PREDETERMINED DEMAND

HOW PAST BIRTH COHORTS PREDICT REGIONAL HOUSE PRICE GROWTH IN SWEDEN

HUGO FALK

JOHAN PERSSON

Bachelor Thesis

Stockholm School of Economics

2026

Predetermined Demand: How Past Birth Cohorts Predict Regional House Price Growth in Sweden

Abstract:

In a panel of 20 Swedish counties observed annually from 1996 to 2022, the share of newborns in a county 27 years earlier is associated with subsequent five-year house price growth: a one standard deviation increase corresponds to about 4.7 percentage points of cumulative growth. The same lagged birth share predicts four housing-market outcomes at different horizons, with apartment prices peaking at lag 24, house prices at lag 27, transactions at lag 30, and construction at lag 32. Tracing the same cohort at successive ages from birth through home-buying age, birth-time location is a more reliable predictor of regional house prices than the cohort's location at any later age, and predictive power weakens sharply through the university years.

Keywords:

Demographic demand, Birth cohorts, House prices, Inelastic housing supply, Regional housing markets

Authors:

Hugo Falk (26095)
Johan Persson (26094)

Tutor:

Paula Roth, Postdoc Fellow, Swedish House of Finance

Examiner:

Ramin Baghai, Professor, Department of Finance

Bachelor Thesis
Bachelor Program in Business and Economics
Stockholm School of Economics

I. Introduction

The demand for homeownership follows a life-cycle pattern (Francke and Korevaar, 2025). People rent during their early adulthood, enter homeownership in their late twenties and early thirties as their income stabilises and they form families, and gradually exit the market later in life. Because housing demand is so concentrated by age, the size of the cohort reaching its prime buying years can itself become a determinant of demand. When an unusually large cohort reaches its late twenties, the number of households seeking to buy housing rises, placing upward pressure on prices.

This paper asks whether demographic changes predict regional house price growth in Sweden. Residential property is the largest asset class for most households, and house price growth is a dominant driver of household wealth accumulation. The demographic demand channel is interesting because it is predetermined. Cohort sizes from 25 to 30 years ago are observable today in public records, and if they predict future regional price growth, this points to a potentially forecastable component of housing returns. Market efficiency would suggest such a signal should already be incorporated into prices and into supply decisions in the years before the cohort arrives. Our findings suggest the signal is not fully incorporated in advance.

Establishing the link between demographics and prices in modern data is difficult because house price growth has many confounders. Real house prices were roughly stable historically, but in modern times they fluctuate substantially in response to economic shocks and policy changes, making the demographic component hard to isolate. Cross-country comparisons are especially fragile because these shocks and changes occur at different times in different countries. A cross-county panel within Sweden addresses this in two ways. Institutionally, all Swedish counties share monetary policy, mortgage market regulation, and tax treatment of housing, so these common factors are absorbed by year fixed effects. The Swedish housing market also concentrates demographic demand sharply. A queue-rationed and shrinking rental sector, early household formation, and slow permitting that limits supply responses together push the demand of an arriving cohort into priced ownership tenures (Section III). Sweden also offers detailed regional records over long time periods, including population by single year of age, a quality-adjusted house price index, transaction registers, construction, and income statistics, all at the county level.

We construct a county-level panel of 20 Swedish counties from 1996 to 2022 and regress centred five-year log house price changes on the share of newborns in each county 27 years prior, controlling for income growth among 25-29 year olds and absorbing county and year fixed effects. We use newborns rather than older age groups because the geographic distribution of a cohort at birth has not yet been reshaped by inter-county migration, making the regressor predetermined relative to current housing market conditions.

In our main specification, a one standard deviation increase in the lagged birth share is associated with approximately 4.7 percentage points higher cumulative five-year house price growth, significant at the 1% level. The coefficient peaks sharply at a lag of 27 years, slightly before the typical Swedish home-buying age. The sign and statistical significance of the effect survive non-overlapping estimation windows, annual price changes, migration controls, and the exclusion of the three major metropolitan counties, though the magnitude is not stable across different time periods. The same demographic shift shows up in four sequential housing market outcomes: apartment prices peak first at lag 24, house prices at lag 27, transaction volumes at lag 30, and new construction at lag

32. The five-year lead of prices over construction is consistent with a market in which short-run housing supply is inelastic, so demographic demand shows up in prices before quantities can adjust.

A separate analysis tests which observation age of the same cohort actually carries predictive power. Holding the target cohort fixed at the year in which it reaches age 27, we measure the cohort's county share at every age from birth through home-buying age and re-estimate the regression at each. The resulting cohort observation age profile traces a W-shape. The demographic signal is strongest at birth, recovers around age 18 after a childhood decline, drops sharply during the university years, and only partially returns by buying age. The dip during the university years and the partial recovery by buying age suggest that many people leave their birth county for studies and early career but end up buying a house in the same county where they were born.

II. Related Literature and Contribution

The closest related work is Francke and Korevaar (2025), who study historical Amsterdam (with Paris rent data as a supplementary case) and find that birth cohorts reaching home-buying age predict higher house prices, without a similar effect observed for rents. They use lagged birth rates rather than the contemporaneous age structure, arguing that birth rates are determined long before any housing-market outcome and so cannot be affected by it. Their internet appendix extends the result to a modern OECD panel using a related specification.

Monnet and Wolf (2017) extend the predetermined-cohort approach to a cross-country panel of OECD countries. Their dependent variable is residential investment relative to GDP, and they instrument current growth of the 20–49 age group with the growth of the same birth cohort observed twenty years earlier. The instrument is strong in their low-migration sub-sample, which includes Sweden, and weak in the high-migration sub-sample, which directly supports the use of predetermined cohort variables in the Swedish setting. They argue against house prices as a cross-country dependent variable on comparability grounds, and their own price regressions deliver no robust demographic effect.

The related literature started with Mankiw and Weil (1989), who construct a U.S. demographic demand index from the age distribution weighted by age-specific housing demand from the Census, and find that this index predicts long-term real house price movements at the national level. Engelhardt and Poterba (1991) replicate this exercise on Canadian national time-series data and obtain insignificant and mostly negative coefficients on demographic demand, casting doubt on the strength of the link in national time series. An important difference in our paper is the exploitation of within-country differences instead of looking at the national level.

More broadly, our study relates to the finance literature on demographic structure and asset returns. Abel (2003) develops an overlapping-generations model in which a baby boom raises the price of capital through higher saving by the young, with the increase subsequently mean-reverting as the capital stock grows. Poterba (2001), testing whether U.S. age structure predicts real returns on T-bills, bonds, and equities, finds only weak evidence on returns but somewhat firmer evidence linking demographic measures to the price-dividend ratio. Empirically, DellaVigna and Pollet (2007) demonstrate that demographic demand shifts predict stock returns in age-sensitive industries, suggesting that markets do not fully incorporate predictable demographic changes. Our paper applies the same logic to housing: if demand from birth cohorts is predictable and not fully priced in, lagged births should predict regional house price growth.

Related research examines how housing markets adjust to demand shocks. Glaeser and Gyourko (2005) show that the durability of housing makes the supply curve asymmetric, with negative demand shocks falling mostly on prices because the housing stock disappears slowly, and positive shocks falling mostly on quantities so long as new construction is elastic. Saiz (2010) measures the regulatory and geographic constraints that limit that elasticity, and shows that the responsiveness of housing supply shapes how demand shocks translate into prices.

Our paper contributes to this literature in three ways. First, we estimate the same regression at every observation age of the same cohort, from birth through buying age, and compare the resulting predictive power. Francke and Korevaar (2025) make the case for predetermined birth rates on identification grounds, and our comparison confirms empirically that birth-time measurement is preferable to contemporaneous demographic structures. Second, we document a four-stage timing chain through the housing market in which apartment prices peak at lag 24, house prices at lag 27, transactions at lag 30, and construction at lag 32. The five-year lead of prices over construction implies short-run supply is inelastic. Third, the within-country panel isolates demographic variation from national-level shocks that limit the comparability of cross-country studies in this literature.

III. Institutional context

The Swedish rental sector is heavily regulated and rationed by queue rather than by price. Landlords and tenant unions set rents through collective bargaining under the so-called utility value system, holding rents below equilibrium levels in high-demand areas. As a result, an unusually large cohort entering the housing market does not show up as higher rents. OECD (2025) classifies Sweden's rent controls as the most restrictive in the OECD, with no immediately available vacancies in any of Sweden's three main cities (Stockholm, Gothenburg, and Malmö). In the Stockholm region the average queuing time for a rental apartment was 9 years in 2025, with inner-city apartments requiring an average of 21 years (Bostadsförmedlingen i Stockholm, 2025). A structural decline in the metropolitan rental stock has intensified this pressure, with the number of rental apartments per 1,000 residents in Sweden's metropolitan areas falling from 244 to 176 between 1990 and 2018, and the Stockholm rental stock contracting by 12 percent over a period when the population grew by 43 percent (Olsén Ingefeldt and Thell, 2019). New rental construction has not made up the difference. The municipal housing companies that dominate the rental segment have limited incentive to build new units, and private developers concentrate their projects in the owner-occupied segment, leaving rental supply largely stagnant (OECD, 2025). With the rental stock contracting, an additional cohort arriving at buying age has effectively no rental segment to take it in, and the demographic demand pressure shows up in the priced ownership segment.

Apartment ownership in Sweden takes the form of the *bostadsrätt*, a tenant-owned cooperative in which the cooperative owns the building and the household holds a share that carries the right to occupy a specific unit (Olsén Ingefeldt and Thell, 2019). Although the legal structure is different, the economic substance is similar to ownership: the share trades at market price, is financed by mortgage credit, and serves as the household's primary housing asset. Olsén Ingefeldt and Thell (2019) treat bostadsrätter and single-family houses as a single category of owned housing in their analysis of household entry into the housing market. The bostadsrätt segment is large and growing fast, especially in cities. Between 1990 and 2018 the number of bostadsrätter per 1,000 residents in Sweden's metropolitan areas grew from 105 to 161, compared with 66 to 80 in the rest of the country (Olsén Ingefeldt and Thell, 2019). Two channels drive this growth. The first is the conversion of existing rental buildings into bostadsrätter, in which the tenants form a cooperative and collectively buy the property from the landlord (Olsén Ingefeldt and Thell, 2019). Conversions have contributed to the decrease in the metropolitan rental share documented above (OECD, 2025). The second is new private construction in the owner-occupied segment, which Swedish developers concentrate on, adding units to the priced side of the market over time (OECD, 2025). While the rental stock contracts, the bostadsrätt stock continues to expand on both fronts and can accommodate additional demand. The housing demand of new cohorts in Sweden therefore flows primarily into the two priced routes into ownership: single-family houses, and bostadsrätter.

Young adults in Sweden form independent households earlier and stay with their parents less commonly than peers in most of Europe. The average age at which young adults leave the parental household was 21.9 years in 2024, compared to an EU27 average of 26.2 years and above 30 years in much of southern Europe (Eurostat, 2024), and only 42 percent of Swedes aged 16 to 29 lived with their parents in 2017, among the lowest shares in Europe (Olsén Ingefeldt and Thell, 2019). For those entering university, student housing temporarily fills the gap. Student housing is allocated through a separate student-specific queue, and in part because units cycle with each graduating cohort rather than being held long-term, it is substantially more accessible than the regular rental segment. After graduation, students lose access to this form of housing, the regular rental segment is limited, and returning to live with parents is less common in Sweden. Demand therefore arrives at the priced ownership segment in a tightly concentrated age band.

When this demand reaches the priced market, supply expands only slowly. Land-use regulation is decentralised, with each municipality holding exclusive authority over zoning and detailed development plans within its borders. Average zoning and permitting lead times for multi-dwelling projects rose by roughly a quarter, from 45 months in 2015 to 56 months in 2024, with the full process now taking three years in the fastest-moving municipalities and as much as six in the slowest (OECD, 2025). New construction of homes in Sweden has averaged only 0.6 percent of the existing housing stock per year over 1990 to 2018 (Olsén Ingefeldt and Thell, 2019), and the OECD documents that residential construction has responded to Swedish population growth with a lag, especially in the largest cities (OECD, 2025). In the locations where construction is permitted, buildable land is also scarce, which further limits the supply response. The value of land underlying Swedish housing has appreciated roughly 27-fold since 1980 (OECD, 2025). We therefore expect the demand of an arriving cohort to show up in prices before it shows up in additional units.

Owner-occupied housing in Sweden is also favoured by the tax system. Recurrent property tax revenues are among the lowest in the OECD, capped at a low flat rate, and mortgage interest payments are tax-deductible. OECD (2025) estimates that the tax system absorbs roughly half of the cost of mortgage-financed homeownership in Sweden, and places the marginal effective tax rate on debt-financed residential property among the most favourable in the OECD. Together with the regulated rental segment described above, these provisions concentrate demographic demand for housing in the priced ownership market.

These institutional features apply identically across Sweden's counties, on both the demand and supply sides. The mortgage market is regulated by Finansinspektionen, the financial supervisory authority, with a national loan-to-value cap and amortization rules. Tax policy is likewise set in national law, and roughly 80 percent of Swedish mortgages carry variable rates (OECD, 2025).

IV. Data

A. *Sources and sample*

We construct a county-level panel of 20 Swedish counties observed annually from 1996 to 2022 by merging seven SCB statistical products on common county and year identifiers: the Fastighetsprisindex (house price index), single-year-of-age population statistics, population-change statistics (for migration flows), median income by age group, granted registrations of title (lagfarter), completed dwellings, and bostadsrätter transaction data. The county is the most granular geography at which SCB reports a quality-adjusted house price index. Raw price statistics are available at the municipality level, but no published index goes below the county level. Kalmar and Gotland are merged throughout because SCB reports their house price index as a single combined series, leaving 20 distinct cross-sectional units.

The sample period is bounded by data availability on each end. SCB single-year-of-age population data begin in 1968, but the 5-year income growth control requires income data from 1991, which sets the first usable year at 1996. The end year is set by the centred price-change window, since prices two years ahead are needed for each observation, making 2022 the last year for which the dependent variable is observed. The resulting panel contains 540 county-year observations.

B. *Variables*

The dependent variable is the centred 5-year log change in SCB's real estate price index (Fastighetsprisindex), computed as $\log p_{i,t+2} - \log p_{i,t-3}$. The Fastighetsprisindex is a chain-linked Laspeyres index for permanently occupied single- and two-family houses, reported annually by county with base year 1990 = 100. Five-year changes reduce noise in annual county-level price movements, following Francke and Korevaar (2025). The centring ensures that the coefficient at lag L captures price growth around age L. The index is nominal, but year fixed effects absorb national inflation.

The main regressor is the birth share, the number of 0-year-olds in a county divided by total county population, lagged 27 years. The contemporaneous share of 27-year-olds in the county-year is used in the predetermined-versus-contemporaneous test reported in Appendix A.3.

The control variable in the main specification is the 5-year log change in median income for the 25–29 age group, available from SCB by county, age group, and year.

In robustness specifications we also include net internal and net international migration rates by county, drawn from SCB's population-change statistics and normalised by population. County-level migration data are available from 2000 onward, which reduces the subsample with migration controls to 460 observations.

Two additional outcomes follow how a demographic demand shock shows up beyond prices. The transaction rate (the per-capita number of granted registrations of title, lagfarter, for ordinary market purchases) tests whether demographic pressure shows up first in prices or in trading volume. The construction rate (the per-capita number of completed dwellings in one- or two-family buildings, småhus) tests how quickly new units enter the market in response to demand. Both series come from SCB and are normalised by total county population.

We also examine apartment prices as a separate test of the same demographic effect in a different price segment. SCB does not publish a county-level price index for apartments comparable to the Fastighetsprisindex, so we construct one from SCB's bostadsrätter transaction data, aggregating sub-regional median prices to the county level using a sales-weighted average and computing the same centred 5-year log change.

C. Data quality

Sweden's administrative registers provide the data quality this design requires. The Fastighetsprisindex is computed by SCB from Lantmäteriet's Fastighetsprisregister, which records every registered title transfer in Sweden (Statistics Sweden, 2026). The lagfarter (transactions) series comes from the same register. Universal coverage removes sampling-frame and selection concerns that affect listing-price or broker indices used elsewhere in the literature. Population data come from Folkbokföring, the residency register, with single-year-of-age and county-of-residence definitions stable since 1968. Income data come from Inkomst- och taxeringsregistret, the tax register. Coverage is the population of taxed residents rather than a survey sample. The migration series derives from the same population-register flows.

The variables themselves carry caveats. The construction outcome covers one- or two-dwelling buildings (småhus) only, not multi-dwelling apartment buildings. We read the construction series as the supply response in the single-family house market specifically. The apartment-price series we construct from SCB bostadsrätter data uses median transaction prices, which reflect both true price changes and shifts in the type of apartments sold in any given year. We treat the apartment results as an additional test rather than as a direct counterpart to the house price index.

D. *Birth shares by county*

Figure 1 traces the share of newborns by county over more than half a century. National fertility swings (peak in the early 1990s, decline through the late 1990s, partial recovery in the 2000s and early 2010s, renewed decline since the mid-2010s) dominate the visual movement, and persistent cross-county level differences are also visible, with Stockholm consistently above Norrbotten. Both kinds of variation are absorbed by the design described in Section V, leaving the residual cohort-specific fluctuation as the variation that identifies the coefficient.

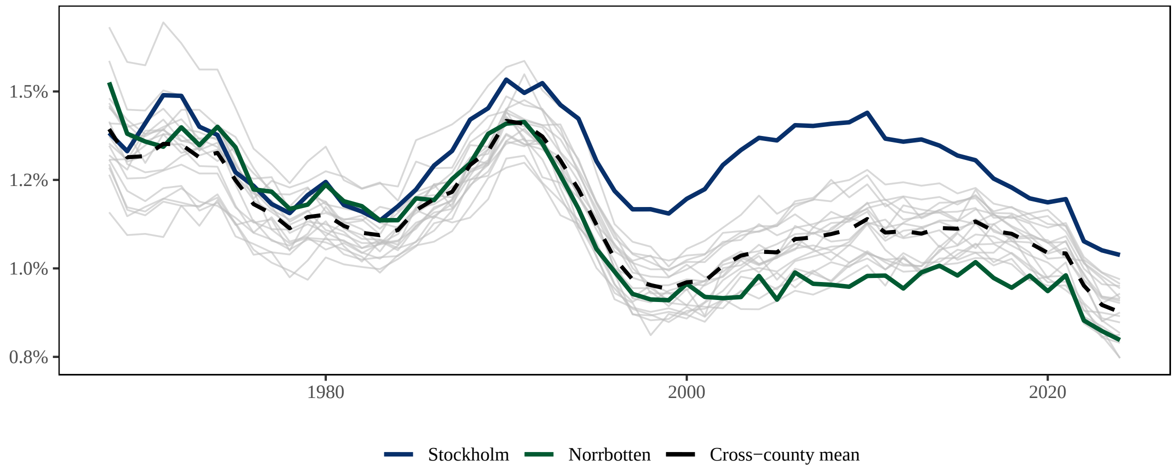

Figure 1. Share of newborns by county, 1968–2024. Each line shows newborns as a fraction of total county population. Stockholm (blue) and Norrbotten (green) are highlighted to illustrate the cross-county dispersion. The cross-county mean is overlaid as a dashed line. The 18 remaining counties are shown in grey. Data: SCB single-year-of-age population by county (Kalmar and Gotland combined).

E. *House prices by county*

Cumulative price growth varies sharply across counties, with Stockholm sitting at roughly five times its 1990 level by 2024, while Norrbotten is around three times. The bulk of this divergence accumulated between the late 1990s and the late 2000s, with cross-county dispersion roughly stable since. Long-run cross-county gaps of this kind are absorbed by the design. The dependent variable is the 5-year change rather than the cumulative level, and the residual within-county-within-year variation identifies the coefficient.

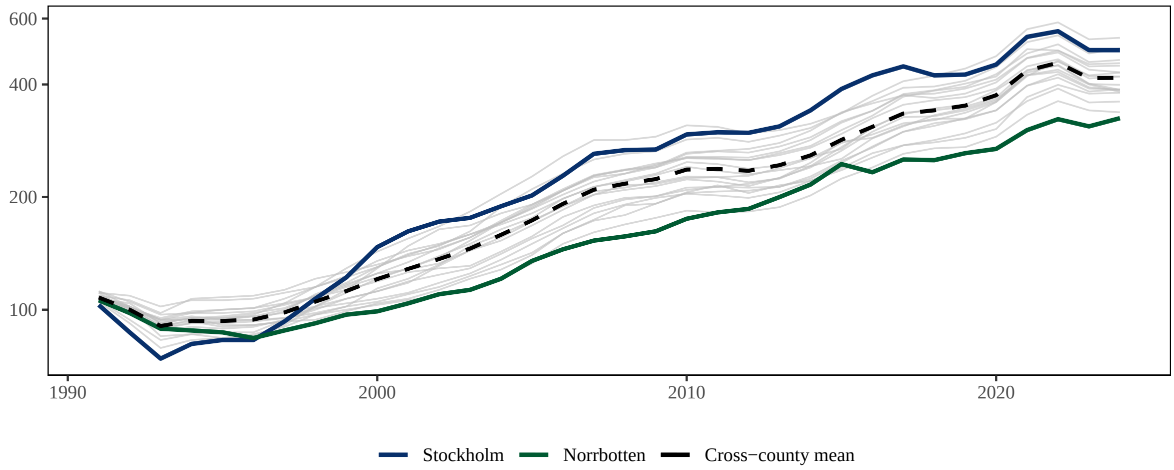

Figure 2. House price index by county, 1990–2024. SCB Fastighetsprisindex (1990 = 100) by county, plotted on a log scale so that proportional differences are visually comparable. Stockholm (blue) and Norrbotten (green) are highlighted. The cross-county mean is overlaid as a dashed line. The 18 remaining counties are shown in grey. Cumulative growth varies substantially across counties, with metropolitan counties growing fastest and rural northern counties slowest.

F. Net buying by age

Figure 3 reports net buying by age group as descriptive evidence about the age structure of housing demand. This series does not enter the regressions. The concentration of net entry in the late twenties and early thirties aligns with the institutional timeline described in Section III: Swedish young adults form independent households early, lose access to student housing at graduation, and concentrate ownership entry in a narrow age band.

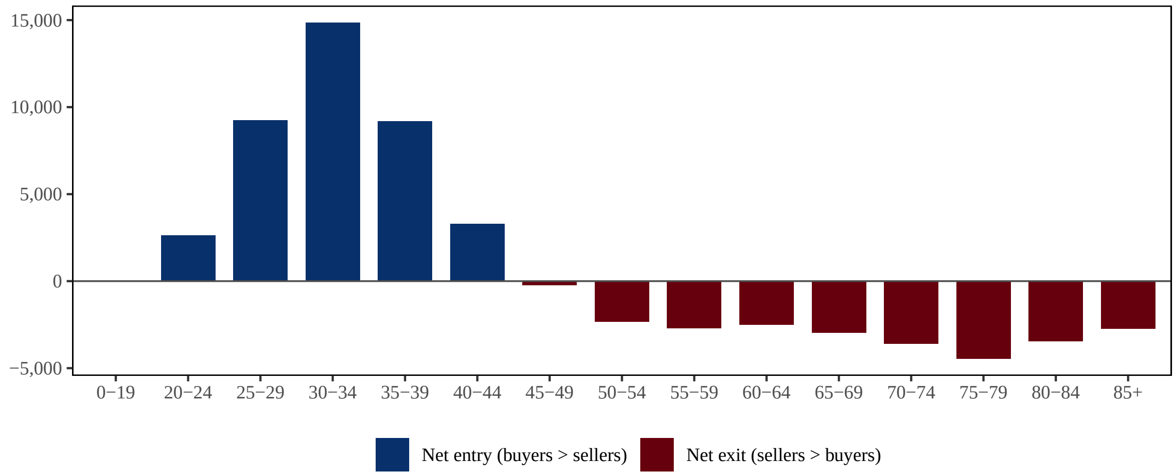

Figure 3. Net buying by age group in the Swedish housing market. For each five-year age bin, the bar shows the number of buyers minus the number of sellers. Positive values (blue) indicate net entry into homeownership. Negative values (red) indicate net exit. Subtracting sellers from buyers nets out within-stock churn and isolates net entry and exit at each age. Net entry peaks at ages 30–34 and turns into net exit around age 50, with the largest net exits among those over 60. Data: SCB lagfartsstatistik (granted registrations of title for ordinary market purchases).

G. Descriptive statistics

Several features of the panel motivate the design (Table I). The 5-year log price change ranges from 0.03 to 0.67 across county-years, with a standard deviation of 0.11. The regression decomposes that dispersion. The lagged birth share has a much narrower range (0.97% to 1.70%). Within any given year, the cross-county dispersion is around 0.09 percentage points, and this is the variation the year fixed effects isolate. The apartment-price series is roughly three times as noisy as the house price index in standard-deviation terms (0.34 versus 0.11), reflecting its construction from median transaction prices. This is why we treat the apartment results as a supplementary test.

Table I. Descriptive statistics. Summary statistics for the main estimation sample of 20 Swedish counties observed annually from 1996 to 2022. The 5-year log price change is centred on year t (running from $t-3$ to $t+2$). The lagged birth share is the share of 0-year-olds in the county 27 years prior to year t. Income growth is the 5-year log change in median earnings for the 25–29 age group. Transaction and construction rates are scaled per capita. The apartment price change and migration rate are reported on their narrower samples. SCB apartment-price data begin in 2000 and inter-county migration data begin in 2000.

Variable	N	Mean	SD	Min	Max
5-year log price change (centered)	540	0.2757	0.1118	0.0324	0.6688
Birth share (age 0), 27 years prior	540	0.0124	0.0013	0.0097	0.0170
5-year log income growth (25-29)	540	0.1318	0.0579	-0.0615	0.2653
Transactions per capita	480	0.0128	0.0032	0.0053	0.0283
New dwellings per capita	540	0.0009	0.0005	0.0001	0.0031
5-year log apt. price change (centered)	400	0.6077	0.3428	-0.0171	1.5353
Net internal migration rate	460	-0.0009	0.0030	-0.0103	0.0082

V. Methodology

A. *Specification*

We estimate how a county's lagged birth cohort size predicts subsequent house price growth on the panel of 20 Swedish counties observed annually from 1996 through 2022 (540 county-year observations). The dependent variable is the centred 5-year log change in the SCB house price index, $\Delta_5 \log p_{i,t} \equiv \log p_{i,t+2} - \log p_{i,t-3}$. The regressor is the share of newborns in county i in year $t-27$. The estimating equation is

$$\Delta_5 \log p_{i,t} \;=\; \beta\, B_{i,t-27} \;+\; \gamma\, \Delta_5 \log y_{i,t} \;+\; \alpha_i \;+\; \delta_t \;+\; \varepsilon_{i,t}$$

with $B_{i,t-27}$ the lagged birth share, $\Delta_5 \log y_{i,t}$ the 5-year log change in median income for the 25–29 age group ending at year t, and α_i and δ_t county and year fixed effects. The coefficient of interest is β, which scales the lagged birth share into the 5-year log price change. Because birth shares are measured as proportions (typically around 0.012), a one-percentage-point increase in the share corresponds to a $0.01\,\beta$ change in the 5-year log price. Income enters as a control to absorb county-specific changes in regional prosperity that would independently affect house prices, with 25–29 income growth as a proxy for the broader desirability shifts that would otherwise contaminate the demographic estimate. The birth share and income growth do not move together within the panel. The variance-inflation factor between them is 1.00, the lowest possible value (the conventional concern threshold is 5).

B. *Identification*

The two-way fixed-effects design absorbs two distinct sources of variation. County fixed effects absorb time-invariant regional differences in urbanisation, industrial structure, geography, and baseline housing demand: whatever drives Stockholm's housing market to grow persistently faster than Norrbotten's over decades drops out before β is computed. Year fixed effects absorb nationwide variation in macroeconomic and demographic conditions: the 2008 financial crisis or the post-2022 monetary tightening enter the regression only insofar as their effects differ across counties. County and year fixed effects leave only the within-county, within-year deviation of birth shares as identifying variation for β. County-specific fertility timing differences 27 years ago still show up today, and we test whether they predict county-specific deviations from the national price trend in the present.

The regressor is predetermined. Birth shares are fixed in the historical record 27 years before the price outcomes we measure. This rules out reverse causality (current price growth cannot affect 27-year-old birth records) and any year-t shock that would simultaneously move the regressor and the outcome. It does not rule out slow-moving county-specific processes that span the 27-year lag and could affect both fertility then and prices now. For this reason we use the share of newborns rather than the contemporaneous share of 27-year-olds. Birth-time geography is fixed in the historical record, while the current share of 27-year-olds has been reshaped by 27 years of inter-county migration that itself responds to local economic and housing-market conditions. We test the migration channel directly by including net internal and international migration rates as controls in robustness specifications, which reduces the sample to 460 observations because SCB inter-county migration data begin in 2000.

The age profile re-estimates the regression at every lag from 20 to 38 (Section VI.B). The shape of the resulting coefficient curve provides a credibility check on the design. If the relationship between lagged birth shares and house prices ran through a lurking variable correlated with both, the coefficient would be similar at every lag. The observed pattern, with the coefficient concentrated near the buying age and small or insignificant at lags away from it, is consistent with the demographic mechanism and inconsistent with confounders that would produce a flat coefficient profile across lags. The age profile also determines the main specification's choice of lag, namely lag 27, the peak of the curve. The cohort observation age analysis (Section VI.C) holds the target cohort fixed (those who are 27 in year t) and varies the age at which we measure their county share. The way predictive power changes with observation age is what we would expect if life-cycle migration weakens the demographic signal over time. A formal regression in Appendix A.3 completes this suite by including both the lagged birth share and the contemporaneous share of 27-year-olds in the same equation, confirming that the lagged measure remains significant while the contemporaneous measure is absorbed.

Some potential confounders survive the two-way fixed-effects design. The threats take a common form, a county-time-varying process correlated with both the lagged birth share and current price growth. Such processes range from a single shock around the time the cohort was born to slowly evolving county-specific trajectories that span the 27-year lag. Regional labour-market booms, broader changes in county prosperity, sustained immigration patterns, and university expansions would each fit this template. Our specifications address such threats directly. Excluding the three counties containing Sweden's largest cities (Appendix A.1) tests for drivers specific to metropolitan counties. Adding migration rates as controls (Appendix A.2) tests whether migration flows mediate the result. The income control in the main specification absorbs the income-growth component of these prosperity channels. Francke and Korevaar (2025) address similar threats in their main specification by adding controls lagged to match the birth-rate lag. SCB income data begin only in 1991, which precludes adding lagged income aligned to our 27-year horizon for our panel.

C. *Standard errors*

The centred 5-year price-change windows overlap across consecutive years, inducing a moving-average autocorrelation of order four in the residuals. We report Newey-West standard errors with bandwidth four, the order of that induced moving-average structure. As a check on this correction, Appendix A.4 re-estimates the main specification on a non-overlapping subsample where the autocorrelation is removed by construction. The demographic effect remains significant at the 5% level under cluster-robust standard errors. We do not cluster standard errors by county. With only 20 counties, cluster-robust standard errors are below the threshold at which they produce reliable test rejection rates, and the Newey-West correction handles the within-county serial dependence that clustering would otherwise capture. We also do not weight by population. The unweighted estimator gives β the interpretation of an average within-county relationship across the 20 counties. Weighting by population would give Stockholm, Västra Götaland, and Skåne disproportionate influence and pull β toward the relationship that holds in those metropolitan markets specifically.

VI. Analysis

A. *Main result*

The lagged birth share predicts house price growth across the panel. The coefficient is 36.2 in the specification with county and year fixed effects only and 35.5 once 25–29 income growth is added as a control, both significant at the 1% level (Table II). Adding the income control changes neither the magnitude materially nor the statistical significance. Because the regressor is predetermined (Section V.B), this association reflects historical fertility patterns rather than current demand shocks.

A one-standard-deviation increase in the lagged birth share (0.135 percentage points) is associated with approximately 4.7 percentage points higher cumulative 5-year house price growth in our main specification, or about 1 percentage point per year. Mean 5-year price growth in the sample is about 27 percent, so a one-standard-deviation demographic shift covers roughly a sixth of typical 5-year price growth. The implied size of the effect is sensitive to the sample window, as the migration subsample in Appendix A.2 and the discussion in Section VII make clear. Cook's distance flags 41 of 540 county-years as having unusual influence on the estimate. The coefficient rises slightly when we refit after dropping them (from 35.5 to 38.0), so these observations are not producing the result. If anything, they are slightly pulling it down.

The closest external benchmark is the OECD-panel evidence in the Internet Appendix of Francke and Korevaar (2025), which regresses changes in house prices on lagged changes in demographic shares across 18 OECD countries from 1970 to 2020. Their headline regressor is the 20-year lag of the change in the under-15 share, which they interpret as the cohort entering the housing market at an average age of about 28. Their main specification finds a positive and significant coefficient on this regressor for house prices. The pattern and the implied buyer-age window match ours. Direct point-estimate comparison is uninformative because the regressors are constructed differently, and the authors themselves caution against it. Our specification differs in the source of variation, with their identification resting on cross-country differences in demographic timing, while the year fixed effect in our equation absorbs every national-level driver of Swedish house prices

Table II. Birth cohort size and house price growth. Two-way fixed-effects OLS regressions of the centred five-year log house price change on the lagged birth share. Column 1 includes the birth share at lag 27 only. Column 2 adds five-year log income growth (25–29) as a control. Both specifications absorb county and year fixed effects. Standard errors are Newey-West with bandwidth four. Sample: 540 county-year observations, 20 Swedish counties, 1996–2022.

Dependent Variable:	5-yr log price change	
Model:	(1)	(2)
Variables		
Birth share (age 0, t-27)	36.22***	35.51***
	(10.05)	(9.930)
5-yr log income growth (25-29)		0.3743**
		(0.1740)
Fixed-effects		
County	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	540	540
R^2	0.71200	0.72092
Within R^2	0.07129	0.10003

Newey-West (L=4) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

and leaves only within-Sweden, between-county variation to identify β.

We add the 25–29 income control to absorb any contemporaneous shift in regional desirability that would independently raise local prices. If a county becomes more attractive through for example new employers, that shift would typically raise both local incomes for prime home-buying ages and local house prices. The coefficient barely changes between columns 1 and 2 of Table II, indicating that this story is not driving the result.

B. Age profile

We chose lag 27 because the coefficient peaks there, found by re-estimating the same regression at every lag from 20 to 38 and reading off the maximum. The resulting curve, plotted in Figure 4, does two things: it picks out lag 27 as the peak, and its shape provides our most informative identification check.

The coefficient is negative or insignificant at lags 20–23, turns positive at lag 24, rises sharply to peak at 27, then declines through the early thirties before returning to zero around lag 35. The 25–30 band is broadly elevated, but the rise into the peak is sharper than the decline away from it. Estimates at lags beyond 35 lose precision as the sample shrinks, since SCB single-year-of-age population data begin in 1968. The centred five-year window smooths the year-to-year variation but does not shift the location of the peak. This is shown by re-estimating the age profile on annual price changes, which again confirms the peak at age 27 (Appendix A.5).

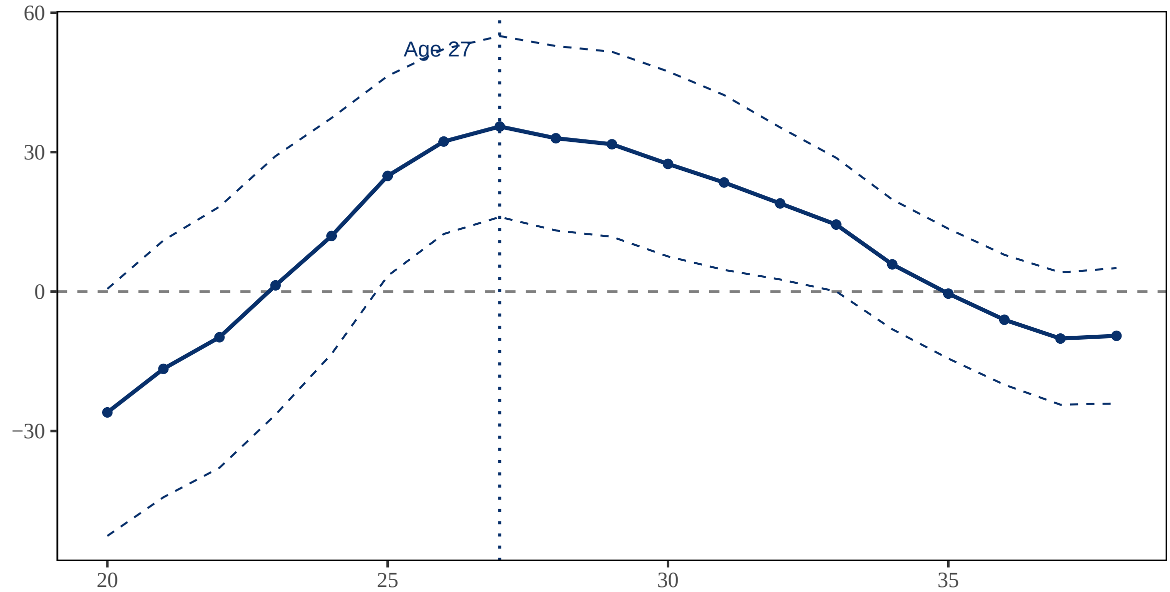

Figure 4. Age profile of cohort effects on house prices. Each point is the OLS coefficient on the birth share at lag L from a separate regression of the centred five-year log house price change on the lagged birth share, controlling for five-year log income growth (25–29) and absorbing county and year fixed effects. The horizontal axis is the lag in years, equal to the approximate age of the cohort at the time of observation. The shaded band is the 95% confidence interval based on Newey-West standard errors with bandwidth four. The coefficient peaks sharply at age 27 and declines through the early thirties. Sample: 20 Swedish counties, 1991–2024 ($N = 540$ at the peak lag).

If the result reflected slowly evolving regional fundamentals such as productivity trends or persistent migration that happened to correlate with cohort sizes 27 years earlier, the coefficient would be roughly flat across lags, since such fundamentals do not localise on a particular birth year. Instead, the effect concentrates in a five- to six-year window matching the housing-demand life cycle, and is small or insignificant outside it. Lags far from buying age serve as placebo lags, and the absence of an effect there argues against confounders that would produce a flat profile.

The location of the peak fits the institutional timeline of Section III. Swedish young adults leave the parental household earlier than peers in most of Europe (21.9 years on average, Eurostat 2024) but face a queue-rationed and shrinking rental segment, making the priced ownership market the main feasible route into independent household formation. Student housing serves as a temporary bridge for those entering university but ends with the loss of student status. Entry into ownership in Sweden therefore happens mostly in the mid-to-late twenties, which matches the sharp rise in the coefficient between lags 24 and 27.

The peak's location is consistent with prior evidence on age and housing demand. Francke and Korevaar (2025) identify a 25–29 year birth-rate lag as the strongest predictor of rent-price ratios and house prices in their 1550–1884 Amsterdam sample, sitting adjacent to our lag 27. Mankiw and Weil (1989), using US Census data, document that age-specific housing demand rises sharply between ages 20 and 30 and is approximately flat thereafter, consistent with the rising portion of our age profile.

C. Cohort observation age

The same cohort can be measured at many different ages, from birth to the present. Mankiw and Weil (1989) use the contemporaneous age structure, weighting each age group by its housing demand. Francke and Korevaar (2025) use lagged birth rates, arguing they are determined long before any housing-market outcome and so cannot be affected by it. The OECD-panel evidence in their Internet Appendix uses the lagged share of children under 15. This span aggregates births with several years of subsequent migration, which motivates the decomposition by age we run here. The choice matters, because measuring the cohort at different ages picks up different amounts of migration.

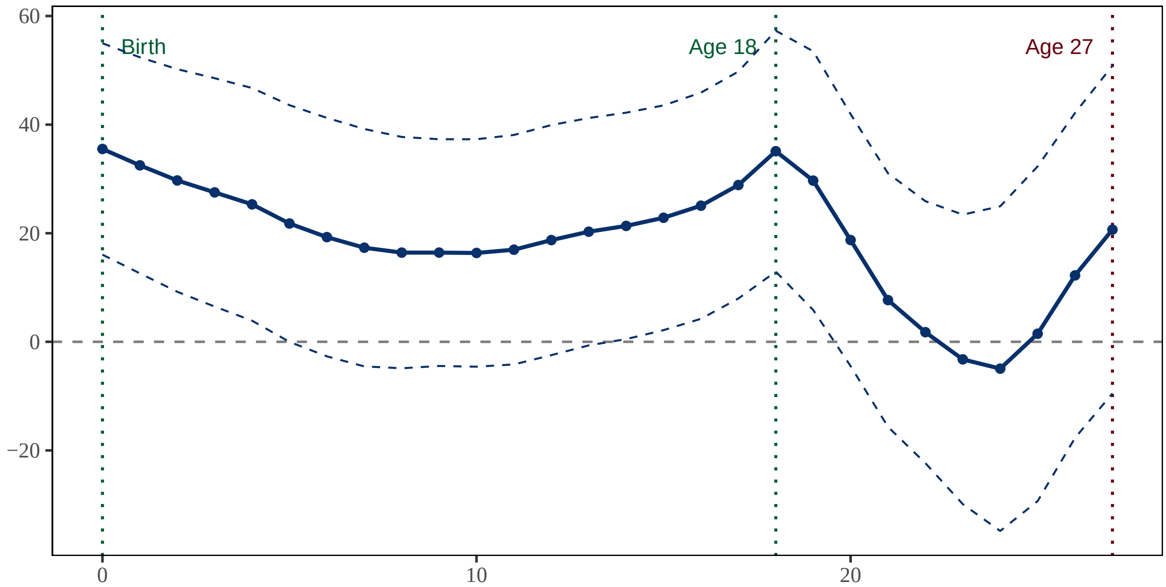

Figure 5. Effect on house prices by age at which the cohort is observed. Each point is the OLS coefficient from a separate regression in which the same cohort, those who are 27 in year t, is measured as a county population share at a different age. The dependent variable is the centred five-year log house price change in year t. Regressions control for five-year log income growth (25–29) and absorb county and year fixed effects. Vertical lines mark birth ($a = 0$), upper secondary school completion ($a = 18$), and the contemporaneous endpoint ($a = 27$). Shaded bands are 95% confidence intervals (Newey-West, bandwidth four). The coefficient peaks at birth and again at age 18 with essentially equal magnitude, falls in middle childhood, and drops sharply during the university years.

We re-run the main regression once for every age from 0 to 27, each time using the share of the cohort that is currently age a and was born in year $t - 27$ (Figure 5). At observation age 0, the regressor is the share of newborns 27 years before t, which is our main specification. At observation age 27, it is the contemporaneous share of 27-year-olds in year t. The intermediate values measure the same cohort at school age, late adolescence, university age, and so on. The dependent variable, controls, and fixed effects are unchanged across the 28 regressions.

The coefficient is large and statistically significant at observation age 0 ($\beta = 35.5$, $p < 0.001$) and declines steadily through the first six years of life. Through middle childhood (ages 7–13), it sits in the range 16–20 and is at most marginally significant, with t-statistics between roughly 1.5 and 1.9. From age 14 the coefficient rises through adolescence and peaks at age 18 ($\beta = 35.1$, $p = 0.002$), where it is essentially equal in magnitude to the coefficient at birth. After 18 it drops sharply, falls outside the 5% significance band by age 20, and then turns slightly negative around ages 23–24, to recover only partially by age 27, where the contemporaneous coefficient is positive but insignificant ($\beta = 20.7$, $p = 0.18$). The profile has two local maxima of comparable size, at 0 and 18, separated by a childhood trough.

The figure's shape reflects how Swedish residential mobility changes over the life cycle. One reading is that the cohort share at age a tracks price growth to the extent that the cohort's geography at that age anticipates where its members will eventually buy housing. The high coefficient at birth fits this reading, since the cohort is measured before later migration has reshaped its geography. A decline through early childhood would follow if some families relocate during the children's first years. The recovery to near the value at birth by age 18 is consistent with families that relocated having largely returned by upper secondary school age. After graduation, young adults disperse for university and early-career moves, and through the early twenties the cohort share appears to contain mostly migration noise. As the cohort approaches buying age, the coefficient moves back toward positive values without fully recovering, which fits a story in which some young adults return to settle in their birth county. The contemporaneous share at age 27 includes both these returners and later in-migrants whose moves likely respond to contemporaneous economic conditions. This may be one reason the contemporaneous measure is a weaker predictor than the birth-time measure.

The cohort share at birth dominates any measurement closer to the present. The closest competitor is the share at age 18, which has essentially the same magnitude but a lower t-statistic (3.10 versus 3.58) and a forecast horizon of nine years rather than twenty-seven. The economic case for using births rather than the contemporaneous age structure is the one Francke and Korevaar (2025) make conceptually, since the contemporaneous count of 27-year-olds in a county reflects 27 years of migration that itself responds to local economic conditions, while births 27 years earlier are predetermined. County-level Swedish data lets us test this directly by measuring the same cohort at every age from birth through buying age. A regression including both the birth-time and contemporaneous measures (Appendix A.3) confirms this in a single specification. When both are included, the contemporaneous measure is absorbed by the birth-time measure.

D. Mechanisms

The age profile shows when prices respond to demographic demand. Transactions provide an independent check on the demand-driven reading. Prices could move for reasons other than demographic demand, but if demand from the arriving cohort is the mechanism, transaction volumes should also respond to the same lagged birth-share regressor. Construction, which adjusts most slowly of the three outcomes, indicates the supply rigidity that lets demand pressure show up first as prices.

D.1. Transactions

Re-estimating the age profile with the per-capita transaction rate as the dependent variable (Figure 6), we find the coefficient is positive but small and not statistically significant at lag 27, where prices peak. It rises through the late twenties, becomes marginally significant around lag 28, and peaks at lag 30, where it is significant at the

1% level.

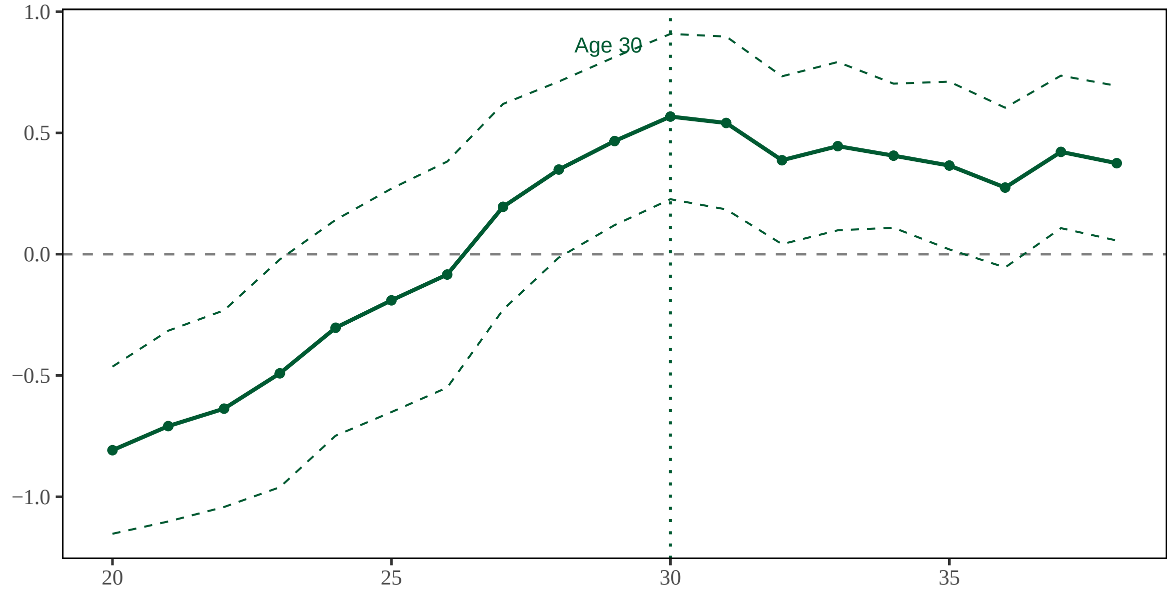

Figure 6. Age profile of cohort effects on transaction volume. Each point is the OLS coefficient on the birth share at lag L from a separate regression of the per-capita transaction rate (granted titles of transfer per resident) on the lagged birth share, controlling for five-year log income growth (25–29) and absorbing county and year fixed effects. The horizontal axis is the lag in years. The shaded band is the 95% confidence interval based on Newey-West standard errors with bandwidth four. Unlike the price profile, the transaction effect rises with age and remains elevated through the late thirties. Sample: 20 Swedish counties, 1991–2024.

The three-year delay between the price peak and the transaction peak is consistent with a supply-constrained market in which prices respond to demand shocks faster than the stock or transaction volume can adjust (Glaeser and Gyourko, 2005; Saiz, 2010). The transaction peak at lag 30 also matches the descriptive evidence in Section IV.F, where net entry into homeownership peaks at ages 30–34. Under a reading in which supply is rigid, demand pressure from a large arriving cohort pushes prices up first, while the stock of homes available for sale does not expand at the same speed. As prices rise, more sellers come to market, and transaction volumes catch up.

Francke and Korevaar (2025) find a similar delay in their historical Amsterdam data, where the probability of sale responds to lagged birth cohorts at horizons beyond the price peak. The lag values are not directly comparable to ours, since their regressor uses five-year birth rates rather than single-year birth shares, but the directional pattern is the same, with prices leading the transaction response.

After the peak, the transaction profile and the price profile diverge. Prices capture a one-time demand shock, and once the cohort is housed, upward pressure on prices fades. Transactions capture ongoing market turnover, which the same cohort generates as it moves between homes within the county. The transaction coefficient therefore stays elevated through the late thirties, while the price coefficient declines after its peak.

D.2. Construction

Re-estimating the age profile with completed dwellings per resident as the dependent variable (Figure 7), we find the construction response emerges only when the cohort reaches buying age. The coefficient turns positive at lag 28, becomes significant by lag 30, and peaks at lag 32 ($\beta = 0.097$, $p < 0.01$), after both the price and transaction peaks.

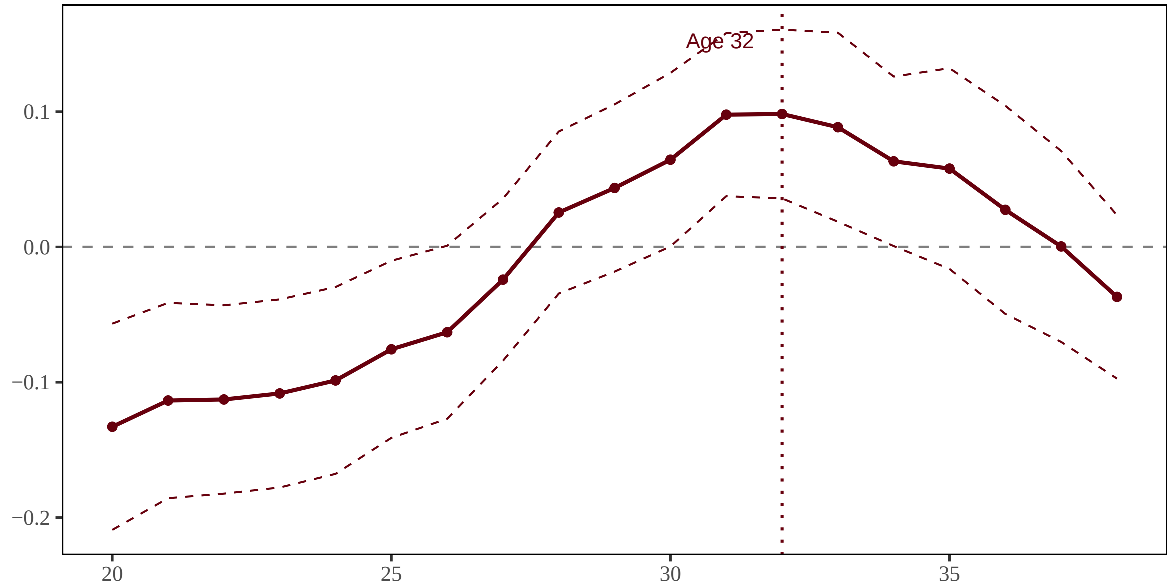

Figure 7. Age profile of cohort effects on residential construction. Each point is the OLS coefficient on the birth share at lag L from a separate regression of the per-capita construction rate (completed dwellings per resident) on the lagged birth share, controlling for five-year log income growth (25–29) and absorbing county and year fixed effects. The shaded band is the 95% confidence interval based on Newey-West standard errors with bandwidth four. The construction response peaks at age 32, five years after the price peak at 27. Sample: 20 Swedish counties, 1991–2024.

The five-year delay between the price peak and the construction peak is consistent with Sweden's permitting and building environment. The dependent variable measures completed dwellings, so a developer responding to demand pressure must first clear municipal zoning and permitting and then proceed through construction itself before the unit appears in the data. Combined zoning and permitting has averaged 45 months in 2015 and 56 months in 2024 for multi-dwelling projects, with the full process taking around three years in the fastest-moving municipalities and around six in the slowest (OECD, 2025). The observed five-year price-to-completion lag falls within this range.

Even with permits in hand, the supply response is small in magnitude. New construction has averaged only 0.6% of the existing housing stock per year in Sweden (Olsén Ingefeldt and Thell, 2019). The combination of slow timing and small magnitude places Sweden, in the standard housing-supply framework, in the case of low supply elasticity. When supply cannot adjust quickly, demand shocks fall mostly on prices rather than quantities (Saiz, 2010), consistent with the asymmetry between prices and quantities that follows from housing durability, as Glaeser and Gyourko (2005) derive in their kinked supply model.

D.3. Combined age profiles

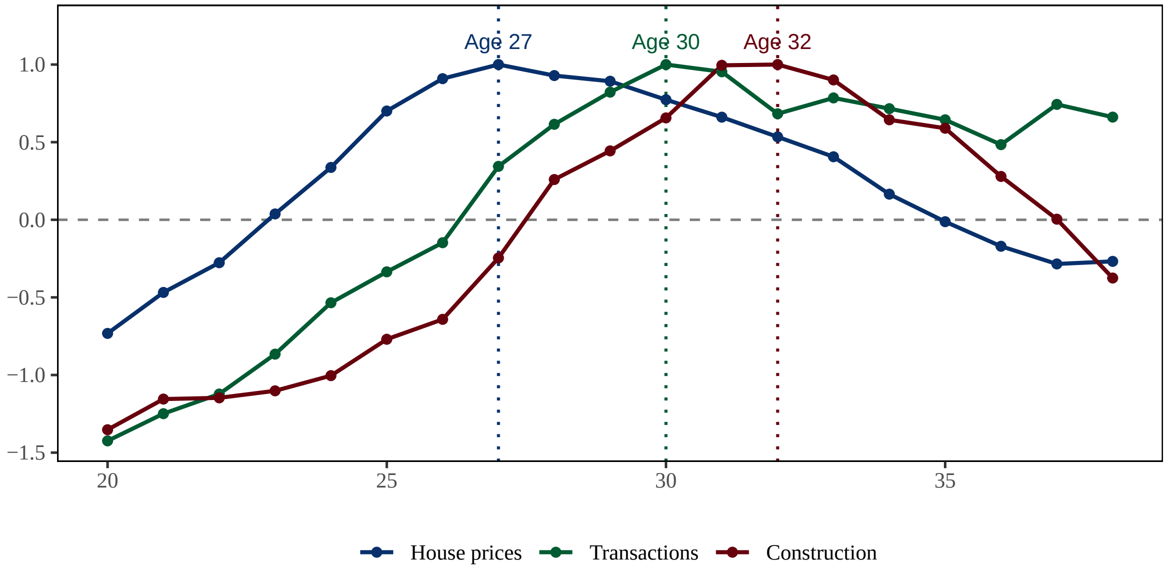

Figure 8. Timing of demographic effects across outcomes. Overlay of three age-profile coefficient curves, one per outcome, rescaled so that each series's peak equals one. The dependent variables are the centred five-year log house price change (blue), the per-capita transaction rate (green), and the per-capita construction rate (red). Each point is the OLS coefficient on the birth share at lag L, with county and year fixed effects, controlling for five-year log income growth (25–29). The vertical dotted lines mark each series's peak. House prices respond first (age 27), followed by transactions (age 30), then new construction (age 32). Sample: 20 Swedish counties, 1991–2024.

The combined profile in Figure 8 shows the cohort effect through three sequential outcomes. Prices respond first at lag 27, transactions follow at lag 30, and construction peaks at lag 32. The five-year span between the price peak and the construction peak is consistent with a housing market in which short-run supply is inelastic. As the limitations section notes, the same gap is also consistent with the physical time required to permit and complete new buildings. If supply is rigid, demographic demand shows up in prices first, then in transaction volumes as more sellers come to market, and finally in completed dwellings once the supply pipeline catches up.

The different shapes of the three profiles are an identification check beyond the timing alignment. A single confounder correlated with lagged birth shares would have to produce these distinct shapes simultaneously: humped responses in prices and construction, sustained elevation in transactions. The demographic mechanism is consistent with each shape arising from a separate process. An entry shock fits the price profile, ongoing turnover fits the transaction profile, and a delayed supply response fits the construction profile.

Francke and Korevaar (2025) present suggestive evidence of a similarly delayed supply response in historical Amsterdam, with construction responding to birth cohorts 25–34 years earlier. The Swedish pattern shares the same interpretation, with prices leading the supply response, consistent with a housing market that adjusts slowly to demographic demand shocks.

E. Apartments

The cohort effect extends to the apartment market and peaks at an earlier lag. Re-estimating the main specification with apartment prices as the dependent variable, the coefficient on the lagged birth share peaks at lag 24, three years before the house-price peak at lag 27 (Table III). Both peak coefficients are statistically significant at the 1% level. The magnitudes are not directly comparable across the two price series.

Table III. Demographic effects on house prices versus apartment prices. Column 1 reproduces the main specification on house prices using the birth share at lag 27 (the house-price peak). Column 2 estimates the same specification on apartment prices using the birth share at lag 24, the lag at which the apartment-price response peaks. Apartment prices are noisier median transaction prices from SCB's bostadsrätter data. The apartment sample begins in 2000. Both specifications absorb county and year fixed effects and use Newey-West standard errors (bandwidth four).

Dependent Variable:	5-yr price change	
	Houses (lag 27)	Apartments (lag 24)
Model:	(1)	(2)
Variables		
Birth share (age 0, t-27)	35.51***	
	(9.930)	
Birth share (age 0, t-24)		85.78***
		(25.77)
5-yr log income growth (25-29)	0.3743**	0.4461
	(0.1740)	(0.3817)
Fixed-effects		
County	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	540	400
R^2	0.72092	0.83732
Within R^2	0.10003	0.05670

Newey-West (L=4) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The three-year gap is consistent with a housing ladder embedded in Swedish institutional structure. With the rental segment queue-rationed and shrinking (Section III), bostadsrätter (priced apartments) serve as the entry segment for young adults forming independent households. The apartment peak at lag 24 corresponds to this entry, while the house peak at lag 27 corresponds to the subsequent upgrade as families form.

Francke and Korevaar (2025) document a similar within-market segmentation in their historical Amsterdam data, though using a different proxy. Their street-level analysis splits Amsterdam by homeownership share: streets with high rental shares respond to younger cohorts, while streets with high homeownership respond at home-buying ages. Our split is by housing type rather than neighbourhood composition. Both proxies capture the same underlying logic that different segments of the housing market respond at different cohort ages.

The pattern across both segments adds an identification check to the main result. The two effects come from completely different data sources: the Fastighetsprisindex for houses (a quality-adjusted index) and bostadsrätter median transaction prices for apartments. A confounder unrelated to demographics would have to produce both the house effect at lag 27 and the apartment effect at lag 24 simultaneously, in two independently constructed datasets, at lags that map onto each segment's typical role in the housing ladder.

VII. Discussion

A. Predictability and anticipation

SCB has reported population data at the county level since 1968. For decades there has been publicly available information that seems to predict house price growth. If buyers acted on this signal, they would enter the market before prices peaked and the price response would flatten. If builders acted on it, they would break ground earlier and meet some of the demographic demand with new supply. Yet the birth share twenty-seven years earlier is positively associated with house price growth, and the construction response peaks five years later.

On the demand side, households generally buy homes to live in rather than to time the market. A 22-year-old who anticipated that prices would rise at 27 would have limited ability to act on that information. Life-cycle stages such as university graduation and employment are not easily moved earlier, and the national loan-to-value cap and amortization rules (Section III) limit the ability of younger applicants to qualify for mortgages. Olsén Ingefeldt and Thell (2019) show that among young adults buying homes in Sweden, the 25–30 age bracket accounts for the largest share and has been growing, supporting the view that households tend to arrive at the market when they need housing rather than when prices are most attractive.

On the supply side, construction peaks about five years after the price peak. This lag is consistent with developers reacting to observed prices rather than to demographic forecasts, since the cohort approaching buying age is in principle observable years in advance, yet the construction response only appears once prices have already moved. The structural barriers described in Section III likely mute the supply response. Even so, we do observe a clear construction response, just one that arrives late. A developer who anticipated the demographic signal could break ground earlier and sell while prices are still rising rather than after the peak. A possible explanation is a genuine information gap, in which demographic data is not among the signals developers weigh. DellaVigna and Pollet (2007) document a similar pattern in equity markets, where demographic demand shifts predict abnormal returns in age-sensitive industries. They attribute the surviving predictability to investors who neglect demographic information at horizons of more than several years ahead.

B. Cohort observation age and migration

Our findings suggest that for predicting future housing demand, older demographic data can be more useful than recent data. Cohort size at birth is fixed when it is recorded.

Cohort size today has been reshaped by 27 years of migration that itself responds to local economic conditions, including the prices we are trying to explain. Francke and Korevaar (2025) argue for lagged birth rates on identification grounds. Our age-by-age comparison shows this is also empirically the right call. The contemporaneous coefficient falls below conventional significance, while measurement at birth gives the strongest signal.

What appears as in-migration to a county at buying age is partly return migration, with people moving back to the county where they were born. The dip in our W-shape during university age, followed by a partial recovery by buying age, fits this pattern. Return migration is well-known in demography, but it also has an implication for housing forecasts. Using current population flows alone treats every move as new, while birth records already capture the share of future demand that returns home. Our migration robustness check in Appendix A.2 supports this reading. Controlling for internal, international, and total migration rates leaves the birth-share coefficient essentially unchanged, suggesting that much of the migration relevant to future housing demand is already determined by historical birth records.

C. External validity

Sweden may be particularly well suited for identifying this effect. The design we use exploits variation in birth rates across counties. The methodology itself could apply in any country with sufficient regional variation, but two features of the Swedish setting support it further. First, SCB has published single-year-of-age population data by county continuously since 1968, with consistent county definitions throughout the panel, giving a long and uniform series that is not always available elsewhere. Second, Sweden's low international migration supports the predetermined-cohort identification, as Monnet and Wolf (2017) find in a cross-country panel of 20 OECD countries, where lagged cohort variables predict current age-group size strongly in low-migration countries (Sweden among them) and weakly in high-migration ones.

The same Swedish features that enable identification may also make the effect larger than it would be in other settings. Living with parents is less common among Swedish young adults than elsewhere in Europe, and the regulated rental sector is allocated by queues that run for years (Section III). When the cohort reaches buying age there is little capacity to take demand into rental tenure or back into the parental household, so demand concentrates into a narrow age band of priced ownership purchases. In countries with more accessible rental markets or more common parental cohabitation, a comparable demographic shock would be partly taken up elsewhere. Young adults could remain with parents, rent for longer, or wait. The supply side is similarly constrained (Section III), so the demand that does appear goes mostly into prices rather than into new construction. Thus, the magnitudes we estimate are likely higher than what would be observed in less constrained settings, although the underlying mechanism should still operate elsewhere. The timing pattern (prices first, then transactions, then construction) is likely more robust across settings than the magnitudes, since the ordering depends only on supply being slow to expand, while the specific lags depend on Sweden's permitting timelines and patterns of household formation.

The apartment evidence in Section VI.E shows the same dynamic in tenant-owned apartments (bostadsrätter), with the cohort effect peaking three years earlier than for houses, reflecting the typically younger buyer age in the apartment segment. The same demographic pattern thus appears in both priced segments of the Swedish ownership market. The mechanism is therefore not specific to one type of housing, and the same logic could extend to other asset types whose demand is concentrated in a specific age range.

D. *Limitations*

The headline coefficient is sensitive to the sample period. On the full 1996–2022 panel the birth share at lag 27 produces a coefficient of 35.51. On the 2000–2022 subsample used for the migration robustness check (Appendix A.2), the same specification produces 21.61. The effect remains statistically significant in both samples, but the magnitude is not stable. The magnitude estimate therefore carries more uncertainty than a single specification's standard error reflects, and is likely to differ on a different sample.

We interpret the timing patterns in our data as a sequence in which cohort entry drives demand, demand raises prices, transactions follow, and supply responds last. We document the timings between the peaks but do not formally test this sequence. The five-year gap between the price peak and the construction peak, for example, could reflect developers responding slowly to the demand signal, or it could reflect the physical time required to plan and complete new dwellings, which is on a similar scale. Similarly, the W-shape in cohort observation age is suggestive of return migration but cannot demonstrate it directly. Our data identify when each variable peaks but cannot provide evidence for the underlying mechanisms.

Two further data constraints bound the scope of our claims. Aggregation is at the county level, so within-county variation, especially within metropolitan counties, is averaged out. And the population series begins in 1968, which is too recent to support the roughly 60-year lag needed to test whether older cohorts entering retirement depress prices, so the exit side of the demographic life cycle is untested.

VIII. Conclusion

We asked whether past birth cohorts predict regional house price growth in Sweden, and our findings suggest that they do. Birth shares twenty-seven years prior predict five-year house price growth at the county level, and the same demographic shift propagates through apartment prices, transactions, and construction in a sequence that places prices ahead of quantities.

The cohort sizes that move prices have been publicly observable since 1968, yet the response only appears when the cohort reaches home-buying age. Neither buyers nor builders price the cohort in advance. Sweden's institutional setup, with a queue-rationed rental sector and slow permitting, likely sharpens the pattern. The underlying logic, that ownership entry concentrates by age and supply expands slowly, is not specific to Sweden.

Measuring the same cohort at every age from birth through buying age, we document that birth-time measurement carries the strongest signal in our regressions, while the contemporaneous measure falls below significance. For studies of demographic effects on housing in other settings, where regional birth records exist, they are preferable to current age structure.

The size of birth cohorts recorded in Sweden during the late 1990s and 2000s is observable today, and if the relationship in our sample continues to hold, this signal is informative about regional housing pressure in the late 2020s and beyond.

REFERENCES

Abel, Andrew B., 2003, The effects of a baby boom on stock prices and capital accumulation in the presence of social security, *Econometrica* 71, 551–578.

Bostadsförmedlingen i Stockholm, 2025, Hur lång tid tar det? [how long does it take?], Accessed March 2026.

DellaVigna, Stefano, and Joshua M. Pollet, 2007, Demographics and industry returns, *American Economic Review* 97, 1667–1702.

Engelhardt, Gary V., and James M. Poterba, 1991, House prices and demographic change: Canadian evidence, *Regional Science and Urban Economics* 21, 539–546.

Eurostat, 2024, Estimated average age of young persons leaving the parental household by sex, Dataset code: yth_demo_030. Accessed May 2026.

Francke, Marc, and Matthijs Korevaar, 2025, Baby booms and asset booms: Demographic change and the housing market, *Journal of Finance* 80, 3021–3056.

Glaeser, Edward L., and Joseph Gyourko, 2005, Urban decline and durable housing, *Journal of Political Economy* 113, 345–375.

Mankiw, N. Gregory, and David N. Weil, 1989, The baby boom, the baby bust, and the housing market, *Regional Science and Urban Economics* 19, 235–258.

Monnet, Eric, and Clara Wolf, 2017, Demographic cycle, migration and housing investment, *Journal of Housing Economics* 38, 38–49.

OECD, 2025, *OECD Economic Surveys: Sweden 2025* (OECD Publishing, Paris).

Olsén Ingefeldt, Niclas, and Viktor Thell, 2019, Young adults and the housing market, FI Analysis 19, Finansinspektionen.

Poterba, James M., 2001, Demographic structure and asset returns, *Review of Economics and Statistics* 83, 565–584.

Saiz, Albert, 2010, The geographic determinants of housing supply, *Quarterly Journal of Economics* 125, 1253–1296.

Statistics Sweden, 2026, Kvalitetsdeklaration: Fastighetspriser och lagfarter [Quality declaration: Real estate prices and registrations of title], Quality declaration BO0501, Statistics Sweden (SCB), Version 1, published 2026-02-05.

Appendix A. Additional results

Appendix A.1. Excluding major urban counties

A natural concern is that the result is driven by large, growing metropolitan counties such as Stockholm, Skåne, and Västra Götaland, which have higher birth shares and stronger cumulative price growth than the other counties. We re-estimate the main specification excluding these three counties.

Table IV. Main result excluding major urban counties. Column 1 reproduces the main specification on the full sample of 20 counties. Column 2 re-estimates excluding the three counties containing Sweden's largest cities (Stockholm, Skåne with Malmö, and Västra Götaland with Gothenburg). The birth-share coefficient is larger and more precisely estimated outside the major cities, indicating that the result does not depend on their inclusion. Newey-West standard errors (bandwidth four).

Dependent Variable:	5-yr log price change	
	Full sample	Excl. major urban counties
Model:	(1)	(2)
Variables		
Birth share (age 0, t-27)	35.51***	39.30***
	(9.930)	(9.280)
5-yr log income growth (25-29)	0.3743**	0.3693*
	(0.1740)	(0.1908)
Fixed-effects		
County	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	540	459
R^2	0.72092	0.77933
Within R^2	0.10003	0.15783

Newey-West (L=4) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The coefficient on the lagged birth share is 39.30 on the 17-county subsample, compared with 35.51 on the full sample. Despite the smaller sample (459 versus 540 observations), the standard error tightens from 9.93 to 9.28 and the within R^2 rises from 0.10 to 0.16. Smaller samples typically produce larger standard errors, so the precision improvement is unlikely to be a mechanical artefact of the sample reduction. The 10% coefficient increase falls within typical sampling variation for the smaller subsample, and the precision improvement is the more informative comparison.

We interpret the precision improvement as a feature of the regressor itself. Metropolitan counties have absorbed the bulk of internal and international migration over the sample period (Section III), and for migrants who were born elsewhere, the birth share at lag 27 in their current county does not include them. The regressor is therefore a noisier proxy for the current population aged 27 in metros than in non-metros, and excluding the metros sharpens the demographic signal. The migration check in Appendix A.2 addresses the same concern from a different angle.

Appendix A.2. Migration controls

Year fixed effects absorb aggregate migration trends, and county fixed effects absorb persistent county-level migration patterns. For migration to bias the demographic estimate, it would have to both correlate with lagged birth shares and independently affect price growth. We control for this directly by adding net internal, net international, and net total migration rates as separate controls. The sample reduces to 460 observations because SCB inter-county migration data begins in 2000.

Table V. Main result with migration controls. Column 1 reproduces the main specification on the migration subsample. The sample begins in 2000 because SCB inter-county migration data is unavailable before, so the baseline coefficient differs slightly from Table II. Columns 2–4 add net internal migration, net international migration, and net total migration rates respectively as additional controls. The birth-share coefficient remains stable, suggesting migration does not confound the result. All specifications absorb county and year fixed effects, with Newey-West standard errors (bandwidth four).

Dependent Variable:	5-yr log price change			
	Baseline	+ Internal	+ International	+ Total
Model:	(1)	(2)	(3)	(4)
Variables				
Birth share (age 0, t-27)	21.61*	22.92**	21.43*	22.82**
	(11.17)	(10.70)	(11.09)	(10.66)
5-yr log income growth (25-29)	0.4038**	0.3479**	0.4332**	0.3762**
	(0.1710)	(0.1675)	(0.1681)	(0.1628)
Net internal migration rate		5.231*		
		(2.870)		
Net international migration rate			5.103*	
			(2.544)	
Net total migration rate				5.615**
				(2.308)
Fixed-effects				
County	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	460	460	460	460
R^2	0.74371	0.74915	0.74758	0.75381
Within R^2	0.06577	0.08562	0.07988	0.10261

Newey-West (L=4) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The relevant comparison is across columns within Table V rather than between the subsample and the full-sample result. The baseline on the 2000–2022 subsample is 21.61

(column 1), substantially lower than the 35.51 obtained on the full 1996–2022 sample, indicating that the coefficient is sensitive to which years are included. Within this subsample, the birth-share coefficient is stable across the three migration controls: 22.92 with internal migration, 21.43 with international migration, and 22.82 with total migration (columns 2–4). The demographic effect is not absorbed by contemporaneous migration controls in this subsample. We cannot, however, separately explain why the coefficient is lower on the 2000–2022 subsample than on the full panel. Migration patterns that differ across the two windows are one of several possibilities, and we discuss the magnitude sensitivity further in Section VII.

The stability of the coefficient under migration controls aligns with the cohort observation age analysis in Section VI.C, where the price response appears to draw largely on people born in the county, including those who left and returned by buying age, rather than on inflows from elsewhere. Appendix A.1 shows the same logic from a different angle. Removing the metropolitan counties with high migration sharpens the demographic signal.

Inflowing migrants might also make fewer house purchases than settling cohorts. Some moves are temporary, international migrants may face barriers to mortgage finance, and some rely on whatever public housing or queue-allocated rental options remain. If so, migration would exert less price pressure than its raw flows suggest, weakening any potential to confound the birth-share estimate.

Francke and Korevaar (2025) flag the same endogeneity concern at a conceptual level: contemporaneous age structure responds to migration, which itself responds to local conditions. Monnet and Wolf (2017) document empirically that the predetermined-cohort identification works best in countries with low migration, including Sweden.

Appendix A.3. Predetermined vs contemporaneous

The cohort-observation-age figure in Section VI.C traces the demographic coefficient as the same target cohort is measured at successive ages, from birth to its contemporaneous endpoint. The endpoint coefficient is positive but insignificant. Here we run a single regression that includes both the predetermined birth share at lag 27 and the contemporaneous share of 27-year-olds. Francke and Korevaar (2025) argue that birth rates are a more credible source of exogenous variation in housing demand than current demographic structure, since the age structure today is reshaped by 27 years of migration that itself responds to local economic conditions. Their historical data do not include contemporaneous age counts, so they cannot run this comparison. Modern Swedish county data make it feasible.

Column 1 reports the predetermined birth share alone ($\beta = 35.51$, $p < 0.01$). Column 2 substitutes the contemporaneous share of 27-year-olds ($\beta = 20.67$, insignificant). Column 3 includes both. The predetermined coefficient is essentially unchanged ($\beta = 33.31$, $p < 0.01$), while the contemporaneous coefficient falls from 20.67 to 9.70 and remains insignificant. Predetermined birth shares absorb the predictive content of the contemporaneous share.

The figure shows this absorption age by age. The regression shows it in one specification. At the lag where prices peak, the demographic signal is carried by predetermined data on births, matching the identification argument in Section V.B.

Table VI. Lagged birth share versus contemporaneous age share. Two-way fixed-effects OLS regressions of the centred five-year log house price change. Column 1 reproduces the main specification using the predetermined birth share at lag 27. Column 2 replaces it with the contemporaneous share of 27-year-olds in the county-year. Column 3 includes both. All regressions control for five-year log income growth (25–29) and absorb county and year fixed effects. Newey-West standard errors (bandwidth four). Predetermined births absorb the contemporaneous share when both are included. The predetermined coefficient barely moves (35.51 to 33.31), while the contemporaneous coefficient more than halves (20.67 to 9.70).

Dependent Variable:	5-yr log price change		
	Lagged births	Contemp. share	Both
Model:	(1)	(2)	(3)
Variables			
Birth share (age 0, t-27)	35.51***		33.31***
	(9.930)		(11.69)
5-yr log income growth (25-29)	0.3743**	0.4130**	0.3852**
	(0.1740)	(0.1765)	(0.1772)
Share aged 27 (contemp.)		20.67	9.703
		(15.39)	(18.48)
Fixed-effects			
County	Yes	Yes	Yes
Year	Yes	Yes	Yes
Fit statistics			
Observations	540	540	540
R^2	0.72092	0.70465	0.72193
Within R^2	0.10003	0.04757	0.10329

Newey-West (L=4) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Appendix A.4. Non-overlapping panels

The main specification uses centred five-year log price changes, where the dependent variable in year t is $\log p_{t+2} - \log p_{t-3}$. Two adjacent observations share four of their five input years, inducing a moving-average autocorrelation of order four in the residuals. We address this with Newey-West standard errors at bandwidth four. A sharper check removes the overlap mechanically. We restrict the panel to price changes centred on 2000, 2005, 2010, 2015, and 2020, leaving 100 county-year observations. The five-year spacing eliminates serial correlation by construction, so cluster-robust standard errors at the county level are sufficient.

The point estimate rises from 35.51 to 44.25 while the standard error roughly doubles (9.93 to 18.97). The reduced precision is expected with only 100 observations, five time points across 20 counties. Despite the much smaller sample, the coefficient remains significant at the 5% level. The result is not an artifact of the overlapping windows in the main specification.

Table VII. Main result with non-overlapping estimation windows. Two-way fixed-effects OLS regressions of the centred five-year log house price change. Column 1 reproduces the main specification (overlapping windows, Newey-West standard errors). Column 2 restricts to non-overlapping windows (price changes centred on 2000, 2005, 2010, 2015, and 2020) with cluster-robust standard errors at the county level. Sample sizes are 540 and 100. The coefficient remains significant at the 5% level despite the fivefold sample reduction.

Dependent Variable:	5-yr log price change	
	Overlapping (NW)	Non-overlapping (clustered)
Model:	(1)	(2)
Variables		
Birth share (age 0, t-27)	35.51***	44.25**
	(9.930)	(18.97)
5-yr log income growth (25-29)	0.3743**	0.3862*
	(0.1740)	(0.2004)
Fixed-effects		
County	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	540	100
R^2	0.72092	0.76725
Within R^2	0.10003	0.12257

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Appendix A.5. One-year price changes

The centred five-year window in the main specification smooths noise from year to year in the price index, which could in principle distort the shape of the age profile or its statistical significance. We re-estimate the age profile using annual log price changes, with one-year income growth (25–29) as the control. Annual changes do not overlap, so heteroskedasticity-robust standard errors replace the Newey-West correction.

Annual changes deliver the same shape as the centred five-year profile. The coefficient is insignificant before age 24, climbs to its peak at age 27, and returns toward zero after the mid-thirties. The five-year smoothing in the main specification reduces year-to-year noise but does not generate the location of the peak or the surrounding significance pattern.

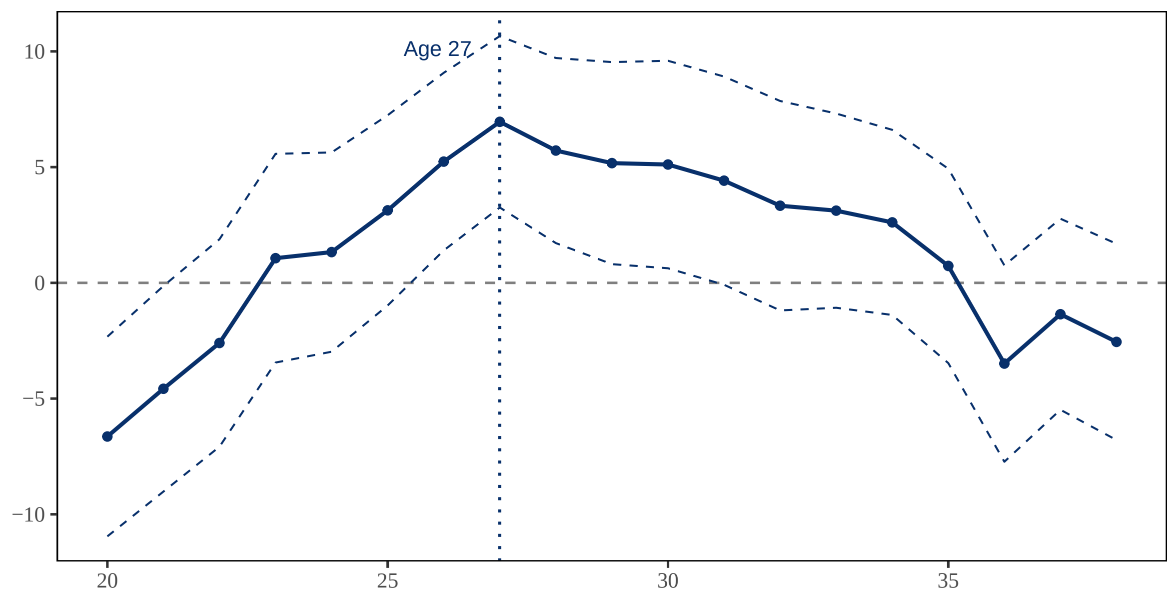

Figure 9. Age profile using annual price changes. Each point is the OLS coefficient on the birth share at lag L from a separate regression of the annual log house price change on the lagged birth share, controlling for one-year log income growth (25–29) and absorbing county and year fixed effects. Standard errors are heteroskedasticity-robust. The shaded band is the 95% confidence interval. Sample: 20 Swedish counties, 1991–2024. The profile peaks at age 27 ($\beta = 6.96$, $p < 0.001$) with the same hump shape as the centred-five-year profile in Figure 4.

Appendix B. AI disclosure

In writing this thesis we used Claude Opus (Anthropic) for several distinct tasks, with particular attention to ensuring the integrity of our citations. We extracted the full source text of every cited paper into a structured local archive and configured an automated verification workflow that ran in the background each time we wrote new material referring to a source. The primary check was accuracy of paraphrasing: whether every claim attributed to a paper faithfully represents what that paper actually argues. As a secondary check we verified originality of language, ensuring our wording was sufficiently distinct from the source's own. The result is that every citation in the final draft has been continuously cross-checked against its underlying source.

We also used Claude Opus as a coding assistant for the R analyses underlying every figure and table reported here. As we came up with hypotheses about the data and the regression specifications to test them, the assistant helped us code and run those regressions faster than we could have alone. It also assisted in writing the export pipeline that produces all figures and tables in LaTeX-compatible form, converted our draft into the structured LaTeX source in this submission, migrated our manually formatted citations into BibTeX entries, and edited the writing for grammar and phrasing.

All analytical decisions, including hypothesis development, the choice of empirical specifications, the interpretation of results, and the conclusions drawn, are our own.