No. DP26-2

RCESR Discussion Paper Series

Too Many Houses, Too Few People: Demographic Optimism, Housing Oversupply, and Falling House Prices across Advanced Economies

March 2026

Yongheng Deng, University of Wisconsin-Madison

Tomoo Inoue, Seikei University

Kiyohiko Nishimura, The University of Tokyo

Chihiro Shimizu, Hitotsubashi University

RCESR

The Research Center for Economic and Social Risks
Institute of Economic Research
Hitotsubashi University

2-1 Naka, Kunitachi, Tokyo, 186-8603 JAPAN
http://risk.ier.hit-u.ac.jp/

Too Many Houses, Too Few People:

Demographic Optimism, Housing Oversupply, and Falling House Prices across Advanced Economies

Yongheng Deng[†] Tomoo Inoue[‡] Kiyohiko Nishimura[§] Chihiro Shimizu[¶]

March 2026

Abstract. In population bonus periods, optimism about future housing demand fuels rapid construction and self-reinforcing price appreciation. In population onus periods, pessimism—amplified by the systematic failure of governments to revise demographic projections downward in a timely manner—drives structural oversupply, rising vacancy rates, and prolonged price stagnation. We formalise this mechanism through a present-value model in which the demographic regime asymmetrically shifts expected rent growth and the discount rate, and test it using annual panel data for 16 advanced economies over 1975–2019. Pooled mean-group estimation shows that a rising old-age share significantly depresses long-run real house prices; the unanticipated ageing component (-8.826) substantially exceeds the government-projected component (-5.005). A rising youth share raises prices. Demographic structure further conditions monetary policy transmission: interest-rate cuts stimulate housing markets far more strongly in young than in ageing economies.

Keywords: demographic optimism; demographic pessimism; population bonus and onus; housing vacancy and oversupply; demographic forecast errors

JEL Codes: E31, J11, R21, R31

[†] Wisconsin School of Business, University of Wisconsin-Madison.

[‡] Faculty of Economics, Seikei University.

[§] School of Economics, The University of Tokyo; Former Deputy Governor, Bank of Japan.

[¶] School of Social Data Science, Hitotsubashi University. Corresponding author: 2-1 Naka, Kunitachi, Tokyo 186-8601, Japan. Tel.: +81-(0)42-580-9211. E-mail: `c.shimizu@r.hit-u.ac.jp`.

Acknowledgements. Shimizu gratefully acknowledges financial support from the Japan Society for the Promotion of Science (JSPS KAKENHI Grant 24H00012). We thank Iichiro Uesugi, Tomohiro Hirano, Robert Hill, Marc Francke, and Miriam Steurer for helpful comments and suggestions; Miriam Steurer also provided valuable comments at a conference presentation of an earlier version. The views expressed are those of the authors and do not represent the views of any institution with which they are affiliated.

1 Introduction

Between 1965 and 1990, Japan built more dwellings per capita than any other advanced economy. Official projections pointed to continuing population growth and sustained household formation, and both public and private actors responded accordingly: housing starts frequently exceeded 1.5 million per year, new towns were planned on urban fringes, and mortgage lending expanded to match the demand. By 2018, the same country counted 8.5 million vacant dwellings— 13.6 percent of its entire housing stock and the highest vacancy rate in the developed world. In many regional municipalities, one in three properties now stands empty. Japan's experience—decades of optimistic overbuilding followed by structural excess supply and persistent price stagnation—is the clearest illustration of a pattern now being replicated across the ageing societies of East Asia and Europe.

Germany provides a parallel European case. Following reunification, the federal government and Länder projected rapid population convergence between east and west, and invested accordingly: more than two million new dwellings were built in eastern Germany during the 1990s. The actual population trajectory fell far short of projections, leaving vacancy rates in cities such as Leipzig, Magdeburg, and Halle above 20 percent by the early 2000s. The federal response—the "Stadtumbau Ost" programme, which subsidised the demolition of nearly 350,000 dwellings between 2002 and 2016—was without precedent in Western European housing administration and illustrates the long-lived fiscal consequences of demographic forecasting errors. Italy has an estimated seven million vacant dwellings—roughly 22 percent of its housing stock—concentrated in the south and inland areas where depopulation has proceeded faster than any official forecast anticipated; as in Germany, the gap between projected and realised population closely tracks the regional distribution of vacancy. South Korea is currently at an earlier stage of the same trajectory: official fertility projections have been revised sharply downward at every five-year review since the early 2000s—the total fertility rate reached 0.72 in 2023, the lowest ever recorded in any country—yet the housing supply system continues to operate on earlier, more optimistic targets, with housing starts in 2022–23 far exceeding the pace consistent with flat or declining household counts.

This paper argues that the central mechanism linking demographic change to housing market outcomes is the systematic *optimism bias* embedded in government population forecasts. Official projections throughout the period of fertility decline have persistently overestimated future populations, anchoring on the high-fertility norms of earlier decades and revising downward only gradually (as shown by successive UN WPP vintages for Japan). Developers, municipalities, and mortgage lenders who calibrate housing investment to these projections therefore build more than realised household demand warrants. The resulting excess supply—manifest as rising vacancy rates and stagnant rents—depresses house prices below the level that even the already-pessimistic realised demographic path would imply.

We frame this argument within the *population bonus* and *population onus* distinction of Nishimura (2016). During the bonus period, when the working-age share is rising and young cohorts are large, demographic optimism is largely self-fulfilling: housing demand grows, supply expands, and prices appreciate. The transition to the onus period, when old-age shares rise and household formation slows, is where the forecast-error mechanism bites most severely. Builders and planners anchored to optimistic projections continue constructing at bonus-era rates, generating a structural mismatch between housing stock and demand. Because housing is durable and location-specific, this excess supply cannot be quickly absorbed: vacant properties accumulate, rents fall, and house prices stagnate for years or decades after the demographic turning point.

We test this argument using annual panel data for 16 advanced economies over 1975–2019, estimating pooled mean-group (PMG) error-correction models following Pesaran et al. (1999). By matching each country–year observation to the relevant vintage of United Nations World Population Prospects, we construct a direct measure of the five-year-ahead forecast error for the old-age and youth population shares.

Our results yield three principal findings. First, age structure has large, statistically robust effects on long-run real house prices: a one-unit rise in the youth-to-working-age ratio raises prices by 2.895 (significant at 1%), while the same rise in the old-age-to-working-age ratio reduces prices by 2.549. Second, the demographic state governs the sensitivity of house prices to interest rate changes: a one-percentage-point cut raises prices by approximately 2 percent in the youngest economies of the sample (Ireland, New Zealand) but has near-zero impact in the oldest (Germany, Italy, Sweden). Third, and most importantly, decomposing the old-age share into its projected and forecast-error components reveals that unanticipated ageing carries a coefficient of -8.826, compared with -5.005 for the projected component—a difference of 76 percent that represents the additional price depression attributable to the excess housing stock induced by forecast optimism.

These findings speak directly to four debates within the housing studies literature. First, they provide a cross-national mechanism for the vacancy and abandonment trends documented in shrinking-city research: the accumulation of unwanted dwellings is not merely a local planning failure but a predictable consequence of national-level demographic forecasting bias (Couch and Cocks, 2013; Gyourko et al., 2013). Second, they contribute to a growing literature—including the recent historical evidence in Francke and Korevaar (2025) that cohort-size shocks predict house prices with a 25–30-year lag—confirming that demographic fundamentals are a primary driver of long-run housing values, while extending this literature by identifying the additional price depression caused by forecast errors. They also resolve the long-standing disagreement between Mankiw and Weil (1989) and the rational-expectations critics (Hamilton, 1991; Hendershott, 1991): both sides are partially correct. Demographics do affect prices, and supply does adjust—but only along the projected demographic path; the *unanticipated* portion of ageing generates additional excess supply that rational supply adjustment cannot prevent. Third, they identify a practical early-warning indicator for housing policy: tracking the gap between official population projections and emerging demographic data provides a direct signal of impending housing market imbalances. Fourth, the finding that monetary policy effectiveness depends on demographic structure has implications for housing-oriented fiscal and macroprudential policy in ageing societies. Taken together, these results suggest that the standard toolkit of housing policy—zoning reform, planning incentives, social housing investment—needs to be supplemented by systematic monitoring of the gap between official demographic projections and emerging demographic reality, and that the size of this gap should trigger automatic reviews of housing supply targets.

The remainder of the paper is structured as follows. Section 2 reviews the relevant literature. Section 3 presents the conceptual framework. Section 4 describes the data and estimation methodology. Section 5 reports the empirical results. Section 6 discusses the implications for housing markets, vacancy, and policy. Section 7 concludes.

2 Related Literature

Four bodies of work inform our analysis: research on housing vacancy and shrinking cities; the demography–house price nexus; the supply-elasticity debate; and the role of demographic expectations.

Housing vacancy, shrinking cities, and demographic decline. A growing body of work within the housing studies literature examines the supply-side consequences of population ageing and decline at the local and regional scale. Couch and Cocks (2013) document how housing vacancy and urban shrinkage interact in the United Kingdom, showing that population decline generates persistent housing surpluses that depress local markets for decades and that conventional market mechanisms are insufficient to clear excess stock without active policy intervention—a finding directly relevant to the irreversibility mechanism central to our argument. Lozano Alcántara and Vogel (2023) examine the intersection of demographic ageing and housing costs in Germany, demonstrating that rising housing expenditure has substantially amplified income poverty among the elderly—evidence that the demographic–housing market nexus has concrete distributional consequences beyond aggregate price effects. Tamai et al. (2017) project municipality-level property prices in Japan through 2040 and find severe declines in the areas with the highest projected dependency ratios, precisely those where earlier overbuilding was most pronounced. The German literature on *Stadtumbau* (urban restructuring) documents how demographic optimism in the post-reunification decade produced a durable supply overhang whose effects persisted for twenty years despite active demolition programmes (Wiechmann and Pallagst, 2012). For southern Europe, Martínez-Gutiérrez et al. (2016) trace the interaction between demographic decline, construction cycles, and spatial price divergence in Spain, noting that official projections systematically underestimated the pace of depopulation in interior regions. In East Asia, emerging evidence from South Korea and China points to the early stages of a similar dynamic: housing completions in Chinese third- and fourth-tier cities already substantially exceed net household formation in those areas, consistent with our theoretical prediction that forecasting bias generates oversupply ahead of the demographic turning point (Glaeser et al., 2017). Gyourko et al. (2013) provide the key theoretical insight for our mechanism: in shrinking cities, housing markets can remain persistently illiquid even at very low prices because the stock cannot be costlessly destroyed or relocated. This irreversibility is what transforms a temporary demand shortfall into a structural equilibrium of chronically depressed rents and prices.

Demographics and house prices. Mankiw and Weil (1989) established the foundational insight that housing demand is strongly age-dependent, and projected a 47 percent real decline in US house prices as the baby-boom cohort aged past its peak home-buying years. Although the prediction proved too stark, the underlying mechanism has been confirmed repeatedly. Takáts (2012) and Nishimura and Takáts (2012) use an overlapping-generations model to show that working-age population growth raises real property prices in the long run. Empirical studies by Saita et al. (2016) for Japan and the United States, and by Tamai et al. (2017) for Japanese municipalities, confirm that population composition significantly affects house prices at both national and regional scales. Knoll et al. (2017) document that virtually all long-run global house price appreciation since 1870 reflects land price appreciation, consistent with inelastic urban land supply amplifying demographic demand shocks. Most recently, Francke and Korevaar (2025) exploit several centuries of historical data for Amsterdam and Paris to show that cohort-size shocks to housing demand—with a lag of 25–30 years as birth cohorts enter the home-buying age—generate substantial real house price appreciation, and that these effects are concentrated in prices rather than rents, reflecting the market frictions and segmentation that prevent immediate supply responses. Their evidence that *realised* demographic shocks drive prices over multi-decade horizons complements our focus on the role of *forecast errors*:

whereas Francke and Korevaar establish that demographic fundamentals systematically predict future prices, we show that the gap between government projections and realised demographics generates an additional layer of price distortion through the oversupply mechanism. We extend this literature by quantifying the price effect arising from the *discrepancy* between realised and projected demographic change across a contemporary cross-country panel.

Supply elasticity and the rational-expectations debate. The principal objection to demographic explanations of house price fluctuations is that demographic change is slow and predictable: Hamilton (1991) and Hendershott (1991) argued that rational developers would adjust supply ahead of any projected demand decline, preventing price collapse. Engelhardt and Poterba (1991) provide supporting evidence from Canada, and Ohtake and Shintani (1996) and Shimizu and Watanabe (2010) find similar results for Japan. Our contribution is to show that this argument holds only for the *projected* component of demographic change—the portion that could in principle be anticipated. The unanticipated component, driven by forecasting bias, generates additional excess supply that supply-side rationality cannot prevent. In this sense, our findings are compatible with, but extend beyond, the rational-expectations critique.

Demographic expectations and housing market misalignment. Nishimura (2016) argues that population bonuses and demographic pessimism have historically conditioned housing booms and prolonged stagnation, and that central banks should account for demographic structure when calibrating monetary policy. The broader role of expectation errors in driving asset-price cycles has been emphasised by Reinhart and Rogoff (2009). Our paper identifies a specific, institutionally grounded source of optimism bias: the anchoring behaviour of national statistical agencies and the United Nations World Population Prospects, which have repeatedly projected fertility reversals that did not materialise. The literature on forecast bias in official projections (Keyfitz, 1981; Keilman, 1998) documents that long-run demographic forecasts systematically overestimate fertility during periods of sustained decline. Our contribution is to link this forecasting failure to housing market outcomes by constructing a direct panel measure of demographic forecast errors from successive UN

WPP vintages and quantifying its independent contribution to long-run real house price dynamics across 16 countries—to our knowledge the first study to do so at the international level.

Panel cointegration of house prices. Our econometric framework builds on the panel ARDL/PMG approach of Pesaran et al. (1999), which allows for heterogeneous short-run dynamics across countries while imposing homogeneous long-run equilibrium coefficients. Badarinza et al. (2021) apply a related panel cointegration framework to international housing and capital markets, providing a methodological precedent for our cross-country estimation strategy.

3 Conceptual Framework

This section develops the argument in three steps. First, we introduce the present-value relation for housing assets and explain how demographic structure enters through both the rental income and the discount rate. Second, we formalise the population bonus–onus distinction and show how optimism in the bonus period and pessimism amplified by forecast errors in the onus period generate characteristic housing supply mismatches. Third, we derive the three empirical specifications that translate the conceptual argument into testable hypotheses.

3.1 The Present-Value Relation for Housing Assets

The long-run equilibrium *real* price of a residential property reflects the discounted present value of the real stream of housing services it delivers. All quantities in this paper are expressed in real (inflation-adjusted) terms: the dependent variable is the log real house price index ($\log P_t^{rppi} - \log P_t^{cpi}$), real income serves as the proxy for housing-service rents, and the interest rate is the real long-term rate. Under standard no-arbitrage conditions, the long-run equilibrium real house price is given by:

$$Q_t \equiv \frac{P_t^{rppi}}{P_t^{cpi}} \approx \frac{R_t}{r_t^* - g_{R,t}}, \qquad r_t^* > g_{R,t}, \tag{1}$$

where R_t is the real flow of housing-service rents, r_t^* is the long-run *real* discount rate, and $g_{R,t}$ is the expected long-run real growth rate of rents. Because all three objects are measured in real terms, inflation affects neither the left-hand side nor the right-hand side of (1): price-level changes cancel in the ratio Q_t and drop out of the real interest rate r_t^*. Demographic structure enters through both terms. On the income side, real rents depend on real working-age per-capita income scaled by an effective capacity ratio κ_t: $R_t \approx \kappa_t Y_t / \text{pop}_t^W$. A younger population raises housing demand relative to supply, increasing κ_t and rents; an ageing population lowers κ_t. On the discount side, a youthful population raises expected real rent growth $g_{R,t}$ and the natural *real* interest rate r_t^*, but the net discount term $r_t^* - g_{R,t}$ is compressed (amplifying prices) when demand growth is strong. In ageing economies the fall in expected real demand growth dominates, widening the net discount term and depressing real house prices. The full structural derivation is provided in Online Appendix A.3.

3.2 From Demographic Bonus to Onus: Optimism, Oversupply, and Vacancy

The PVR in equation (1) helps explain not only the *level* of house prices but also the pronounced *cycles* associated with demographic transitions. We formalise the population bonus–onus framework (Nishimura, 2016) and connect it directly to housing supply dynamics.

Bonus period: optimism and the seeds of oversupply. During the bonus period—when the working-age share is rising and youth cohorts are large—housing demand grows visibly and persistently. Developers, municipalities, and mortgage lenders extrapolate this trend and calibrate investment to official demographic projections. At this stage the projections tend to be approximately correct, or even conservative. Supply expands rapidly to meet genuine demand, and the compression of the net discount term $r_t^* - g_{R,t}$ supports high prices. The large, geographically dispersed housing stock built during this period is—critically—durable. Once constructed, it cannot easily be removed if future demand falls short of expectations.

Onus period: amplified pessimism, forecast errors, and accumulating vacancies. As the old-age share rises and household formation slows, demand for new housing weakens. However, the adjustment of housing supply is slow and incomplete for two reinforcing reasons. First, the irreversibility of the existing stock means that vacant properties accumulate rather than disappear. Second—and this is the key mechanism of this paper—government demographic projections anchor on previous, more optimistic assumptions and revise downward only gradually (as documented by successive UN WPP vintages for Japan). Developers and planners acting on these projections therefore continue building at rates calibrated to a population that is not materialising. The resulting excess stock depresses rents, reduces R_t in equation (1), and lowers expected rent growth $g_{R,t}$, causing house prices to fall further than the projected demographic path alone would imply.

This mechanism predicts that the *unanticipated* component of ageing— the realised deviation from official five-year-ahead projections—should depress prices more than the *projected* component, because only the unanticipated portion triggers additional excess supply that could not have been avoided even by a developer acting rationally on official forecasts. Formally, letting $\hat{n}^O_{jt}$ denote the five-year-ahead projected old-age share and $\varepsilon^O_{jt} = n^O_{jt} - \hat{n}^O_{jt}$ the forecast error, the oversupply mechanism implies $|\beta_7| > |\beta_5|$ in equation (4) below.

Asymmetry and irreversibility. The mechanism is asymmetric. During the bonus period, supply can expand rapidly in response to demand. During the onus period, contraction is slow: demolition is costly, property rights are fragmented, and owners resist writing down asset values. Moreover, as falling rents erode the collateral value of residential property, credit supply to the housing sector contracts, reinforcing the price decline through a mechanism analogous to the credit-cycle amplification documented by Kiyotaki and Moore (1997). This asymmetry amplifies the duration and severity of price stagnation in the onus period relative to the appreciation in the bonus period—a pattern consistent with the long periods of vacancy accumulation observed in Japan, Germany, and Italy.

3.3 Empirical Specifications

The framework above maps into three empirical long-run specifications, each estimated as a pooled mean-group (PMG) error-correction model.

Model 1 (Baseline age-structure effects). The baseline specification tests whether youth and old-age dependency ratios have independent long-run effects on real house prices:

$$\log\left(\frac{P_{jt}^{rppi}}{P_{jt}^{cpi}}\right) = \beta_{0j} + \beta_1 \log\left(\frac{Y_{jt}}{\mathrm{pop}_{jt}^{W}}\right) + \beta_2\, r_{jt} + \beta_3 \log \mathrm{pop}_{jt}^{total} + \beta_4(n_{jt}^{Y} - n_{jt}^{W}) + \beta_5(n_{jt}^{O} - n_{jt}^{W}) + \varepsilon_{jt}, \tag{2}$$

where the dependent variable is the log real house price index (`lrrppi` $\equiv \log P_{jt}^{rppi} - \log P_{jt}^{cpi}$), r_{jt} is the real long-term interest rate (nominal rate minus CPI inflation), $Y_{jt}/\mathrm{pop}_{jt}^{W}$ is real GDP per working-age person, and demographic shares enter as differences from the working-age share, which imposes the constraint $\beta_Y + \beta_W + \beta_O = 0$ implied by the identity $n_{jt}^{Y} + n_{jt}^{W} + n_{jt}^{O} = 1$, thereby avoiding collinearity. The disturbance ε_{jt} represents the deviation from long-run equilibrium; equations (2)–(4) specify the long-run equilibrium relationship embedded in the error-correction model (5). The theory predicts $\beta_4 > 0$ and $\beta_5 < 0$.

Model 2 (Interest rate × age-structure interactions). To test whether monetary policy transmission to housing markets depends on demographic regime:

$$\log\left(\frac{P_{jt}^{rppi}}{P_{jt}^{cpi}}\right) = \beta_{0j} + \beta_1 \log\left(\frac{Y_{jt}}{\mathrm{pop}_{jt}^{W}}\right) + \beta_2\, r_{jt} + \beta_3 \log \mathrm{pop}_{jt}^{total} + \beta_4(n_{jt}^{Y} - n_{jt}^{W}) + \beta_5(n_{jt}^{O} - n_{jt}^{W}) + \beta_6\, r_{jt}\cdot(n_{jt}^{Y} - n_{jt}^{W}) + \beta_7\, r_{jt}\cdot(n_{jt}^{O} - n_{jt}^{W}) + \varepsilon_{jt}. \tag{3}$$

The marginal effect of the real interest rate is $\partial \log(P/P^{cpi})/\partial r_{jt} = \beta_2 + \beta_6(n_{jt}^{Y} - n_{jt}^{W}) + \beta_7(n_{jt}^{O} - n_{jt}^{W})$. The theory predicts $\beta_6 < 0$ (youth amplifies the real interest-rate channel) and $\beta_7 > 0$ (ageing attenuates it), reflecting the state-dependence of the net discount term $r^*_t - g_{R,t}$.

Model 3 (Forecast-error decomposition). The central test decomposes the old-age and youth shares into government-projected and forecast-error components:

$$\log\left(\frac{P^{rppi}_{jt}}{P^{cpi}_{jt}}\right) = \beta_{0j} + \beta_1 \log\left(\frac{Y_{jt}}{\text{pop}^W_{jt}}\right) + \beta_2\, r_{jt} + \beta_3 \log \text{pop}^{total}_{jt} + \beta_4(\hat{n}^Y_{jt} - n^W_{jt}) + \beta_5(\hat{n}^O_{jt} - n^W_{jt}) + \beta_6\, \varepsilon^Y_{jt} + \beta_7\, \varepsilon^O_{jt} + \varepsilon_{jt}, \tag{4}$$

where $\hat{n}^g_{jt}$ is the five-year-ahead UN projection, $\varepsilon^g_{jt} = n^g_{jt} - \hat{n}^g_{jt}$ is the realised forecast error, and r_{jt} is the real interest rate. The oversupply hypothesis predicts $\beta_5 < 0$, $\beta_7 < 0$, and $|\beta_7| > |\beta_5|$: unanticipated ageing depresses real house prices more than anticipated ageing.

4 Data

4.1 Sample and Variables

We use annual panel data for 16 advanced economies over 1975–2019: Australia, Belgium, Canada, Switzerland, Germany, Denmark, France, the United Kingdom, Ireland, Italy, Japan, the Netherlands, Norway, New Zealand, Sweden, and the United States.

Residential property prices. The dependent variable is the log real residential property price index (`lrrppi`), constructed by deflating the BIS Residential Property Price Index (nominal, local currency) by the CPI. Because the BIS index is published at quarterly frequency, we convert it to annual frequency by taking a simple average of the four quarterly observations.

Income. Log real GDP per working-age person (`ly2wpop`) is our proxy for the long-run rental service flow. Real GDP is taken from the World Bank's World Development Indicators (WDI); working-age population (ages 15–64) from the United Nations World Population Prospects (WPP).

Interest rates. The real long-term interest rate (`rint`) is the key interest rate variable in the empirical models, consistent with the real present-value framework in Section 3.1. It is constructed as the nominal long-term interest rate minus the contemporaneous CPI inflation rate (static expectations). The nominal rate is taken from OECD Statistics, with gaps filled from the IMF's International Financial Statistics (IFS); all rates are converted to continuously compounded form before deflation.

Population age structure. Age-group population shares—youth (`ny`, ages 0–14), working-age (`nw`, ages 15–64), and old-age (`no`, ages $\geq$ 65)—are computed from the UN WPP country-level data. The dependency ratios entering the regressions are `ny_nw` $\equiv n^Y_{jt}/n^W_{jt}$ and `no_nw` $\equiv n^O_{jt}/n^W_{jt}$.

Demographic forecast errors. We construct the projected demographic shares $\hat{n}^k_{jt}$ $(k = Y, O)$ using the sequence of official UN World Population Prospects revisions published between 1982 and 2012. For each country and year, we take the five-year-ahead projection published by the most recent available revision and interpolate between revision years using linear interpolation. The full construction procedure is described in Online Appendix A.2. The forecast error components are defined as the difference between realised and projected shares: `ny_py` $\equiv n^Y_{jt} - \hat{n}^Y_{jt}$ and `no_po` $\equiv n^O_{jt} - \hat{n}^O_{jt}$. Because UN country-level projections are available only from 1982, the forecast-error model (Model 3) is estimated on the shorter sample beginning in 1986.

4.2 Descriptive Statistics

The log real house price index (`lrrppi`) has a mean of 4.303 and a standard deviation of 0.410, reflecting substantial variation both across countries and over time. Log real GDP per working-age person (`ly2wpop`) has a mean of 6.971 and a standard deviation of 1.350, capturing the wide dispersion in income levels across our 16-country sample. The real interest rate (`rint`) averages 0.025, ranging from negative values to a maximum of approximately 0.098, documenting the sharp swings in monetary conditions over our four-decade sample.

The population structure variables display the secular demographic transition in our sample. The average youth share (`ny`) is 0.203 and declining; the average working-age share (`nw`) is 0.655 and roughly stable; the average old-age share (`no`) is 0.142 and rising. The smaller standard deviation of nw (0.026) compared with ny (0.052) and no (0.039) reflects the well-known smoothness of the working-age share relative to the youth and old-age shares.

The forecast-error variables (ny_py and no_po) are available for 561 observations. Their sample means are close to zero, confirming that official forecasts are approximately unbiased on average over short horizons. However, the non-trivial ranges of both variables document that forecast errors of economically meaningful magnitude arise at specific country–year combinations, consistent with the pattern of accumulated optimism documented for Japan through successive UN WPP vintage comparisons.

4.3 Econometric Methodology

Because the main series are integrated of order one (see Online Appendix A.1 for unit root tests), we estimate the models using the panel ARDL(2,1) framework of Pesaran et al. (1999), following the application to international housing markets in Badarinza et al. (2021). The error-correction representation is:

$$\Delta y_{it} = \theta_i(y_{i,t-1} - \beta' \mathbf{X}_{i,t-1}) + \lambda_i \Delta y_{i,t-1} + \delta_i' \Delta \mathbf{X}_{i,t-1} + \mu_i + \varepsilon_{it}, \tag{5}$$

where $\theta_i < 0$ is the country-specific speed of adjustment to long-run equilibrium, and β is the vector of long-run coefficients. We estimate the *pooled mean group* (PMG) version of equation (5), which imposes homogeneity of the long-run coefficients β across countries while allowing country-specific short-run dynamics $(\theta_i, \lambda_i, \delta_i)$. This restriction reflects the theoretical prediction from the no-arbitrage condition in equation (1): the fundamental determinants of house prices should be common across well-integrated capital markets.

The validity of the homogeneity restriction is tested via a Hausman test comparing PMG and mean-group (MG) estimates. In all three models, the Hausman test fails to reject homogeneity ($p > 0.35$ in all cases), supporting the PMG specification.

5 Empirical Results

5.1 Model 1: Age Structure and the Long-Run House Price Equilibrium

Table 1 reports PMG estimates of equation (2) for four alternative start years (1975, 1976, 1977, and 1978), all ending in 2019. The robustness of the estimates across start years is one of our key identification checks.

The error-correction coefficient θ_i is negative and significant at the 1% level in all specifications, averaging -0.180 in column (1). This implies that roughly 18 percent of any deviation from long-run equilibrium is corrected within a single year, consistent with the well-known sluggishness of housing markets.

Among the long-run coefficients, log real GDP per working-age person (`ly2wpop`) carries a coefficient of 2.040 (s.e. 0.219), positive and significant at the 1% level across all columns. This confirms that income growth is a powerful long-run driver of house prices, consistent with the PVR. Log total population (`ltpop`) is also significantly positive (0.792, s.e. 0.181), reflecting the independent role of market size in supporting demand. The real interest rate (`rint`) enters with the expected negative sign (-0.301) but is not statistically significant, a finding that motivates the interaction-term specification in Model 2.

The central finding of Model 1 is the symmetric, statistically significant role of the youth and old-age dependency ratios. *Youth share:* a one-unit increase in `ny_nw` (the ratio of youth to working-age population) raises the log real house price index by **2.895** (s.e. 0.770), significant at the 1% level. *Old-age share:* a one-unit increase in `no_nw` (the old-age-to-working-age ratio) reduces the log real house price index by **2.549** (s.e. 0.650), significant at the 1% level.

These estimates are highly robust: the old-age coefficient remains between -1.919 and -2.549 across all four start years, and the youth coefficient ranges from 1.096 to 2.895. The pattern is precisely what the demand-driven PVR model predicts: a more youthful population generates more household formation and higher housing demand, while a more elderly population depresses demand for new housing.

The Hausman test p-value of 0.427 in column (1) confirms that the PMG homogeneity restriction on long-run coefficients cannot be rejected, validating our pooling assumption.

Table 1: Model 1: PMG Estimates of the Long-Run House Price Equilibrium, Baseline and Robustness Across Start Years

	(1) (1975–2019)		(2) (1976–2019)		(3) (1977–2019)		(4) (1978–2019)	
ec								
ly2wpop	2.040***	(0.129)	1.909***	(0.125)	1.937***	(0.124)	1.981***	(0.112)
rint	-0.301	(0.587)	-0.200	(0.606)	0.087	(0.598)	0.065	(0.570)
ltpop	0.792***	(0.262)	0.859***	(0.266)	0.777***	(0.272)	0.623**	(0.265)
ny_nw	2.895***	(0.770)	2.388***	(0.813)	1.879**	(0.825)	1.096	(0.822)
no_nw	-2.549***	(0.650)	-2.427***	(0.649)	-2.248***	(0.639)	-1.919***	(0.608)
SR								
ec	-0.180***	(0.039)	-0.182***	(0.039)	-0.193***	(0.040)	-0.202***	(0.042)
LD.lrrppi	0.486***	(0.041)	0.479***	(0.038)	0.482***	(0.037)	0.464***	(0.036)
D.ly2wpop	0.628***	(0.177)	0.657***	(0.173)	0.591***	(0.173)	0.555***	(0.159)
D.rint	0.404**	(0.198)	0.376**	(0.180)	0.288	(0.199)	0.282	(0.217)
D.ltpop	2.142*	(1.250)	2.377**	(1.165)	1.858	(1.301)	1.488	(1.489)
D.ny_nw	-7.555***	(1.930)	-7.418***	(2.080)	-7.506***	(2.460)	-7.299***	(2.715)
D.no_nw	7.947***	(2.650)	8.245***	(2.711)	7.976**	(3.225)	7.476**	(3.431)
_cons	-4.256***	(0.931)	-4.392***	(0.954)	-4.458***	(0.957)	-4.267***	(0.951)
Observations	688		672		656		640	
Countries	16		16		16		16	
Years	43		42		41		40	
Hausman MG vs PMG chi2	4.911		3.589		4.040		4.587	
Hausman MG vs PMG p-value	0.427		0.610		0.544		0.468	
Spec ID	S0004		S0005		S0006		S0007	

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. EC = long-run (error-correction) equation; SR = short-run equation. Dependent variable: log real residential property price index (`lrrppi` $= \log P^{rppi} - \log P^{cpi}$). Interest rate: real long-term rate (`rint` $=$ nominal rate $-$ CPI inflation). `ny_nw` = youth-to-working-age ratio; `no_nw` = old-age-to-working-age ratio. Lag order ARDL(2,1) selected by Schwarz criterion.

5.2 Model 2: Age Structure and the Interest Rate Channel

Table 2 reports PMG estimates of the interaction-term model (3). The core result is that the demographic composition of the population significantly moderates the sensitivity of house prices to interest rate changes.

In the long-run equation, the main effect of the real interest rate becomes statistically indistinguishable from zero once the interaction terms are included, while the interaction term between the interest rate and the old-age dependency ratio (`r_no_nw`) is significantly positive (29.367, s.e.

10.333), and the corresponding youth interaction term (`r_ny_nw`) is negative (-12.587), consistent with the theory that youth amplifies and ageing attenuates the interest-rate channel. The positive coefficient on `r_no_nw` means that in countries with a higher old-age share, the *sensitivity* of house prices to interest rate changes is attenuated: ageing populations are less responsive to monetary stimulus.

To quantify this heterogeneity, Table 3 and Figure 1 report the marginal effect of the real interest rate on house prices, evaluated at each country's sample-mean age composition:

$$\frac{\partial \log(\hat{P}_j^{rppi}/P_j^{cpi})}{\partial r_j} = \beta_2 + \beta_6\,(\bar{n}_j^Y - \bar{n}_j^W) + \beta_7\,(\bar{n}_j^O - \bar{n}_j^W), \tag{6}$$

A negative marginal effect means that a one-percentage-point rise in the real interest rate reduces house prices; by symmetry, a one-percentage-point fall in the real interest rate raises them by the same amount. The cross-country pattern is striking:

For the **full-sample average** (ALL), the marginal effect is -0.740 (95% CI: $[-1.404, -0.076]$), which is significantly negative. A one-percentage-point reduction in the real interest rate is therefore associated with an average house price increase of 0.74%.

Countries with a **relatively young population** during the sample period show much larger marginal effects: Ireland (-2.234, 95% CI significantly negative), New Zealand (-1.969), Australia (-1.929), Canada (-1.850), and the United States (-1.641). In these countries, a one-percentage-point interest rate cut is associated with a house price increase of approximately 2 percent.

By contrast, countries with older average populations—Germany, Italy, Sweden, Denmark—show marginal effects close to zero or even positive (Germany: $+0.245$; Italy: $+0.299$; Sweden: $+0.519$). The positive estimates for these countries, while not statistically significant at conventional levels, suggest that the interest rate transmission channel for house prices is substantially diminished, or even reversed, in the most aged economies of our sample.

These results are consistent with Nishimura (2016), who argues that demographic structure is a key determinant of how monetary policy transmits to asset prices. They also have a direct implication for policy: a uniform interest rate cut will generate highly unequal house price responses across countries depending on their demographic structure, a concern particularly relevant for the design of monetary policy in currency unions such as the euro area.

Table 2: Model 2: PMG Estimates with Interest Rate × Age-Structure Interactions

	(1) (1975–2019)		(2) (1976–2019)		(3) (1977–2019)		(4) (1978–2019)	
ec								
ly2wpop	2.127***	(0.126)	2.170***	(0.126)	1.934***	(0.135)	1.925***	(0.128)
rint	8.352	(6.958)	-10.801	(7.766)	-2.908	(7.683)	0.405	(7.321)
ltpop	0.478	(0.291)	0.633**	(0.294)	0.729**	(0.293)	0.566**	(0.278)
ny_nw	3.097***	(0.807)	5.164***	(0.779)	3.121***	(0.885)	3.176***	(0.898)
no_nw	-2.632***	(0.676)	-3.161***	(0.573)	-2.670***	(0.613)	-2.355***	(0.588)
r_ny_nw	-12.587	(12.914)	-58.975***	(14.408)	-23.842*	(13.995)	-25.021*	(14.640)
r_no_nw	29.367***	(10.333)	34.084***	(11.042)	17.782*	(10.726)	26.591***	(10.148)
SR								
ec	-0.174***	(0.039)	-0.163***	(0.036)	-0.180***	(0.035)	-0.186***	(0.037)
LD.lrrppi	0.467***	(0.042)	0.457***	(0.038)	0.468***	(0.037)	0.445***	(0.037)
D.ly2wpop	0.698***	(0.181)	0.749***	(0.167)	0.691***	(0.164)	0.713***	(0.144)
D.rint	1.429	(6.749)	0.181	(6.295)	2.770	(6.692)	1.151	(5.997)
D.ltpop	2.633**	(1.171)	2.677**	(1.047)	2.519**	(1.147)	2.683**	(1.163)
D.ny_nw	-6.602***	(2.012)	-6.056***	(1.815)	-5.976***	(1.958)	-5.674***	(1.953)
D.no_nw	6.447**	(2.866)	6.025***	(2.318)	5.899**	(2.557)	5.265**	(2.506)
D.r_ny_nw	-4.526	(7.344)	-0.493	(7.005)	-3.144	(8.619)	-8.483	(7.868)
D.r_no_nw	4.751	(10.256)	-1.468	(8.657)	5.622	(9.479)	7.421	(9.440)
_cons	-3.319***	(0.773)	-3.475***	(0.769)	-3.946***	(0.795)	-3.540***	(0.740)
Observations	688		672		656		640	
Countries	16		16		16		16	
Years	43		42		41		40	
Hausman MG vs PMG chi2	7.713		11.061		7.901		5.304	
Hausman MG vs PMG p-value	0.359		0.136		0.341		0.623	
Spec ID	S0040		S0041		S0042		S0043	

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. `r_ny_nw` = real interest rate (`rint`) × `ny_nw`; `r_no_nw` = real interest rate (`rint`) × `no_nw`. All specifications use the real (inflation-adjusted) interest rate and the log real house price index as dependent variable.

5.3 Model 3: Government Forecast Errors and House Prices

Table 4 reports the central results of the paper: PMG estimates of the forecast-error model (4) over the sample 1986–2019 (and alternative start years 1987–1989 for robustness). This model decomposes the age-structure variables into government-projected components and realised forecast errors.

Table 3: Marginal Effect of the Real Interest Rate on House Prices by Country, Evaluated at Sample-Mean Age Composition (Model 2)

Country	ny	nw	no	CI (low,2SD)	CI (low,1SD)	ME(rint)	CI (high,1SD)	CI (high,2SD)
ALL	0.1913	0.6595	0.1492	-2.0487	-1.4041	-0.7402	-0.0762	0.5684
AU	0.2144	0.6643	0.1213	-3.3446	-2.6474	-1.9293	-1.2113	-0.5141
BE	0.1796	0.6584	0.1621	-1.6681	-0.9429	-0.1960	0.5509	1.2761
CA	0.1930	0.6820	0.1250	-3.5152	-2.6950	-1.8501	-1.0053	-0.1851
CH	0.1705	0.6747	0.1549	-2.3253	-1.4591	-0.5670	0.3252	1.1913
DE	0.1552	0.6710	0.1738	-1.8320	-0.8089	0.2450	1.2988	2.3219
DK	0.1822	0.6586	0.1592	-1.7445	-1.0421	-0.3187	0.4048	1.1072
FR	0.1970	0.6457	0.1573	-1.6928	-1.0286	-0.3444	0.3398	1.0041
GB	0.1889	0.6499	0.1611	-1.5728	-0.8964	-0.1996	0.4971	1.1736
IE	0.2457	0.6424	0.1119	-4.2612	-3.2627	-2.2343	-1.2059	-0.2075
IT	0.1616	0.6639	0.1744	-1.5798	-0.6545	0.2985	1.2515	2.1767
JP	0.1685	0.6652	0.1664	-1.7590	-0.9143	-0.0443	0.8257	1.6704
NL	0.1870	0.6731	0.1399	-2.6806	-1.9446	-1.1865	-0.4284	0.3076
NO	0.1961	0.6483	0.1556	-1.7453	-1.0947	-0.4246	0.2455	0.8961
NZ	0.2287	0.6521	0.1192	-3.5585	-2.7756	-1.9692	-1.1628	-0.3799
SE	0.1806	0.6419	0.1775	-1.1101	-0.3077	0.5187	1.3452	2.1476
US	0.2122	0.6600	0.1278	-2.9818	-2.3211	-1.6406	-0.9601	-0.2994

Notes: The marginal effect (ME) of a one-percentage-point increase in the real long-term interest rate (`rint`) on the log real house price index (`lrrppi`), computed from equation (6) evaluated at each country's sample mean of `ny_nw` and `no_nw`. CI (1SD) and CI (2SD) denote 68% and 95% confidence intervals. Country codes: AU = Australia; BE = Belgium; CA = Canada; CH = Switzerland; DE = Germany; DK = Denmark; FR = France; GB = United Kingdom; IE = Ireland; IT = Italy; JP = Japan; NL = Netherlands; NO = Norway; NZ = New Zealand; SE = Sweden; US = United States.

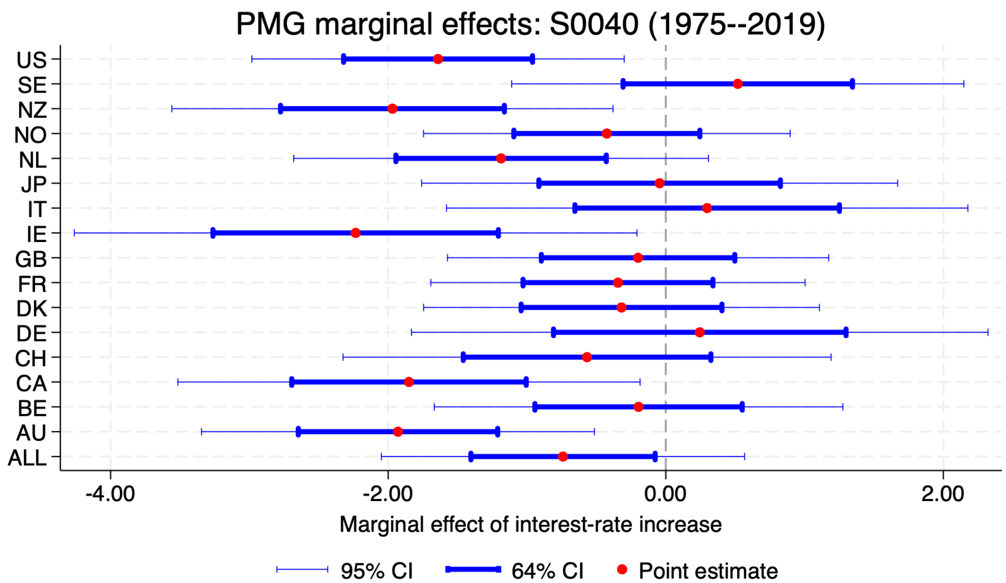

Figure 1: Marginal Effect of the Real Interest Rate on House Prices: Isoline Plot by Youth and Old-Age Dependency Ratios (Model 2)

Notes: Isolines of the marginal effect of the real interest rate on the log real house price index, as a function of the youth-to-working-age ratio (horizontal axis) and the old-age-to-working-age ratio (vertical axis). Country-specific sample-mean evaluation points are overlaid as labeled markers. Darker shading indicates more negative (larger) effects of interest rate increases in depressing house prices.

The error-correction coefficient is -0.187 (s.e. 0.038), negative and significant at the 1% level, confirming long-run equilibrium reversion. Income (`ly2wpop`: 1.048) and population size (`ltpop`: 2.214) retain their expected signs and significance. The real interest rate coefficient of -3.119 is now statistically significant at the 1% level, consistent with the prediction from the PVR.

The central finding of Model 3 concerns the decomposed old-age share variables:

Projected old-age share (`po_nw`): the coefficient is -5.005 (s.e. 0.687), significant at the 1% level. This is the *anticipated* component of ageing. Even the portion of ageing that was foreseeable from official forecasts depresses house prices significantly, indicating that supply adjustment did not fully offset the known trajectory of demographic decline.

Forecast error on old-age share (`no_po`): the coefficient is -8.826 (s.e. 3.158), significant at the 1% level. This is the *unanticipated* component—the degree to which ageing exceeded official projections. Its absolute value is 76 percent larger than that of the projected component (8.826 versus 5.005), precisely as our central hypothesis predicts.

The interpretation is straightforward: when the old-age population grows faster than government forecasts predicted, housing supply—which was calibrated to the (overly optimistic) official projection—is excessive relative to realised demand. The resulting oversupply exerts an additional downward pressure on prices, on top of the demand-reduction effect that the projected ageing alone would have caused.

For the youth-share variables, the results are weaker. The projected youth share (`py_nw`) is positive but not significant, and the forecast error (`ny_py`) is negative but not significant. This asymmetry between youth and old-age variables is consistent with the demographic reality of our sample: it is ageing, not the youth share, that is the dominant source of demographic forecast errors in advanced economies over this period. The finding that old-age forecast errors drive house prices more powerfully than projected old-age shares—and that youth-share errors are largely insignificant—suggests that policymakers should focus especially on the quality of long-run projections for the elderly population when assessing housing market risks.

The Hausman test *p*-value of 0.642 again confirms that the PMG homogeneity restriction cannot be rejected, supporting the pooling of long-run coefficients across countries.

Table 4: Model 3: PMG Estimates with Decomposed Demographic Variables, 1986–2019 and Robustness Across Start Years

	(1) (1986–2019)		(2) (1987–2019)		(3) (1988–2019)		(4) (1989–2019)	
ec								
ly2wpop	1.048***	(0.176)	1.048***	(0.176)	0.235	(0.223)	0.158	(0.199)
rint	-3.119***	(0.813)	-3.119***	(0.813)	-1.940**	(0.933)	0.352	(0.808)
ltpop	2.214***	(0.409)	2.214***	(0.409)	3.761***	(0.437)	2.404***	(0.281)
py_nw	1.284	(1.167)	1.284	(1.167)	1.300	(1.321)	-5.270***	(1.592)
po_nw	-5.005***	(0.687)	-5.005***	(0.687)	-7.029***	(1.133)	-0.739	(0.861)
ny_py	-3.336	(2.685)	-3.336	(2.685)	-6.723**	(2.783)	-3.472	(3.198)
no_po	-8.826***	(3.158)	-8.826***	(3.158)	-11.829***	(3.314)	0.229	(3.034)
SR								
ec	-0.187***	(0.038)	-0.187***	(0.038)	-0.134***	(0.048)	-0.144***	(0.046)
LD.lrrppi	0.419***	(0.051)	0.419***	(0.051)	0.387***	(0.052)	0.441***	(0.050)
D.ly2wpop	0.892***	(0.189)	0.892***	(0.189)	1.059***	(0.219)	0.932***	(0.233)
D.rint	0.982***	(0.210)	0.982***	(0.210)	0.896***	(0.221)	0.801***	(0.191)
D.ltpop	1.347	(1.156)	1.347	(1.156)	3.226*	(1.659)	2.413	(1.508)
D.py_nw	-2.554	(3.349)	-2.554	(3.349)	-2.935	(3.859)	-1.942	(3.339)
D.po_nw	0.806	(2.882)	0.806	(2.882)	0.591	(2.423)	-0.668	(2.919)
D.ny_py	-2.774	(3.892)	-2.774	(3.892)	-2.847	(4.375)	-1.849	(3.989)
D.no_po	2.330	(3.486)	2.330	(3.486)	2.024	(2.646)	0.564	(3.093)
_cons	-7.841***	(1.582)	-7.841***	(1.582)	-8.252***	(2.975)	-5.749***	(1.839)
Observations	496		496		480		464	
Countries	16		16		16		16	
Years	31		31		30		29	
Hausman MG vs PMG chi2	5.145		5.145		2.320		6.886	
Hausman MG vs PMG p-value	0.642		0.642		0.940		0.441	
Spec ID	S0321		S0322		S0323		S0324	

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable: log real residential property price index ($\texttt{lrrppi} = \log P^{rppi} - \log P^{cpi}$). Interest rate: real long-term rate (`rint`). `py_nw` = projected youth-to-working-age ratio (5-year-ahead UN forecast); `po_nw` = projected old-age-to-working-age ratio; `ny_py` = forecast error on youth share ($n^Y_{jt} - \hat{n}^Y_{jt}$); `no_po` = forecast error on old-age share ($n^O_{jt} - \hat{n}^O_{jt}$).

6 Discussion

6.1 From Forecast Optimism to Structural Oversupply

The central finding of Model 3—that the unanticipated component of old-age population growth (ε^O_{jt}) carries a price coefficient of -8.826, compared with -5.005 for the projected component ($\hat{n}^O_{jt}$)—is consistent with the oversupply mechanism developed in Section 3.2. The difference of

3.821 log-index points represents the additional price depression attributable to excess housing stock built because official forecasts overestimated future household demand. The magnitude is economically substantial: across our sample, average annual forecast errors for the old-age share are small in percentage-point terms but cumulate over years, generating persistent structural excess capacity in the housing sector.

Japan's experience provides the clearest illustration. Between the 1992 revision of Japan's official population projections and the actual TFR outcome in 2010, government experts projected a TFR of approximately 1.8; the realised value was 1.4, a gap of 0.4 that accumulated into significant overestimates of future household counts over the intervening two decades. The old-age share exceeded its projected values throughout the 1990s and 2000s. Applying our Model 3 estimates, these forecast errors imply an additional real house price depression of roughly $8.826 \times \bar{\varepsilon}^O$ beyond the $5.005 \times \hat{n}^O$ effect of the projected ageing trajectory—a combined force consistent with the approximately 40 percent real decline in Japanese residential property prices between 1991 and 2015.

The transmission runs directly through the housing stock. Japan's vacancy rate rose from approximately 9 percent in 1993 to 13.6 percent in 2018, with the pace of increase broadly tracking the gap between projected and realised population. The vacant dwellings are not evenly distributed: they are concentrated in suburban areas and regional cities that were developed during the 1970s–1990s in anticipation of population growth that did not materialise. This spatial pattern—large, standardised, medium-density residential developments now falling empty—is the physical manifestation of the forecast-error mechanism documented in our panel data.

Germany's eastern Länder provide a direct European parallel. The post-reunification housing programme was predicated on projections of population convergence that proved far too optimistic: by the early 2000s, vacancy rates in Halle, Magdeburg, and Leipzig had risen above 20 percent, triggering the "Stadtumbau Ost" demolition programme. The cumulative forecast error for the eastern German old-age share over the 1990s— reflecting faster-than-projected ageing as young migrants left for western cities—is, by our estimates, associated with an additional real price depression of the order of $8.826 \times \bar{\varepsilon}^O \approx 5$–8 percent relative to the trajectory implied by the projected demographic path alone. Italy's estimated seven million vacant dwellings are concentrated in the south and in Apennine interior regions, precisely the areas where the discrepancy between projected and realised population was largest: the National Institute of Statistics (ISTAT) repeatedly projected stabilisation of southern Italy's population in the 1990s and 2000s, while emigration to northern Italy and Germany continued to exceed natural increase, leaving a housing stock calibrated to a population that had long since departed.

A further implication of the forecast-error mechanism is that its effects are not symmetric across space within countries. Regions that received the largest infrastructure and housing investments during the bonus period— often based on explicitly quantified regional population projections—tend to accumulate the largest excess stock during the onus period. In Japan, the suburban "new towns" (shin-toshi) developed around major cities in the 1970s–1990s now account for a disproportionate share of the national akiya stock, precisely because they were built to accommodate a projected wave of household formation that attenuated as fertility fell. This within-country spatial heterogeneity, while not directly testable in our country-level panel, is consistent with the mechanism and points toward productive extensions using regional data.

6.2 Why Does Even Anticipated Aging Depress Prices?

A natural question arises: why does the *projected* old-age share ($\hat{n}^O$) also significantly depress prices? Under the rational-expectations benchmark, developers acting on accurate projections should pre-emptively adjust supply, preventing price declines. Our findings suggest two complementary explanations.

First, even the "projected" component of ageing is itself contaminated by forecast optimism. As successive TFR projection vintages illustrate, each successive vintage of Japanese official projections anchored on the previous vintage's TFR assumption and revised downward only slowly. The projected trajectory that enters our Model 3 therefore reflects an *optimistic* official path—not a rational expectation of realised demographic outcomes. Developers acting on this projected path still overbuilt relative to what a truly unbiased forecast would have warranted.

Second, even with accurate projections, the irreversibility of the housing stock means that some supply mismatch is inevitable. Land-use regulations, construction finance, and the spatial fixity of existing dwellings mean that supply contraction in response to projected demand decline is slow, partial, and costly (Knoll et al., 2017; Gyourko et al., 2013). Excess dwellings built during the bonus era remain part of the stock for decades, continuously exerting downward pressure on rents and prices regardless of what developers knew at the time of construction.

6.3 Monetary Policy Transmission in Aging Economies

The results of Model 2 reveal that the effectiveness of interest rate policy for housing markets is substantially conditioned by demographic structure. In the youngest economies of our sample—Ireland (average youth share 24.6 percent) and New Zealand (22.9 percent)—a one-percentage-point decline in the real interest rate raises real house prices by approximately 2.2 and 2.0 percent respectively. In the oldest economies—Italy (old-age share 17.4 percent), Germany (17.4 percent), and Sweden (17.8 percent)—the same rate cut produces effects close to zero or even slightly positive.

This asymmetry has important implications for housing policy. Central banks operating across demographically heterogeneous environments—most obviously the European Central Bank—cannot rely on uniform interest rate adjustments to achieve symmetric housing market outcomes across member states. Younger member states may experience destabilising house price booms in response to monetary easing, while ageing ones remain largely unresponsive. The Hausman test p-value of 0.359 in the baseline column confirms that long-run coefficient homogeneity cannot be rejected, supporting the PMG pooling strategy. This suggests that housing-specific instruments—loan-to-value limits, debt-service ratios, supply incentives—need to be calibrated to the demographic conditions of individual countries or regions, rather than set uniformly.

For monetary policy more broadly, the results confirm that the "lost decades" in Japan cannot be fully explained by demand-side factors alone. Even at near-zero interest rates, the housing market failed to recover because the structural excess supply induced by demographic forecast errors held down rents and expected rent growth $g_{R,t}$, keeping the net discount term $r_t^* - g_{R,t}$ wide and prices depressed. Monetary easing can stimulate housing markets only when there is genuine unmet demand; when the market is structurally oversupplied, the interest-rate channel is blocked.

6.4 The Lost Decades as a Housing Supply Problem

Japan's three-decade property price stagnation is conventionally attributed to the bursting of the 1980s asset bubble and the subsequent demand contraction as the working-age population declined. Our analysis identifies an additional structural layer: the systematic overbuilding induced by official demographic optimism created a persistent excess stock that kept downward pressure on prices even during periods of economic recovery.

The mechanism is self-reinforcing. Falling prices reduce construction incentives, but they do not reduce the existing stock of vacant dwellings. Landlords facing declining rents have limited ability to withdraw supply: demolition costs money, tax rules penalise vacant land over vacant buildings, and the decentralised structure of Japanese housing ownership makes coordinated withdrawal of supply practically impossible. The result is a market permanently in excess supply, where each cyclical recovery in demand is met by a large pool of idle stock rather than by new investment—a pattern consistent with Japan's persistently low housing investment share relative to other advanced economies throughout the post-bubble period.

South Korea, China, and several Southern European countries are at various stages of entering a similar trajectory. South Korea's total fertility rate reached 0.72 in 2023—the lowest ever recorded in any country—yet housing construction has continued at rates calibrated to the government's 2020-round projections, which assumed a TFR floor near 1.0. On our estimates, the cumulative forecast error that will accumulate over the next decade implies a structural excess of housing stock whose price effects could rival those seen in Japan after 1991. China's situation is more complex: officially reported vacancy rates vary widely across city tiers, but satellite-based measures of housing occupancy in third- and fourth-tier cities suggest vacancy rates in some areas already above

20 percent, consistent with construction decisions based on migration projections that overestimated urbanisation-driven household formation (Glaeser et al., 2017).

Our findings suggest that the key policy variable to monitor is not the house price level or the interest rate environment in isolation, but the divergence between official demographic projections and the emerging demographic reality. Specifically, a simple early-warning indicator can be constructed as the accumulated deviation of the realised old-age share from its official five-year-ahead projection—the same object that, in our Model 3, carries the largest price coefficient in the panel. When this indicator turns persistently positive, it signals that the housing stock is likely to be in excess supply relative to realised demand, and that downward pressure on rents and prices will persist even after cyclical recoveries in economic activity or easing of monetary conditions. Real-time publication of this indicator by national statistical agencies or international organisations such as the UN or the OECD would represent a low-cost but potentially high-value contribution to housing market surveillance.

7 Conclusion

Population ageing depresses house prices. This finding, first formalised by Mankiw and Weil (1989) and subsequently contested, is confirmed in our 16-country panel over nearly five decades. But the *mechanism* matters as much as the *direction*: housing markets do not simply weaken as populations age; they accumulate structural excess supply, in the form of rising vacancy rates and stagnant rents, because the institutions responsible for planning and financing residential construction have been systematically overoptimistic about future demographic trajectories.

Our core finding is that the unanticipated component of population ageing— the degree to which realised old-age shares exceeded official five-year-ahead projections—carries a long-run price coefficient of -8.826, substantially larger than the -5.005 coefficient on the projected component. This difference is not merely statistical: it represents the additional price depression caused by housing stock built for a population that did not arrive. Developers, municipalities, and mortgage lenders that relied on government demographic projections were not irrational; they were acting on the best available official information. The problem lies upstream, in the systematic reluctance of national statistical agencies—and of the UN World Population Prospects—to revise fertility and population forecasts downward as quickly as the underlying demographic reality required.

Three implications follow for housing policy and research. First, housing supply planning should be formally linked to demographic forecast quality. Where official projections have a documented track record of optimism bias—measurable as the historical average of the forecast errors constructed in this paper—planning authorities should discount projected housing demand by a corresponding margin. The simplest implementation would be to require that housing supply targets in national spatial plans be validated against scenarios in which the official demographic projection is revised downward by the historical average bias; any supply target that generates excess stock under this scenario should be automatically revisited. The construction of new dwellings in areas of accelerating demographic decline—as is currently occurring in parts of South Korea, China, and Southern Europe—will otherwise generate the same structural oversupply that has characterised Japan for three decades. Second, housing finance regulation should incorporate demographic trajectory risk into mortgage underwriting standards. Properties in areas whose population projections have been repeatedly revised downward carry structural oversupply risk that standard credit models— which treat dwelling values as permanently supported by current demand— substantially underestimate; loan-to-value limits and stress-testing frameworks should be calibrated to the probability that local population will fall short of official projections, not merely to interest rate and income shocks. Third, the demographic dependence of monetary transmission documented in Model 2 has direct implications for macroprudential policy design in ageing societies and heterogeneous currency unions. Interest rate cuts that stimulate supply-constrained, younger markets may fail entirely in structurally oversupplied ageing markets; policymakers in such environments should rely more heavily on demand-side instruments—mortgage guarantee schemes, housing vouchers, renovation and adaptation subsidies—that can generate genuine incremental demand rather than simply re-pricing a stock that is already in excess.

Our panel approach operates at the country level and therefore cannot capture the sharp within-country variation in demographic trajectories that is a distinctive feature of advanced ageing societies. City and regional-level analyses that link local vacancy rates, local demographic forecast errors, and local house price dynamics represent an important extension. The role of international migration—increasingly the dominant source of demographic uncertainty in many advanced economies—deserves separate treatment, as migration flows are subject to different forecasting biases from those that affect natural population change. Finally, the interaction between demographic forecast errors, zoning and planning regulations, and financial sector behaviour deserves systematic investigation: the countries in our sample differ substantially in the degree to which planning systems could in principle accommodate supply adjustments to revised demographic projections, and these differences may explain some of the cross-country variation in our estimates that the PMG framework attributes to heterogeneous short-run dynamics.

References

Badarinza, C., Ramadorai, T., & Shimizu, C. (2021). Gravity, counterparties, and foreign investment. *Journal of Financial Economics*, *145*(2), 532–555.

Couch, C., & Cocks, M. (2013). Housing vacancy and the shrinking city: Trends and policies in the UK and the city of Liverpool. *Housing Studies*, *28*(4), 499–519.

Engelhardt, G. V., & Poterba, J. M. (1991). House prices and demographic change. *Regional Science and Urban Economics*, *21*(4), 539–546.

Francke, M. K., & Korevaar, M. (2025). Baby booms and asset booms: demographic change and the housing market. *Journal of Finance*, *80*(5), 3021–3056.

Glaeser, E., Huang, W., Ma, Y., & Shleifer, A. (2017). A real estate boom with Chinese characteristics. *Journal of Economic Perspectives*, *31*(1), 93–116.

Gyourko, J., Mayer, C., & Sinai, T. (2013). Superstar cities. *American Economic Journal: Economic Policy*, *5*(4), 167–199.

Hamilton, B. W. (1991). The baby boom, the baby bust, and the housing market: a second look. *Regional Science and Urban Economics*, *21*(4), 547–552.

Hendershott, P. H. (1991). Are real house prices likely to decline by 47 percent? *Regional Science and Urban Economics*, *21*(4), 553–563.

Keilman, N. (1998). How accurate are the United Nations world population projections? *Population and Development Review*, *24*(Supplement), 15–41.

Keyfitz, N. (1981). The limits of population forecasting. *Population and Development Review*, *7*(4), 579–593.

Kiyotaki, N., & Moore, J. (1997). Credit cycles. *Journal of Political Economy*, *105*(2), 211–248.

Knoll, K., Schularick, M., & Steger, T. (2017). No price like home: global house prices, 1870–2012. *American Economic Review*, *107*(2), 331–353.

Lozano Alcántara, A., & Vogel, C. (2023). Rising housing costs and income poverty among the elderly in Germany. *Housing Studies*, *38*(7), 1220–1238.

Mankiw, N. G., & Weil, D. N. (1989). The baby boom, the baby bust, and the housing market. *Regional Science and Urban Economics*, *19*(2), 235–258.

Martínez-Gutiérrez, A., Maza, A., & Rubiera-Morollón, F. (2016). Demographic change, housing investment and regional divergence in Spain. *Papers in Regional Science*, *95*(S1), S101–S122.

Nishimura, K. G. (2016). Three "seismic shifts" in the global economy and the policy challenges they pose. *International Finance*, *19*(2), 219–229.

Nishimura, K. G., & Takáts, E. (2012). Ageing, property prices and money demand (BIS Working Papers No. 385). Bank for International Settlements.

Ohtake, F., & Shintani, M. (1996). The effect of demographics on the Japanese housing market. *Regional Science and Urban Economics*, *26*(2), 189–201.

Pesaran, M. H., Shin, Y., & Smith, R. P. (1999). Pooled mean group estimation of dynamic heterogeneous panels. *Journal of the American Statistical Association*, *94*(446), 621–634.

Reinhart, C. M., & Rogoff, K. S. (2009). *This time is different: eight centuries of financial folly*. Princeton University Press.

Saita, Y., Shimizu, C., & Watanabe, T. (2016). Aging and real estate prices: evidence from Japanese and US regional data. *International Journal of Housing Markets and Analysis*, *9*(1), 66–87.

Shimizu, C., & Watanabe, T. (2010). Housing bubble in Japan and the United States. *Public Policy Review*, *6*(3), 431–472.

Takáts, E. (2012). Aging and house prices. *Journal of Housing Economics*, *21*(2), 131–141.

Tamai, Y., Shimizu, C., & Nishimura, K. G. (2017). Aging and property prices: a theory of very-long-run portfolio choice and its predictions on Japanese municipalities in the 2040s. *Asian Economic Papers*, *16*(3), 48–74.

Wiechmann, T., & Pallagst, K. M. (2012). Urban shrinkage in Germany and the USA: a comparison of transformation patterns and local strategies. *International Journal of Urban and Regional Research*, *36*(2), 261–280.

Online Appendix

A.1 Unit Root Tests

Prior to estimation, we verify the integration properties of the main series using the Im–Pesaran–Shin (IPS) panel unit root test (Im et al., 2003) and the Fisher-type ADF test. The test equation is:

$$\Delta y_{it} = \rho_i y_{i,t-1} + \sum_{L=1}^{p_i} \theta_{iL} \Delta y_{i,t-L} + \alpha_i d_t + \varepsilon_{it}, \tag{7}$$

where d_t includes a constant (and optionally a trend) and the null is $H_0 : \rho_i = 0$ for all i. Results are reported in Tables A1–A3 below.

Appendix: Variable-by-Variable Summary of Unit Root Test Results

The time-series ADF test results for the 16 countries in the sample (Australia, Belgium, Canada, Switzerland, Germany, Denmark, France, the United Kingdom, Ireland, Italy, Japan, the Netherlands, Norway, New Zealand, Sweden, and the United States) exhibit several common features across variables. In general, the unit root null cannot be rejected in levels for many countries, while rejections become more frequent in first differences, suggesting that the main variables exhibit strong persistence. Because each country's time series spans only 43 years, the power of the ADF test is inherently limited, and the results below should be interpreted as supplementary evidence.

lrrppi (log real residential property price index) **Time-series ADF tests.** The level-series ADF statistics for `lrrppi` are uniformly weak, providing almost no evidence in favor of level stationarity for Australia, Belgium, Canada, Switzerland, Germany, Denmark, France, the United Kingdom, Japan, the Netherlands, New Zealand, Sweden, and the United States, among others. In first differences, the test statistics become substantially more negative for Canada, Germany, France, Ireland, Italy, Japan, the Netherlands, Sweden, and the United States, pointing more clearly toward stationarity. Overall, it is appropriate to treat the real house price index as an $I(1)$ series.

Panel unit root tests. Results for the level series are mixed. In the constant-only specification, LLC, IPS, and Breitung cannot reject the unit root null, while Fisher-ADF rejects it. In the specification with a constant and trend, LLC, IPS, and Fisher-ADF reject the null, but the Hadri test rejects the stationarity null. In first differences, however, the unit root null is clearly rejected by the majority of tests, and it is appropriate overall to treat `lrrppi` as an $I(1)$ series.

ly2wpop (log real GDP per working-age person) **Time-series ADF tests.** Non-stationarity in levels is particularly pronounced for this variable. The level-series ADF statistics take large positive values for Australia, Belgium, Canada, Switzerland, Germany, Denmark, France, the United Kingdom, the Netherlands, New Zealand, and the United States, leaving very little evidence for level stationarity. In first differences, results improve toward rejection for Australia, Belgium, Canada, Switzerland, Germany, Denmark, France, the United Kingdom, Ireland, Italy, Japan, the Netherlands, Norway, Sweden, and the United States. `ly2wpop` is therefore among the variables most clearly consistent with $I(1)$ behavior.

Panel unit root tests. There is some inconsistency across tests for the level series. In the constant-only specification, LLC and Fisher-ADF reject the unit root null, while IPS and Fisher-PP do not. In the specification with a trend, LLC and Breitung reject but IPS and Fisher-PP do not uniformly do so. In first differences, the unit root null is strongly rejected in all specifications, and it is appropriate to classify `ly2wpop` as an $I(1)$ series.

rint (real interest rate) **Time-series ADF tests.** The unit root null cannot be rejected in levels for many countries, including Australia, Belgium, Canada, Switzerland, Germany, Denmark, France, the United Kingdom, Ireland, Italy, Japan, the Netherlands, New Zealand, Sweden, and the United States, suggesting non-stationarity in levels. In first differences, clearer rejections are obtained for Australia, Belgium, Canada, Switzerland, Germany, Denmark, France, the United Kingdom, Norway, New Zealand, Sweden, and the United States. Overall, the real interest rate is consistent with the non-stationary-in-levels, stationary-in-differences pattern, although cross-country heterogeneity is somewhat greater than for house prices and income.

Panel unit root tests. A relatively large share of tests already suggests stationarity in levels. In the constant-only specification, IPS, Breitung, Fisher-ADF, and Fisher-PP reject the unit root null, and similar patterns hold in the specification with a trend for LLC, IPS, Fisher-ADF, and Fisher-PP. The Hadri test rejects the stationarity null, leaving some inconsistency across tests. In first differences, stationarity is strongly supported by all major tests. Because the possibility of $I(0)$ behavior in levels cannot be ruled out, `rint` is best treated cautiously as a series on the boundary between $I(0)$ and $I(1)$.

ltpop (log total population) Time-series ADF tests. This variable exhibits particularly strong persistence in levels. ADF statistics are uniformly weak across most countries—Australia, Belgium, Canada, Switzerland, Germany, Denmark, France, the United Kingdom, Ireland, Italy, Japan, the Netherlands, Norway, New Zealand, Sweden, and the United States—reflecting the strong influence of long-run demographic trends. In first differences, some improvement is observed for France, Japan, the Netherlands, and the United States, but the evidence is not as uniform as for `lrrppi` or `ly2wpop`. It is therefore natural to treat `ltpop` as an $I(1)$ series, bearing in mind that its variation is largely driven by gradual and persistent demographic trends rather than short-run cyclical fluctuations.

Panel unit root tests. Results for the level series predominantly suggest non-stationarity. In the constant-only specification, LLC, IPS, and Breitung cannot reject the unit root null, and only Fisher-type tests do so. In the specification with a trend, LLC and Fisher-ADF reject but IPS, Fisher-PP, and Breitung do not, leaving a mixed picture. In first differences, however, the unit root null is rejected consistently in both specifications, making it natural to treat `ltpop` as an $I(1)$ series.

ny_nw (youth population ratio n^Y/n^W) Time-series ADF tests. A substantial number of countries—Australia, Belgium, Canada, Switzerland, Germany, Denmark, the United Kingdom, Ireland, Japan, the Netherlands, Sweden, and the United States—fail to reject the unit root null in levels, suggesting high persistence. In France, Italy, Norway, and New Zealand, the level-series statistics are somewhat more negative than in other countries. First-difference results show some improvement but are less clear-cut than for house prices or income. While the variable exhibits high statistical persistence, `ny_nw` is an inherently bounded ratio that is not straightforwardly comparable to a standard macroeconomic level variable.

Panel unit root tests. Results for the level series relatively often support stationarity. In the constant-only specification, LLC, IPS, Fisher-ADF, Fisher-PP, and Breitung strongly reject the unit root null; in the specification with a trend, LLC, IPS, and Fisher-ADF yield similar rejections. Although the Hadri test rejects the stationarity null, the preponderance of panel evidence for stationarity in levels—consistent with the bounded nature of ratio variables—makes it appropriate to interpret `ny_nw` as likely exhibiting $I(0)$ behavior.

no_nw (old-age population ratio n^O/n^W) **Time-series ADF tests.** Like `ny_nw`, this variable displays high persistence. The unit root null cannot be rejected in levels for Australia, Belgium, Canada, Switzerland, Denmark, France, the United Kingdom, Ireland, Japan, the Netherlands, New Zealand, and the United States, providing limited evidence for level stationarity. Some improvement in first differences is observed for Italy, Norway, and Sweden, but the evidence remains mixed. Although strong persistence is confirmed statistically, the interpretation of the stochastic process for `no_nw` requires caution in light of its nature as a bounded compositional share.

Panel unit root tests. Several test results suggest stationarity in levels. In the constant-only specification, LLC and IPS cannot reject the unit root null, but Fisher-ADF and Breitung reject it. In the specification with a trend, LLC, IPS, and Fisher-ADF reject the null, providing some evidence for level stationarity. The Hadri test rejects the stationarity null, and inconsistency across tests remains. Nevertheless, given the bounded nature of ratio variables, `no_nw` may also exhibit $I(0)$-type behavior, as does `ny_nw`.

Table A1: Time-Series ADF Tests: Level and First-Difference Series (No Constant)

Country	lrrppi	ly2wpop	rint	ltpop	ny_nw	no_nw	lrrppi	ly2wpop	rint	ltpop	ny_nw	no_nw
	Levels						First Differences					
AU	2.92 [2]	7.48 [0]	-1.62* [0]	32.63 [0]	-0.55 [1]	-1.13 [1]	-4.16*** [0]	-3.27*** [0]	-8.06*** [0]	-0.48 [0]	-1.13 [0]	0.44 [0]
BE	1.06 [2]	6.86 [0]	-0.86 [1]	1.58 [1]	-0.50 [2]	-1.09 [2]	-3.18*** [1]	-3.73*** [0]	-8.38*** [0]	-0.77 [0]	-2.26** [1]	-2.14** [1]
CA	1.43 [1]	3.95 [0]	-0.80 [1]	2.83 [1]	-0.04 [2]	-1.98** [2]	-3.78*** [0]	-4.39*** [0]	-9.31*** [0]	-0.35 [0]	-3.44*** [1]	-0.95 [1]
CH	1.24 [2]	4.03 [0]	-1.50 [1]	3.21 [2]	-0.08 [2]	-0.70 [1]	-2.49** [1]	-4.09*** [0]	-11.04*** [0]	-0.99 [1]	-1.86* [1]	-0.48 [0]
DE	0.48 [1]	5.64 [0]	-0.95 [0]	1.22 [1]	-0.99 [1]	-0.75 [2]	-2.27** [0]	-3.90*** [0]	-7.03*** [0]	-3.46*** [0]	-1.58 [0]	-1.99** [1]
DK	0.53 [1]	5.71 [0]	-1.02 [0]	2.12 [2]	-0.34 [1]	-0.91 [2]	-3.49*** [0]	-3.81*** [0]	-8.85*** [0]	-0.46 [0]	-0.91 [0]	-1.26 [1]
FR	0.71 [1]	3.02 [1]	-0.87 [0]	2.19 [2]	-0.20 [1]	-1.33 [2]	-2.41** [0]	-2.86*** [0]	-8.07*** [0]	-0.80 [1]	-1.13 [0]	-2.43** [1]
GB	1.22 [1]	2.48 [1]	-0.80 [3]	2.08 [3]	-0.08 [2]	-0.85 [2]	-3.21*** [0]	-2.70*** [0]	-1.82* [4]	0.81 [2]	-2.20** [1]	-1.98** [1]
IE	0.84 [4]	5.49 [0]	-2.06** [0]	2.44 [2]	0.06 [1]	-0.61 [1]	-3.14*** [3]	-3.08*** [0]	-6.16*** [0]	-1.41 [1]	-0.80 [0]	-0.56 [0]
IT	0.26 [1]	1.93 [1]	-1.62 [0]	1.21 [1]	0.01 [2]	-1.38 [2]	-2.81*** [0]	-3.89*** [0]	-1.40 [4]	-2.54** [0]	-1.61 [1]	-1.97** [1]
JP	0.24 [1]	7.58 [0]	-1.69* [0]	-1.24 [2]	0.14 [1]	-1.47 [2]	-2.71*** [0]	-2.91*** [0]	-7.97*** [0]	-3.21*** [1]	-1.11 [0]	-0.40 [1]
NL	1.16 [4]	2.26 [1]	-0.78 [0]	2.62 [3]	-0.93 [1]	-1.58 [2]	-2.85*** [2]	-2.64*** [0]	-6.55*** [0]	-0.94 [2]	-2.22** [1]	-0.79 [1]
NO	1.71 [1]	1.89 [1]	-1.69* [0]	1.77 [2]	-0.02 [1]	-0.46 [1]	-3.73*** [0]	-2.56** [0]	-8.68*** [0]	-0.47 [1]	-1.09 [0]	-0.56 [0]
NZ	1.46 [1]	3.06 [2]	-2.14** [0]	8.67 [0]	-0.31 [1]	-1.07 [1]	-3.10*** [0]	-3.73*** [0]	-7.68*** [0]	-2.35** [0]	-1.16 [0]	0.39 [0]
SE	0.72 [1]	5.07 [0]	-1.72* [0]	1.70 [2]	-0.20 [1]	-0.93 [2]	-2.81*** [0]	-3.66*** [0]	-9.60*** [0]	0.52 [0]	-1.42 [0]	-1.67* [1]
US	0.72 [2]	6.54 [0]	-1.51 [0]	-0.41 [1]	-0.22 [1]	-1.77* [1]	-4.07*** [1]	-3.10*** [0]	-6.56*** [0]	-0.99 [0]	-2.41** [0]	1.06 [0]

Notes: ADF statistics; lag order in brackets (AIC). Critical values (no constant): 1% = -2.6, 5% = -1.95, 10% = -1.62. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A2: Time-Series ADF Tests: Level and First-Difference Series (Constant Only)

	lrrppi	ly2wpop	rint	ltpop	ny_nw	no_nw	lrrppi	ly2wpop	rint	ltpop	ny_nw	no_nw
Country	Levels						First Differences					
AU	0.37 [2]	0.18 [0]	-3.11** [0]	2.32 [0]	-2.88* [1]	-0.52 [1]	-5.64*** [1]	-5.92*** [0]	-7.99*** [0]	-4.30*** [0]	-0.72 [0]	-0.11 [0]
BE	-0.28 [2]	-1.34 [0]	-2.55 [0]	0.85 [1]	-6.50*** [1]	-0.19 [2]	-3.36** [1]	-5.95*** [0]	-8.27*** [0]	-1.72 [0]	-1.98 [1]	-2.36 [1]
CA	0.12 [1]	-0.36 [0]	-2.13 [0]	0.19 [1]	-1.20 [2]	1.10 [2]	-4.08*** [0]	-5.36*** [0]	-9.19*** [0]	-2.84* [0]	-2.93** [1]	-1.73 [1]
CH	-0.18 [2]	-0.24 [0]	-1.74 [1]	1.34 [2]	-5.26*** [1]	-0.68 [1]	-2.80* [1]	-5.30*** [0]	-10.97*** [0]	-3.34** [1]	-1.63 [1]	-0.72 [0]
DE	-1.86 [1]	-0.94 [0]	-0.59 [0]	-1.17 [1]	-4.48*** [2]	-0.49 [2]	-2.26 [0]	-5.51*** [0]	-7.09*** [0]	-3.67*** [0]	-1.16 [0]	-2.10 [1]
DK	-1.28 [1]	-1.27 [0]	-1.01 [0]	3.01 [2]	-6.42*** [1]	-3.13** [1]	-3.52*** [0]	-5.37*** [0]	-8.84*** [0]	-2.51 [1]	-0.80 [0]	-1.52 [1]
FR	-1.01 [1]	-0.68 [1]	-1.53 [0]	-1.51 [2]	-4.75*** [1]	1.09 [2]	-2.52 [0]	-4.37*** [0]	-7.98*** [0]	-2.35 [1]	-1.02 [0]	-2.79* [1]
GB	-1.38 [1]	-1.04 [1]	-4.45*** [0]	-1.74 [3]	-5.34*** [1]	-4.03*** [1]	-3.54*** [0]	-3.84*** [0]	-1.77 [4]	-1.25 [2]	-2.00 [1]	-2.12 [1]
IE	-1.05 [4]	0.79 [0]	-3.13** [0]	0.90 [2]	-3.53*** [1]	-4.20*** [1]	-3.27** [3]	-4.44*** [0]	-6.08*** [0]	-2.85* [1]	-1.05 [0]	-0.50 [0]
IT	-2.37 [1]	-1.84 [1]	-3.04** [0]	-0.57 [1]	-10.57*** [1]	0.59 [2]	-2.80* [0]	-4.47*** [0]	-1.41 [4]	-2.78* [0]	-1.48 [1]	-2.46 [1]
JP	-1.99 [1]	-3.21** [0]	-2.70* [0]	-3.00** [1]	-4.65*** [1]	-0.68 [2]	-2.65* [0]	-4.87*** [0]	-7.87*** [0]	-1.65 [1]	-1.15 [0]	-1.39 [1]
NL	-0.82 [4]	-0.06 [1]	-0.99 [0]	-2.64* [2]	-6.18*** [1]	0.29 [2]	-3.37** [3]	-3.58*** [0]	-6.49*** [0]	-2.80* [2]	-1.51 [0]	-1.41 [1]
NO	-0.91 [1]	-1.89 [1]	-2.73* [0]	1.41 [2]	-3.33** [1]	-4.86*** [1]	-4.25*** [0]	-3.26** [0]	-8.58*** [0]	-1.80 [1]	-0.93 [0]	-0.57 [0]
NZ	-0.28 [1]	1.17 [0]	-3.12** [0]	2.09 [0]	-4.46*** [1]	-1.83 [1]	-3.47** [0]	-4.66*** [1]	-7.66*** [0]	-4.53*** [0]	-0.85 [0]	0.07 [0]
SE	-0.53 [1]	-0.01 [0]	-2.51 [0]	2.23 [2]	-4.15*** [1]	-2.65* [1]	-2.90** [0]	-5.07*** [0]	-9.49*** [0]	-1.72 [1]	-1.39 [0]	-1.89 [1]
US	-1.15 [2]	-1.41 [0]	-2.48 [0]	-2.18 [1]	-2.11 [1]	-3.36** [1]	-4.11*** [1]	-4.68*** [0]	-6.48*** [0]	0.15 [0]	-1.96 [0]	0.46 [0]

Notes: ADF statistics; lag order in brackets (AIC). Critical values (constant): 1% = -3.51, 5% = -2.89, 10% = -2.58. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A3: Panel Unit Root Tests: p-values

Test	`lrrppi`	`ly2wpop`	`rint`	`ltpop`	`ny_nw`	`no_nw`	`lrrppi`	`ly2wpop`	`rint`	`ltpop`	`ny_nw`	`no_nw`
	Levels						First Differences					
LLC t^*	0.458	0.002***	0.072*	0.735	<0.001***	1.000	<0.001***	<0.001***	<0.001***	0.013**	<0.001***	0.015**
IPS *arW*	0.990	0.999	0.002***	1.000	<0.001***	0.996	<0.001***	<0.001***	<0.001***	<0.001***	0.015**	0.122
Fisher-ADF (P)	<0.001***	0.016**	<0.001***	0.026**	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***
Fisher-PP (P)	1.000	0.569	<0.001***	<0.001***	<0.001***	1.000	<0.001***	<0.001***	<0.001***	<0.001***	0.998	1.000
Breitung λ	0.236	1.000	0.003***	0.079*	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	0.012**	0.138	0.002***
Hadri–Tzavalis z	1.000	0.999	<0.001***	1.000	0.485	1.000	<0.001***	<0.001***	<0.001***	<0.001***	0.885	0.984
Hadri z	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	0.002***	<0.001***	<0.001***	<0.001***

Notes: p-values reported. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. LLC = Levin–Lin–Chu; IPS = Im–Pesaran–Shin; Fisher tests are inverse chi-squared. Hadri test has stationarity as the null; all other tests have unit root as the null.

A.2 Construction of Five-Year Population Projections

The United Nations publishes revisions to the World Population Prospects every two years. Revision years used in this study: 1982, 1984, 1988, 1990, 1992, 1994, 1996, 1998, 2000, 2002, 2004, 2006, 2008, 2010, 2012. Each revision publishes projections for years ending in 0 or 5. We construct annual five-year-ahead projections $\hat{n}^k_{jt}$ by linear interpolation between published quinquennial values, selecting the vintage whose projection date is exactly five years prior to the observation year. The forecast error is $\varepsilon^k_{jt} = n^k_{jt} - \hat{n}^k_{jt}$.

1. *Revision years ending in 0 or 5:* use the UN's published five-year projection directly.
2. *Other revision years:* interpolate linearly between the two adjacent quinquennial projections.
3. *Non-revision years:* interpolate linearly between the two adjacent revision-year projections.

A.3 Structural Derivation of the Demographic Present-Value Relation

A.3.1 Production structure and the rental price of housing

We model housing services as a productive input in a Cobb–Douglas economy: $Y_t = A_t K_t^{\alpha} L_t^{\beta} \tilde{H}_t^{1-\alpha-\beta}$, where $\tilde{H}_t \equiv H_t \phi(D_t)$ is effective housing capacity, H_t is the physical stock, and $\phi(D_t)$ is a demographic efficiency factor. Under profit maximisation, equilibrium rents satisfy $R_t = \kappa_t Y_t$, $\kappa_t \equiv (1-\alpha-\beta)/[H_t \phi(D_t)]$. Because H_t and $\phi(D_t)$ are slow-moving, expected long-run rent growth is

$$g_{R,t} \approx \mathbb{E}_t[\Delta \ln Y_{t+1}] - \mathbb{E}_t[\Delta \ln H_{t+1}] - \mathbb{E}_t[\Delta \ln \phi(D_{t+1})].$$

A.3.2 Asset pricing and the long-run PVR

Under no-arbitrage, all quantities in real terms: $Q_t \equiv P_t^{rppi}/P_t^{cpi} = \mathbb{E}_t[\sum_{s=1}^{\infty} M_{t,t+s}R_{t+s}]$. Applying a Gordon-growth approximation yields the real demographic PVR: $Q_t \approx \kappa_t(Y_t/\mathrm{pop}_t^W)/(r_t^* - g_{R,t})$, where r_t^* is the real natural interest rate.

A.3.3 Demographic channels and comparative statics

Let n_t^Y, n_t^W, n_t^O be youth, working-age, and old-age shares. The sign restrictions $\partial g_R/\partial n_t^Y > 0$, $\partial g_R/\partial n_t^O < 0$, $\partial r^*/\partial n_t^Y > 0$, $\partial r^*/\partial n_t^O < 0$ follow from overlapping-generations logic. Differentiating $\ln P_t$:

$$\frac{\partial \ln P_t}{\partial n_t^g} = \frac{\partial \ln \kappa_t}{\partial n_t^g} + \frac{\partial \ln Y_t}{\partial n_t^g} - \frac{1}{r_t^* - g_{R,t}}\left(\frac{\partial r_t^*}{\partial n_t^g} - \frac{\partial g_{R,t}}{\partial n_t^g}\right).$$

These signs imply $\partial \ln Q_t/\partial n_t^Y > 0$ and $\partial \ln Q_t/\partial n_t^O < 0$, where $Q_t = P_t^{rppi}/P_t^{cpi}$ is the real house price. The real interest-rate elasticity $\partial \ln Q_t/\partial r_t^* = -1/(r_t^* - g_{R,t})$ is state-dependent: it is larger in absolute value in young economies (small net discount term) and attenuated in ageing ones—the comparative static behind Model 2.

References (Online Appendix)

Breitung, J. (2000). The local power of some unit root tests for panel data. In B. Baltagi (Ed.), *Advances in econometrics, Vol. 15: Nonstationary panels, panel cointegration, and dynamic panels* (pp. 161–177). JAI Press.

Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, *74*(366), 427–431.

Dickey, D. A., & Fuller, W. A. (1981). Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica*, *49*(4), 1057–1072.

Engle, R. F., & Granger, C. W. J. (1987). Co-integration and error correction: representation, estimation, and testing. *Econometrica*, *55*(2), 251–276.

Granger, C. W. J., & Newbold, P. (1974). Spurious regressions in econometrics. *Journal of Econometrics*, *2*(2), 111–120.

Hadri, K. (2000). Testing for stationarity in heterogeneous panel data. *Econometrics Journal*, *3*(2), 148–161.

Im, K. S., Pesaran, M. H., & Shin, Y. (2003). Testing for unit roots in heterogeneous panels. *Journal of Econometrics*, *115*(1), 53–74.

Kao, C. (1999). Spurious regression and residual-based tests for cointegration in panel data. *Journal of Econometrics*, *90*(1), 1–44.

Levin, A., Lin, C.-F., & Chu, C.-S. J. (2002). Unit root tests in panel data: asymptotic and finite-sample properties. *Journal of Econometrics*, *108*(1), 1–24.

Maddala, G. S., & Wu, S. (1999). A comparative study of unit root tests with panel data and a new simple test. *Oxford Bulletin of Economics and Statistics*, *61*(S1), 631–652.

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, *55*(3), 703–708.

Pedroni, P. (1999). Critical values for cointegration tests in heterogeneous panels with multiple regressors. *Oxford Bulletin of Economics and Statistics*, *61*(S1), 653–670.

Pedroni, P. (2004). Panel cointegration: asymptotic and finite sample properties of pooled time series tests with an application to the PPP hypothesis. *Econometric Theory*, *20*(3), 597–625.

Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence. *Journal of Applied Econometrics*, *22*(2), 265–312.

Pesaran, M. H., & Shin, Y. (1999). An autoregressive distributed lag modelling approach to cointegration analysis. In S. Storm (Ed.), *Econometrics and economic theory in the 20th century: the Ragnar Frisch centennial symposium* (pp. 371–413). Cambridge University Press.

Pesaran, M. H., Shin, Y., & Smith, R. P. (1999). Pooled mean group estimation of dynamic heterogeneous panels. *Journal of the American Statistical Association*, *94*(446), 621–634.

Phillips, P. C. B., & Hansen, B. E. (1990). Statistical inference in instrumental variables regression with $I(1)$ processes. *Review of Economic Studies*, *57*(1), 99–125.

Phillips, P. C. B., & Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika*, *75*(2), 335–346.

Said, S. E., & Dickey, D. A. (1984). Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika*, *71*(3), 599–607.

Stock, J. H., & Watson, M. W. (1993). A simple estimator of cointegrating vectors in higher order integrated systems. *Econometrica*, *61*(4), 783–820.

Westerlund, J. (2007). Testing for error correction in panel data. *Oxford Bulletin of Economics and Statistics*, *69*(6), 709–748.