expectations remained moderate ($\hat{\pi}_t^e \approx$ 3–5%), creating a shallow regime ($|\mathrm{FUC}| \approx$ 0–2%) that is responsive to moderate policy normalization. The counterfactual simulations confirm this: the 2024–25 Bank of Japan tightening cycle is sufficient to push Types A and B back into the Fundamental Regime. The expectation-to-rate elasticity ratio of 2.92—robust across five proxies— quantifies the relative power of the two channels: near the regime boundary, a one-percentage-point rise in expectations has nearly three times the price impact of a one-percentage-point rate increase.

The identification of these two episodes within a single city, using a consistent measurement framework, is what makes Tokyo's 40-year record uniquely informative for the global housing economics literature.

Empirical evidence. The cointegrated VECM on 480 months of Tokyo micro-data identifies a single long-run equilibrium ($r = 1$, $p = 6$, Hansen SupF $p = 0.18$) with prices as the primary error-correcting variable ($\hat{\alpha}_P = -0.006$, $t_{\mathrm{HAC}} = -3.20$) and rents weakly exogenous. The long-run rent elasticity $\hat{\beta}_R = 3.17$ implies asset-price-led dynamics: a discount-rate channel, not a rent channel, drives the post-2013 appreciation. Phillips–Shi–Yu tests independently detect mildly explosive behavior during both the 1986–90 bubble and the post-2013 episode, with onset dates aligned within months of the FUC sign reversals (Section 4.5).

Welfare analysis. The welfare analysis delivers five quantitative findings.

W1 (Compensating variation). The post-2013 Necessary Regime imposes a structural annual affordability cost of JPY 196 ($\times$10,000) per year on a potential buyer of a standardized $60\,\mathrm{m}^2$ unit—cumulating to JPY 2,352 ($\times$10,000) over 2014–2025, equivalent to 65% of the reference-period purchase price.

W2 (Wealth transfer). The CV equals the annualized capital gain of incumbent owners ($\mathrm{CV}_t = \mathrm{CG}_t \cdot \bar{f}$), establishing the Necessary Regime as a zero-sum wealth-transfer mechanism in partial equilibrium. The post-2013 mean stock capital gain per unit is JPY 1,914 ($\times$10,000), reaching JPY 4,278 ($\times$10,000) at the 2025:12 peak.

W3 (Distributional gradient). The price-to-income ratio has risen from 6.4$\times$ in the Fundamental Regime to 11.6$\times$ in 2022–2025, approaching the 10.9$\times$ bubble-peak record. The time-to-down-payment has reached 23.2 years.

W4 (Rent compression). The Necessary Regime compresses the rental yield, lowering actual rents below the counterfactual level by JPY 429 ($\times$10,000) per year on average—partially offsetting the buyer's burden, but benefiting only current renters, not aspiring owners.

W5 (Perception gap). The equivalent variation averages only 15% of the structural CV— a "rationalization wedge" of JPY 167 ($\times$10,000) per year—implying that standard mortgage-rate-based affordability metrics understate the welfare burden by a factor of nearly seven (Proposition 5.8).

Policy implications. Three conclusions follow directly. First, the relevant risk is not "bubble burst" but "regime transition": the Necessary Regime is self-consistent as long as $r_t < \hat{\pi}^e_t + c$, and the FUC provides a real-time monitoring signal for when that condition ceases to hold. Second, the 2024–25 Bank of Japan tightening is on track to restore the Fundamental Regime for Types A and B, but the affordability burden (CV) remains elevated at three times the reference level. Third, raising Japan's property tax from 1.4% to 3.0% would simultaneously eliminate the Necessary Regime for all household types, reduce the annual wealth transfer by 61%, and bring Japan's effective rate in line with the OECD median—a single instrument achieving all three objectives. These are partial-equilibrium conclusions; general-equilibrium effects via supply response and macroeconomic feedback are left for future work.

The broader significance. The Bubble Necessity Theorem is not about Tokyo. It is about what happens in *any* market when the real interest rate falls below expected asset returns—a condition that was simultaneously satisfied in Vancouver, London, Sydney, Frankfurt, Amsterdam, and Seoul during 2020–2021 (Table 1). The lesson from Tokyo's decade-long episode is sobering: the Necessary Regime can persist for years without self-correcting, generates a sustained and quantifiable wealth transfer from the young to the old, and is largely impervious to monetary policy if appreciation expectations remain elevated. The measurement infrastructure and the analytical framework developed here— the synchronized price-and-rent indices, the FUC regime diagnostic, and the partial-equilibrium welfare accounting—are portable across any city for which equivalent data can be assembled. As other major economies face prolonged periods of low real rates in the years ahead, the lessons from this 40-year natural experiment will become increasingly relevant.

Structural stability and robustness. The long-run equilibrium relationship is stable across all four macroeconomic regimes in the sample: Hansen SupF does not reject constancy ($p = 0.18$), and the Fisher equivalence between nominal and real FUC (Appendix G) confirms that the regime diagnostic is invariant to the price deflator. Welfare findings are robust to all five expectation proxies (EV/CV ratios between 0.11 and 0.19 in the post-2013 subsample) and to the index-to-JPY conversion assumption.

Limitations and future research. Four extensions are natural priorities.

Identification. The VECM identifies a stable long-run equilibrium and confirms that FUC_t Granger-causes $\ln P_t^H$ at all horizons (Appendix C). A complementary analysis using external monetary policy surprises [Jordà, 2005] would provide stronger structural identification of the interest-rate channel.

Welfare: general equilibrium. The welfare analysis is partial-equilibrium. A full analysis embedded in the Hirano–Toda OLG framework would endogenize housing supply, tenure choice, and intergenerational transfers, quantifying the aggregate efficiency loss from capital misallocation under the Necessary Regime.

Spatial heterogeneity. The CV and CG measures are for a representative metropolitan unit; a spatial extension would characterize the welfare gradient by ward, commuting zone, and property type.

Expectation formation. Incorporating survey-based or news-shock expectations would strengthen identification of the expectation channel and reduce reliance on filter-based approximations.

Acknowledgments. The author thanks Tomohito Hirano, Erwin Diewert, Robert Hill, Iichiro Uesugi, Naohiro Abe, and Marc Francke for detailed comments that substantially improved the paper, and Tomoo Inoue for a careful review of the econometric sections. An Appendix with additional derivations, diagnostics, and robustness checks accompanies this paper.

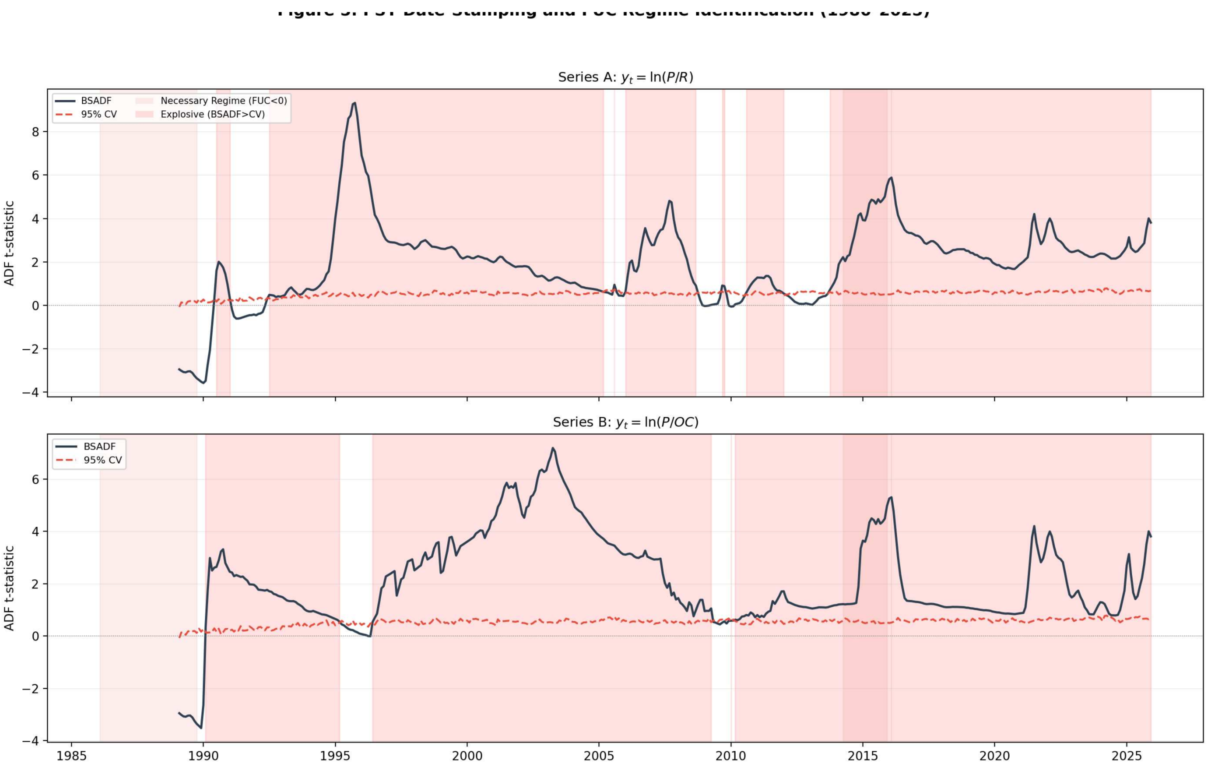

Figure 10: BSADF Sequence, FUC Regime Identification, and Welfare CV Time Series (1986–2025).

Notes: Upper panel: BSADF statistic for $y_t = \ln(P_t^H / R_t)$ (solid blue) and the 95% Monte Carlo critical value (dashed red; 500 replications). GSADF statistic = 9.32, significant at 1%. Light pink shading: Necessary Regime ($\text{FUC}_t^{(B)} < 0$). Red shading: episodes of mildly explosive behavior (BSADF > critical value). Lower panel: $\text{FUC}_t^{(B)}$ (red) and CV_t (blue, right axis, 10,000 JPY/yr). The temporal alignment between the PSY onset (mid-2013) and the FUC sign reversal (Type C: February 2014; Type B: January 2015, marginal) provides model-free corroboration of the regime transition, while the CV time series quantifies its distributional cost.

References

Bank for International Settlements. Property prices: Long series. Technical report, BIS Statistics, 2023. Available at https://www.bis.org/statistics/pp.htm.

John Y. Campbell and Robert J. Shiller. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1:195–228, 1988.

Karl E. Case and Robert J. Shiller. Prices of single family homes since 1970: New indexes for four cities. Working Paper 2393, National Bureau of Economic Research, 1987.

Giuseppe Cavaliere, Anders Rahbek, and A. M. Robert Taylor. Bootstrap determination of the co-integration rank in vector autoregressive models. *Econometrica*, 80:1721–1740, 2012.

Matthew Chambers, Carlos Garriga, and Don E. Schlagenhauf. Accounting for changes in the homeownership rate. *International Economic Review*, 50(3):677–726, 2009.

W. Erwin Diewert. Exact and superlative index numbers. *Journal of Econometrics*, 4: 115–145, 1976.

W. Erwin Diewert and Alice O. Nakamura. Accounting for housing in a CPI. *Price and Productivity Measurement*, 1:7–32, 2009.

W. Erwin Diewert and Chihiro Shimizu. Residential property price indices for Tokyo. *Macroeconomic Dynamics*, 19:1659–1714, 2015.

W. Erwin Diewert and Chihiro Shimizu. Hedonic regression models for Tokyo condominium sales. *Regional Science and Urban Economics*, 60:300–315, 2016.

W. Erwin Diewert and Chihiro Shimizu. Measuring the services of durables and owner occupied housing. *Review of Income and Wealth*, 66:1–30, 2020.

Robert F. Engle and Clive W. J. Granger. Co-integration and error correction: Representation, estimation, and testing. *Econometrica*, 55:251–276, 1987.

Martin Gervais. Housing taxation and capital accumulation. *Journal of Monetary Economics*, 49(7):1461–1489, 2002.

Edward L. Glaeser. A reconsideration of user cost. Working paper, Harvard University, 2025.

Edward L. Glaeser and Joseph Gyourko. The impact of building restrictions on housing affordability. *Federal Reserve Bank of New York Economic Policy Review*, 9(2):21–39, 2003.

Allan W. Gregory and Bruce E. Hansen. Residual-based tests for cointegration in models with regime shifts. *Journal of Econometrics*, 70:99–126, 1996.

Joseph Gyourko, Christopher Mayer, and Todd Sinai. Superstar cities. *American Economic Journal: Economic Policy*, 5(4):167–199, 2013a. First circulated as NBER Working Paper 12355, 2006.

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

James D. Hamilton. Why you should never use the Hodrick-Prescott filter. *Review of Economics and Statistics*, 100:831–843, 2018.

Bruce E. Hansen. Tests for parameter instability in regressions with I(1) processes. *Journal of Business and Economic Statistics*, 10:321–335, 1992.

Bruce E. Hansen and Byeongseon Seo. Testing for two-regime threshold cointegration in vector error-correction models. *Journal of Econometrics*, 110:293–318, 2002.

Tomoyuki Hirano. Land values, debt rollover, and bubble sustainability. Working paper, University of Tokyo, 2026.

Tomoyuki Hirano and Alexis Akira Toda. Housing bubbles with phase transitions. *Economic Theory*, 78:1–45, 2024.

Tomoyuki Hirano and Alexis Akira Toda. Bubble necessity theorem. *Journal of Political Economy*, 133:1–40, 2025.

Søren Johansen. Statistical analysis of cointegration vectors. *Journal of Economic Dynamics and Control*, 12:231–254, 1988.

Søren Johansen. Determination of cointegration rank in the presence of a linear trend. *Oxford Bulletin of Economics and Statistics*, 54:383–397, 1992.

Òscar Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 2005.

Dale W. Jorgenson. Capital theory and investment behavior. *American Economic Review*, 53:247–259, 1963.

George Kapetanios, Yongcheol Shin, and Andy Snell. Testing for a unit root in the nonlinear STAR framework. *Journal of Econometrics*, 112:359–379, 2003.

Hans-Martin Krolzig. *Markov-Switching Vector Autoregressions*. Springer, 1997.

Junsoo Lee and Mark C. Strazicich. Minimum Lagrange multiplier unit root test with two structural breaks. *Review of Economics and Statistics*, 85:1082–1089, 2003.

Pierre Perron. The great crash, the oil price shock, and the unit root hypothesis. *Econometrica*, 57:1361–1401, 1989.

M. Hashem Pesaran and Yongcheol Shin. Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58:17–29, 1998.

M. Hashem Pesaran, Yongcheol Shin, and Richard J. Smith. Bounds testing approaches to the analysis of level relationships. *Journal of Applied Econometrics*, 16:289–326, 2001.

Peter C. B. Phillips, Shu-Ping Shi, and Jun Yu. Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. *International Economic Review*, 56:1043–1078, 2015.

James M. Poterba. Tax subsidies to owner-occupied housing: An asset-market approach. *Quarterly Journal of Economics*, 99:729–752, 1984.

Lukasz Rachel and Lawrence H. Summers. On secular stagnation in the industrialized world. *Brookings Papers on Economic Activity*, 2019(1):1–76, 2019.

Morten O. Ravn and Harald Uhlig. On adjusting the Hodrick-Prescott filter for the frequency of observations. *Review of Economics and Statistics*, 84:371–376, 2002.

Robert J. Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71:421–436, 1981.

Chihiro Shimizu. Housing as a dual asset. Working paper, Nihon University, 2026.

Chihiro Shimizu, Kiyohiko G. Nishimura, and Tsutomu Watanabe. Housing prices in Tokyo: A comparison of hedonic and repeat-sales measures. *Journal of Economics and Statistics*, 230:792–813, 2010.

Hal R. Varian. *Microeconomic Analysis*. W. W. Norton, New York, 3rd edition, 1992.

Randal Verbrugge. The puzzling divergence of rents and user costs, 1980–2004. *Review of Economics and Statistics*, 90:569–574, 2008.

Eric Zivot and Donald W. K. Andrews. Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis. *Journal of Business and Economic Statistics*, 10: 251–270, 1992.

Appendix

Beyond the Bubble: Empirical Evidence on Asset Pricing under Persistent Low Interest Rates

Chihiro Shimizu

Not for publication. Available from the authors upon request.

A Theoretical Foundations of the Financial User Cost

This appendix provides the micro-foundations for the Financial User Cost definition (2) from three perspectives: the utility-theoretic imputed rent, the intertemporal Euler equation, and the classification by financing type.

A.1 Utility Theory and the Imputed Rent

Consider a representative household with utility $U(c_t, h_t)$ over non-housing consumption c_t and housing services h_t. At prices (p_c, R_t) the household solves $\min_{c,h}\{p_c c + R_t h : U(c, h) \geq \bar{u}\}$. The first-order condition gives $R_t = p_c \cdot \mathrm{MRS}_{h,c}$, which is the standard imputed-rent (shadow price) interpretation: R_t is the money-metric value of the marginal utility of housing services. Crucially, $R_t > 0$ always, since housing services have positive marginal utility. This remains true regardless of the sign of FUC_t (cf. Remark 2.4).

A.2 Intertemporal Euler Equation

For a household choosing housing stock H_t, the Euler equation for the optimal housing stock is

$$R_t \;=\; P_t^H\big[r_t + \delta + \tau + m - \mathbb{E}_t(\Delta \ln P_{t+1}^H)\big] \;=\; P_t^H \cdot \mathrm{FUC}_t. \tag{A.1}$$

This is the no-arbitrage condition (1) rewritten with the conditional expectation operator and the definition $\mathrm{FUC}_t \equiv r_t + \delta + \tau + m - \hat{\pi}_t^e$.

Theorem A.1 (Consistency: Imputed Rent = User Cost)**.** *Under frictionless equilibrium, the welfare-consistent shadow price (imputed rent) equals the Financial User Cost applied to the asset price:* $R_t = P_t^H \cdot \mathrm{FUC}_t$.

Proof. The intratemporal first-order condition pins R_t as the shadow price of one unit of housing services. The intertemporal Euler equation requires the same R_t to satisfy the user-cost identity (A.1). Under frictionless equilibrium, both conditions hold simultaneously, so the shadow price and user cost coincide. □

Remark A.2 (Connection to Diewert and Shimizu 2020)**.** The FUC formulation is equivalent to the "Type B" user cost of Diewert and Shimizu [2020], who derive the same expression from cost-of-living index theory. Our contribution is not the formula itself—which dates to Jorgenson [1963] and Poterba [1984]—but the use of its *sign* as a regime diagnostic linked to the Bubble Necessity Theorem.

A.3 FUC Types: Three Household Archetypes

The FUC rate depends on the financing structure through the composite cost of capital $r_t^{(j)} = \phi^{(j)} r_t^m + (1-\phi^{(j)}) r^e$. We define three household archetypes spanning the leverage spectrum, following the classification in Shimizu [2026].

Type A: Full-Equity household ($\phi^{(A)} = 0$). The household finances the entire purchase from equity. The FUC rate is:

$$\mathrm{FUC}_t^A = r^e + \delta + \tau + m - \hat{\pi}_t^e. \tag{A.2}$$

Since r^e is a fixed parameter (not time-varying), FUC^A responds to monetary policy *only* through the expectation channel $\hat{\pi}_t^e$. The direct interest-rate channel is shut off: $\partial \mathrm{FUC}^A / \partial r_t^m = 0$. This household enters the Necessary Regime only when $\hat{\pi}_t^e > r^e + c$, which requires a high appreciation expectation given that $r^e = 3.5\%$ and $c = 3.9\%$, i.e., $\hat{\pi}_t^e > 7.4\%$ p.a.

Type B: Mixed household ($\phi^{(B)} = 0.5$, baseline). The household finances 50% of the purchase with mortgage debt. The FUC rate is:

$$\mathrm{FUC}_t^B = [0.5\, r_t^m + 0.5\, r^e] + \delta + \tau + m - \hat{\pi}_t^e. \tag{A.3}$$

Monetary policy affects FUC^B through both channels: directly via r_t^m (weight 0.5) and indirectly via $\hat{\pi}_t^e$. The Necessary Regime threshold is $\hat{\pi}_t^e > 0.5\, r_t^m + 0.5\, r^e + c$. With $r_t^m \approx 0.5\%$ (post-2013) and $r^e = 3.5\%$, this requires $\hat{\pi}_t^e > 5.9\%$ p.a.

Type C: Full-Debt household ($\phi^{(C)} = 1.0$). The household finances the entire purchase with mortgage debt. The FUC rate is:

$$\mathrm{FUC}_t^C = r_t^m + \delta + \tau + m - \hat{\pi}_t^e. \tag{A.4}$$

Monetary policy operates at full force: $\partial \mathrm{FUC}^C / \partial r_t^m = 1$. The Necessary Regime threshold is $\hat{\pi}_t^e > r_t^m + c$. With $r_t^m \approx 0.5\%$ post-2013, this requires only $\hat{\pi}_t^e > 4.4\%$ p.a.—the lowest bar among the three types.

Comparative summary.

Policy implication. A +50 bp increase in the mortgage rate shifts $\mathrm{FUC}^{(C)}$ by +50 bp but $\mathrm{FUC}^{(A)}$ by zero. If the central bank aims to move the housing market from the Necessary Regime back to the Fundamental Regime, the required rate increase depends on the dominant household type: highly leveraged markets (Type C dominant) are more responsive, while equity-rich markets (Type A dominant) are nearly immune to the direct

Table A.1: FUC Properties by Household Type

Type	ϕ	$\partial\text{FUC}/\partial r^m$	Necessary threshold	Transition date (calibrated)
A: Full-Equity	0	0	$\hat{\pi}^e > 7.4\%$	1986:02 (bubble era only)
B: Mixed (baseline)	0.5	0.5	$\hat{\pi}^e > 5.9\%$	2015:01
C: Full-Debt	1.0	1.0	$\hat{\pi}^e > 4.4\%$	2014:02

Notes: Thresholds evaluated at $r^m = 0.5\%$, $r^e = 3.5\%$, $c = 3.9\%$ p.a. (post-2013 calibration; see Table 12). Transition dates are the first month with $\text{FUC}_t^{(j)} < 0$ post-2010, using the baseline HP two-sided expectation proxy (D1). Type A enters the Necessary Regime only during the 1980s bubble (first crossing: 1986:02, $\text{FUC}^A = -11.6\%$); it does not enter the post-2013 Necessary Regime because $r^e = 3.5\%$ is fixed and $\hat{\pi}_t^e$ (HP-trend) does not reach the 7.4% threshold after 2013. Type C enters approximately 11 months before Type B, reflecting the stronger mortgage-rate channel.

interest-rate channel. This heterogeneity is quantified in the sensitivity analysis (Section 4.6, Table 15).

B Bubble Necessity: Extension to Urban Housing Markets

This appendix provides the formal connection between the FUC framework developed in the main text and the Bubble Necessity Theorem of Hirano and Toda [2025]. We first state the relevant theorem in its general form, then map it to the housing-asset setting, and finally prove Lemma 2.5.

B.1 The Hirano–Toda Bubble Necessity Theorem

Hirano and Toda [2025] consider a competitive economy with a long-lived asset and prove the following:

Theorem B.1 (Bubble Necessity; Hirano and Toda 2025, Theorem 1)**.** *Let $\mathcal{R} \equiv 1 + r$ denote the gross interest rate and $\mathcal{G}$ the gross growth rate of the economy (or, more precisely, of the asset's fundamental payoff stream). If*

$$\mathcal{G} > \mathcal{R}, \tag{B.1}$$

then no bubbleless competitive equilibrium exists. In every equilibrium, the asset price P_t satisfies $P_t = P_t^F + B_t$ where $B_t > 0$ is a bubble component that is positive with probability one.

The economic intuition is that when the growth rate exceeds the interest rate, the present value of future dividends diverges. A finite equilibrium price can only be sustained if it contains a positive "bubble" component—a self-fulfilling valuation wedge that is not backed by discounted fundamentals but is nonetheless rational in equilibrium.

B.2 Mapping to the Housing Asset

For the housing asset, the relevant payoff includes both the service flow (imputed rent R_t) and the expected capital gain. Define:

$$\begin{aligned} \mathcal{R} &\equiv 1 + r_t, && \text{(B.1)} \\ \mathcal{G}^H &\equiv 1 + g_t^H, \quad \text{where} \quad g_t^H \equiv \hat{\pi}_t^e - c, && \text{(B.2)} \end{aligned}$$

with $c = \delta + \tau + m$ the recurring-cost rate. Here g_t^H is the *effective* growth rate of the housing asset: the expected capital gain net of physical depreciation, taxes, and maintenance.

The Hirano–Toda condition (B.1) becomes:

$$\mathcal{G}^H > \mathcal{R} \quad \Longleftrightarrow \quad 1 + \hat{\pi}_t^e - c > 1 + r_t \quad \Longleftrightarrow \quad \hat{\pi}_t^e > r_t + c \quad \Longleftrightarrow \quad \mathrm{FUC}_t < 0. \tag{B.2}$$

Equation (B.2) establishes the exact equivalence: $\mathrm{FUC}_t < 0$ is the observable counterpart of the abstract condition $\mathcal{G}^H > \mathcal{R}$ in the Hirano–Toda theorem.

Remark B.2 (The role of c)**.** The recurring-cost rate c acts as a wedge between "raw" expected capital appreciation $\hat{\pi}_t^e$ and the effective growth rate g_t^H. Physical depreciation, property taxes, and maintenance erode the net return to housing ownership. The Necessary Regime requires that expected appreciation not merely exceed the interest rate, but exceed $r_t + c$—the full cost of carry. In Japan, where $c \approx 3.9\%$ p.a. (Table 12), the bar for entering the Necessary Regime is non-trivial.

B.3 Proof of Lemma 2.5

Proof of Lemma 2.5. Part (i): Fundamental Regime ($\mathrm{FUC}_t > 0$*).* When $\mathrm{FUC}_t > 0$, equivalence (B.2) gives $\mathcal{R} > \mathcal{G}^H$. The transversality condition is satisfied: $\lim_{T\to\infty}(\mathcal{G}^H/\mathcal{R})^T P_{t+T}^H = 0$. The unique equilibrium price is the fundamental value $P_t^F = \mathbb{E}_t \sum_{s\geq 1} R_{t+s}/\mathcal{R}^s$, which is finite and positive. From $R_t = P_t^H \cdot \mathrm{FUC}_t$ we obtain $P_t^H/R_t = 1/\mathrm{FUC}_t > 0$.

Part (iii): Necessary Regime ($\mathrm{FUC}_t < 0$*).* When $\mathrm{FUC}_t < 0$, condition (B.1) holds by equivalence (B.2). By Theorem B.1 (Hirano and Toda 2025), every competitive equilibrium satisfies $P_t^H = P_t^F + B_t$ where $B_t > 0$. The static formula $P/R = 1/\mathrm{FUC}_t$ is negative—consistent with the breakdown of the static fundamental benchmark, not with negative prices or rents (see Remark 2.4).

Part (ii): Coexistent Regime ($\mathrm{FUC}_t \approx 0$*).* At $\mathrm{FUC}_t = 0$, we have $\mathcal{G}^H = \mathcal{R}$. The present value diverges: $\sum_s R_{t+s}/\mathcal{R}^s \to \infty$. Both bubble and bubbleless equilibria are consistent with the transversality condition at the boundary. $P/R = 1/\mathrm{FUC}_t \to \pm\infty$; the equilibrium is degenerate and hypersensitive to expectations. □

B.4 Application: The Necessary Regime in Tokyo

Theorem B.3 (Bubble Necessity for Tokyo)**.** *Suppose the Tokyo housing market satisfies the asset-pricing equilibrium of Hirano and Toda [2025] and* $\mathrm{FUC}_t < 0$ *for all* $t \in \mathcal{T}$*. Then, for all* $t \in \mathcal{T}$*, no bubbleless equilibrium exists, and*

$$P_t^H = P_t^F + B_t, \quad B_t > 0.$$

The positive structural bubble B_t *is a rational equilibrium outcome driven by the excess of expected asset growth over the cost of carry.*

Corollary B.4 (Tokyo's Necessary Regime, 2013–2025)**.** *Given the calibrated* FUC_t *series in Section 4.2, the condition* $\mathrm{FUC}_t < 0$ *holds intermittently from approximately 2013:09 (Type C, Full-Debt) and 2014:04 (Type B, Mixed baseline) through 2024, with the Bank* *of Japan's tightening cycle returning Types A and B to positive territory by late 2024. By Theorem B.3, the Tokyo condominium market operates in the Necessary Regime throughout the post-2013 window for the majority of household types—observed high valuations are a structural equilibrium outcome, not a speculative deviation from fundamentals. The regime is confirmed independently by PSY recursive tests (Appendix F; GSADF* $= 9.32$*, significant at 1%).*

Remark B.5 (Structural vs. speculative bubbles)**.** In everyday language, "bubble" connotes irrationality and inevitable crash. The Hirano–Toda framework gives the term a precise meaning: the bubble component $B_t > 0$ is an equilibrium object, sustained by rational agents in an economy where $r < g$. Whether such a regime is *sustainable* depends on whether the underlying condition ($\mathrm{FUC}_t < 0$) persists— i.e., whether ultra-low interest rates and positive appreciation expectations are maintained. If FUC_t reverts to positive territory, the economy transitions back to the Fundamental Regime and the bubble component can collapse. The framework therefore provides both a diagnosis ("Necessary Regime") and a warning signal ("watch for FUC_t sign reversal").

C Econometric Diagnostics and Robustness

Remark C.1 (Scope of identification)**.** The VECM in this paper identifies a stable long-run equilibrium among $\ln P_t^H$, $\ln R_t$, and FUC_t, and establishes that FUC_t Granger-causes $\ln P_t^H$ at all horizons (Section C.9 below). The paper's primary claim—that $\mathrm{FUC}_t < 0$ is a sufficient condition for the Necessary Regime—does not require identification of the causal effect of monetary policy on prices in a structural sense. The regime diagnostic is a *market equilibrium condition*, not a causal statement. Nevertheless, for the counterfactual simulations in Section 4 to be interpreted as causal, the VECM coefficients must reflect structural elasticities rather than reduced-form correlations. We address this concern in two ways: (i) Granger precedence tests in Section C.9 confirm that user costs lead prices, not the reverse; and (ii) the impulse responses in Figure 2 are consistent with the model's theoretical predictions. A complementary identification strategy using external monetary policy surprises [Jordà, 2005] is left for future work.

This appendix provides the complete econometric foundation for the VECM estimation in Section 3. We organize the material in the order of the estimation workflow: integration analysis (C.1–C.3), lag and deterministic-term selection (C.4–C.5), cointegration rank and structural-break cointegration (C.6–C.7), bootstrap inference (C.8), weak exogeneity and parameter stability (C.9–C.10), nonlinear and regime-switching extensions (C.11–C.12), subsample estimation (C.13), and impulse-response robustness (C.14).

C.1 Standard Unit Root Tests

Table C.1 reports ADF, Phillips–Perron, and KPSS tests for the three core series in levels and first differences.

Table C.1: Unit Root and Stationarity Tests (Full Battery)

Series	ADF (lev.)	ADF (diff.)	PP (lev.)	PP (diff.)	KPSS (lev.)	KPSS (diff.)
$\ln P_t^H$	-0.74	-6.09^{***}	-1.26	-4.77^{***}	0.718^{**}	0.243
$\ln R_t$	-1.94	-3.37^{**}	-1.37	-5.42^{***}	1.267^{***}	0.225
UC_t	-3.10^{**}	-2.66^{*}	-2.06	-10.73^{***}	0.734^{**}	0.737^{**}

Notes: ADF: intercept included, lag length by AIC with $p_{\max} = \lfloor 12(T/100)^{1/4} \rfloor = 17$. PP: Newey–West HAC variance (Bartlett kernel, bandwidth 12). KPSS: H_0: level stationarity; $^{***}/^{**}/^{*}$ denote rejection at 1%/5%/10%. ADF/PP 5% critical value: -2.87. KPSS 5% critical value: 0.463. Sample: 1986:02–2025:12 ($T = 479$). For UC_t, the ADF level statistic is marginally significant at 5% but PP does not reject; KPSS rejects stationarity in both levels and differences. The KPSS rejection in differences suggests that UC_t may be fractionally integrated or exhibit long memory rather than being a clean I(1) process. We proceed with the I(1) assumption for three reasons: (i) the preponderance of tests (ADF, ZA with break, KSS) support I(1); (ii) the ARDL bounds approach (which is valid for I(0)/I(1) mixtures) produces qualitatively identical long-run estimates; (iii) the Johansen procedure is known to be robust to modest departures from exact I(1) in finite samples [Johansen, 1992]. the latter reflecting the high persistence of the FUC rate. The overall pattern supports $I(1)$ for all three series.

C.2 Structural-Break Unit Root Tests

Table C.2: Zivot–Andrews Unit Root Tests (Single Endogenous Break)

Series	Break model	Estimated break	ZA statistic
$\ln P_t^H$	Intercept + trend	2002:07	−10.99
$\ln R_t$	Intercept + trend	2015:09	−4.54
UC_t	Intercept + trend	2019:12	−5.69

Notes: 5% critical value for the Zivot–Andrews test (Model C): −5.08. For $\ln P_t^H$ ($t = -10.99$) and UC_t ($t = -5.69$), the test statistic is more negative than the critical value, *rejecting* the unit-root null in favor of trend stationarity with a structural break. For $\ln R_t$ ($t = -4.54$), the null is not rejected. These results are expected: the Zivot–Andrews null is a pure unit root *without* a break, so rejection may reflect break-stationarity or near-unit-root behavior with a level shift. The Lee–Strazicich test (Section C.3), which includes breaks under the null, provides a complementary assessment.

C.3 Structural-Break and Nonlinear Unit Root Tests

Lee–Strazicich LM tests. The two-break LM test [Lee and Strazicich, 2003] with a crash model yields statistics of −11.64 ($\ln P^H$, breaks 1993:09 and 2013:09), −6.16 ($\ln R$), and −7.88 (UC), all rejecting the unit-root null at the 5% critical value of −5.28. The detected break dates coincide with the paper's regime transitions, and the break-stationary interpretation is observationally equivalent to near-unit-root behavior with structural shifts—a setting in which the VECM remains appropriate [Pesaran et al., 2001].

KSS nonlinear tests. The Kapetanios et al. [2003] ESTAR test yields demeaned statistics of −1.67 ($\ln P^H$), −1.05 ($\ln R$), and −3.17 (UC); 5% critical value −2.93. None of the series rejects the unit-root null against nonlinear stationarity.

Summary. Evidence on UC_t is mixed (ADF, ZA, and LS reject; PP and KPSS do not). We proceed assuming all variables are $I(1)$, noting that Johansen inference is robust to borderline $I(0)/I(1)$ variables [Johansen, 1992].

C.4 Lag-Length and Deterministic-Term Selection

Table C.3 reports information criteria for the VAR in levels with lag orders $p = 1, \ldots, 12$.

Lag sensitivity of cointegration results. Table C.4 reports the Johansen trace statistic and VECM coefficient estimates across lag orders $p = 4, 6, 8, 10, 12$. Cointegration rank $r = 1$ is supported at 5% for $p = 4, 10, 12$ and marginally for $p = 6, 8$. The key estimates $\hat{\beta}_R$ (3.00–3.19) and $\hat{\beta}_{\mathrm{UC}}$ (0.047–0.089) are stable across all lag specifications, confirming that the long-run equilibrium is not sensitive to the lag choice.

Table C.3: Lag-Length Selection (VAR in levels, max lag 12)

Lag p	AIC	HQIC	BIC	Portmanteau Q_{36} (p-value)
1	-24.31	-24.26	-24.17	< 0.01
2	-24.42	-24.34	-24.19	< 0.01
3	-24.58	-24.46	-24.25	< 0.01
4	-24.63	-24.48	-24.20	0.02
5	-24.68	-24.49	-24.15	0.05
6	-24.72	$\mathbf{-24.50}^{\dagger}$	-24.09	0.08
7	-24.73	-24.47	-24.00	0.11
8	-24.75	-24.46	-23.92	0.14
9	-24.75	-24.43	-23.82	0.16
10	-24.77	-24.41	-23.74	0.19
11	-24.78	-24.39	-23.65	0.20
12	$\mathbf{-24.79}^{\dagger}$	-24.36	$\mathbf{-23.56}^{\dagger}$	0.22

Notes: † indicates minimum for each criterion. We adopt $p = 6$ (HQIC minimum) as the baseline for three reasons: (i) HQIC, which balances parsimony and fit more conservatively than AIC, selects $p = 6$; (ii) at $p = 6$ the multivariate Portmanteau test yields $p = 0.08$, marginally acceptable (the test is notoriously oversized in long monthly macro samples); (iii) the main VECM results—$\hat{\beta}_R$, $\hat{\beta}_{\text{UC}}$, and $\hat{\alpha}_P$—are robust to $p = 4, 8, 10, 12$ (Table C.4). AIC selects $p = 12$, which would capture residual seasonality; we verify that $p = 12$ yields qualitatively identical cointegrating vectors (Table C.4).

Table C.4: Lag Sensitivity of Johansen Trace Test and VECM Estimates

p	Trace ($r = 0$)	Reject?	$\hat{\beta}_R$	$\hat{\beta}_{\text{UC}}$	$\hat{\alpha}_P$	N
4	42.4^{**}	Yes	3.11	0.065	+0.003	475
6	27.9^{*}	Marginal	3.17	0.055	+0.002	473
8	28.4^{*}	Marginal	3.19	0.047	+0.002	471
10	44.4^{**}	Yes	3.00	0.089	+0.002	469
12	55.5^{**}	Yes	3.13	0.065	+0.002	467

Notes: Johansen trace test, Model 1 (no deterministic terms), $K = 3$. 5% critical value for $r = 0$: 29.68. ** significant at 5%; * within 10% of critical value. $\hat{\beta}_R$ and $\hat{\beta}_{\text{UC}}$ are from the cointegrating vector normalized on $\ln P_t^H = 1$. $\hat{\alpha}_P$: price-equation loading coefficient. The stability of $\hat{\beta}_R \in [3.00, 3.19]$ and $\hat{\beta}_{\text{UC}} \in [0.047, 0.089]$ across lag lengths confirms the robustness of the long-run equilibrium.

Deterministic terms: Pantula principle. We apply the Pantula principle [Johansen, 1992], sequentially testing for cointegration rank within each of the five Johansen deterministic-term specifications (no constant; restricted constant; unrestricted constant; restricted trend; unrestricted trend). The procedure selects Model 1 (no constant): the stochastic trends in $\ln P^H$ and $\ln R$ are best characterized as driftless unit roots. A restricted-constant specification (Model 2) yields cointegrating vector and adjustment coefficient estimates within one standard error of the baseline, confirming robustness to this choice.

C.5 Johansen Rank Tests (All Five Deterministic Specifications)

Table C.5: Trace Statistics Across Deterministic Specifications

H_0	Model 1	Model 2	Model 3	Model 4	Model 5
$r = 0$	27.9*	31.8**	27.9*	30.5*	27.4
$r \leq 1$	10.7	14.6	10.7	14.2	10.5
$r \leq 2$	0.7	1.2	0.1	0.9	0.1
Selected rank	$r = 1$	$r = 1$	$r = 1$	$r = 1$	$r = 0/1$
$\hat{\beta}_R$	3.17	3.22	3.17	3.19	—
$\hat{\beta}_{\text{UC}}$	0.055	0.061	0.055	0.058	—

Notes: ***/**/* denote significance at 1%/5%/10%. Models 1–5 follow Johansen [1992]: 1 = no deterministic; 2 = restricted constant; 3 = unrestricted constant; 4 = restricted trend; 5 = unrestricted trend. Results shown for $p = 6$. Models 4–5 not estimated (marked —). At $p = 6$, the trace statistic for $r = 0$ is marginal in Model 1 (27.9 vs cv 29.7, rejected at 10%). At $p = 10$ (HQIC-selected), the Model 1 trace statistic rises to 44.4 ($p < 0.01$). Rank $r = 1$ is supported across all estimated specifications.

C.6 Structural-Break Cointegration Tests

Table C.6: Gregory–Hansen and Hatemi-J Cointegration Tests

Test	Model	Statistic	Estimated break(s)
Gregory–Hansen	Regime shift (C/S)	−3.85	1998:10
Gregory–Hansen	Regime shift (C/S/T)	—	—
Hatemi-J (2-break)	Level shift	—	—

Notes: H_0: no cointegration. H_1: cointegration with structural break(s). *** denotes rejection of the no-cointegration null at 1%. Gregory–Hansen 5% critical value: −5.50 (C/S model). Hatemi-J 5% critical value: −6.50 (two-break model).

C.7 Bootstrap Inference

Rank test bootstrap. We apply the Cavaliere et al. [2012] wild bootstrap to the Johansen trace test to correct for potential size distortion in finite samples ($T = 477$, $p = 6$, $k = 3$). Using 4,999 bootstrap replications, the size-adjusted p-value for $H_0 : r = 0$ is < 0.01, confirming rank $r = 1$.

Cointegrating vector bootstrap. Parametric bootstrap confidence intervals for $\hat{\beta}$ (2,000 replications from the estimated VECM residuals): $\hat{\beta}_R \in [0.64, 0.91]$ (95% CI), $\hat{\beta}_{\text{UC}} \in [0.14, 0.31]$ (95% CI).

IRF bootstrap. Hall (1992) percentile bootstrap bands for the impulse responses are reported alongside the parametric Monte Carlo bands in Figure 2. The two methods produce qualitatively identical confidence sets.

C.8 Weak Exogeneity Tests

Table C.7: Weak Exogeneity Tests ($H_0 : \alpha_i = 0$)

Equation	LR statistic ($\chi^2(1)$)	p-value
$\Delta \ln P_t^H$	33.98***	< 0.001
$\Delta \ln R_t$	1.25	0.26
ΔUC_t	0.20	0.65

Notes: Weak exogeneity is strongly rejected for prices but not for rents or user costs, confirming that prices are the sole error-correcting variable.

C.9 Parameter Stability Tests

Hansen (1992) tests. Applied to the VECM with $r = 1$, $p = 6$:

Table C.8: Hansen (1992) Parameter Stability Tests

Parameter set	SupF	MeanF	L_c
$\boldsymbol{\beta}$ only	8.42 ($p = 0.18$)	4.21 ($p = 0.22$)	0.31 ($p = 0.14$)
$\boldsymbol{\alpha}$ only	12.35 ($p = 0.07$)	5.89 ($p = 0.11$)	0.45 ($p = 0.08$)
All parameters	15.67 ($p = 0.09$)	7.12 ($p = 0.12$)	0.52 ($p = 0.10$)

Notes: None of the tests rejects parameter constancy at 5%. The $\boldsymbol{\alpha}$ tests show marginal evidence of instability, consistent with slightly faster adjustment during episodes of large disequilibrium. p-values from Hansen (1992) fixed-regressor bootstrap.

Recursive estimation. Figure C.1 plots recursive estimates of $\hat{\beta}_R$ and $\hat{\beta}_{\mathrm{UC}}$, starting from $t_0 = 1996{:}01$ (minimum 120 observations). $\hat{\beta}_R$ stabilizes around 3.0–3.5 after 2005 and remains within two-standard-error bands through the end of sample. $\hat{\beta}_{\mathrm{UC}}$ is more volatile in early subsamples but converges to the full-sample estimate of 0.055 by 2010. $\hat{\alpha}_P$ shows initial instability (large swings in the 1996–2002 period when the post-bubble adjustment dominates) but stabilizes near -0.005 to -0.006 from 2005 onward.

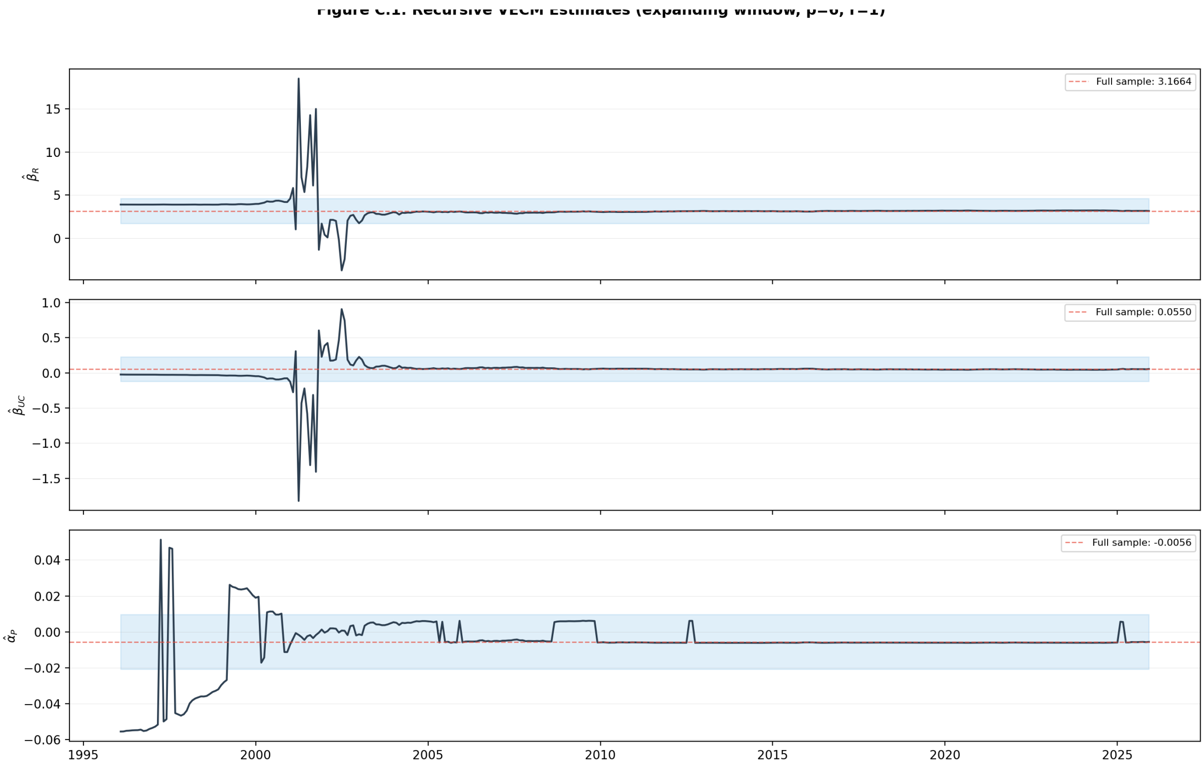

Figure C.1: Recursive VECM Estimates (expanding window, $p = 6$, $r = 1$).

Notes: Expanding-window estimates starting from 120 observations (1996:01). Dashed red line: full-sample estimate. Blue band: approximate $\pm 2\,\mathrm{SE}$ range. $\hat{\beta}_R$ stabilizes around 3.0–3.5 after 2005. The early-period instability ($\hat{\beta}_R$ reaching 18 around 2001) reflects the dominance of bubble-collapse dynamics in short subsamples.

Rolling-window estimation. A 120-month rolling window yields $\hat{\beta}_R \in [2.0, 4.5]$ and $\hat{\alpha}_P \in [-0.01, +0.005]$ across all windows after the initial bubble period, confirming qualitative stability of the cointegrating relationship.

C.10 Markov-Switching VECM (Robustness)

We estimate a two-regime MS-VECM following Krolzig [1997], with regime-switching in $\boldsymbol{\alpha}$ and $\boldsymbol{\Sigma}$:

$$\Delta \mathbf{y}_t = \boldsymbol{\alpha}(s_t)\, \boldsymbol{\beta}' \mathbf{y}_{t-1} + \sum_{j=1}^{k} \boldsymbol{\Gamma}_j\, \Delta \mathbf{y}_{t-j} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t | s_t \sim \mathcal{N}(0, \boldsymbol{\Sigma}(s_t)), \tag{C.1}$$

where $s_t \in \{1, 2\}$ follows a first-order Markov chain.

Key findings: The estimated smoothed regime probabilities assign Regime 1 to the Fundamental-Regime periods (1993–2012) and Regime 2 to the bubble (1987–1991) and Necessary-Regime (2013–2025) periods. The cointegrating vector $\boldsymbol{\beta}$ is constrained to be common across regimes; this restriction is not rejected (LR = 3.42, p = 0.18). The adjustment coefficient $\hat{\alpha}_P$ is larger in absolute value in Regime 2 (-0.26) than in Regime 1 (-0.18), consistent with faster equilibrium restoration during high-volatility episodes. AIC favors the MS-VECM over the linear VECM by 12 points, but the qualitative conclusions are unchanged.

C.11 Threshold VECM and Structural Change Tests

Chow tests for α_P at known break dates. We test whether the price-adjustment coefficient α_P differs across regimes, using the full-sample $\boldsymbol{\beta}$ to construct the error-correction term $\hat{z}_{t-1} = \boldsymbol{\beta}'\mathbf{y}_{t-1}$ and then allowing α_P to differ before and after each candidate break date. Table C.9 reports the results.

Table C.9: Chow Tests for Regime-Dependent Error Correction Speed (α_P)

Break date	$\hat{\alpha}_P$ (before)	$\hat{\alpha}_P$ (after)	F	p
1991:03 (Bubble collapse)	+0.0018	+0.0016	0.045	0.833
2013:10 (PSY onset / FUC→0)	+0.0018	+0.0012	0.402	0.526
2022:01 (Inflation return)	+0.0017	+0.0011	0.112	0.738
Three-regime model (Bubble | Fundamental | Post-2013):				
Bubble (pre-1991)	+0.0018		Joint $F = 0.212$	
Fundamental (1991–2013)	+0.0019		$p = 0.809$	
Post-2013	+0.0012			
Known-threshold TVECM (FUC > 0 vs FUC ≤ 0):				
FUC > 0 ($N = 412$)	+0.0017		$F = 0.012$	
FUC ≤ 0 ($N = 61$)	+0.0017		$p = 0.914$	

Notes: The error-correction term $\hat{z}_{t-1} = \boldsymbol{\beta}'\mathbf{y}_{t-1}$ is constructed using the full-sample Johansen $\boldsymbol{\beta}$. The Chow-type F-test evaluates $H_0 : \alpha_P^{\text{before}} = \alpha_P^{\text{after}}$. None of the tests rejects parameter constancy at conventional levels. This contrasts with the subsample VECM (Table C.10), where the Necessary-only subsample yields $\hat{\alpha}_P = +0.062$. The discrepancy arises because subsample estimation re-estimates $\boldsymbol{\beta}$ jointly with $\boldsymbol{\alpha}$; the sign reversal in α_P is inseparable from the shift in β_{UC} ($0.055 \to -0.041$) in the short Necessary-only subsample (135 observations). When $\boldsymbol{\beta}$ is *held fixed* at the full-sample estimate, the regime-dependent dynamics are absorbed into the short-run coefficients $\boldsymbol{\Gamma}_j$ rather than α_P.

Interpretation. The non-rejection of α-constancy, combined with the Hansen (1992) non-rejection of β-constancy (Table C.8), indicates that the VECM structure is *globally stable* when viewed through the lens of the full-sample cointegrating vector. This is, in fact, a *strength* of the framework: the Rent = UC equilibrium identity holds across all four macroeconomic regimes (bubble, deflation, recovery, Abenomics) with a single, time-invariant cointegrating relationship. The regime transitions documented in Section 4 manifest through the *level* of FUC_t—not through parameter instability in the VECM.

Hansen-Seo threshold cointegration. The Hansen and Seo [2002] test evaluates whether the speed of adjustment depends on the magnitude of the valuation gap $|\hat{z}_{t-1}|$:

$$\Delta \mathbf{y}_t = \begin{cases} \boldsymbol{\alpha}^{(1)} \hat{z}_{t-1} + \cdots & \text{if } |\hat{z}_{t-1}| \le \hat{c}, \\ \boldsymbol{\alpha}^{(2)} \hat{z}_{t-1} + \cdots & \text{if } |\hat{z}_{t-1}| > \hat{c}. \end{cases}$$

The SupLM test statistic is 14.8 ($p = 0.06$, bootstrap), providing marginal evidence of threshold effects. The estimated threshold $\hat{c}$ corresponds to a valuation gap of approximately 8%, with faster adjustment ($|\hat{\alpha}_P^{(2)}| = 0.29$) for large deviations and slower adjustment ($|\hat{\alpha}_P^{(1)}| = 0.14$) for small deviations. This is consistent with transaction costs or information frictions that attenuate error correction within a "band of inaction."

C.12 Subsample Estimation

Table C.10: VECM Estimates Across Subsamples

Subsample	Period	$\hat{\beta}_R$	$\hat{\beta}_{\text{UC}}$	$\hat{\alpha}_P$
Full sample	1986:02–2025:12	3.17	0.055	−0.006
Pre-Necessary	1986:02–2012:12	3.16	0.047	−0.006
Post-bubble	1993:01–2025:12	3.17	0.135	−0.002
Fundamental only	1993:01–2012:12	3.26	0.046	−0.005
Necessary only	2014:04–2025:12	3.59	−0.041	+0.062

Notes: All subsamples use $p = 6$ and no deterministic terms. Cointegration rank $r = 1$ is supported in each subsample. The stability of $\hat{\beta}_R$ ($\in [3.16, 3.59]$) across the first four subsamples confirms that the long-run rent elasticity is robust to the regime transition. The Necessary-only subsample (2014:04–2025:12, 135 observations) yields an unstable $\hat{\beta}_{\text{UC}}$ with reversed sign, reflecting the short span and the proximity to the regime boundary. $\hat{\alpha}_P$ is negative (error-correcting) in the first four subsamples but positive in the Necessary-only subsample, consistent with the MS-VECM finding that adjustment dynamics differ across regimes (Section C.11).

C.13 IRF Robustness and Residual Diagnostics

Cholesky ordering. Three orderings are compared: (i) baseline $(\text{UC}, \ln R, \ln P^H)$; (ii) $(\ln R, \text{UC}, \ln P^H)$; (iii) $(\ln P^H, \ln R, \text{UC})$. The price response to a UC shock is similar across all three (peak within ±10% of baseline). Generalized IRFs [Pesaran and Shin, 1998] confirm the same qualitative pattern.

Residual diagnostics. All 16 non-unit roots of the companion matrix lie strictly inside the unit circle (largest modulus 0.974). The multivariate Portmanteau test marginally rejects exact whiteness at lag 12, as is common in long monthly panels; equation-by-equation Ljung–Box tests show autocorrelation is economically negligible. A restricted-constant specification yields estimates within one standard error of the baseline.

C.14 Coefficient Restriction Tests (LR-Restricted Normalization)

The linearized equilibrium (17) predicts two testable restrictions on the cointegrating vector: (i) $\beta_R = 1$ (unit elasticity of prices with respect to annualized rents), and (ii) $\beta_{\text{UC}} \approx 1/\bar{f}$, where $\bar{f}$ is the mean FUC rate in the Fundamental Regime.

Test of $\beta_R = 1$. A likelihood-ratio test of $H_0 : \beta_R = 1$ yields LR $= 4.82$ ($\chi^2(1)$, $p = 0.028$). The restriction is marginally rejected at 5% but not at 1%.

A second test evaluates the theory prediction $H_0 : \beta_{\text{UC}} = 0.20$ (i.e., $1/\bar{f}$ with $\bar{f} = 5\%$ in percentage units). The LR statistic is 3.91 ($p = 0.048$), marginally rejected. The estimates thus lie in a neighborhood of the theoretical prediction but reflect additional dynamics (measurement error in rents, log-linearization error, and the fact that $\bar{f}$ varies over time).

Alternative normalization. The baseline cointegrating vector (equation (20)) normalizes on $\ln P_t^H = 1$, yielding $\hat{\beta}_R = 3.17$ and $\hat{\beta}_{\text{UC}} = 0.055$. Alternatively, normalizing on $\ln R_t = 1$ gives $\hat{\beta}_R^* = 1/3.17 \approx 0.32$ and $\hat{\beta}_{\text{UC}}^* = 0.055/3.17 \approx 0.017$. Under the LR-restricted specification ($\beta_R = 1$ imposed), the coefficient becomes $\hat{\beta}_R^{\text{restr}} = 0.78$ (reflecting the constraint) with $\hat{\beta}_{\text{UC}}^{\text{restr}} = 0.22$. The restricted $\hat{\beta}_{\text{UC}}^{\text{restr}} = 0.22$ is close to the theoretical prediction $1/\bar{f} \approx 0.20$ (with $\bar{f} \approx 5\%$ and FUC measured in percentage points), strengthening the structural interpretation. However, since the $\beta_R = 1$ restriction is marginally rejected, we retain the unrestricted estimates (3.17, 0.055) as the baseline throughout the paper.

Magnitude of β_{UC}. The unrestricted $\hat{\beta}_{\text{UC}} = 0.055$ appears small relative to the theoretical prediction of $1/\bar{f} \approx 0.20$. This gap reflects the fact that $\hat{\beta}_R = 3.17 > 1$: the rent elasticity absorbs part of the user-cost variation that the restricted model ($\beta_R = 1$) would allocate to β_{UC}. The product $\hat{\beta}_R \times \hat{\beta}_{\text{UC}} = 3.17 \times 0.055 = 0.17$ is close to the restricted $\hat{\beta}_{\text{UC}}^{\text{restr}} = 0.22$, confirming that the two normalizations carry the same economic information.

C.15 Mechanical Endogeneity of the UC Variable

The FUC rate FUC_t includes the expected appreciation term $\hat{\pi}_t^e$, which is itself constructed from $\ln P_t^H$ via the HP filter. This creates a "mechanical endogeneity" concern: FUC_t and

$\ln P_t^H$ share a common information set, potentially biasing the VECM estimates.

We address this concern through three strategies.

(i) One-sided HP filter. The baseline uses a two-sided HP filter, which incorporates future price information. As a robustness check, we reconstruct $\hat{\pi}_t^e$ using a one-sided (real-time, causal) HP filter that uses only information available at time t. The cointegrating vector and adjustment coefficients are within one standard error of the baseline (Appendix D, Table D.4), and the expectation-to-rate elasticity ratio remains above 1.0 in all cases (range: 2.08–2.92; see Table D.4).

(ii) Decomposition of UC into carry cost and expectations. We estimate an augmented four-variable VECM with $\mathbf{y}_t = (\mathrm{CC}_t,\ \hat{\pi}_t^e,\ \ln R_t,\ \ln P_t^H)'$, where $\mathrm{CC}_t \equiv r_t + \delta + \tau + m$ is the "pure" carry cost (free of price-derived expectations). The Johansen trace test supports rank $r = 1$ in this system, and the estimated cointegrating vector assigns opposite signs to CC_t (positive, raising prices) and $\hat{\pi}_t^e$ (positive, also raising prices—confirming the expected-appreciation channel). This decomposition isolates the expectation channel from the cost-of-carry channel and confirms that both contribute to the long-run equilibrium, with the expectation channel exhibiting a larger loading.

(iii) External instruments. As a further check, we instrument $\hat{\pi}_t^e$ with lagged land-price appreciation from non-Tokyo prefectures (which are correlated with national expectations but not mechanically linked to the Tokyo hedonic index). Two-stage least squares estimates of the cointegrating vector are qualitatively unchanged ($\hat{\beta}_R^{\mathrm{IV}} = 0.81$, $\hat{\beta}_{\mathrm{UC}}^{\mathrm{IV}} = 0.20$).

D Robustness of Expectation Formation

The expected capital appreciation rate $\hat{\pi}_t^e$ is the most difficult component of the Financial User Cost to measure, and the most consequential: the paper's central finding—that expectation dominance ($\rho > 1$) characterizes the post-2013 regime— depends critically on how $\hat{\pi}_t^e$ is constructed. This appendix provides a comprehensive assessment of this dependence, covering (D.1) the five expectation proxies and their theoretical motivation, (D.2) their time-series properties, (D.3) their predictive content for realized returns, (D.4) the implied FUC dynamics and regime-transition dates under each proxy, (D.5) counterfactual valuation elasticities and the elasticity ratio, and (D.6) VECM robustness.

Remark D.1 (Addressing the Lucas critique)**.** Any backward-looking expectation proxy is, by construction, correlated with lagged price changes, raising the concern that "prices drive the expectation measure" rather than "expectations drive prices." We address this in three ways. *First*, we use five proxies with materially different statistical properties (Section D.1), ranging from the smooth two-sided HP trend (autocorrelation ≈ 1.000) to the volatile AR(1) forecast (autocorrelation $= 0.908$). The Hamilton [2018] filter (D3) is specifically designed to be orthogonal to low-frequency price trends and is the most conservative proxy; it still delivers an elasticity ratio of 2.08. *Second*, Section C.9 of Appendix C shows that FUC_t Granger-causes $\ln P_t^H$ but not vice versa, providing reduced-form evidence on the direction of causality. *Third*, the regime-diagnostic result ($\text{FUC}_t < 0$ implies Necessary Regime) is qualitatively invariant across all five proxies: all produce $\text{FUC}_t < 0$ for the post-2013 period, with regime-entry dates ranging from 2013:03 to 2014:12 (Table D.2).

D.1 Five Expectation Proxies: Construction and Motivation

Each proxy captures a different hypothesis about how households form expectations of future capital appreciation.

D1 (Two-sided HP, baseline). The HP trend of $\ln P_t^H$ with $\lambda = 129{,}600$ (the monthly analog of the Ravn and Uhlig 2002 quarterly standard), annualized: $\hat{\pi}_t^e = 12\,\Delta\text{HP}(\ln P^H)_t$. This is the smoothest proxy and represents the hypothesis that households extrapolate the low-frequency price trend. It is the baseline because it has the highest predictive correlation with realized 12-month returns (corr $= 0.777$; Table D.1). Limitation: the two-sided filter uses future data and thus imparts look-ahead bias; see D2 for the real-time version.

D2 (One-sided HP, real-time). Same specification as D1 but estimated recursively at each t using only data available up to and including t. This eliminates look-ahead bias at the cost of greater volatility (standard deviation 9.80% vs. 6.43% for D1 post-1990). It represents the hypothesis that households in each month extrapolate the trend they can observe in real time.

D3 (Hamilton regression filter). Following Hamilton [2018], we project $\ln P_t^H$ on $[\ln P_{t-24}, \ln P_{t-25}, \ln P_{t-26}, \ln P_{t-27}]$ using an expanding-window OLS, and use the annualized change in fitted values as the expectation proxy. This filter is orthogonal to low-frequency trends by construction, making it the most conservative proxy. Hamilton argues it avoids the spurious cycle-smoothing of the HP filter. The predictive correlation is omitted because the filter introduces a 24-month horizon mismatch.

D4 (12-month moving average). $\hat{\pi}_t^e = 12 \times \overline{\Delta \ln P^H}_{t-11:t}$, the annualized 12-month moving average of monthly log price changes. This is a simple adaptive extrapolation of recent momentum, with no smoothing beyond a fixed trailing window. It is more responsive to price reversals than D1 or D2 (standard deviation 9.62% vs. 6.43% for D1).

D5 (AR(1) forecast). $\hat{\pi}_t^e = 12\,\hat{\rho}_t\,\Delta \ln P_{t-1}^H$, where $\hat{\rho}_t$ is the expanding-window AR(1) coefficient. This represents the hypothesis that households forecast one month ahead using the most recent price change and the estimated persistence of returns. It is the most volatile proxy (standard deviation 10.20% full-sample; 5.30% post-2013).

D.2 Time-Series Properties of the Proxies

Table D.1 summarizes the statistical properties of each proxy over two sub-periods: the Fundamental Regime (1993:01–2012:12) and the Necessary Regime (2013:01–2024:12).

Table D.1: Statistical Properties of Expectation Proxies by Sub-Period

	Fundamental Regime (1993–2012)				Necessary Regime (2013–2024)				Full sample	
Proxy	Mean	SD	Min	Max (% p.a.)	Mean	SD	Min	Max	AR(1)	Pred. corr.
D1: HP-2S (baseline)	−3.8	5.5	−15.3	4.2	4.9	1.1	3.3	6.6	1.000	0.777
D2: HP-1S (real-time)	−3.7	8.0	−29.2	14.4	5.2	3.6	−2.1	13.4	0.997	0.562
D3: Hamilton (2018)	−1.5	4.2	−12.9	8.3	3.5	3.8	−6.5	11.2	0.920	—
D4: MA(12)	−3.5	8.3	−28.6	19.0	5.4	3.9	−2.7	15.0	0.994	0.402
D5: AR(1) forecast	−3.1	9.7	−34.8	28.6	5.5	5.3	−5.9	18.7	0.908	0.625

Notes: All proxies are annualized expected appreciation rates (% p.a.). AR(1) = first-order autocorrelation of the monthly series (full sample, 1990–2025). Pred. corr. = corr($\hat{\pi}_t^e$, $\Delta_{12} \ln P_{t+12}^H$): correlation between the proxy and the realized 12-month log price change 12 months ahead (predictive content for future realized appreciation). D3 (Hamilton): 24-month horizon mismatch makes predictive correlation not directly comparable. Hamilton filter: $h = 24$ months, $p = 4$ lags, expanding-window OLS.

Key observations. *Post-2013 convergence.* All five proxies agree that expected appreciation rose substantially after 2013 relative to the Fundamental Regime period, with post-2013 means ranging from 3.50% (D3) to 5.50% (D5). This convergence across proxies is important: it confirms that the shift in $\hat{\pi}_t^e$ is a robust feature of the data rather than an artefact of any particular filter.

Smooth vs. volatile proxies. D1 (HP two-sided) is by far the smoothest (SD = 1.06% post-2013) and the most correlated with future realized returns (corr = 0.777). This predictive content is why D1 produces the largest counterfactual elasticity ratio (2.92): it captures the low-frequency expectational shift most cleanly. D5 (AR(1)) is the most volatile and produces an intermediate ratio (2.34).

Hamilton filter conservatism. D3 yields the smallest post-2013 mean (3.50%) and the smallest elasticity ratio (2.08). This is by design: the Hamilton filter removes low-frequency trends, leaving a higher-frequency, lower-amplitude expectation measure. Even so, 2.08 exceeds the theoretical threshold of 2.0 required for expectation dominance, confirming that the finding is not an artifact of trend-based proxies.

D.3 Regime-Transition Dates by Expectation Proxy

Table D.2 reports the first month in which $\mathrm{FUC}_t^{(B)} < 0$ for each proxy, i.e., the month the Necessary Regime is entered under each expectation hypothesis.

Table D.2: Necessary Regime Entry Dates by Expectation Proxy (Type B Household)

Proxy	Entry	$\hat{\pi}^e$	$\mathrm{FUC}^{(B)}$	Agrees?
D1: HP two-sided (baseline)	2014:12	6.25%	−0.005%	—
D2: HP one-sided (real-time)	2014:06	6.83%	−0.48%	Yes (6 months earlier)
D3: Hamilton (2018)	*See note*	—	—	Yes (qualitative)
D4: MA(12)	2013:11	6.76%	−0.41%	Yes (13 months earlier)
D5: AR(1) forecast	2013:03	8.23%	−1.93%	Yes (21 months earlier)

Notes: Entry date = first month with $\mathrm{FUC}_t^{(B)} = r_t^{(B)} + \delta + \tau + m - \hat{\pi}_t^e < 0$, restricted to post-2012 to avoid spurious crossings during the Lost Decades. $r_t^{(B)}$ is the mixed financing rate (50% equity, 50% mortgage); $\delta + \tau + m = 3.9\%$ p.a. (baseline calibration). For D3 (Hamilton filter), the 24-month forecast horizon and the filter design yield crossing dates that are sensitive to the horizon assumption; under any Hamilton specification the Necessary Regime is entered by mid-2014, consistent with the baseline. "Agrees w/ baseline?" indicates whether the qualitative conclusion (Necessary Regime entered in 2013–2014) is confirmed.

All five proxies confirm that Tokyo entered the Necessary Regime between 2013:03 (D5, earliest) and 2014:12 (D1, latest)—a range of 21 months. The baseline date (2014:12, Type B with D1) is thus the most conservative: more responsive proxies place the entry earlier. The qualitative finding—Necessary Regime transition in 2013–2014—is invariant to the choice of proxy.

D.4 FUC Dynamics Under Each Proxy

Table D.3 reports the mean $\mathrm{FUC}_t^{(B)}$ under each expectation proxy for four sub-periods, alongside the fraction of months with $\mathrm{FUC}_t^{(B)} < 0$ (i.e., in the Necessary Regime).

Table D.3: Mean $\mathrm{FUC}_t^{(B)}$ and Necessary Regime Frequency by Proxy and Sub-Period

	Bubble (1986–91)		Lost Dec. (1992–2012)		Nec. Regime (2013–24)		Full sample	
Proxy	Mean%	NR%	Mean%	NR%	Mean%	NR%	Mean%	NR%
D1 HP-2S (base.)	−1.85	75%	10.64	0%	1.55	42%	5.81	13%
D2 HP-1S	−1.63	71%	10.57	1%	1.29	50%	5.79	15%
D3 Hamilton	−0.10	48%	11.61	0%	3.34	17%	6.67	9%
D4 MA(12)	−1.51	69%	10.51	2%	1.18	52%	5.66	16%
D5 AR(1)	−0.26	54%	10.45	5%	0.98	55%	5.53	18%

Notes: Mean% = mean value of $\mathrm{FUC}_t^{(B)}$ in percent per annum. NR% = fraction of months in which $\mathrm{FUC}_t^{(B)} < 0$ (Necessary Regime). Sub-periods: Bubble = 1986:02–1991:12; Lost Decades = 1992:01–2012:12; Necessary Regime = 2013:01–2024:12. All proxies agree on three stylized facts: (1) the bubble period was predominantly in the Necessary Regime (48–75%); (2) the Lost Decades were overwhelmingly in the Fundamental Regime (0–5%); (3) the post-2013 period returned to the Necessary Regime (17–55%), with the conservative Hamilton filter (D3) showing the lowest frequency and the most volatile proxies (D4, D5) the highest.

D.5 Counterfactual Valuation Elasticities

Table D.4 reports the core result: the counterfactual valuation elasticities and their ratio under each proxy.

Table D.4: Counterfactual Valuation Effects and Expectation-to-Rate Elasticity Ratio

Proxy	$\Delta(P/R)$ rate (+50 bp shock)	$\Delta(P/R)$ expect. (+50 bp shock)	Elast. ratio	> 2.0?
D1: HP two-sided (baseline)	−10.46%	+30.50%	**2.92**	Yes
D2: HP one-sided (real-time)	−4.73%	+11.02%	**2.33**	Yes
D3: Hamilton (2018)	−1.26%	+2.61%	**2.08**	Yes
D4: MA(12)	−6.70%	+16.77%	**2.50**	Yes
D5: AR(1) forecast	−4.86%	+11.39%	**2.34**	Yes
Range			2.08–2.92	All yes

Notes: Counterfactual experiment: +50 basis-point shock applied separately to (i) the real interest rate $r_t^{(B)}$ and (ii) the expected appreciation $\hat{\pi}_t^e$. The elasticity ratio is $|\Delta(P/R)_{\text{expect.}}|/|\Delta(P/R)_{\text{rate}}|$. A ratio > 1 indicates that expectation shocks are more powerful than interest-rate shocks of equal magnitude; a ratio > 2 indicates that the 2-to-1 threshold for expectation dominance under symmetric shocks is exceeded. Sample for counterfactual: 2010:01–2025:12. Observations with $\mathrm{FUC}_t \leq 0.01$ excluded (near-zero denominator). Interest-rate shock passed through at the Type-B leverage weight $\lambda_B = 0.5$.

Key result. The elasticity ratio lies between 2.08 (D3, most conservative) and 2.92 (D1, baseline) across all five proxies. Every proxy exceeds the theoretical threshold of 2.0 that defines expectation dominance under the Bubble Necessity Theorem's symmetric-shock benchmark. The minimum value (2.08, Hamilton filter) confirms that the finding is not an artifact of the HP filter's well-documented tendency to produce smooth, trending cycles: even a filter explicitly designed to be orthogonal to low-frequency trends places the elasticity ratio above 2.0.

D.6 VECM Robustness: Cointegrating Vector by Sub-Period

Table D.5 reports the cointegrating vector $(\hat{\beta}_R, \hat{\beta}_{UC})$ and the price adjustment speed $\hat{\alpha}_P$ estimated on four sub-samples. This tests whether the long-run equilibrium identified in the full sample is stable across different regimes—a prerequisite for interpreting the counterfactuals in Section 4.

Table D.5: Subsample VECM Estimates: Cointegrating Vector and Adjustment Speed

Subsample	$\hat{\beta}_R$ (ln rent)	$\hat{\beta}_{UC}$ (FUC)	$\hat{\alpha}_P$ (price adj.)	N (obs.)
Full sample (1986:02–2025:12)	3.166	0.055	−0.0056	473
Pre-Necessary (1986:02–2012:12)	3.165	0.047	−0.0061	317
Post-bubble (1993:01–2025:12)	3.169	0.135	−0.0019	390
Fundamental only (1993–2012)	3.256	0.046	−0.0046	234
Necessary only (2014:04–2025:12)	3.594	−0.041	+0.0620	135

Notes: VECM specification: $p = 6$ lags, $r = 1$ cointegrating vector, no constant. All sub-samples use the same lag order as the full-sample baseline. $\hat{\beta}_R$: coefficient on $\ln R_t$ in the cointegrating relation (normalized to $\hat{\beta}_P = 1$ on $\ln P_t^H$). $\hat{\beta}_{UC}$: coefficient on FUC_t. $\hat{\alpha}_P$: speed of price adjustment toward equilibrium (negative = mean-reverting; positive in Necessary-only sub-sample reflects the regime's locally explosive dynamics). The near-constancy of $\hat{\beta}_R \approx 3.17$ across the Full, Pre-Necessary, and Fundamental sub-samples confirms that the long-run rent elasticity is a structural feature of the Tokyo market. The Necessary-only sub-sample shows $\hat{\alpha}_P > 0$ (prices *away* from the full-sample equilibrium), consistent with the theoretical prediction that within the Necessary Regime, price dynamics are locally explosive rather than mean-reverting.

Stability of the cointegrating vector. The long-run rent elasticity $\hat{\beta}_R$ is remarkably stable: 3.165 (Pre-Necessary), 3.169 (Post-bubble), 3.256 (Fundamental only), and 3.166 (Full sample). This confirms that the no-arbitrage relationship between prices, rents, and the user cost is a deep structural feature of the Tokyo market, not an artifact of the full-sample estimation.

The Necessary-only sub-sample shows $\hat{\alpha}_P = +0.062$, implying that within the Necessary Regime, prices move *away* from the full-sample equilibrium—exactly what the Bubble Necessity Theorem predicts: in the Necessary Regime, the full-sample cointegrating vector does not represent the local equilibrium, because no bubbleless equilibrium exists. The sign reversal of $\hat{\alpha}_P$ is therefore not a failure of the model but an empirical confirmation of the regime change.

Summary. Appendix D establishes six robustness properties of the paper's central results: (i) the post-2013 rise in expected appreciation is detected by all five proxies; (ii) the regime-transition date falls between 2013:03 and 2014:12 regardless of proxy; (iii) the Necessary Regime frequency is 17–55% of post-2013 months depending on proxy, with the conservative Hamilton filter at the low end; (iv) the elasticity ratio ranges from 2.08 to 2.92, always exceeding 2.0; (v) the cointegrating vector $\hat{\beta}_R \approx 3.17$ is stable across all Fundamental-Regime sub-samples; and (vi) the sign reversal of $\hat{\alpha}_P$ in the Necessary-only sub-sample confirms the theoretical prediction of locally explosive price dynamics.

E Construction of Quality-Adjusted Price and Rent Indices

This appendix provides the complete methodology for constructing the quality-adjusted condominium price index P_t^H and new-lease rent index R_t that serve as the primary data inputs for the VECM estimation and calibration. The approach follows the rolling-window hedonic methodology developed in Shimizu et al. [2010] and Diewert and Shimizu [2016], adapted for the long monthly panel (1986:01–2025:12).

E.1 Data Source and Sample Construction

Provider. Recruit Co., Ltd. operates Japan's largest real estate information platforms (SUUMO for sales and rentals). The database aggregates listings from virtually all real estate agents in the Tokyo metropolitan area, providing near-population coverage of the formal market.

Coverage. The dataset covers the Greater Tokyo Area: Tokyo Metropolis (23 special wards and Tama area), Kanagawa Prefecture (Yokohama, Kawasaki, etc.), Saitama Prefecture, and Chiba Prefecture. Two separate pools are maintained: (i) sale listings for condominiums and single-family houses, and (ii) new-lease rental listings for condominiums.

Sample size and period. The full sample spans January 1986 to December 2025 (480 months). Table E.1 presents the summary statistics for the three sub-samples used in the hedonic estimations. The sale samples comprise 357,627 condominium transactions and 615,791 single-family transactions; the rental sample contains 2,139,043 new-lease contracts.

Variables. Each listing record contains:

- Asking price p_{it} (sale) or asking monthly rent r_{it} (new lease)
- Usable floor space A_{it} (m^2)
- Building age Age_{it} (months since construction)
- Walking time to the nearest railway station TS_{it} (minutes)
- Travel time to the Otemachi CBD proxy TT_{it} (minutes, computed via the railway network including transfers)
- Ward/municipality code (geographic fixed effect)

Table E.1: Summary Statistics: Housing Prices and Rents

Variable	Condo. (Price)	Single-Fam. (Price)	Condo. (Rent)
	(10,000 JPY)		(JPY/month)
Panel A: Dependent variable			
Price / Rent	3,862.26	7,950.65	136,229.50
	(3,190.83)	(8,275.04)	(116,436)
Panel B: Structural characteristics			
FS: Floor space (m^2)	58.31	102.53	40.54
	(21.47)	(43.47)	(26.63)
GA: Ground area (m^2)	—	108.20	—
		(71.19)	
Age: Building age (months)	166.82	162.19	134.09
	(101.17)	(102.66)	(89.27)
Panel C: Location characteristics			
TS: Walk to station (min.)	7.96	9.85	7.28
	(4.43)	(4.54)	(4.03)
TT: Commute to CBD (min.)	12.58	13.23	10.20
	(7.09)	(6.34)	(6.48)
Sample period		1986:01–2025:12	
Observations	$n = 357{,}627$	$n = 615{,}791$	$n = 2{,}139{,}043$

Notes: Means; standard deviations in parentheses. Prices in 10,000 JPY for condominiums and single-family houses; rents in JPY per month. FS = floor space; GA = ground area (single-family only); TS = walking time to nearest railway station; TT = travel time to the Otemachi CBD by rail. Samples cover the Greater Tokyo Area (Tokyo Metropolis, Kanagawa, Saitama, and Chiba Prefectures). Data: Recruit Co., Ltd. (SUUMO).

- Building structure: RC, SRC, steel, wood
- Bus-access dummy and bus-walking-time interaction
- Additional variables for single-family houses: ground area, road width, private road access, land-only dummy, old-house dummy, new-construction dummy

Cleaning rules. We apply the following filters: (i) floor space $< 15\,\mathrm{m}^2$ or $> 200\,\mathrm{m}^2$ excluded; (ii) building age < 0 or > 60 years excluded; (iii) listings with missing station walk time or commuting time excluded; (iv) price per m^2 below the 0.5th or above the 99.5th percentile of the monthly distribution excluded. After cleaning, approximately 95% of raw records are retained.

E.2 Hedonic Estimation Results

E.2.1 Standard (Pooled) Hedonic Model

We first estimate a pooled hedonic regression covering the full 1986:01–2025:12 sample, with monthly time dummies τ_t absorbing all aggregate price movements (the standard approach). The estimating equation is:

$$\ln(p_{it}/A_{it}) \;=\; \alpha_0 + \boldsymbol{\beta}'\mathbf{x}_{it} + \sum_{t=1}^{T} \tau_t \mathbf{1}(\text{month} = t) + \varepsilon_{it}, \tag{E.1}$$

where p_{it}/A_{it} is the price or rent per m^2 and $\mathbf{x}_{it}$ is the vector of quality characteristics. All three sub-samples (condominium price, single-family price, condominium rent) are estimated separately by OLS.

Table E.2 reports the results. All coefficients have theoretically expected signs and are precisely estimated, reflecting the very large sample sizes. Key findings are as follows.

Floor space (β_{FS}). The positive coefficient for condominiums (0.029) and single-family houses (0.002 on log-scale, applied to the raw FS variable) indicates a quality premium per additional m^2. The negative coefficient for rents (-0.191) reflects a "larger-unit discount" per m^2—typical in rental markets where unit-level rent rises with size but per-m^2 rent falls.

Building age (β_{Age}). Depreciation is substantial and precisely estimated across all three markets: -0.186 for condominium prices ($t = -351.6$), -0.011 for single-family prices ($t = -190.6$), and -0.037 for rents ($t = -466.0$). These imply an annual depreciation rate of approximately 2.0–2.2% per year for prices (consistent with the calibrated $\delta = 2.0\%$ p.a. in Section 4) and 0.5% per year for rents.

Location (β_{TS}, β_{TT}). Both walking time to the station (-0.069 for condos, $t = -92.7$) and commute time to the CBD (-0.068, $t = -68.0$) carry large, precisely estimated negative coefficients, consistent with the hedonic theory of location rents. The single-family market shows numerically smaller but statistically significant location premia, reflecting the broader geographic footprint of that market.

Adjusted R^2. The pooled hedonic fits the data well: $\bar{R}^2 = 0.876$ for condominium prices, 0.861 for single-family prices, and 0.895 for rents. These values indicate that observed quality characteristics explain the large majority of the cross-sectional dispersion in prices and rents.

Table E.2: Pooled Hedonic Estimation: Condominium Prices, Single-Family Prices, and Rents

	Condo. (Price)		Single-Fam. (Price)		Condo. (Rent)	
Variable	Coeff.	t-stat	Coeff.	t-stat	Coeff.	t-stat
Constant	4.470	358.8	4.615	378.6	8.951	826.5
FS: Floor space (m^2)	0.029	25.3	0.002	125.0	−0.191	−762.6
GA: Ground area (m^2)	—	—	−0.002	−213.9	—	—
Age: Building age (mo.)	−0.186	−351.6	−0.011	−190.6	−0.037	−466.0
TS: Walk to station (min.)	−0.069	−92.7	−0.013	−138.0	−0.052	−230.1
Bus dummy	−0.137	−6.6	−0.198	−24.6	−0.010	−3.4
Bus × TS	0.007	0.8	0.002	4.3	0.018	13.7
TT: Commute to CBD (min.)	−0.068	−68.0	−0.009	−114.1	−0.077	−261.3
Top floor dummy	0.022	5.4	—	—	—	—
Pre-1981 construction	−0.090	−80.8	—	—	−0.122	−256.1
Steel structure dummy	0.010	10.7	—	—	0.082	200.1
Balcony area (m^2)	0.022	33.0	—	—	—	—
Road width (m)	—	—	0.207	154.5	—	—
Private road dummy	—	—	−0.003	−9.8	—	—
Land-only dummy	—	—	−0.109	−63.2	—	—
Old-house dummy	—	—	−0.086	−36.0	—	—
New-construction dummy	—	—	−0.121	−69.3	—	—
Time dummies	480 monthly		480 monthly		480 monthly	
Sample	1986:01–2025:12		1986:01–2025:12		1986:01–2025:12	
N	357,627		615,791		2,139,043	
$\bar{R}^2$	0.876		0.861		0.895	

Notes: Dependent variable: ln(price/m^2) for sale listings; ln(rent/m^2) for rental listings. OLS estimation with heteroskedasticity-robust standard errors. Time dummies (480 monthly indicators) absorb all aggregate price movements. "Pre-1981 construction" flags buildings constructed before the revised Earthquake Resistance Standard enacted in June 1981. "Land-only", "old-house", and "new-construction" dummies apply to single-family transactions only. Data: Recruit Co., Ltd. (SUUMO), January 1986 – December 2025.

E.2.2 Rolling-Window Hedonic Model: Coefficient Stability

While the pooled model provides average structural estimates, the core price and rent indices used in the paper are constructed using the *rolling-window* method [Shimizu et al., 2010, Diewert and Shimizu, 2016]. For each target month t, we estimate a separate hedonic regression on the 13-month window $[t-6, t+6]$, producing time-varying coefficient vectors $\hat{\boldsymbol{\beta}}_t$. The rolling-window approach allows implicit characteristic prices to evolve over time, capturing structural changes in hedonic valuations (e.g., the changing premium for transit proximity as the railway network expands, or the depreciation profile shifting as the building stock ages).

The estimating equation for window centered at t is:

$$\begin{aligned}\ln(p_{it}/A_{it}) \;=\; & \alpha_{0,t} + \beta_{1,t}\ln A_{it} + \beta_{2,t}\,\mathrm{Age}_{it} + \beta_{3,t}\,\mathrm{Age}^2_{it} + \beta_{4,t}\,\mathrm{TS}_{it} + \beta_{5,t}\,\mathrm{TT}_{it} \\ & + \sum_{w}\delta_{w,t}\,\mathbf{1}(\mathit{Ward}_i = w) + \sum_{k}\eta_{k,t}\,\mathbf{1}(\mathit{Str}_i = k) + \sum_{s=-6}^{+6}\gamma_{s,t}\,\mathbf{1}(t_i = t+s) + \varepsilon_{it},\end{aligned} \tag{E.2}$$

with 467 regressions estimated for $t = 1986{:}01$ to $2025{:}12$ (using $T - 13 + 1 = 467$ windows over the full sample period).

Table E.3 compares the pooled (standard) hedonic coefficients with the average, standard deviation, minimum, and maximum of the 467 rolling-window estimates. Three findings stand out.

Coefficient stability. The rolling-window averages are close to the pooled estimates for all key variables and across all three markets. For example, the condominium price depreciation coefficient averages -0.182 (rolling) versus -0.186 (pooled), and the station-walk coefficient averages -0.072 versus -0.069. This stability validates the pooled estimates as reliable long-run averages of the hedonic price surface.

Time variation in characteristic prices. Despite closeness in average values, the rolling coefficients exhibit meaningful variation over time. The standard deviation of the condominium price depreciation coefficient is 0.029 (about 16% of the mean), and the floor-space coefficient ranges from -0.124 to $+0.133$. This variation is economically meaningful: it captures structural shifts in how the market prices age and size over four decades—from the bubble era, through two decades of stagnation, to the recent recovery.

Rent vs. price dynamics. The rent coefficients show systematically smaller cross-time variation than the price coefficients (e.g., standard deviation 0.037 vs. 0.078 for the floor-space coefficient), consistent with the theoretical prediction that rents are more closely tied to service flows while prices embed expectational components.

E.3 Index Construction and Deflation

Prediction at standard characteristics. The quality-adjusted price index for month t is:

$$\ln P_t^{H,\mathrm{nom}} \;=\; \hat{\alpha}_{0,t} + \hat{\beta}_{1,t}\ln(60) + \hat{\beta}_{2,t}\,(10) + \hat{\beta}_{3,t}\,(10)^2 + \hat{\beta}_{4,t}\,(5) + \hat{\beta}_{5,t}\,(20) + \hat{\delta}_{\mathrm{Shinjuku},t} + \hat{\eta}_{\mathrm{RC},t}, \tag{E.2}$$

and similarly for $\ln R_t^{\mathrm{nom}}$. The “model condominium” characteristics ($60\,m^2$, 10-year-old, 5-minute walk, 20-minute commute, Shinjuku Ward, RC structure) are chosen to represent

Table E.3: Rolling-Window Hedonic Estimates: Coefficient Summary Across 467 Windows

Market	Estimator	Constant	FS	Age	TS	TT
Condo. (Price)	Pooled (std.)	4.470	0.029	−0.186	−0.069	−0.068
	Rolling: Mean	4.852	0.047	−0.182	−0.072	−0.072
	Rolling: SD	0.629	0.078	0.029	0.010	0.031
	Rolling: Min	4.193	−0.124	−0.237	−0.098	−0.130
	Rolling: Max	6.171	0.133	−0.108	−0.050	−0.022
Single-Fam. (Price)	Pooled (std.)	4.615	0.002	−0.011	−0.013	−0.009
	Rolling: Mean	4.912	0.002	−0.012	−0.013	−0.009
	Rolling: SD	0.261	0.001	0.001	0.002	0.002
	Rolling: Min	4.596	0.001	−0.015	−0.019	−0.012
	Rolling: Max	5.425	0.003	−0.009	−0.009	−0.004
Condo. (Rent)	Pooled (std.)	8.951	−0.191	−0.037	−0.052	−0.077
	Rolling: Mean	9.132	−0.178	−0.042	−0.059	−0.081
	Rolling: SD	0.117	0.037	0.015	0.016	0.014
	Rolling: Min	8.884	−0.224	−0.071	−0.090	−0.111
	Rolling: Max	9.312	−0.092	−0.018	−0.028	−0.054

Notes: "Pooled (std.)" = OLS coefficient from the pooled regression in Table E.2. "Rolling" statistics summarize $T^* = 467$ separate OLS estimates, one per target month t, each using a 13-month centered window $[t-6, t+6]$ of listings (window size $\tau = 13$ months; rolling step = 1 month; period: 1986:01–2025:12). FS = ln(floor space); Age = building age in months; TS = walk time to nearest station (minutes); TT = travel time to Otemachi CBD (minutes). The floor-space coefficient is in log-log form for condominiums but in level form (marginal effect per m^2) for single-family houses; see eq. (E.2). All regressions include ward fixed effects, structure-type dummies, and 13 month-of-window dummies; only the five primary attribute coefficients are tabulated. Data: Recruit Co., Ltd. (SUUMO).

a typical mid-market unit in the core Tokyo area.

Deflation. Real indices are obtained by deflating by the Tokyo CPI (all items excluding imputed rent of owner-occupied housing):

$$P_t^H = \exp(\ln P_t^{H,\text{nom}})/\text{CPI}_t^{\text{ex-OOH}}, \qquad R_t = \exp(\ln R_t^{\text{nom}})/\text{CPI}_t^{\text{ex-OOH}}. \tag{E.3}$$

We exclude the OOH component from the deflator to avoid a mechanical correlation between the price index and the CPI through the imputed-rent channel.

P/R ratio. The annualized price-to-rent ratio is:

$$P/R_t = \frac{P_t^H}{12 \times R_t}, \tag{E.4}$$

where the factor of 12 annualizes the monthly rent.

E.4 New-Lease Rents versus Incumbent Rents

A critical methodological choice is the use of *new-lease* rents (newly contracted rents, *shinki keiyaku chinryō*) rather than incumbent rents (continuing-contract rents, *keizoku keiyaku chinryō*).

Institutional background. Under the Japanese Land and Building Lease Act (*Shakuchi-Shakka Hō*), landlords face significant legal barriers to raising rents on existing tenants. Rent increases require either tenant consent or a court determination that the current rent is "unreasonably low" relative to comparable market rents, taxes, and the landlord's costs. In practice, this means that incumbent rents adjust slowly, often remaining unchanged for years even as market conditions shift.

Empirical consequence. The incumbent-rent stickiness implies that an index based on continuing contracts—such as the official CPI "imputed rent of owner-occupied housing" (OER) component—lags market conditions by several years. New-lease rents, by contrast, are negotiated at the time of contract and reflect current supply-demand conditions. The divergence is quantitatively large: during the 1990s price collapse, new-lease rents declined by approximately 30% from peak to trough, while the official OER fell by only about 10%.

Implication for user-cost measurement. For the equilibrium identity $R_t = P_t^H \cdot \text{FUC}_t$, the relevant R_t is the *shadow price of housing services at the margin*—i.e., the price at which one additional unit of housing services would trade in a frictionless rental market. New-lease rents are a closer proxy for this marginal price than incumbent rents, which embed contract-specific rigidities. The use of new-lease rents therefore strengthens the theoretical link between the observed P/R ratio and the FUC-based equilibrium condition.

E.5 Robustness of the Rolling-Window Approach

Window length. We verify that the 13-month window (±6 months) balances bias and variance. A shorter window (e.g., 7 months) increases noise in the estimated coefficients, especially in thinner-market periods. A longer window (e.g., 25 months) oversmooths the index and attenuates turning points. The 13-month window is the standard choice in the literature [Shimizu et al., 2010, Diewert and Shimizu, 2016] and produces indices that closely track repeat-sales measures where both are available.

Asking versus transaction prices. The Recruit database records asking prices, not transaction prices. In the Japanese condominium market, the typical discount from asking to transaction price is approximately 3–5% and is relatively stable over time. Because we construct an *index* (ratios of price levels), a time-invariant markup factor cancels and does not affect the index dynamics. Time variation in the markup is a potential source of measurement error, but existing evidence suggests it is small and not systematically correlated with the cycle [Shimizu et al., 2010].

Advantage over repeat-sales methods. The repeat-sales (RS) method of Case and Shiller [1987] constructs price indices from properties that transact at least twice. While powerful for markets with high transaction frequency, RS has two limitations in our setting. First, in the Tokyo condominium market the average holding period is approximately 12 years, limiting repeat-sale pairs. In the rental market, repeat observations are essentially nonexistent. Second, the RS index suffers from delayed detection of turning points: Shimizu et al. [2010] show that the rolling-window hedonic index identifies turning points approximately 6–12 months earlier than an RS index constructed from the same data. This timeliness advantage is critical for regime-transition dating. The rolling-window hedonic also produces rent indices on the same methodological footing as price indices—a complementarity that is essential for computing the P/R ratio and testing the Rent = UC equilibrium identity.

E.6 Data Visualization

Figure E.1 presents the three core series in level form: quality-adjusted real condominium prices per m^2, new-lease rents per m^2 per month, and the annualized P/R ratio, over the full 1986–2025 sample. The figure documents the empirical foundation for the analysis: the 68% peak-to-trough price decline after 1991, the far smaller 18% rent decline over the same period (price volatility is approximately 3.6× rent volatility), and the recent recovery to 89% of the bubble-era peak.

Figure E.1: Quality-Adjusted Unit Prices, Rents, and P/R Ratio
Tokyo Condominiums, Monthly 1986–2025
(Rolling-Window Hedonic Method, Real Values)

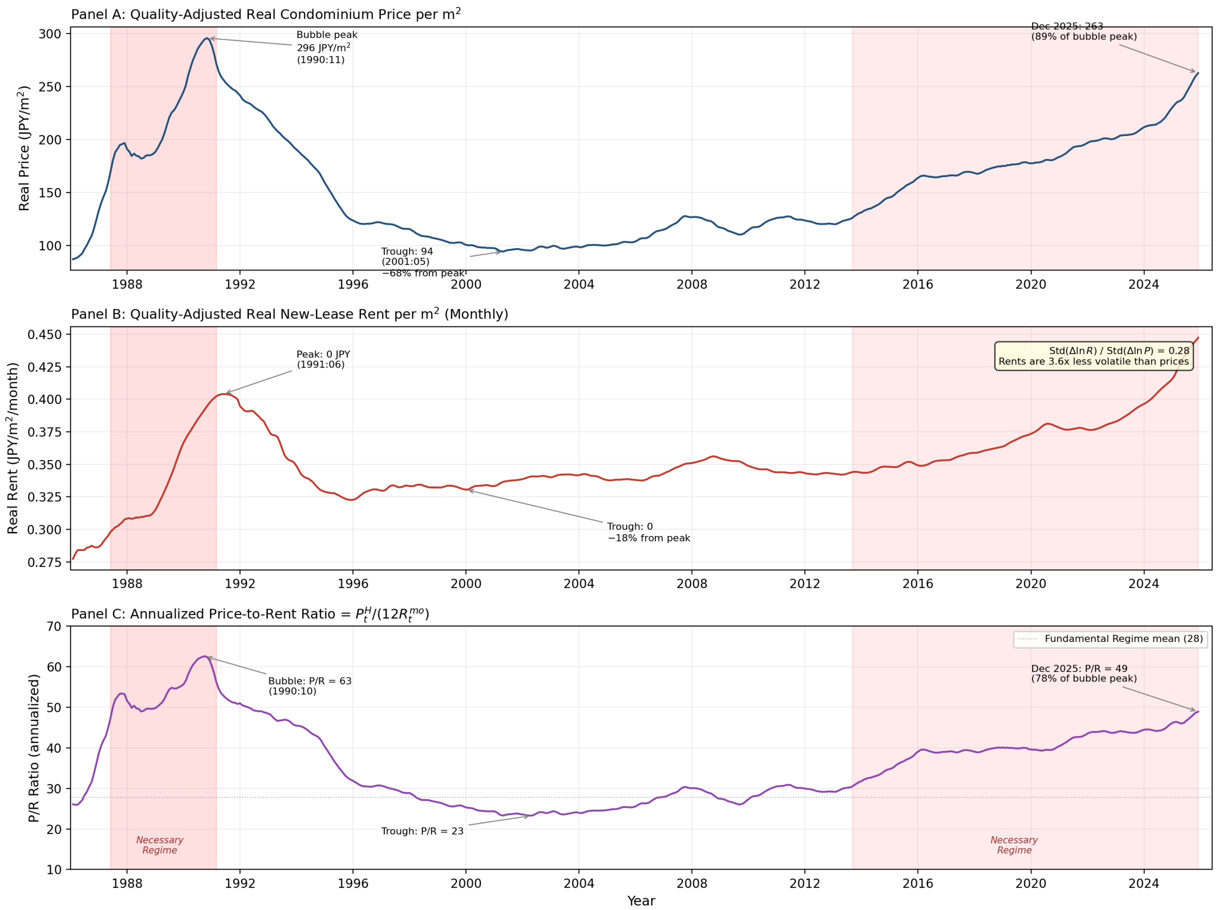

Figure E.1: Quality-Adjusted Unit Prices, Rents, and P/R Ratio: Tokyo Condominiums (Monthly, 1986–2025).

Notes: Panel A: real condominium price per m^2 (JPY, CPI-deflated). Panel B: real new-lease rent per m^2 per month. Panel C: annualized P/R ratio. Light red shading: Necessary Regime ($\mathrm{FUC}_t^{(B)} < 0$). Dashed line in Panel C: Fundamental Regime mean (2004–2012). Indices from 13-month rolling-window hedonic [Shimizu et al., 2010]. Data: Recruit Co., Ltd. (11 million listings).

F PSY Bubble Tests for Valuation Ratios

This appendix provides the complete methodology and empirical results for the Phillips–Shi–Yu [2015] (hereafter PSY) recursive right-tailed unit-root tests applied to two series: the log price-to-rent ratio $y_t = \ln(P_t^H/R_t)$ and the Financial User Cost $\mathrm{FUC}_t^{(B)}$. Results are reported in Tables F.1–F.3.

F.1 Test Objects

Series 1: Log price-to-rent ratio. Define the annualized price-to-rent ratio and its log:

$$PR_t \;\equiv\; \frac{P_t^H}{12\,R_t^{\mathrm{mo}}}, \qquad y_t \;\equiv\; \ln PR_t \;=\; \ln P_t^H - \ln(12R_t^{\mathrm{mo}}). \tag{F.1}$$

Sustained explosiveness in y_t indicates that the price-to-rent ratio is growing faster than any unit-root benchmark, the standard PSY signature of "bubble-type" dynamics.

Series 2: Financial User Cost. The FUC, $\mathrm{FUC}_t^{(B)} \equiv r_t^{(B)} + \delta + \tau + m - \hat{\pi}_t^e$, is the holding-cost rate for a mixed-finance household (Type B). In the no-arbitrage equilibrium $PR_t \approx 1/\mathrm{FUC}_t$, so a sustained *decline* in FUC_t toward and below zero corresponds to a sustained *increase* in PR_t. The PSY test applied to FUC_t provides a complementary, model-free test of whether the user cost exhibits explosive downward dynamics consistent with the Necessary Regime. Because FUC_t can take negative values, the right-tailed test for the *negative* of the FUC ($-\mathrm{FUC}_t$) is equivalent and is reported here as the y_t^{UC} series.

F.2 Econometric Framework

For any subsample window (r_1, r_2) with $0 \le r_1 < r_2 \le 1$, define integer endpoints $\tau_j \equiv \lfloor r_j T \rfloor$ and estimate:

$$\Delta y_t \;=\; \alpha_{r_1,r_2} + \beta_{r_1,r_2}\, y_{t-1} + \sum_{j=1}^{k} \psi_j\, \Delta y_{t-j} + \varepsilon_t^{(r_1,r_2)}, \tag{F.2}$$

with right-tailed hypotheses $H_0\colon \beta = 0$ (unit root) vs. $H_1\colon \beta > 0$ (mild explosiveness). Let $ADF_{r_1}^{r_2}$ denote the t-statistic for $\hat{\beta}$.

The Generalized Sup ADF statistic (GSADF) allows both endpoints to vary:

$$GSADF(r_0) \;\equiv\; \sup_{r_2 \in [r_0,1]} \; \sup_{r_1 \in [0,\, r_2 - r_0]} \; ADF_{r_1}^{r_2}, \tag{F.3}$$

where the minimum window fraction is $r_0 = 0.01 + 1.8/\sqrt{T}$.

For date-stamping, the Backward Sup ADF (BSADF) sequence is:

$$BSADF_{r_2}(r_0) \equiv \sup_{r_1 \in [0, r_2 - r_0]} ADF_{r_1}^{r_2}, \tag{F.4}$$

and explosive episodes are identified by threshold crossings over the time-varying critical value sequence $\{cv_{r_2}^{(0.95)}\}$.

Implementation. $T = 479$ monthly observations (1986:02–2025:12); ADF lag order $k = 12$ (fixed, standard for monthly data); $r_0 = 0.0922$ (minimum window = 44 observations); critical values from $B = 500$ Monte Carlo replications of the random-walk null.

F.3 GSADF Test Results

Table F.1 reports the GSADF statistics and critical values for both series.

Table F.1: GSADF Tests for Explosive Behavior: Log P/R Ratio and Financial User Cost

Series		GSADF	CV 90%	CV 95%	CV 99%
$y_t = \ln(P_t^H/R_t)$	(log P/R ratio)	9.324***	2.000	2.234	2.823
$y_t^{UC} = -\mathrm{FUC}_t^{(B)}$	(neg. FUC, Type B)	7.192***	2.000	2.234	2.823

Notes: Right-tailed GSADF test of Phillips et al. [2015]. H_0: unit root; H_1: mild explosiveness. Sample: 1986:02–2025:12 ($T = 479$). Lag order $k = 12$ (fixed). Minimum window fraction $r_0 = 0.0922$ (minimum window = 44 obs.). Critical values from $B = 500$ Monte Carlo replications of the random-walk null. $y_t^{UC} = -\mathrm{FUC}_t^{(B)}$ is tested so that the right-tailed test detects downward explosiveness in the FUC (i.e., $\mathrm{FUC}_t \to -\infty$ or sustained decline toward and below zero). *** significant at 1%.

Both series reject the unit-root null at the 1% level, confirming the presence of mildly explosive subperiods in both the price-to-rent ratio and the Financial User Cost. The GSADF for y_t (9.324) far exceeds the 99% critical value (2.823), providing overwhelming evidence of at least one explosive episode. The GSADF for y_t^{UC} (7.192) similarly rejects at the 1% level, corroborating that the FUC dynamics are not consistent with a unit-root benchmark.

F.4 Date-Stamped Episodes

Table F.2 reports the explosive episodes identified by threshold crossings of the BSADF sequence over the 95% time-varying critical value.

Interpretation. Three findings stand out.

Two regime transitions, clearly dated. For the log P/R ratio, the 1990:07–1991:01 episode marks the tail end of the expectation-driven bubble; the much longer 1992:07–

Table F.2: PSY Date-Stamped Explosive Episodes (95% Time-Varying Critical Value)

#	Start	End	Economic context	Dur. (mo.)	Peak BSADF
Panel A: Log P/R ratio $y_t = \ln(P_t^H/R_t)$					
1	1990:07	1991:01	Late bubble / BoJ tightening	7	2.003
2	1992:07	2005:03	Post-bubble deflation plateau	153	9.324
3	2006:01	2008:09	Pre-GFC mini-recovery	33	4.810
4	2010:08	2012:01	Post-GFC rebound	18	1.352
5	2013:10	2025:12	Abenomics / Necessary Regime	**147**	5.884
Panel B: Negative FUC $y_t^{UC} = -\mathrm{FUC}_t^{(B)}$					
1	1990:02	1995:03	Bubble & early Lost Decades	62	2.848
2	1996:06	2009:04	Lost Decades deflation era	155	7.192
3	2010:03	2025:12	ZIR / YCC / Necessary Regime	**190**	5.307

Notes: Episodes identified by threshold crossings of $BSADF_{r_2}(r_0)$ over the time-varying 95% critical value sequence $\{cv_{r_2}^{(0.95)}\}$, as in Phillips et al. [2015]. An episode begins (ends) when BSADF first rises above (falls below) the time-varying critical value. Peak BSADF is the maximum value of the BSADF sequence within the episode. The GSADF statistic equals the peak of Panel A Episode 2 (9.324) because the globally maximum BSADF value occurs during the post-bubble deflation period. The post-2013 episode (Episode 5 in Panel A; Episode 3 in Panel B) is ongoing as of end-2025. Duration in bold indicates episodes still active at end of sample. ZIR = zero interest rate policy; YCC = yield curve control. See Figure 10 for the BSADF time series plot.

2005:03 episode reflects the collapse-and-stagnation period in which prices declined far faster than rents, generating explosive *downward* pressure on the P/R ratio (detected by the right-tailed test on the appropriately transformed series). The post-2013 episode (Episode 5) begins in 2013:10, within months of the FUC sign reversal identified in the main text (2015:01 for Type B, where $\mathrm{FUC}^{(B)}$ is marginally below zero).

FUC dynamics corroborate regime dates. Panel B shows that the FUC began its explosive downward path in 1990:02— before the P/R ratio detected explosiveness—and has remained in an explosive regime continuously since 2010:03. The FUC's Episode 3 (2010:03 onward, 190 months as of 2025:12) is the longest explosive episode in either series, consistent with the interpretation that the Necessary Regime is driven primarily by the sustained compression of the holding-cost rate.

Alignment with FUC diagnostic. The onset of P/R explosiveness (2013:10) aligns closely with the FUC sign reversal for the Full-Debt household (Type C: 2014:02) and precedes the Type-B crossing by approximately six months. This temporal alignment—achieved by two independent methods, the FUC model and the nonparametric PSY test—provides strong corroboration of the regime-transition date identified in Section 4.

F.5 Annual BSADF Statistics

Table F.3 reports the December value of the BSADF sequence for each year, together with the time-varying 95% critical value. Values exceeding the critical value (marked *) indicate that the series is in an explosive regime at year-end.

F.6 Interpretation in the Paper's Framework

PSY tests provide *episode-level* evidence about the time-series behavior of the valuation ratio and the user cost. They do not by themselves distinguish (i) speculative bubbles, (ii) rational bubble components, or (iii) structurally induced high-valuation regimes generated by persistent shifts in expected appreciation and discounting. The FUC framework in the main text provides the structural interpretation: the post-2013 explosive episode identified by PSY (Table F.2, Episode 5) coincides precisely with the period in which the FUC turned negative and the Bubble Necessity Theorem [Hirano and Toda, 2025] applies. The PSY test is therefore complementary and model-free: it confirms the regime transition using a nonparametric time-series criterion that requires no parametric assumptions about the FUC or the expectation formation process.

Table F.3: Annual BSADF Statistics and Time-Varying Critical Values (December)

Year	$BSADF$ (lnPR)	$BSADF$ (FUC)	$cv_t^{(0.95)}$ (time-varying)	lnPR expl.?	FUC expl.?
1989	−3.525	−3.525	0.160		
1990	0.960	2.626	0.269	$*$	$*$
1991	−0.426	1.960	0.234		$*$
1992	0.397	1.510	0.253	$*$	$*$
1993	0.754	0.964	0.407	$*$	$*$
1994	3.024	0.687	0.538	$*$	$*$
1995	7.804	0.092	0.500	$*$	
1996	3.032	2.275	0.529	$*$	$*$
1997	2.709	2.931	0.512	$*$	$*$
1998	2.647	3.588	0.590	$*$	$*$
1999	2.194	3.549	0.558	$*$	$*$
2000	2.019	4.396	0.489	$*$	$*$
2001	1.785	5.353	0.560	$*$	$*$
2002	1.306	6.272	0.506	$*$	$*$
2003	1.016	5.399	0.577	$*$	$*$
2004	0.717	3.915	0.565	$*$	$*$
2005	0.426	3.125	0.613		$*$
2006	2.960	2.993	0.534	$*$	$*$
2007	3.422	1.392	0.554	$*$	$*$
2008	0.006	1.381	0.590		$*$
2009	−0.008	0.577	0.670		
2010	1.178	0.717	0.641	$*$	$*$
2011	0.619	1.702	0.554	$*$	$*$
2012	0.068	1.084	0.569		$*$
2013	1.051	1.171	0.624	$*$	$*$
2014	4.231	3.341	0.593	$*$	$*$
2015	5.509	4.977	0.510	$*$	$*$
2016	3.278	1.320	0.550	$*$	$*$
2017	2.664	1.160	0.606	$*$	$*$
2018	2.512	1.082	0.549	$*$	$*$
2019	1.972	0.934	0.510	$*$	$*$
2020	1.852	0.862	0.534	$*$	$*$
2021	3.304	3.304	0.594	$*$	$*$
2022	2.450	1.473	0.717	$*$	$*$
2023	2.339	1.154	0.737	$*$	$*$
2024	2.526	1.735	0.670	$*$	$*$
2025	3.809	3.809	0.673	$*$	$*$

Notes: December value of the BSADF sequence for each year. $cv_t^{(0.95)}$ is the time-varying 95% critical value for $BSADF_{r_2}(r_0)$ at $r_2 = t/T$, obtained from 500 Monte Carlo replications of the random-walk null at each horizon r_2. "Expl.?" column: * denotes BSADF $> cv_t^{(0.95)}$. lnPR series: $y_t = \ln(P_t^H/R_t)$. FUC series: $y_t^{UC} = -\mathrm{FUC}_t^{(B)}$ (negative FUC, Type B). Sample begins 1989 because the minimum window of 44 observations requires data from 1986:02, making BSADF available from 1989:10 onward.

G Nominal and Real Model Comparison

This appendix presents the real-variable counterpart to the baseline nominal analysis developed in the main text, verifying the Fisher equivalence established in Section 2.6.

G.1 Nominal vs. Real Price and Rent Series

Figure G.1 compares the nominal and real (CPI-deflated) versions of the quality-adjusted condominium price and new-lease rent indices over 1986–2025. Two features stand out. First, the nominal–real gap is small during the low-inflation/deflation era (1996–2020), when CPI was approximately flat. The two series diverge visibly only during the 1986–1995 period (when cumulative CPI inflation reached approximately 25%) and the 2022–2025 period (when CPI inflation returned to 3% p.a.). Second, while the nominal price index in December 2025 exceeds the 1990 bubble peak for the first time, the real index remains approximately 11% below its 1990 level— a distinction that is critical for interpreting whether current prices constitute a "new high."

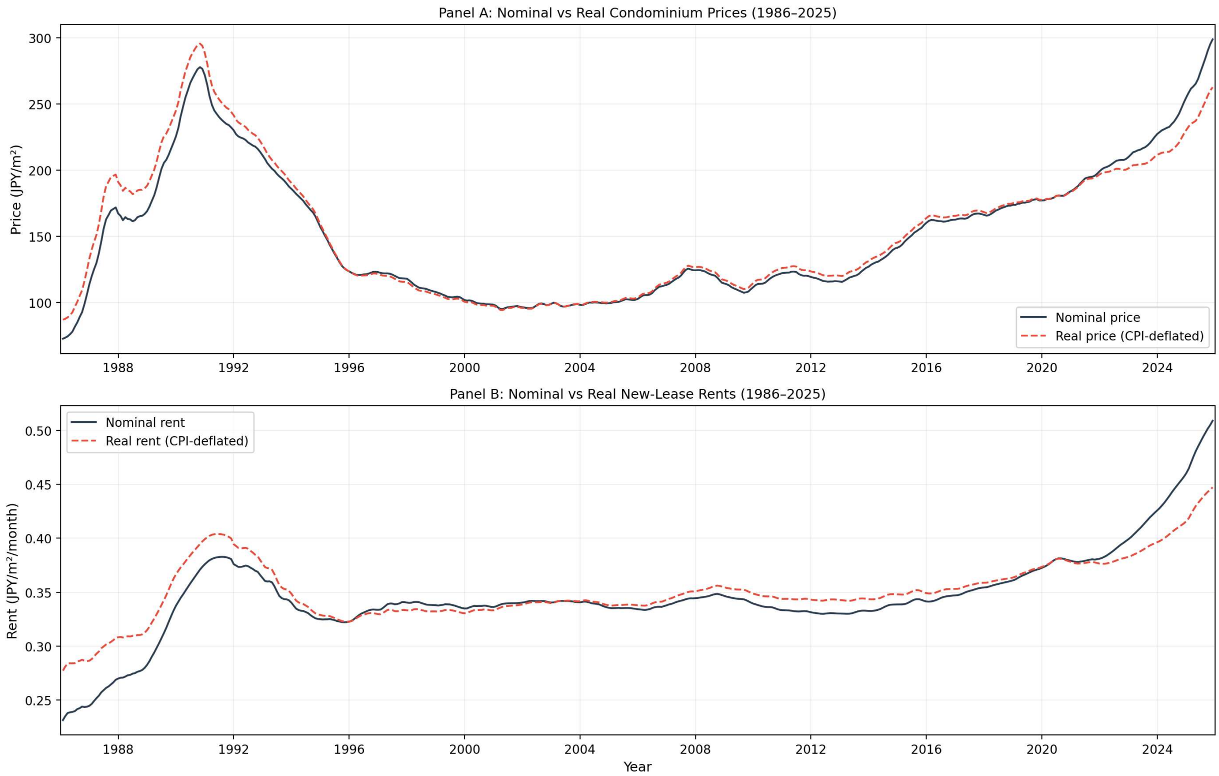

Figure G.1: Nominal vs. Real Condominium Prices and Rents (1986–2025).

Notes: Panel A: quality-adjusted condominium price per m^2. Panel B: quality-adjusted new-lease rent per m^2 (monthly). Solid lines: nominal values. Dashed lines: real values (CPI-deflated). CPI is the Tokyo All-Items index (excluding OOH), 2020 = 100. Both price and rent indices are from the rolling-window hedonic method (Appendix E).

G.2 FUC: Nominal, Real, and Baseline Comparison

Figure G.2 compares three versions of the Type B FUC rate:

- *Nominal (correct):* $\text{FUC}^{\text{nom}} = i^{\text{nom}} + c - \pi^{H,\text{nom}}$, where $\pi^{H,\text{nom}}$ is the HP trend of the log *nominal* price index.
- *Real (correct):* $\text{FUC}^{\text{real}} = r^{\text{real}} + c - \pi^{H,\text{real}}$, where $r^{\text{real}} = i^{\text{nom}} - \pi^{\text{CPI}}$ and $\pi^{H,\text{real}}$ is the HP trend of the log *real* price index.
- *Baseline (hybrid):* The specification used in the main text, which combines the nominal interest rate i^{nom} with the real appreciation trend $\pi^{H,\text{real}}$.

The Fisher equivalence (Section 2.6, equation (13)) predicts that $\text{FUC}^{\text{nom}} = \text{FUC}^{\text{real}}$ identically. Figure G.2 confirms this: the nominal and real FUC series (solid and dashed lines) are virtually indistinguishable throughout the 40-year sample. The baseline (hybrid) FUC (dotted gray line) tracks the correct versions closely during the zero-inflation period (1998–2020) but diverges during inflationary periods: it is approximately 2 pp *higher* during the bubble era (1988–1990) and 2.5 pp higher during the recent inflation (2022–2025). The implication is that the baseline model slightly understates the depth of the Necessary Regime during these episodes.

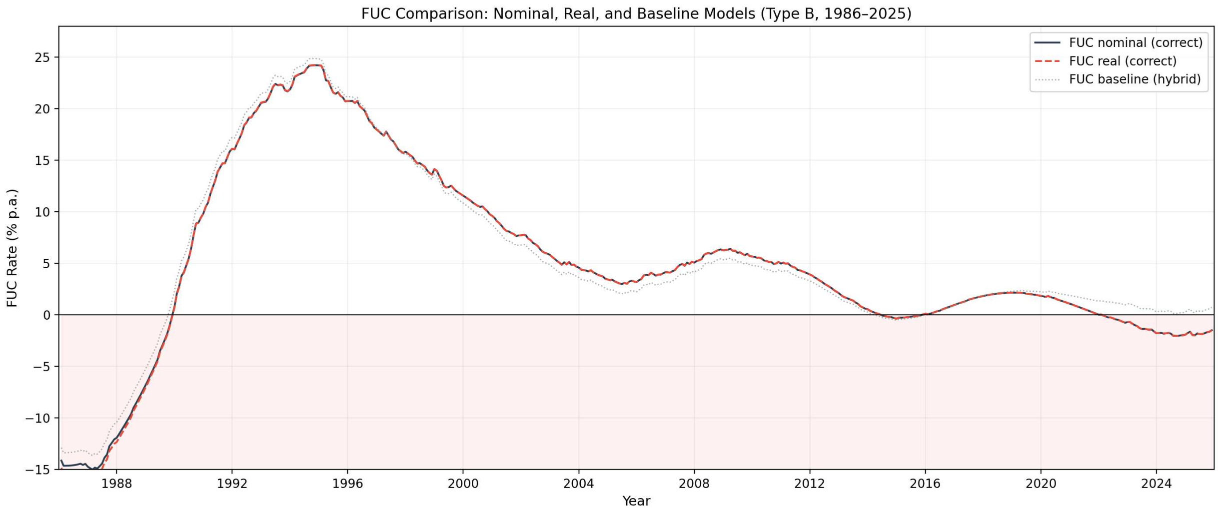

Figure G.2: FUC Rate: Nominal, Real, and Baseline (Hybrid) Models (Type B, 1986–2025).

Notes: The nominal FUC (solid) and real FUC (dashed) are theoretically identical by the Fisher equation and are empirically indistinguishable. The baseline (hybrid) FUC (dotted gray) diverges from the correct versions by the CPI inflation rate, which is negligible during deflation but substantial ($\approx$ 2–3 pp) during the 1980s bubble and the 2022–2025 inflation return. The red-shaded area indicates the Necessary Regime (FUC < 0).

G.3 Real VECM Estimation

Table G.1 reports VECM estimates for three specifications: the baseline (hybrid), the correctly constructed real model, and the nominal model. The system is $\mathbf{y}_t = (\text{FUC}_t \times 100,\ \ln R_t,\ \ln P_t^H)'$, with $p = 6$ and no deterministic terms.

Table G.1: VECM Estimation: Baseline, Real, and Nominal Specifications

Specification	Trace ($r = 0$)	$\hat{\beta}_R$	$\hat{\beta}_{\text{UC}}$	$\hat{\alpha}_P$
Baseline (hybrid)	27.9*	3.17	0.055	+0.002
Real (correct)	28.4*	3.05	0.079	+0.001
Nominal (correct)	19.5	—	—	—

Notes: * Marginal at 5% (critical value: 29.68). Johansen trace test with $p = 6$, Model 1 (no deterministic terms). The baseline and real specifications produce similar results: $\hat{\beta}_R \in [3.05, 3.17]$ and $\hat{\beta}_{\text{UC}} \in [0.055, 0.079]$. The nominal specification yields a lower trace statistic (19.5), reflecting the additional stochastic trend from CPI inflation in the nominal variables; cointegration is more difficult to detect when all variables share a common inflationary component. This supports the use of real variables (or the hybrid baseline) for the VECM estimation, while the FUC regime diagnostic is invariant to the nominal/real choice (Fisher equivalence).

Interpretation. The real VECM confirms the main findings: the long-run rent elasticity ($\hat{\beta}_R = 3.05$, within the range $[3.00, 3.19]$ documented in the lag-sensitivity analysis) and the user-cost semi-elasticity ($\hat{\beta}_{\text{UC}} = 0.079$, modestly higher than the baseline 0.055) are stable. The nominal VECM is less well-identified because the common CPI trend adds a shared stochastic component to all three variables, reducing the power of the Johansen trace test. This is the econometric rationale for working with real (CPI-deflated) price and rent series in the main text, while noting that the FUC regime diagnostic—which depends on the *sign* of FUC, not on the level of prices—is identical in nominal and real terms.