Global house prices since 1950*

Haroon Mumtaz[†] and Roman Šustek[‡§]

October 1, 2024

Abstract

What drives house prices? Applying a parsimonious model to house prices in 12 countries since 1950, we show that expectations about future fundamentals were the key driver behind major house price movements. In the model, house prices depend on the expected future housing stock, population, income per capita, and age distribution. The growth rates of these fundamentals contain persistent components, estimated from the data, affecting expectations. The model accounts for the spectacular boom and bust in Japan, the boom starting in many countries in early 1990s, and the house price cycles in Switzerland. A decomposition into the contributing factors is carried out.

JEL Classification Codes: E20, G50, J11, R21, C11.

Keywords: House prices, international, growth shocks, fundamentals, demographics.

*We are grateful for helpful comments and suggestions to Carlos Garriga, Morris Davis, Víctor Ríos-Rull, Tatsuro Senga, Winfried Koeniger, Burhan Kuruscu, and Toni Braun, as well as seminar and conference participants at the City University, Masaryk University, Queen Mary, Keio, Durham, St Gallen, and the Vienna Macro Café. We are also grateful to Allen Head for his discussion in Vienna. Jason Hilton and Jakub Bijak from the Department of Social Statistics and Demography at the University of Southampton and Sara Hertog from the UN Population Division were very helpful with our enquiries about the demographic aspects of the paper.

[†]Queen Mary University of London.

[‡]Queen Mary University of London and the Centre for Macroeconomics, LSE.

[§]Correspondence: School of Economics and Finance, Queen Mary University of London, Mile End Road, London, E1 4NS, UK, tel: +44 20 7882 8831, e-mail: r.sustek@qmul.ac.uk.

1 Introduction

What drives house prices? The boom and bust in the US housing market in the 2000s has firmly placed this question in mainstream macroeconomics. During the boom (2000-2006), US real house prices grew on average by 5.4% per annum, while during the bust (2006-2012) the market suffered average annual decline in real house prices of −5.1%. The tremendous volume of academic work that followed has taught us important lessons about housing and mortgage markets and their effects on different households. However, generating large and persistent movements in house prices has been challenging and the question of what drives house prices is still not settled.[1]

While the focus of the literature on the US boom-bust period is understandable, in terms of house prices alone, the long-term international experience (documented in the next section) is arguably even more interesting. For instance, in the mid-1990s many advanced economies embarked on a path of real house price appreciation that was roughly at par with the growth rates observed during the US boom and more than twice as high as in the previous decades. Many commentators at the time speculated that house price bubbles were forming around the world (see Case and Shiller, 2003). The most dramatic changes in house prices in the post-WWII history occurred, however, in Japan. After WWII, Japan experienced an unprecedented four-decade-long house price bonanza, with house prices growing at staggering 9.3% per year, on average, in real terms. That is almost double the growth rate during the US boom. In 1991, however, the boom in Japan suddenly turned into two decades of a sustained house price bust, with the average growth rate of −3.2% per year in real terms. The house price movement in Japan is so stark and persistent that successfully accounting for it must inevitably bring important insights into the factors driving house prices at the aggregate level. At the other end of the spectrum is Switzerland, where house prices have modest long-run growth, but exhibit recurrent cycles comparable in magnitude to the US boom-bust period.

The goal of the paper is to advance our understanding of the main drivers of aggregate house prices by accounting for the above patterns since WWII in a common theoretical framework. Unlike the recent literature, the approach taken here abstracts from the details of mortgage and housing markets, as well as government policies and regulations, which evolve over time and differ across

[1] For the US period, see Favilukis, Ludvigson and Van Nieuwerburgh (2017), Garriga, Manuelli and Peralta-Alva (2019), Justiniano, Primiceri and Tambalotti (2019), Kaplan, Mitman and Violante (2020), Garriga and Hedlund (2020), Albanesi, Giorgi and Nosal (2022), Arslan, Guler and Kuruscu (2023), and Greenwald and Guren (2024). For a review of the literature on housing in macroeconomics see Davis and Van Nieuwerburgh (2015), Guerrieri and Uhlig (2016) and Piazzesi and Schneider (2016).

countries.[2] While such considerations matter for house prices, as the research on the US boom-bust period has shown, the goal is to explore how far a common set of a few fundamentals can go in accounting for historical house price movements in different countries. In other words, to present a parsimonious theory that, while not explaining every house price swing, would be informative about observed house prices, irrespective of a country's institutional and market environments.

Our sample consists of 12 countries and annual data for the period 1950-2019.[3] Viewing the housing stock as an asset providing shelter services over time, we embed a parsimonious optimizing model of demand for housing services into a standard asset pricing framework. The *price of housing services* depends on the current housing stock, total population, income per capita, and the age distribution of the population, jointly referred to as 'fundamentals'.[4] *House prices* are the present value of the prices of housing services expected to prevail in the future. The assumed time-varying stochastic discount factor is based on the real interest rate taken from the data. Importantly, we allow the stochastic processes for the growth rates of the fundamentals to contain *persistent random components*, building on Barsky and DeLong (1993), Bansal and Lundblad (2002) and Bansal and Yaron (2004). Such specifications capture persistent demographic changes as well as the notion that long-run economic growth is not constant. For instance, in the 1970s GDP growth was lower than during the IT revolution of the 1990s, as established by the growth literature. The parameters, and the persistent components, of the stochastic processes are estimated from the data on the respective fundamentals by Bayesian state space methods.[5] We emphasize that the estimates of the key drivers are obtained outside of the asset pricing model. The estimated persistent components are then fed into the model, where they determine the rational expectations of future fundamentals, thus affecting the deviations of house prices from current fundamentals. A few parameters—related mainly to people's preferences—are estimated by Bayesian methods within the asset pricing model, conditional on the estimated parameters of the exogenous processes. Existing studies, including micro-level studies, are used to set priors on these parameters.

[2]We also abstract from the possibility of bubbles.

[3]The sample consists of Australia, Belgium, Canada, Denmark, Finland, France, Japan, the Netherlands, Sweden, Switzerland, the United Kingdom, and the United States.

[4]We do not model optimal housing supply. The evolution of the housing stock is taken as exogenous. Knoll, Schularick and Steger (2017) argue that zoning and land control restrictions have been the main constraints on housing supply in the post-WWII period. In line with this view, their work shows that across countries the bulk of house price movements after WWII is due to land prices, not construction costs; see also Davis and Heathcote (2007) and Braun and Lee (2021).

[5]To achieve a better precision of the estimates, we estimate the parameters on a panel of the 12 countries. Monte Carlo simulations are used to test whether the estimation technique can uncover the persistent components.

To summarize the findings, when the observed fundamentals and their estimated persistent components are fed into the model, the model accounts well for the three aforementioned house price patterns since 1950.[6] Specifically, the most striking finding is that the model reproduces almost exactly the unprecedented boom and bust in Japan. The model also accounts for the fact that between 1993 and 2007 most countries experienced house price growth twice as fast as in the preceding decades. Finally, the model generates the recurrent fluctuations in house prices, in the presence of weak long-run growth, observed in Switzerland. The findings withstand a number of robustness checks, noted in the paper and contained in the on-line Appendix. We are not aware of another study that would successfully account for these patterns. The results are also novel by bringing fundamentals back into the fore of house price determination. While the earlier work on house prices emphasized fundamentals (Topel and Rosen, 1988), the more recent quantitative studies found it difficult to generate volatile and persistent movements in house prices from such factors (eg, Davis and Heathcote, 2005).

Expectations are crucial in accounting for the three house price patterns. Expectations in our model are rational expectations about future fundamentals, derived from the estimated state space. Expected future house prices are endogenous.[7] Without changes in expectations, house prices would exhibit a relatively stable stochastic trend. For instance, in the case of Japan, the boom would be nowhere near as strong as in the data and there would be no bust, only a leveling-off of house prices. The changes in expectations, driven by the persistent random components in the growth rates of the fundamentals, generate large and persistent deviations of house prices from the stochastic trend, determined by the current levels of the fundamentals. In addition to helping us understand the historical patters in house prices, these findings convey a broader message to the macro-housing literature: models aimed at explaining volatile and persistent house prices should feature persistent shocks to the growth rates of fundamentals (most models work with mean-reverting shocks to the levels of fundamentals).

We provide a decomposition of the findings into the marginal contribution of the individual factors. In the case of Japan, the most important driver was expected future per-capita income (GDP) growth, followed by expected future population growth. These two factors reproduce the boom-bust

[6] We have results for all 12 individual countries. However, to draw general conclusions, we organize the paper in terms of the three house price patterns, which summarize the historical experience of the countries in the sample.

[7] Some authors work with survey-based house price expectations (eg, Landvoigt, Piazzesi and Schneider, 2015). This strategy is more suitable for more recent periods and specific markets, for which the relevant surveys are available, unlike our international sample from 1950.

period almost exactly. Japan experienced a phenomenal growth in GDP per capita and population after WWII. According to the model, expectations of such advances in prosperity continuing into the future got reflected in the rapid growth of house prices. As the growth in the fundamentals gradually stalled (and in the case of population even reversed), the resulting adjustments in expectations turned the boom into a sustained bust.

In the case of the countries that experienced the acceleration of house price growth from around 1993 (all countries in the sample except Japan and Switzerland), the main driver until 2000 was expected future per-capita GDP growth. Afterwards, between 2000 and 2007, the fast house price growth was sustained by expectations of fast population growth. These expectations reflected, respectively, fast underlying growth in GDP per capita during the 1990s, which was followed by a surge in population growth in the 2000s. Finally, in the case of Switzerland, the cyclical nature of house prices in this country is mainly due to recurrent shocks to expected population growth, reflecting net migration tied to the business cycle.[8]

A pertinent question concerns the role of interest rates, especially since the global financial crises. Has loose monetary policy inflated house prices? Our findings support this view. In all countries, house prices would be lower between 2009 and 2019 if interest rates stayed at their post-WWII average. The gap in 2019 is about 12%.[9] However, with the exception of a few short periods, for most of the 70 years, the marginal contribution of interest rates is relatively weak.

Finally, the model is used to gauge the marginal contribution to house prices of the changes in housing demand driven by the age composition of the population.[10] This is an extension of the classic exercise of Mankiw and Weil (1989), taking into account the effect on expectations. All countries in the sample experienced significant population aging between 1950 and 2019. For five broad age groups, the largest losses were in the age group 0-24 and the largest gains in the age groups 55-69 and 70+. In Japan, *expectations* of future population ageing had a positive effect on house prices until the mid-1990s, as the mass of the distribution was slowly moving towards the age categories

[8] In the case of the US (see the on-line Appendix), the model cannot fully account for the boom and bust in the 2000s, thus indirectly confirming the findings of the literature that financial factors and subjective expectations were important during this period. Similarly, the model does not account for sharp short-term boom and busts during different financial crisis in Denmark and Finland.

[9] The effect of interest rates is weaker in Japan and Switzerland than in the other countries. The difference between Japan and Switzerland on one hand and the rest of the countries on the other is in the size of the decline in the real interest rate between 2009 and 2019.

[10] The exercise estimates the marginal contribution of the changes in the age composition of the population, keeping the observed path of the population numbers unchanged.

40-54 and 55-69, which (according to the estimated parameters) are the largest consumers of housing services. After the mid-1990s, expectations of population ageing progressing further into the category 70+ started to weight down on house prices.[11] The rest of the countries appear to be in 2019 where Japan was before the mid-1990s, with expectations of population ageing still having a positive effect on house prices. Taking into account also the effect of the age composition on the *current* demand for housing services, population ageing had, so far, a positive effect on house prices in all countries in the sample. If the age structure of the population stayed at the 1950 distribution, house prices in 2019 would be 14 percent lower (cross-country median).

In terms of the literature, as noted above, most of the recent work has focused on the US boom-bust period and the role of different frictions in mortgage and housing markets. Our paper is related to three (overlapping) strands of the earlier literature. The distinguishing feature of our model, with respect to both the recent and earlier work, is the central role of the persistent random components of the fundamentals and their effects on expectations.[12] First, the paper follows a tradition, going back to Swan (1984) and Topel and Rosen (1988), that ties house prices to the demand for housing services and its deeper determinants. Davis and Heathcote (2005) carry out such an analysis in general equilibrium with sectoral productivity shocks calibrated to the US economy.[13] Case and Shiller (2003) establish a regression-based relationship between house prices and per-capita income in a cross-section of US states, 1985-2002.[14] Knoll et al. (2017) collect and analyze house price data for a number of countries going back to 1870. They describe a 'hockey stick' pattern, whereby house prices were approximately flat until the 1950s, before embarking on an upward trend, and relate the change to residential land becoming a scarce factor. We focus in more detail on the post-1950 period.[15]

Second, the paper is related to an asset valuation approach to house prices. This literature has focused on two aspects: a variance decomposition into expected future returns and rent growth

[11] Relative to the two main drivers of the house price boom and bust in Japan (the growth rates of real GDP per capita and population) this effect is less important.

[12] Kaplan et al. (2020) have a flavor of our mechanism in their paper. They consider a shock to expected future preference for housing.

[13] See also Grossmann, Larin and Steger (2023).

[14] See also Poterba (1991), McCarthy and Peach (2004), and Glaeser and Gyourko (2007).

[15] In addition, a number of authors approach the relationship between house prices and fundamentals (typically income per capita) from the perspective of a time series cointegrating relationship, using US national, state and city data (eg, Gallin, 2006; Holly, Paseran and Yamagata, 2010). Arestis and González (2014) and Geng (2018) carry out such an exercise at the aggregate level for OECD countries for the periods 1970-2011 and 1990-2016, respectively. As cointegration only detects a relationship between the variables along a stochastic trend, given the length of the available data, persistent deviations of house prices from the trend led to inconclusive findings across the different studies.

(Campbell, Davis, Gallin and Martin, 2009; Plazzi, Torous and Valkanov, 2010) and deriving house prices based on the Poterba (1984) user cost theory. In the latter case, authors typically make various ad-hock assumptions about expected future house prices (see, eg, André, 2010, for OECD countries, 1970-2009).[16]

Finally, as we take the age distribution of the population into account, the paper is related to Mankiw and Weil (1989), Hamilton (1991), Green and Hendershott (1996), and Martin (2005). We contrast our results with the classic Mankiw and Weil (1989) paper.

The paper proceeds as follows, Section 2 documents the three patterns of international real house prices, Section 3 introduces the model, Section 4 describes its estimation, Section 5 presents the findings, and Section 6 concludes. Online Appendix provides further details.

2 Three house price patterns

This section documents the observed patterns in international house prices in the post-WWII period. Whenever we speak of house prices, we mean *real* house prices. The sample consists of 12 countries: Australia, Belgium, Canada, Denmark, Finland, France, Japan, the Netherlands, Sweden, Switzerland, the United Kingdom, and the United States. The sample period is 1950-2019, except Canada, which is 1957-2019. The data are annual. For the period 1970-2019, the data come from the OECD database. For 1950-1969, the source is Knoll et al. (2017).[17] The data are expressed as an index (set to 100 in 1957) and are plotted in Figure 1. Visual inspection, confirmed formally below, reveals three broad patterns.

First, the time path of house prices in Japan stands out. It exhibits spectacular four-decade-long growth between 1950 and 1991 (with a small boom-bust period in the early 1970s). As a result, in 1991, house prices in Japan were 39 times higher than in 1950. That is equivalent to an annual growth rate of 9.3%. From 1991, however, house prices steadily declined for almost two decades until 2009, at an average rate of -3.2% per annum. In 2009, house prices were at the same level as in 1978. No other country in the sample experienced such a long-lasting decline.[18] In terms of the

[16] In their application to the US national and metropolitan area data, 1975-2007, Campbell et al. (2009) highlight the importance of expected excess returns (a housing premium) for house price variation over time. However, for the sub-period 1997-2007, they find that expected future rent growth was the dominant factor in house price movements. This would be consistent with shifts in expectations about fundamentals.

[17] The house price data of Knoll et al. (2017) coincide with the OECD data in the post-1970 period.

[18] The boom and bust in Japan was uniform across different regions and reflected mainly the price of residential land.

post-WWII house price boom, France comes somewhat close to Japan. But even in France, which had tremendous house price growth between 1950 and 1967, house prices in 1967 were 'only' 7.9 times higher than in 1950, compared with 16.4 times higher in Japan. And in 1991, house prices in France were only 12.5 times higher than in 1950, compared with the aforementioned 39 fold increase in Japan. Both the boom and bust in house prices in Japan are unprecedented in the post-WWII history and dwarf the US boom-bust period in the 2000s.

Second, most countries seem to exhibit a house price pattern characterized by a faster average growth in the period after the mid-1990s then in the previous decades after WWII. This is more apparent in the bottom chart of Figure 1, which is a zoom-in of the upper chart by removing France and Japan. To confirm and formalize this common pattern, we carry out principal component (PC) analysis of the 12 data series, using the method of Barigozzi, Lippi and Luciani (2021), which extends the PC decomposition to non-stationary data. The 1st PC of the 12 data series is plotted in both charts of Figure 1 as the thick blue line.[19] The loadings of the individual series on the 1st PC reveal that Japan and Switzerland are different from the other countries; see Table 1. The loading for Japan is essentially zero while the loading for Switzerland, although positive, is much smaller than for the other countries. Re-doing the PC decomposition without Japan and Switzerland produces essentially the same 1st PC (not plotted).[20] We will refer to the group of countries excluding Japan and Switzerland as the 'G10 countries' and summarize their common house price experience by the 1st PC of the 10 data series.[21] Based on the 1st PC, these countries experienced house price growth of 1.7% per year on average during 1951-1993 and more than twice as fast growth of 4.8% per year on average during 1993-2007. The growth in the latter period is comparable to the average growth rate witnessed during the US boom in the 2000s.[22] The difference in the growth rates led many commentators to speculate that house price bubbles were forming in many countries around the world (Case and Shiller, 2003).[23]

Finally, Switzerland has the lowest long-run growth rate of house prices among the 12 countries, of only 1.1% per year on average (house prices in 2019 were only slightly above their 1989 levels).

[19] The 1st PC is quantitatively the most important common component of the 12 data series.

[20] Computing the unweighted average or median of the ten countries reveals the same pattern as the 1st PC.

[21] A common component in international house prices has also been detected by Hirata, Kose, Otrok and Terrones (2012) and Jackson, Kose, Otrok and Owyang (2016).

[22] For the purposes of describing the pattern of the G10 countries, we choose 1993 as the starting point of the fast growth period, as 1993 is the starting point of uninterrupted growth of house prices until 2007, based on the 1st PC.

[23] Including the global financial crisis in the second period, the average growth rate (1993-2019) based on the 1st PC is 3.1%. This is still almost double the average growth rate prior to 1993.

House prices in Switzerland, however, exhibit recurrent fluctuations, with three complete 'cycles' in the post WWII period. In terms of a change from peak to trough, the magnitudes of the cycles are: 1960-1962 a decline of 12%, 1973-1977 a decline of 27%, and 1989-1997 a decline of 36%. In terms of the magnitudes, the last two declines are comparable to the US bust in the 2000s.

Although we have results for all 12 countries (contained in the online Appendix), to draw general lessons, we organize the paper in terms of the house price patterns for Japan, the 1st PC of the G10 countries, and Switzerland. Our goal is to account for the three patterns in a common theoretical framework and understand the main driving forces.

3 The model

We view the housing stock as an asset providing a flow of housing services over time and embed a parsimonious model of demand for housing services in a standard asset pricing framework. In contrast to the recent literature, the model abstracts from the details of mortgage and housing markets, as well as government policies and regulations. However, it allows for persistent random components in the exogenous stochastic processes of housing market fundamentals, building on Barsky and DeLong (1993), Bansal and Lundblad (2002), and Bansal and Yaron (2004).[24]

3.1 Pricing the aggregate housing stock

Consider an economy in which people consume a numeraire good and homogenous housing services provided by the existing aggregate housing stock. People differ in their preferences for housing based on their age $j \in J$. We approach the problem of the allocation of the numeraire good and housing services across the agents as a planner's problem (here we provide only a brief description, relegating the details to the online Appendix). Each period, the planner solves a static problem of maximizing equally weighted utility of the different people in the economy, subject to an aggregate per-capita endowment of the numeraire good and the housing stock. We work with the utility function

$$u_j(a_j, c_j) = \psi_j \frac{a_j^{1-\varepsilon_1}}{1-\varepsilon_1} + \frac{c_j^{1-\varepsilon_2}}{1-\varepsilon_2}, \quad \varepsilon_1, \varepsilon_2 \geq 0,$$

[24] In Barsky and DeLong (1993) and Bansal and Lundblad (2002), persistent random components affect only expectations. Bansal and Yaron (2004) extend this idea by modeling the effect of persistent random components also on risk premia. These papers focus on equity pricing. In our model the channel operates only through expectations, as discussed below.

where a_j and c_j are consumption of housing services and the numeraire good of agent j, respectively, and $\psi_j > 0$ is a preference parameter. This utility function is one of the frequently used utility functions in the macro-housing literature (eg, Chambers, Garriga and Schlagenhauf, 2009; Gervais and Fisher, 2011). Amongst its advantages is the property that it allows for different elasticities of consumption and housing. An additional benefit, for our purposes, is that it admits explicit aggregation, whereby the preference heterogeneity is subsumed in a single variable in the planner's indirect utility function, while the indirect utility function preserves the functional form of the individual utility functions (an application of the results of Maliar and Maliar, 2001, 2003).

By the nature of the planner's problem, there is no distinction in the model between renting vs. owning and the scarcity of housing is captured by a single shadow price, given by the marginal rate of substitution between consumption of the numeraire good and housing, based on the indirect utility function of the planner. In logs, the shadow *price of housing services* is

$$\log d_t = \log \Psi_t + \varepsilon_2 \log c_t - \varepsilon_1 \log H_t + \varepsilon_1 \log N_t, \tag{1}$$

where c_t is the per-capita endowment of the numeraire good and the per-capita consumption of housing services is $a_t = H_t/N_t$, where H_t is the existing aggregate housing stock and N_t is the size of the population. Finally,

$$\Psi_t \equiv \left[\int_J \psi_j^{1/\varepsilon_1} d\mu_{j,t} \right]^{\varepsilon_1}, \tag{2}$$

where $d\mu_{j,t}$ is the measure of the agents of age j in the population in period t. The variable Ψ_t subsumes the effects of age heterogeneity on housing demand. It captures the idea that people at different stages of their lifecycle have different preferences/needs for housing and the aggregate effects of these lifecycle considerations on housing demand change as the age distribution of the population changes over time.[25] We refer to the four variables on the right-hand side of equation (1) as 'fundamentals'.

The period-t price of a unit of the aggregate housing stock—the *house price*—is the present value of the prices of housing services expected to prevail in the future. That is, the house price satisfies

[25] To keep the number of parameters to estimate in check, the preference parameters, ψ_j, are time-invariant. The model thus does not capture changes in people's attitudes towards marriage and the start of family (and thus housing), considered, for instance, by Gervais and Fisher (2011).

a standard asset pricing condition

$$q_t = E_t \left[m_{t+1}(q_{t+1} + d_{t+1}) \right],$$

where q_t is the period-t house price, d_{t+1} is the period $t+1$ shadow price of housing services, m_{t+1} is a pricing kernel, and E_t is the expectation operator conditional on the state space in period t. We use a log-normal pricing kernel

$$-\log m_{t+1} = \delta + r_t,$$

where r_t is a (continuously compounded) exogenous real interest rate, taken from the data, and δ is a parameter picking up a 'housing premium'.[26] The pricing kernel simply assumes that the housing stock can be financed at the exogenous rate of return equal to, in expectations, $\exp(\delta + r_t)$. This is similar to, among others, Kaplan et al. (2020) and Greenwald and Guren (2024). For the purposes of our application, treating the real interest rate as being determined outside of the national borders seem to be a sensible approach on empirical grounds. There is evidence that during the period under investigation, the real interest rates of the individual countries in the sample were strongly influenced by global conditions. Specifically, the interest rates were strongly positively correlated with each other, while only weakly correlated with domestic economic conditions.[27] Any possible time-varying departures of the pricing kernel from the real interest rate, resulting from various mortgage market frictions, are abstracted from, as noted at the outset of the paper.

To deal with nonstationarity, the above asset pricing condition is expressed as

$$x_t = E_t \left[m_{t+1} \exp(v_{t+1})(x_{t+1} + 1) \right], \tag{3}$$

where $x_t \equiv q_t/d_t$ and $v_{t+1} \equiv \log d_{t+1} - \log d_t$ is the continuously compounded growth rate of the shadow price of housing services. The term $(x_{t+1} + 1)$ on the right-hand side of equation (3) makes

[26] The annual frequency, dictated by the availability of the required data, precludes estimation of a process for time-varying prices of risk, or time-varying volatility, with an acceptable degree of precision. Consequently, we abstract from time-varying risk premia and estimate only δ, which subsumes constant risk premia, together with other factors, such as maintenance costs.

[27] For instance, as one possible metric, the average correlation between a country's real interest rate and the average real interest rate in the cross-section is 0.8. As another metric, the 1st PC of the 12 real interest rates accounts for 70 percent of their joint variance. The strong correlation holds even for Japan, which in other respects considered in this paper looks different from the other countries. At the same time, real interest rates do not seem to be related to domestic conditions: the average correlation between a country's growth rate of real GDP and its real interest rate is 0.14. In the case of Japan, the correlation is 0.17; in the case of Switzerland it is 0.03. For the persistent components of GDP growth, the correlations are similarly low.

the equation unsuitable for a closed-form solution for x_t in terms of the model's state variables. The standard way to proceed is to rewrite the equation in logs

$$\log x_t = \log E_t \left[\exp(\log m_{t+1} + \upsilon_{t+1} + \log(x_{t+1} + 1) \right]$$

and adopt the Campbell and Shiller (1988) approximation

$$\log x_t \approx \log E_t \left[\exp(\log m_{t+1} + \upsilon_{t+1} + \kappa_0 + \kappa_1 \log x_{t+1}) \right], \tag{4}$$

where $\kappa_0 \equiv \log(\overline{x}+1) - \kappa_1 \log \overline{x}$ and $\kappa_1 \equiv \overline{x}/(\overline{x}+1)$, with $\overline{x}$ being the unconditional mean of x_t. The growth rate υ_{t+1} is obtained as a first difference of equation (1)

$$\upsilon_{t+1} = z_{t+1} + \varepsilon_2 g_{t+1} - \varepsilon_1 h_{t+1} + \varepsilon_1 n_{t+1}, \tag{5}$$

where $z_{t+1} \equiv \log \Psi_{t+1} - \log \Psi_t$, $g_{t+1} \equiv \log c_{t+1} - \log c_t$, $h_{t+1} \equiv \log H_{t+1} - \log H_t$, and $n_{t+1} \equiv \log N_{t+1} - \log N_t$. Equation (4) can be solved analytically for $\log x_t$ by the method of undetermined coefficients, whereby the solution for $\log x_t$ is a linear function of the model's state variables.[28]

Endogenous house prices are given as $\log q_t = \log d_t + \log x_t$, where $\log d_t$ is given by (1) and $\log x_t$ satisfies (4). As $\log d_t$ depends only on the current levels of the fundamentals, whereas $\log x_t$ depends only on expectations of the growth rates of the fundamentals (and the interest rate), we can think of the log of house prices as the sum of a log stochastic trend ($\log d_t$), determined by current fundamentals, and log deviations from trend ($\log x_t$), driven by expectations and the interest rate.[29] In the empirical literature noted in the Introduction, a long-term cointegrating relationship between house prices and fundamentals (a stochastic trend) would refer to an equation like equation (1), while $\log x_t$ would be a part of an error correction mechanism picking up short-term deviations from the stochastic trend. Our model ties both elements together in a theoretically coherent way.

[28] An alternative approach would be to work with equation (3) and approximate the state space by discretization. For our purposes, the advantage of the Campbell-Shiller approximation is that it allows more efficient estimation of the model in terms of computing time.

[29] This interpretation is valid, strictly speaking, up to a scale, as the unconditional mean of $\log x_t$ is in general not equal to zero. The stochastic trend of house prices is thus equal to $\log d_t + \log x$, where $\log x$ is the unconditional mean of $\log x_t$.

3.2 Exogenous processes

To close the model, we specify the stochastic processes for the real interest rate and the growth rates of the variables making up υ_{t+1}. The processes are motivated by the existing literature and the properties of the data. We also carry out a number of specification tests, which are discussed in Section 5.3. In the specification below, the innovations are mutually independent across the processes. However, as a part of the robustness checks, we consider a specification with mutually correlated innovations.

As in Bansal and Lundblad (2002), the real interest rate is assumed to follow an AR(1) process

$$r_{t+1} = \nu_r + \phi_r r_t + \sigma_r \xi_{r,t+1}, \tag{6}$$

where $\phi_r \in (0,1)$ and $\xi_{r,t+1}$ is iid $N(0,1)$. For g_t we use a specification similar to those considered by Bansal and Lundblad (2002) and Bansal and Yaron (2004)

$$g_{t+1} = \nu_g + s_{g,t} + \sigma_g \xi_{g,t+1}, \tag{7}$$

$$s_{g,t+1} = \theta_g s_{g,t} + \varsigma_g \zeta_{g,t+1},$$

where $\theta_g \in (0,1)$ and $\xi_{g,t+1}$ and $\zeta_{g,t+1}$ are iid $N(0,1)$. This process allows for a stochastic autocorrelated component, $s_{g,t}$, which captures the idea that economic growth exhibits random long-run changes, in addition to short-run variation, as established by the growth literature (see, eg, Jones, 2016; Antolin-Diaz, Drechsel and Petrella, 2017, and the references therein). In the data, the unconditional autocorrelation of g_t is close to zero, making the growth rate look like a white noise. The process (7) is consistent with this property if σ_g is sufficiently larger than ς_g. The insight of Bansal and Lundblad (2002) is that even small changes in $s_{g,t}$ can have a large effect on expectations and asset prices if θ_g is sufficiently large. Small changes in the expected long-run growth rate can thus have large effects on house prices.

Unlike growth rates of per-capita consumption (or income or GDP), which resemble a white noise, population growth rates are relatively smooth and persistent. In fact, both population growth rates and the *changes* in population growth rates are smooth and persistent. That is, increases in the growth rate tend to be followed by further increases and declines tend to be followed by further declines.[30] We capture this behavior in a parsimonious way as

$$n_{t+1} = \nu_n + \phi_n n_t + s_{n,t} + \sigma_n \xi_{n,t+1}, \tag{8}$$

$$s_{n,t+1} = \theta_n s_{n,t} + \varsigma_n \zeta_{n,t+1},$$

where $\phi_n, \theta_n \in (0,1)$ and $\xi_{n,t+1}$ and $\zeta_{n,t+1}$ are iid $N(0,1)$.[31] Observe that re-writing the first equation as

$$\Delta n_{t+1} = \nu_n + (\phi_n - 1) n_t + s_{n,t} + \sigma_n \xi_{n,t+1},$$

and for $\phi_n \to 1$, the shock $s_{n,t}$ can be interpreted as a stochastic conditional mean of Δn_{t+1}.

The properties of the process for z_t, the growth rate of the distributional variable Ψ_t, depend on the properties of the time series of the age distribution. Movements in the age distribution, and its first difference, are relatively smooth and persistent but the aggregation in (2) may not translate these properties into Ψ_t. With this in mind, we adopt the general specification

$$z_{t+1} = \nu_z + \phi_z z_t + s_{z,t} + \sigma_z \xi_{z,t+1}, \tag{9}$$

$$s_{z,t+1} = \theta_z s_{z,t} + \varsigma_z \zeta_{z,t+1},$$

where $\phi_z, \theta_z \in (0,1)$ and $\xi_{z,t+1}$ and $\zeta_{z,t+1}$ are iid $N(0,1)$.

Data on the housing stock for the countries and periods under investigation are available only for the United States and, under some assumptions, can be constructed for Japan.[32] For our main results we therefore proceed under the assumption that the growth rate of the housing stock is constant. However, we carry out robustness checks for the United States and Japan using the available housing stock data for these two countries. In that case, the growth rate of the housing stock follows the process

$$h_{t+1} = \nu_h + s_{h,t} + \sigma_h \xi_{h,t+1}, \tag{10}$$

[30] Statistical tests of stationarity of population growth rates are inconclusive and in the samples of the length used here have low power. However, it is theoretically difficult to justify population growth rates that can grow (or decline) without bounds.

[31] This process has a representation as ARMA(2,2). ARMA processes have been used by demographers to model population growth since at least Pflaumer (1992).

[32] Data on residential investment that could potentially be used to construct housing stock data are available for most countries in our sample at best only from the 1970s and in many cases only the 1990s (see the supplementary material to Kydland, Rupert and Šustek, 2016).

$$s_{h,t+1} = \theta_h s_{h,t} + \varsigma_h \zeta_{h,t+1},$$

where $\theta_h \in (0,1)$ and $\xi_{h,t+1}$ and $\zeta_{h,t+1}$ are iid $N(0,1)$. It turns out that the results are not particularly sensitive to using this process instead of a constant growth rate, which is a special case of the above process for $\sigma_h = 0$ and $\varsigma_h = 0$. This is because the effects of the demand factors on house prices (including their future expectations) turn out to be substantially stronger than the supply effects (including their future expectations). The results for the specification with the process (10) for Japan and the United States are therefore contained only in the Online Appendix, with a discussion provided in Section 5.3. For the rest of the main text we treat the housing stock as growing at a constant growth rate.

In theory, the fundamentals should be, to some extent, correlated. For instance, gains in longevity affect z_t by increasing the fraction of the elderly in the population. At the same time they increase n_t by reducing mortality rates. And possibly reduce g_t by increasing the number of economically inactive people in the population. As noted at the outset of this Section, as a robustness check, we capture these effects by estimating the processes under the assumption that the innovations are mutually correlated across the processes.[33]

3.3 Model-implied house prices

Endogenous house prices are given by

$$\log q_t = \underbrace{\log \Psi_t + \varepsilon_2 \log c_t - \varepsilon_1 (\log H_0 + ht) + \varepsilon_1 \log N_t}_{\log d_t} + \log x_t, \tag{11}$$

where the process for $\log x_t$ satisfies equation (4). Exploiting the log-normal structure of the asset pricing model, the method of undetermined coefficients yields

$$\log x_t = \gamma + \gamma_z z_t + \gamma_{sz} s_{z,t} + \gamma_n n_t + \gamma_{sn} s_{n,t} + \gamma_{sg} s_{g,t} + \gamma_r r_t, \tag{12}$$

where

$$\gamma_z = \frac{\phi_z}{1 - \kappa_1 \phi_z}, \quad \gamma_{sz} = \frac{1 + \kappa_1 \gamma_z}{1 - \kappa_1 \theta_z}, \tag{13}$$

[33] Working with the primitives (shocks to birth, longevity, and migration rates) is not operational due to inconsistencies in the published UN data on the age distribution of the population on one hand and birth, longevity, and migration rates on the other. We thank Sara Hertog from the UN Population Division for clarifying this issue to us.

$$\gamma_n = \frac{\varepsilon_1 \phi_n}{1 - \kappa_1 \phi_n}, \quad \gamma_{sn} = \frac{\varepsilon_1 + \kappa_1 \gamma_n}{1 - \kappa_1 \theta_n}, \tag{14}$$

$$\gamma_{sg} = \frac{\varepsilon_2}{1 - \kappa_1 \theta_g}, \quad \gamma_r = -\frac{1}{1 - \kappa_1 \phi_r}. \tag{15}$$

Observe that only a few parameters govern the responses of $\log x_t$ to shocks. The equilibrium coefficients (13)-(15) depend only on the two elasticities ε_1 and ε_2 and the persistence parameters of the exogenous processes. The coefficients are increasing, in absolute value, in the persistence of the shocks and (for κ_1 close to one) the responses to $s_{z,t}$ and $s_{n,t}$ are larger than to z_t and n_t, respectively. Further, note that the two elasticities affect both the responses of $\log x_t$ to the growth rates and the responses of $\log d_t$ to the levels of the same variables (the elasticity of Ψ_t is equal to one). A variable that has a weak effect on the stochastic trend can thus nonetheless have a strong effect on the deviations from the trend if its growth rate contains persistent random components.

The coefficient γ subsumes all constants and conditional variances and is given by

$$\gamma = \frac{-\delta - \varepsilon_1 h + (1 + \kappa_1 \gamma_z)\nu_z + (\varepsilon_1 + \kappa_1 \gamma_n)\nu_n + \varepsilon_2 \nu_g + \kappa_1 \gamma_r \nu_r + \kappa_0 + 0.5\Sigma}{1 - \kappa_1} \tag{16}$$

where

$$\Sigma = (1 + \kappa_1 \gamma_z)^2 \sigma_z^2 + (\varepsilon_1 + \kappa_1 \gamma_n)^2 \sigma_n^2 + (\varepsilon_2)^2 \sigma_g^2 + (\kappa_1 \gamma_r)^2 \sigma_r^2 + (\kappa_1 \gamma_{sz})^2 \varsigma_z^2 + (\kappa_1 \gamma_{sn})^2 \varsigma_n^2 + (\kappa_1 \gamma_{sg})^2 \varsigma_g^2.$$

The fact that the variances increase γ reflects the standard Jensen's inequality effect of variance terms on asset prices.

In the quantitative exercise, the model-implied house prices are generated by equation (11) as the sum of the log of the stochastic trend and the log of the deviations. The log of the stochastic trend is constructed by feeding in data on c_t and N_t and the variable Ψ_t, constructed from data on the age composition of the population. The log of the deviations is generated by equation (12) by feeding in the estimates of the latent state variables $s_{z,t}$, $s_{n,t}$, and $s_{g,t}$, the growth rates z_t and n_t (constructed as log differences from the respective data on levels), and data on the interest rate r_t.[34]

[34] In the extended version with the stochastic process for the housing stock, house price data for Japan and the United States, and the estimated components $s_{h,t}$, are also fed into the model.

4 Estimation

4.1 Data

The house price data are complemented with annual data on real GDP per capita, total population, the age distribution of the population, and the real interest rate (obtained as a difference between the 10-year nominal interest rate and the inflation rate for the same year).[35] Real GDP data are from Penn World Table, version 9.1. The PWT data end in 2017. The last two years are taken from the St Louis Fed FRED database. The GDP data are converted into per capita terms by dividing by total population. Population data come from the United Nations, World Population Prospects 2019. The growth rates of GDP per capita and population are derived from the respective levels as log differences. Data on the age distribution (available by year for ages 0 to 100) also come from the United Nations, World Population Prospects 2019. Data on long-term nominal interest rates and CPI inflation come from the dataset accompanying Jordá, Schularick and Taylor (2017) and, where necessary, are complemented with data from the St Louis Fed FRED database. For the extended version of the model with the process for the housing stock (10), the housing stock data for Japan and the United States are discussed in the Online Appendix. For the results in the main text, we calibrate h, the constant growth rate of the housing stock, to 0.02 for all countries.[36]

To proceed, the theoretical construct Ψ_t needs to be made operational. This is done by splitting the population into J groups. The operational Ψ_t is then given by

$$\log \Psi_t = \varepsilon_1 \log \left(\sum_{j=1}^{J} \psi_j^{1/\varepsilon_1} \mu_{j,t} \right), \tag{17}$$

where $\mu_{j,t}$ is the fraction of the age group j in a country's population in period t, as reported in World Population Prospects 2019. In the application, we opt for $J = 5$: ages 0-24, 25-39, 40-54, 55-69, 70+. Note that Ψ_t and its growth rate are observable only conditionally on the estimates of the parameters entering equation (17). This point is summarized by the notation: $\Psi_t = \Psi(\varepsilon_1, \psi; \mu_t)$ and $z_t = \log \Psi_t - \log \Psi_{t-1} = z(\varepsilon_1, \psi; \mu_t, \mu_{t-1})$, where $\psi = \{\psi_j\}_{j=1}^J$ and $\mu_t = \{\mu_{j,t}\}_{j=1}^J$. Given ε_1 and

[35] For the countries in our sample, data on long-term nominal interest rates are longer in coverage than data on short-term nominal interest rates.

[36] This calibration is based on the average growth rate of the aggregate housing stock during the available samples for Japan and the United States. A long-run average growth rate of about 2% is also implied by the quantity index for shelter consumption in Denmark and Finland (OECD data) and the United States (BEA data) and by available data on the housing stock for Ireland (Central Statistics Office Ireland).

ψ, the data on μ_t determine Ψ_t and z_t.

4.2 Estimation method

The parameters are split into two sets, whereby parameters in set Υ_1 are estimated outside of the asset pricing model. Parameters in set Υ_2 are then estimated within the asset pricing model, conditionally on Υ_1. The set Υ_1 concerns the stochastic processes for the observable exogenous variables, specifically g_t, n_t, and r_t: $\Upsilon_1 = \{\nu_g, \theta_g, \sigma_g, \varsigma_g; \nu_n, \phi_n, \theta_n, \sigma_n, \varsigma_n; \nu_r, \phi_r, \sigma_r\}$. Note that Υ_1 includes four of the eight parameters showing up in the equilibrium coefficients (13)-(15) determining the responses of $\log x_t$ to shocks: the persistence parameters θ_g, ϕ_n, θ_n, and ϕ_r. In particular, θ_g, ϕ_n, and θ_n govern the expectations of the future values of g_t and n_t, the two variables whose dynamics turn out to be quantitatively the most important in accounting for the three house price patterns.

We estimate Υ_1 using Bayesian state-space methods, with the latent state variables $s_{g,t}$ and $s_{n,t}$ obtained by the Kalman filter. As we are attempting to estimate potentially highly persistent processes, to improve the precision of the estimates, the parameters are estimated on a panel of all 12 countries. In the panel, the persistence and variance parameters of a given process are common across all 12 countries but the intercept is allowed to be country specific; the latent states $s_{g,t}$ and $s_{n,t}$ are country specific. The details of the panel estimation are contained in the Online Appendix. The panel is for 1951-2019, except Canada, which is for 1958-2019.[37]

The set Υ_2 consists of the remaining parameters: the housing premium, the two elasticities, the age-dependent preference parameters for housing, and the parameters of the stochastic process for the growth rate of the distributional variable Ψ_t. That is, $\Upsilon_2 = \{\delta, \varepsilon_1, \varepsilon_2, \psi, \nu_z, \phi_z, \theta_z, \sigma_z, \varsigma_z\}$, where $\psi = \{\psi_1, \psi_2, \psi_3, \psi_4, \psi_5\}$, with $\psi_2 = 1$ as a normalization. Four of the parameters in Υ_2 show up in the equilibrium coefficients (13)-(15): ε_1, ε_2, ϕ_z, and θ_z. The parameters in Υ_2 are estimated for each country separately, using the limited information Bayesian approach described in detail by Chernozhukov and Hong (2003); see also Christiano, Trabandt and Walentin (2010). Specifically, given (i) a country's data on c_t (proxied by real GDP per capita), N_t, n_t, r_t, and $\{\mu_t\}_{j=1}^{5}$, (ii) the estimated latent states $s_{g,t}$ and $s_{n,t}$ obtained in the previous step, and (iii) the parameters in Υ_1 estimated in the previous step, the parameters in Υ_2 are chosen so that the distance between the time path of the model-implied house prices, $q(\Upsilon_2)$, and the actual house price data, q^{data}, is as small

[37] Monte Carlo simulations are used to test whether the estimation method can uncover the parameters and latent states. This is discussed in Section 5.3.

as possible. The data are for 1951-2019 (Canada for 1958-2019); the model-implied house prices are given by equation (11). For each country, the data points give us 69 moments for the quasi-likelihood function.[38] The quasi-likelihood function is given by

$$F(q^{data}|\Upsilon_2) = \left(\frac{1}{2\pi}\right)^{\frac{T}{2}} |V|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\left(q^{data} - q\left(\Upsilon_2\right)\right)' V^{-1}\left(q^{data} - q\left(\Upsilon_2\right)\right)\right),$$

where T denotes the number of elements in q^{data} and V is a weighting matrix. In our application, V is chosen to be the identity matrix. The quasi-posterior distribution is defined as

$$F\left(\Upsilon_2|q^{data}\right) \propto F(q^{data}|\Upsilon_2)p\left(\Upsilon_2\right)$$

where $p\left(\Upsilon_2\right)$ denotes the prior distribution. In the presence of a potentially persistent unobserved state $s_{z,t}$, the limited information approach is better behaved in finite samples than a full-information likelihood approach.

We use a random walk Metropolis-Hastings algorithm to approximate the posterior distribution. In each iteration, the algorithm consists of the following steps. First, we draw a candidate parameter vector from the normal density $\Upsilon_2^{new} \sim N\left(\Upsilon_2^{old}, \Omega\right)$, where $\Omega = \lambda \times \text{var}\left(\Upsilon_2\right)$, $\text{var}\left(\Upsilon_2\right)$ is an estimate of the variance of the parameters, and λ is a scaling factor. Second, we accept the draw with probability

$$\alpha = \min\left(1, \frac{F\left(\Upsilon_2^{new}|q^{data}\right)}{F\left(\Upsilon_2^{old}|q^{data}\right)}\right).$$

The total number of iterations is set to 505,000 and we save every 50th draw after a burn-in of 5,000. The unobserved state $s_{z,t}$ is obtained by the Kalman filter.[39]

4.3 Parameter estimates

Table 2, panel A, contains the results of the panel estimation of parameters Υ_1. For space constraints, we report only the common persistence and variance parameters; the country-specific intercepts,

[38] In addition to the house price data, the vector q^{data} includes, as an additional moment, the long-run ratio of house prices to the prices of housing service (ie, the long-run x), which is set equal to 20 for all countries in the sample. This roughly corresponds to the average price-rent ratio reported by various studies in the literature (either in a cross-section or time series). The model vector $q\left(\Upsilon_2\right)$ contains its model counterpart. This ratio pins down the housing premium parameter δ. There are thus in total 70 moments in the quasi-likelihood function.

[39] The starting values of the parameters are obtained by maximizing the log posterior using the covariance matrix adaption algorithm (CMA-ES). Then, an initial run of the Metropolis-Hastings algorithm is used to approximate $\text{var}\left(\Upsilon_2\right)$. We choose the scaling factor λ so that the acceptance rate is about 20%.

which do not affect house price dynamics, are not reported. The 90% error bands are reported in the parentheses. The estimates show that the latent state variable $s_{g,t}$ in the process for g_t is persistent, with the median of the posterior distribution of θ_g equal to 0.9387. We therefore interpret $s_{g,t}$ as the long-run component of the growth rate of GDP per capita. The process for the population growth rate has the medians of the posteriors of ϕ_n and θ_n equal to 0.8694 and 0.9852, respectively. Referring back to the discussion in Section 3.2, the estimates imply that $s_{n,t}$ can be approximately interpreted as the expected change in the population growth rate and that the change is very persistent. The process for population growth is thus characterized by prolonged waves. For the real interest rate, the estimated process has a median posterior autocorrelation of 0.66.

Panel B of Table 2 contains the results of the estimation of parameters Υ_2 based on the country-specific quasi-likelihood function. In the table we report the priors (common across the countries) and the cross-country median and standard deviation of the medians of the country-specific posterior distributions. The individual medians of the country-specific posterior distributions, and the 90% error bands, are contained in the Online Appendix.

The mean of the prior distribution of δ is set equal to 0.06 to reflect that long-run total housing return in many countries is about 7% per year (Jordá et al., 2017) and that the long-run real interest rate is about 1%. The means of the prior distributions of ε_1 and ε_2 are set equal to one, implying a log utility function. The mean of the prior for ψ approximates the typical housing consumption profile over the life-cycle (Mankiw and Weil, 1989; Eichholtz and Lindenthal, 2014; Lisack, Sajedi and Thwaites, 2017); with ψ_2 normalized to equal to one.[40] Gamma distribution is assumed for the priors of δ, ε_1, ε_2, and Ψ to ensure positive values. The priors for the parameters of the stochastic process for z_t are based on a simple pre-estimation of the process; beta distribution is assumed to ensure they lie between zero and one.[41]

In the following discussion of the posterior estimates, we always refer to the cross-country median

[40] This is a profile for housing consumption, not necessarily home ownership (for instance, the estimates of Eichholtz and Lindenthal, 2014, are based on floor space). Although in our model housing consumption over the life-cycle is strictly speaking determined by preferences, it is more appropriate to think about the parameter vector ψ as a projection of housing consumption on age that picks up various other factors, such as credit constraints. This is the interpretation of Mankiw and Weil (1989). To keep the estimation manageable, we assume that ψ is time-invariant. The model thus cannot speak to some of the socio-demographic changes studied by Gervais and Fisher (2011). As in Mankiw and Weil (1989) and others, consumption of young dependents is treated separately. To map it into a household consumption, housing consumption of the category 0-24 would have to be assigned in some proportions to the "adult" categories.

[41] In the pre-estimation, Ψ_t is constructed at the priors of ε_1 and ψ. Then, z_t is derived as a log difference of Ψ_t and $s_{z,t}$ is constructed as a moving average of z_t. The parameter estimates of the process used as priors are then obtained by OLS.

of the medians of the country-specific distributions. The estimate of δ is equal to 0.052. The change relative to the common prior reflects differences in the country-specific real interest rates and housing capital gains. For the elasticities, we obtain $\varepsilon_1 = 0.67$ and $\varepsilon_2 = 1.17$. These values imply increasing share of housing expenditures in total consumption as income increases. Thus, as countries get richer, the share of housing expenditures increases. Knoll et al. (2017) report that this has been happening in many countries since 1950.[42] The estimates of ψ yield a life-cycle pattern with a peak at the age category 55-69. This is similar to Eichholtz and Lindenthal (2014) and Lisack et al. (2017) and could reflect, for instance, inheritance received by people in this age category. Based on European data, Wind, Dewilde and Doling (2020) report that a non-negligible fraction of households of such age have a secondary property not used for rental purposes.[43] Finally, the estimates of the parameters of the process for z_t imply that the latent component $s_{z,t}$ is persistent, although z_t itself is not. The latent component thus has an interpretation as a long-run component of the growth rate of housing demand related to the age structure of the population.

5 Results

This section starts by presenting summary statistics of the model-implied house prices for the 12 individual countries. Second, it compares the model with the three patterns of international house prices established in Section 2. Third, it presents the results of specification tests and robustness checks concerning the ability of the model to account for the three house price patterns. Fourth, it shows the importance of expectations about future fundamentals in accounting for the three house price patterns and carries out a decomposition into the contributing factors. Finally, it considers a counterfactual experiment intended to gauge the effect of population ageing on house prices since 1950.

[42] As in the case of the age-dependent preference parameters, the elasticities may be picking up factors not modeled explicitly. Interestingly, the cross-country median of ε_1 is in the ballpark of a regression-based estimate of Takáts (2012) obtained on a panel of 22 OECD countries 1970-2009. The relative magnitudes of the two elasticities also conform with the structural model of Chambers et al. (2009).

[43] For the United States, the medians of the posterior distributions for the categories 40-54 and 55-69 are similar (see the Online Appendix), which conforms with Mankiw and Weil (1989).

5.1 Summary statistics of model-implied house prices

For each of the 12 countries, the chart on the left hand side of Figure 2 plots the standard deviation of the percentage deviations of house prices from trend. The standard deviation is calculated using the model-implied house prices, 1951-2019 (1958-2019 for Canada), at the posterior distribution of the parameter estimates. The chart thus plots a whole distribution of the standard deviation for each country. The chart on the right hand side of Figure 2 does the same for the first-order autocorrelation of the percentage deviations. In summary, the results (discussed below) imply that house prices are characterized by large and persistent deviations from trend, easily lasting for almost a decade.

Regarding the standard deviation, the cross-country median of the medians of the country-specific posterior distributions is 15% and three standard deviations correspond to almost 50% change in house prices relative to trend. Housing market analysts and commentators often calculate statistics such as the price-to-income ratio as an indicator of house price under- or overvaluation or as a measure of housing affordability. Our results indicate that movements in such measures of 15% to 50% should not be surprising, as they lie in one to three standard deviations of the departures of house prices from trend. The lowest median of the standard deviation, 8%, is found for the United States, whereas the highest median, 48%, is found for Japan, followed by France, 25%. The high values for Japan and France reflect the house price boom in these countries in the decades after WWII, as well as the house price bust after 1991 in the case of Japan (refer back to Figure 1). Observe also that in terms of the standard deviation, Switzerland looks just like the majority of the countries, even though it exhibits only a modest long-run trend.

For the autocorrelation, the cross-country median of the medians of the country-specific posterior distributions is 0.91. This implies half-life of seven years and eight months. The lowest median of the autocorrelation, 0.69, is found for the United States, while the highest, 0.97, is found for Japan.

5.2 The three house price patterns in the model

Figure 3 compares the three house price patterns established in Section 2 with their model counterparts. Specifically, it plots house prices in Japan, the G10 countries, and Switzerland against the respective model-implied house prices. In the case of the model, we plot the median and the 90% error bands based on the posterior distributions. For the G10, both the data and the model are represented by the 1st PC of the ten countries. The loadings on the 1st PC in the model are comparable to those in the data. At the medians, the loadings are: 0.30 (AUS), 0.34 (BEL), 0.34 (CAN), 0.31 (DNK), 0.26 (FIN), 0.32 (FRA), 0.23 (NLD), 0.34 (SWE), 0.37 (GBR), 0.33 (USA). See Table 1 for the data counterparts.

Starting with the pattern for Japan, the model succeeds in generating both the long-lasting boom and bust and tracks the data closely with narrow 90% error bands. A discrepancy between the model and the data occurs in the period after the global financial crises, during which Japan finally experienced a mild recovery in house prices. The model generates faster recovery (and a subsequent decline).

For the G10, the model succeeds in generating the pick up in house price growth in the early 1990s. The timing of the fast growth period is slightly different, starting in 1992 and ending in 2005, compared with 1993-2007 in the data. In the model, based on the median path of the 1st PC, the average growth rates are 2.2% during 1951-1992 and 4.2% during 1992-2005. This compares with 1.7% during 1951-1993 and 4.8% during 1993-2007 in the data. The model also generates the double-dip in house prices in the early and late 1980s, as well as a decline in house prices (and subsequent recovery) around the global financial crisis in 2007, although the magnitudes in the model and the data somewhat differ.[44]

Finally, the model generates the recurrent fluctuations in house prices in Switzerland, in the absence of a pronounced long-run growth. The model tracks the fluctuations in the data, except during 2000-2010.

5.3 Specification tests and robustness

Before proceeding further, we carry out the following tests to check the robustness of our specifications and findings.

Housing supply. For Japan and the United States, we utilize the data on the aggregate housing stock and replace the assumption of a constant growth rate of the housing stock with process (10). This process is estimated on a mini panel of the two countries using the standard Bayesian state space method. The housing stock data and the estimates of the latent state variable $s_{h,t}$ are then fed into the model together with the other exogenous variables. This extension has only a limited effect on the results for these two countries (see the Online Appendix). The housing stock simply does not

[44] The fast-growth period starts off from a higher level of house prices in the model than in the data because the double-dip in the 1980s in the model is not as large as in the data.

vary enough to significantly change the results, providing support for the hypothesis of Knoll et al. (2017) that in the post-WWII period housing supply has been relatively unresponsive to demand factors.[45]

Restriction on elasticities. We have re-estimated the model under the restriction $\varepsilon_1 = \varepsilon_2 = 1$ for all countries. This implies log preferences and a constant share of housing expenditures in total consumption. Observe that under this restriction the coefficients determining the responses of $\log x_t$ to the population growth rate, and its latent state, and to the latent state of the growth rate of GDP per capita are the same across the 12 countries (recall that the persistence parameters are the same across the countries, estimated in the panel); the coefficients γ_z, γ_{zs}, and γ_r are the same across the countries even before imposing the restriction on the elasticities. In addition, under the restriction $\varepsilon_1 = \varepsilon_2 = 1$, the coefficients of the response of $\log d_t$ to the level of population and GDP per capita are also the same across the 12 countries. The differences across countries in the simulated $\log q_t$ thus come only from different realizations of the exogenous variables, not from the differences in the respective coefficients. The impact of this restriction, however, is small (see the Online Appendix), implying that the differences in house prices across the countries are due to different histories of the exogenous variables.

Real interest rate process. We have experimented with a more general specification than the AR(1) process used by Bansal and Lundblad (2002). Namely, $r_{t+1} = \nu_r + \phi_r r_t + s_{r,t} + \sigma_r \xi_{r,t+1}$ and $s_{r,t+1} = \theta_r s_{r,t} + \varsigma_r \zeta_{r,t+1}$, where $\xi_{r,t+1}$ and $\zeta_{r,t+1}$ are normally distributed innovations. Like the other processes, this process is estimated on the panel of the 12 countries using the Bayesian state space method. The posterior estimates of θ_r and ς_r, however, are close to zero, while ϕ_r and σ_r are significantly positive. The data thus prefer the AR(1) process.

Recovery of the estimated processes. We test if the Bayesian state space method used to estimate the processes with the latent state variables can recover such processes from the data. To this end, we use the process (7) as the data generating process (DGP) and test the method on the artificial data (an artificial panel of 12 countries, with each series of the length of 70 periods). The method recovers the DGP. We also test if the method wrongly estimates process (7), when in fact the DGP for the growth rate is an iid process. Again, the method recovers the correct DGP process. We

[45] In Japan, for example, the growth rate of the housing stock is strongly positively correlated with the growth rate of house prices but quantitatively the changes in the housing stock and the latent state variable in its growth rate are not sufficient to offset the effects of the other factors on house prices.

also carry out this test for the more general process (8). Again, the method recovers the DGP. The details of the tests are contained in the Online Appendix.

Mutually correlated innovations. Instead of assuming mutually independent innovations ξ_r, ξ_g, ζ_g, ξ_n, ζ_n, we let these innovations to be mutually correlated and re-estimate the exogenous processes under this assumption. This exercise is meant to capture some of the interdependence between the exogenous processes discussed in Section 3.2. This generalization has, however, only a small effect on the estimates of the persistence parameters. The medians for ϕ_r, θ_g, ϕ_n, and θ_n become, respectively, 0.6911, 0.929, 0.8547, and 0.9818. These values are similar to those in Table 2.

5.4 Decomposition

Referring back to the discussion at the end of Section 3.1, the deviations from the stochastic trend depend only on the interest rate and expectations about future fundamentals, whereas the stochastic trend depends only on the current levels of the fundamentals. To demonstrate the role of the interest rate and expectations in the results of Section 5.2, Figure 4 plots the model-implied house prices against the model stochastic trend. The plots are based on the median paths. It is clear that the stochastic trends alone are unable to account for the key features of the three house price patterns. Based on the trend alone, the boom in Japan would be nowhere near as strong as in the data and there would be no bust, only flattening out of house prices. In the case of the G10, the trend misses the period of the fast house price growth from the early 1990s till the global financial crisis. And in the case of Switzerland, on the basis of the stochastic trend alone, there would be no recurrent house price swings.

We turn next to a decomposition of the three house price patterns into the contributions of the interest rate and expectations about the individual fundamentals. The decomposition is based again on the median paths. In the decomposition, we start off with only the state variables z_t and $s_{z,t}$. All other variables affecting the deviations are held at their post-WWII average. Thus, in the first step, the deviations are determined only by the state variables driving the forecasts of the growth rate of the distributional variable Ψ_t (ie, forecasts of the growth rate of aggregate demand for housing related to changes in the age structure of the population). We then add the state variables n_t and $s_{n,t}$, which determine expectations about population growth. The four state variables z_t, $s_{z,t}$, n_t, and $s_{n,t}$ thus characterize forecasts about the overall demographic changes. After that we add $s_{g,t}$, the state variable driving expectations about the growth rate of income per capita. And finally we add the interest rate r_t, which reproduces the model-implied path of house prices reported in Section 5.2. These experiments thus evaluate the marginal contributions of the variables to the deviations of house prices from the stochastic trend.[46]

5.4.1 Japan

Figure 5 shows the decomposition for Japan. When only expectations about the growth rate of the distributional variable Ψ_t affect the deviations of house prices from the stochastic trend (the upper-left chart), the figure shows that until about 1996, these expectations had a mild positive effect on house prices. However, after that, the effect turned negative, although not sufficiently enough to account for a significant share of the bust. These effects in the model reflect the dynamics of population ageing in the data. Japan experienced significant population ageing since 1950, with the distribution gradually shifting towards older categories, as discussed in more detail in Section 5.5. Until about 1996, expectations of population ageing had a positive effect on house prices, as the share of the middle aged kept increasing. After that, however, expectations of population ageing going forward started to weight down on house prices, as the share of the middle aged started to decline and the increase in the share of the 70+ accelerated.[47]

The upper-right chart shows the effect of expectations about both demographic variables, the age structure and the population growth rate, together. Expectations about future population growth had a significant positive effect on house prices between 1962 and 1976 and, together with expectations about population ageing, account for about half of the bust after 1991. The expectations in the model reflect the dynamics of the population growth rate in the data. After a decline in the aftermath of WWII, population growth gradually increased throughout the 1960s, reaching a peak in the early 1970s. After that, however, population growth embarked on a sustained decline lasting until the end of the sample period. In fact, since 2010 it has been negative (see the plot of the data

[46] The decomposition evaluates the marginal contribution of each of the factors. As noted at the end of Section 3.2, the factors can be driven by common deeper determinants. For instance, population numbers and the age composition of the population can be driven by fertility, longevity, and migration rates. Our decomposition evaluates only the marginal effects of any combination of these deeper determinants propagating through either total population or the age composition. In other words, we answer the question whether total population or the age composition was the more important margin of the effects of demographic changes on house prices.

[47] Recall that according to the life-cycle demand for housing in Table 2, demand for housing at the age category 70+ is lower than at the categories 40-54 and 55-69. Although the Table shows only the cross-country medians, the lower demand at the 70+ category reported in the Table is representative for Japan (see the Online Appendix).

in Figure 8).

The lower-left chart of Figure 5 shows the marginal contribution of expectations about the growth rate of income per capita. The chart shows that this effect essentially closes the gap between the full model and the effect of the two demographic variables. Expectations about future income growth were the single most important factor behind the house price boom and make up about half of the bust. These expectations reflect fast underlying growth in GDP per capita, picked up by the state variable $s_{g,t}$, that reached its peak in the second half of the 1960s. By the early 1970s, the growth rate dropped and stayed roughly constant until the second half of the 1980s. By 1991, however, it dropped further to close to zero and stayed in that region until the end of the sample period; see Figure 8. We return below in more detail to the model's ability to account for the turning point in house prices.

The real interest rate played a substantial role only during a couple of years in the early 1970s, when a drop in the real rate pushed temporarily house prices up, and after 2013, when low interest rates also contributed positively to house prices.

The model misses the turning point in 1991 only by two years (Figure 3). How is it possible that the model gets the turning point almost right? This is the result of two effects that growth rates of fundamentals have on house prices: by affecting the level of the respective fundamental, and thus the level of the stochastic trend of house prices, and by affecting the expectations that generate deviations of house prices from the stochastic trend. The growth rate of population in Japan has been declining since early- to mid-1970s, while staying positive until 2009. The decline in the growth rate had an increasingly negative effect on expectations, while still contributing positively to the level of the stochastic trend until 2009. Before 1989 the positive effect on the trend was dominating. After 1989, however, the negative effect on expectations started to dominate, turning the boom into a bust. A similar mechanism applies also to the effect on the turning point of the growth rate of GDP per capita.

5.4.2 G10 countries

Figure 6 carries out the decomposition for the G10. The upper-left chart shows the contribution of expectations about the growth rate of aggregate demand for housing related to the age structure of the population. The chart shows that since the 1980s, the contribution has been mildly positive.

This is similar to the situation in Japan prior to 1996 noted above.

The upper-right chart shows the joint contribution of expectations about the age structure and population growth. It shows that expectations about population growth had a positive marginal effect on house prices for most of the period since the mid-1980s. In particular, the expectations are important in accounting for the second half of the fast-growth period, having a strong positive effect on house prices between 2000 and 2007. This reflects a rebound in the population growth rate in the data that in most of the G10 countries reached a peak around 2009. At the peak, the growth rate was at the highest level since the late 1960s (see Figure 8).

The lower-left chart of Figure 6 shows the marginal contribution of expectations about the growth rate of income per capita. This variable closes most of the gap between the full model and the path generated by expectations about the two demographic variables. In particular, expectations about income growth generate the double-dip in house prices in the 1980s and account for the first half, 1992-1999, of the fast-growth period. In addition, the effect of these expectations around the global financial crisis is strongly negative. The positive contribution to the fast-growth period in house prices is due to a 1990s rebound in the underlying growth rate of GDP per capita, picked up by the state variable $s_{g,t}$, that in most of the G10 countries reached a peak around 1997. At the peak, the growth rate was at its highest level since the end of the 1960s (see Figure 8).

Regarding the marginal contribution of the interest rate, the most interesting part is the period after the global financial crisis. Many commentators posed the question whether the increase in house prices since the start of the recovery from the global financial crisis is due to loose monetary policy. To the extent that the real interest rate used in our analysis reflects monetary policy, the decomposition supports this view. At the end of our sample period, the gap between the full model and the version without the interest rate for the G10 countries is about 12%.

5.4.3 Switzerland

Finally, Figure 7 carries out the decomposition for Switzerland. Expectations about the growth rate of aggregate demand for housing related to the age structure of the population have essentially no effect on house prices. However, expectations about population growth have a strong effect. In fact, these expectations account for a bulk of the recurrent house price swings in Switzerland. They reflect cyclical movements in the population growth rate in the data that are unconditionally positively correlated (0.80) with net migration. Fluctuations in house prices due to these expectations are quantitatively interesting, reaching close to a 20% departure from the stochastic trend in 1960 and 1975, 12% in 1990, and 14% in 1995. Expectations about growth of income per capita close most of the remaining gap between the full model and the path generated by the two demographic variables. In appears that the recurrent fluctuations in house prices in Switzerland are related to the general business cycle, but the key factor is not as much the resulting expectations about income growth as expectations about net migration (see Figure 8 for the population growth rate data and the persistent component in the growth rate of GDP per capita).[48] The interest rate played a quantitatively interesting role in the early 1980s and towards the end of the sample period. At the end of the sample period, it contributes 6.5% to house prices implied by the full model.

5.5 A demographic counterfactual

In the previous section we have studied the contribution to house prices of expected future demand for housing driven by changes in the age structure of the population. This factor affects only the deviations of house prices from the stochastic trend. This section explores the contribution of population ageing to house prices working through *both* the deviations and the stochastic trend.

In this experiment, we generate model-implied house prices with the age structure of the population fixed at the 1950 distribution. A consequence of fixing the age distribution is that the level of aggregate demand related to the age composition of the population is constant and z_t and $s_{z,t}$ are equal to zero.[49] The results of this experiment, carried out at the medians of the individual-country posterior distributions, are contained in Table 3.

First, for each country, the table shows data on the change in the share of each age group in the population between 1950 and 2019. Population ageing in the 12 countries is immediately apparent. The largest losses in all countries are in the age category 0-24, while the largest gains are in the categories 55-69 and 70+. In addition, between these two age groups, only in Canada is the gain in the group 55-69 larger than in the group 70+. The shares of the categories 25-39 and 40-54 remained relatively unchanged in all countries. The most dramatic change occurred in Japan, which experienced a decline of 33 percentage points in the category 0-24 and an increase of 11 and 18

[48] Using an event study, a significant effect of net migration on house prices in Switzerland has been established also by Helfer, Grossmann and Osikominu (2023).

[49] Feeding in the data on the age distribution but assuming homogenous consumption across the age groups ($\psi_j = 1$ $\forall j$) would, up to a constant, have the same effect.

percentage points, respectively, in the categories 55-69 and 70+.

The right-most column in Table 3 reports the change in house prices—which, according to the model—results from the change in the population distribution between 1950 and 2019, relative to the constant 1950 distribution. In all 12 countries the contribution of aging population to house prices is positive. When the age structure of the population is kept at the 1950 distribution, house prices in 2019 are, in the sense of the cross-country median, 14% lower. The largest positive contribution is obtained for Canada, where at the 1950 distribution house prices would be 28% lower.[50] As noted above, Canada is the only country that had, across the five age groups, the biggest gain in the category 55-69, which is the largest consumer of housing services according to the estimated parameters. Japan has the second largest contribution of population ageing to house prices, where house prices would otherwise be 24% lower. The contribution in Japan is lower than in Canada as the biggest change in the distribution has already shifted further into the category 70+, which consumes less housing services than the category 55-69. The smallest effect of population ageing on house prices is obtained for Sweden and the United States, where house prices at the 1950 distribution would be only 9% and 10% lower, respectively.

For the United States, the classic paper by Mankiw and Weil (1989) predicts a peak in the growth rate of housing demand due to changes in the age distribution in the late 1970s to early 1980s. Afterwards, their model predicts a continuous decline in the growth rate of housing demand and, given a historical reduced-form relationship between housing demand and house prices, a substantial decline in house prices starting in the 1990s. The measure of aggregate demand for housing in their paper is equivalent to our Ψ_t with $\varepsilon_1 = 1$. The authors split the population by year and estimate the age-dependent ψ's from micro data. We have five age groups and estimate the coefficients from the aggregate time series, as described in Section 4.2, however using micro-level estimates as a prior. The resulting posterior estimates remain comparable to the micro-level estimates of Mankiw and Weil (1989).[51] Our model, however, generates a peak in the growth rate of Ψ_t (ie, a peak in z_t) in 2000, two decades later than predicted by Mankiw and Weil (1989). As the age-dependent housing consumption profile is similar across our and their studies, where does the difference in results come from? It comes from the United States experiencing a less dramatic shift towards the category

[50] To be precise, for Canada the starting period is 1957, due to the limited data span.

[51] See the Online Appendix for the estimates for the United States. The corresponding estimates in Mankiw and Weil (1989) are in their Figure 3. Unlike in the cross-country median in Table 2, the posterior for the United States exhibits similar values for the age groups 40-54 and 55-69, as in Mankiw and Weil (1989).

70+ than predicted by the authors on the basis of the 1983 Census Bureau's fertility and mortality forecasts. Population ageing in the United States has been slowed down largely by net migration, which was not predicted in 1983.

6 Conclusion

A parsimonious model was used to account for country-level house prices during the period 1950-2019 in a sample of 12 advanced economies. The model ties house prices to a small number of fundamentals and the real interest rate. In previous research, the role of fundamentals in explaining house prices has been found surprisingly weak. The key elements of our model are persistent random components in the stochastic processes for the growth rates of the fundamentals. In the model, shocks to these components result in large and persistent changes in expectations about future fundamentals, generating large and persistent house price swings around a stochastic trend, easily lasting for as long as a decade. At times, such deviations from current fundamentals may appear as 'affordability crises', as house prices substantially increase above current incomes.

When the observable fundamentals and the estimated persistent components are fed into the model, the model accounts well for the three patterns of house prices in the post-WWII period that, separately, characterize Japan, a group of 10 advanced economies (the G10), and Switzerland. The most remarkable result is that the model reproduces almost exactly the spectacular decades-long boom and bust in Japan, by far the biggest house price swing in the post-WWII history. The model also generates the boom that started in the G10 in the early 1990s, as well as the large cyclical fluctuations around a weak trend in Switzerland. An important aspect of these results is that the driving factors are estimated outside of the asset pricing model.

According to the model, the three historical patterns result from large and persistent deviations of house prices from respective stochastic trends. What drives the deviations? Expectations about future growth of income per capita and population are the two most important factors accounting for the three house price patterns. Expectations about future population ageing (demand for housing services due to changes in the age structure of the population) also play a role, but less important one than the other two factors. In Japan, such expectations are already having a negative effect on house prices, while in the G10 the effect is still moderately positive. When also the effect of population ageing on the stochastic trend is taken into account, in all countries in the sample, ageing population had, so far, a significantly positive effect on house prices. Finally, the decline in interest rates after the global financial crisis had an important positive effect on house prices in the last ten years of the sample.

The model intentionally abstracts from institutional details, the structure of mortgage markets, various housing market frictions, and government policies, which differ across countries. While the recent macro-housing literature has shown that such considerations are important, our goal was to propose a theory that could be informative about house prices irrespective of a country's institutional and market environments. It appears that the few factors considered in our model are sufficient to explain many of the major house price developments over the past 70 years.

Our focus has been on country-level house prices. However, there is no reason why the same mechanism could not be applied to a panel of locations within a single country. Such an application would be a further test of the theory.

A broader message of our analysis regards the role of growth shocks. The recent macro-housing literature has paid special attention to various housing and mortgage market frictions. Nonetheless, the literature had a difficult time generating house price movements of the magnitudes and persistence observed in the data. Most of the models rely on mean-reverting shocks to the levels of exogenous variables. Our analysis points to growth shocks as promising and realistic sources of house price swings that could be incorporated into macro-housing models. Richer and more structural models than the parsimonious model used for the purposes of this study could be used to investigate the interaction of growth shocks with mortgage market frictions and their effects on different types of households.

References

ALBANESI, S., G. D. GIORGI AND J. NOSAL, "Credit Growth and the Financial Crisis: A New Narrative," *Journal of Monetary Economics* 132 (2022), 118–139.

ANDRÉ, C., "A Bird's Eye View of OECD Housing Markets," Economics Department Working Papers 746, OECD, 2010.

ANTOLIN-DIAZ, J., T. DRECHSEL AND I. PETRELLA, "Tracking the Slowdown in Long-Run GDP Growth," *Review of Economics and Statistics* 99 (2017), 343–56.

ARESTIS, P. AND A. R. GONZÁLEZ, "Modelling the Housing Market in OECD Countries," *International Review of Applied Economics* 28 (2014), 131–53.

ARSLAN, Y., B. GULER AND B. KURUSCU, "Credit Supply, Housing Demand, and Boom-Bust Cycles," Unpublished manuscript, 2023.

BANSAL, R. AND C. LUNDBLAD, "Market Efficiency, Asset Returns, and the Size of the Risk Premium in Global Equity Markets," *Journal of Econometrics* 109 (2002), 195–237.

BANSAL, R. AND A. YARON, "Risks for the Long-Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance* 59 (2004), 1481–1509.

BARIGOZZI, M., M. LIPPI AND M. LUCIANI, "Large-dimensional Dynamic Factor Models: Estimation of ImpulseResponse Functions with cointegrated factors," *Journal of Econometrics* 221 (2021), 455–82.

BARSKY, R. B. AND J. B. DELONG, "Why Does the Stock Market Fluctuate?," *Quarterly Journal of Economics* 108 (1993), 291–311.

BRAUN, S. AND G. S. LEE, "The Prices of Residential Land in German Counties," *Regional Science and Urban Economics* 89 (2021), 1–36.

CAMPBELL, J. Y. AND R. J. SHILLER, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies* 1 (1988), 195–228.

CAMPBELL, S. D., M. A. DAVIS, J. GALLIN AND R. F. MARTIN, "What Moves Housing Markets: A Variance Decomposition of the Rent-Price Ratio," *Journal of Urban Economics* 66 (2009), 90–102.

CASE, K. E. AND R. J. SHILLER, "Is There a Bubble in the Housing Market?," *Brookings Papers on Economic Activity* 2 (2003), 299–342.

CHAMBERS, M., C. GARRIGA AND D. SCHLAGENHAUF, "Accounting for Changes in the Homeownership Rate," *International Economic Review* 50 (2009), 677–726.

CHERNOZHUKOV, V. AND H. HONG, "An MCMC approach to classical estimation," *Journal of Econometrics* 115 (2003), 293–346.

CHRISTIANO, L. J., M. TRABANDT AND K. WALENTIN, "DSGE Models for Monetary Policy Analysis," in B. M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*volume 3 (Elsevier Science, 2010), 285–367.

DAVIS, M. A. AND J. HEATHCOTE, "Housing and the Business Cycle," *International Economic Review* 46 (2005), 751–84.

———, “The Price and Quantity of Residential Land in the United States,” *Journal of Monetary Economics* 54 (2007), 2595–620.

Davis, M. A. and S. Van Nieuwerburgh, “Housing, Finance, and the Macroeconomy,” in G. Duranton, J. V. Henderson and W. C. Strange, eds., *Handbook of Regional and Urban Economics* (Science Direct, 2015).

Eichholtz, P. and T. Lindenthal, “Demographics, Human Capital, and the Demand for Housing,” *Journal of Housing Economics* 26 (2014), 19–32.

Favilukis, J., S. C. Ludvigson and S. Van Nieuwerburgh, “The Macroeconomic Effects of Housing Wealth, Housing Finance and Limited Risk Sharing in General Equilibrium,” *Journal of Political Economy* 125 (2017), 140–223.

Gallin, J., “The Long-Run Relationship between House Prices and Income: Evidence from Local Housing Markets,” *Real Estate Economics* 34 (2006), 417–38.

Garriga, C. and A. Hedlund, “Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession,” *American Economic Review* 110 (2020), 1603–34.

Garriga, C., A. Hedlund, Y. Tang and P. Wang, “Rural-Urban Migration, Structural Transformation, and Housing Markets in China,” *AEJ: Macroeconomics* 15 (2023), 413–40.

Garriga, C., R. Manuelli and A. Peralta-Alva, “A Macroeconomic Model of Price Swings in the Housing Market,” *American Economic Review* 109 (2019), 2036–72.

Geng, N., “Fundamental Drivers of House Prices in Advanced Economies,” Working Paper WP/18/164, International Monetary Fund, 2018.

Gervais, M. and J. Fisher, “Why Has Home Ownership Fallen Among the Young?*,” *International Economic Review* 52 (2011), 883–912.

Glaeser, E. L. and J. Gyourko, “Arbitrage in Housing Markets,” NBER Working Paper 13704, 2007.

Green, R. and P. H. Hendershott, “Age, Housing Demand, and Real House Prices,” *Regional Science and Urban Economics* 26 (1996), 465–80.

Greenwald, D. L. and A. Guren, “Do Credit Conditions Move House Prices?,” Unpublished manuscript, 2024.

Grossmann, V., B. Larin and T. Steger, “Das House Kapital: A Long-Run Theory of House Prices and Housing Wealth,” Unpublished manuscript, 2023.

Guerrieri, V. and H. Uhlig, “Housing and Credit Markets,” in J. B. Taylor and H. Uhlig, eds., *Handbook of Macroeconomics* (Elsevier, 2016).

Hamilton, B. W., “The Baby Boom, the Baby Bust, and the Housing Market: A Second Look,” *Regional Science and Urban Economics* 21 (1991), 547–52.

Helfer, F., V. Grossmann and A. Osikominu, “How Does Immigration Affect Housing Costs in Switzerland,” *Swiss Journal of Economics and Statistics* 159 (2023), 1–31.

Hirata, H., M. A. Kose, C. Otrok and M. E. Terrones, “Global House Price Fluctuations: Synchronization and Determinants,” NBER Working Paper 18362, 2012.

Holly, S., M. H. Paseran and T. Yamagata, "A Spatio-Temporal Model of House Prices in the USA," *Journal of Econometrics* 158 (2010), 160–73.

Jackson, L. E., M. A. Kose, C. Otrok and M. T. Owyang, "Specification and Estimation of Bayesian Dynamic Factor Models: A Monte Carlo Analysis with an Application to Global House Price Comovement," in *Dynamic Factor Models (Advances in Econometrics, Vol. 35)* (Emerald Group Publishing Limited, Bingley, 2016).

Jones, C., "The Facts of Economic Growth," in J. Taylor and H. Uhlig, eds., *Handbook of Macroeconomics* (Amsterdam: North Holland, 2016).

Jordá, O., M. Schularick and A. M. Taylor, "Macrofinancial History and the New Business Cycle Facts," in M. Eichenbaum and J. A. Parker, eds., *NBER Macroeconomics Annual 2016* (University of Chicago Press, 2017).

Justiniano, A., G. E. Primiceri and A. Tambalotti, "Credit Supply and the Housing Boom," *Journal of Political Economy* 127 (2019), 1317–50.

Kaplan, G., K. Mitman and G. L. Violante, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy* 128 (2020), 328545.

Knoll, K., M. Schularick and T. Steger, "No Price Like Home: Global House Prices, 1870-2012," *American Economic Review* 107 (2017), 331–53.

Kydland, F. E., P. Rupert and R. Šustek, "Housing Dynamics over the Business Cycle," *International Economic Review* 57 (2016), 1149–77.

Landvoigt, T., M. Piazzesi and M. Schneider, "The Housing Market(s) of San Diego," *American Economic Review* 105 (2015), 1371–1407.

Lisack, N., R. Sajedi and G. Thwaites, "Demographic Trends and the Real Interest Rate," Working Paper 701, Bank of England, 2017.

Maliar, L. and S. Maliar, "Heterogeneity in Capital and Skills in a Neoclassical Stochastic Growth Model," *Journal of Economic Dynamics and Control* 25 (2001), 1367–97.

———, "The Representative Consumer in the Neoclassical Growth Model with Idiosyncratic Shocks," *Review of Economic Dynamics* 6 (2003), 362–80.

Mankiw, G. and D. N. Weil, "The Baby Boom, the Baby Bust, and the Housing Market," *Regional Science and Urban Economics* 19 (1989), 235–58.

Martin, R. F., "The Baby Boom: Predictability in House Prices and Interest Rates," International Finance Division Paper 847, Federal Reserve Board, 2005.

McCarthy, J. and R. W. Peach, "Are Home Prices the Next Bubble?," *FRBNY Economic Policy Review* December (2004), 1–17.

Pflaumer, P., "Forecasting U.S. Population Totals with the BoxJenkins Approach," *International Journal of Forecasting* 8 (1992), 329–338.

Piazzesi, M. and M. Schneider, "Housing and Macroeconomics," in J. B. Taylor and H. Uhlig, eds., *Handbook of Macroeconomics* (Elsevier, 2016).

Plazzi, A., W. Torous and R. Valkanov, "Expected Returns and Expected Growth in Rents of Commercial Real Estate," *Review of Financial Studies* 23 (2010), 3469–519.

Poterba, J. M., "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach," *Quarterly Journal of Economics* 99 (1984), 729–52.

———, "House Price Dynamics: The Role of Tax Policy and Demography," *Brookings Papers on Economic Activity* 1991 (1991), 143–83.

Swan, C., "A Model of Rental and Owner-Occupied Housing," *Journal of Urban Economics* 16 (1984), 297–316.

Takáts, E., "Ageing and House Prices," *Journal of Housing Economics* 21 (2012), 131–141.

Topel, R. and S. Rosen, "Housing Investment in the United States," *Journal of Political Economy* 96 (1988), 718–40.

Wind, B., C. Dewilde and J. Doling, "Secondary Property Ownership in Europe: Contributing to Asset-Based Welfare Strategies and the Really Big Trade-Off," *International Journal of Housing Policy* 20 (2020), 25–52.

Table 1: Loadings of house price data on the 1st principal component, 1950-2019

AUS	BEL	CAN	DNK	FIN	FRA	JPN	NLD	SWE	CHE	GBR	USA
0.21	0.32	0.32	0.35	0.22	0.35	-0.04	0.23	0.35	0.12	0.41	0.34

Notes: Based on the method of Barigozzi et al. (2021) for non-stationary data.

Table 2: Parameter estimates

A. Panel estimation of the exogenous processes (Υ_1)

ϕ_r	σ_r	θ_g	σ_g	ς_g	ϕ_n	θ_n	σ_n	ς_n
0.6620	0.0233	0.9387	0.0188	0.0085	0.8694	0.9852	2.71e-5	3.14e-4
[0.6211, 0.7036]	[0.0224, 0.0243]	[0.9038, 0.9639]	[0.0176, 0.0200]	[0.0070 0.0103]	[0.8645, 0.8771]	[0.9724, 0.9975]	[2.23e-05, 3.35e-05]	[3.02e-04 3.28e-04]

B. Country-specific quasi-likelihood estimation (Υ_2)

	Common prior distribution					Median of the posterior dist.	
	Type	Mean	Variance	LB	UB	Cross-country Median	Std
δ	Gamma	0.06	0.001	0.01	0.1	0.052	0.0058
ε_1	Gamma	1.0	0.3	0.01	10.0	0.67	0.131
ε_2	Gamma	1.0	0.3	0.01	10.0	1.17	0.370
ψ_1	Gamma	0.5	0.1	0.01	10.0	0.40	0.049
ψ_2 *		1.0				1.0	
ψ_3	Gamma	2.0	0.5	0.01	10.0	1.76	0.75
ψ_4	Gamma	2.0	0.5	0.01	10.0	2.15	0.61
ψ_5	Gamma	1.5	0.5	0.01	10.0	1.46	0.32
ϕ_z	Beta	0.3	0.01	0.01	0.999	0.28	0.012
θ_z	Beta	0.9	0.01	0.01	0.999	0.95	0.025
ν_z	Normal	0	0.1	-1.0	1.0	0.014	0.0077
σ_z	Log-normal	0.004	3e-5	1e-9	7.34	0.0026	0.00087
ς_z	Log-normal	0.0006	5e-7	1e-9	7.34	0.00029	0.00011

*ψ_2 is normalized to equal to one.

Notes: In panel A, the 90% error bands are reported in the parentheses. The constant terms in the stochastic processes (ν_r, ν_g, ν_n) are allowed to be country-specific. They are unimportant for the dynamics of the model and for space constraints are not reported. In panel B, only the cross-country median and standard deviation of the medians in the country-specific posterior distributions are reported. The medians and the 90% error bands of the country-specific posterior distributions are reported in the Online Appendix. The age categories are: 0-24 (group 1), 25-39 (group 2), 40-54 (group 3), 55-69 (group 4), 70+ (group 5). The growth rate of the housing stock, h, is calibrated to 0.02 for all countries.

Table 3: The effect of ageing population on house prices, 1950-2019

	1950-2019 change in the share of					Ratio of 2019 house prices under 1950 distribution to under actual distribution
	0-24	25-39	40-54	55-69	70+	
AUS	-0.10	-0.02	0.01	0.04	0.06	0.86
BEL	-0.07	-0.01	-0.02	0.04	0.07	0.87
CAN	-0.20	0.00	0.03	0.10	0.07	0.72
DNK	-0.11	-0.04	0.01	0.05	0.09	0.86
FIN	-0.19	-0.02	0.00	0.09	0.12	0.79
FRA	-0.08	-0.01	-0.01	0.04	0.07	0.88
JPN	-0.33	-0.03	0.07	0.11	0.18	0.76
NLD	-0.18	-0.03	0.03	0.09	0.10	0.80
SWE	-0.07	-0.03	-0.02	0.03	0.09	0.91
CHE	-0.13	-0.01	0.00	0.05	0.08	0.89
GBR	-0.06	-0.02	-0.02	0.03	0.07	0.87
USA	-0.10	-0.02	0.00	0.06	0.06	0.90

Notes: The table shows the combined effect of the age distribution on the stochastic trend and the deviations of house prices. The results are computed at the medians of the posterior distributions. The change in the age distribution may not add up to zero due to rounding.

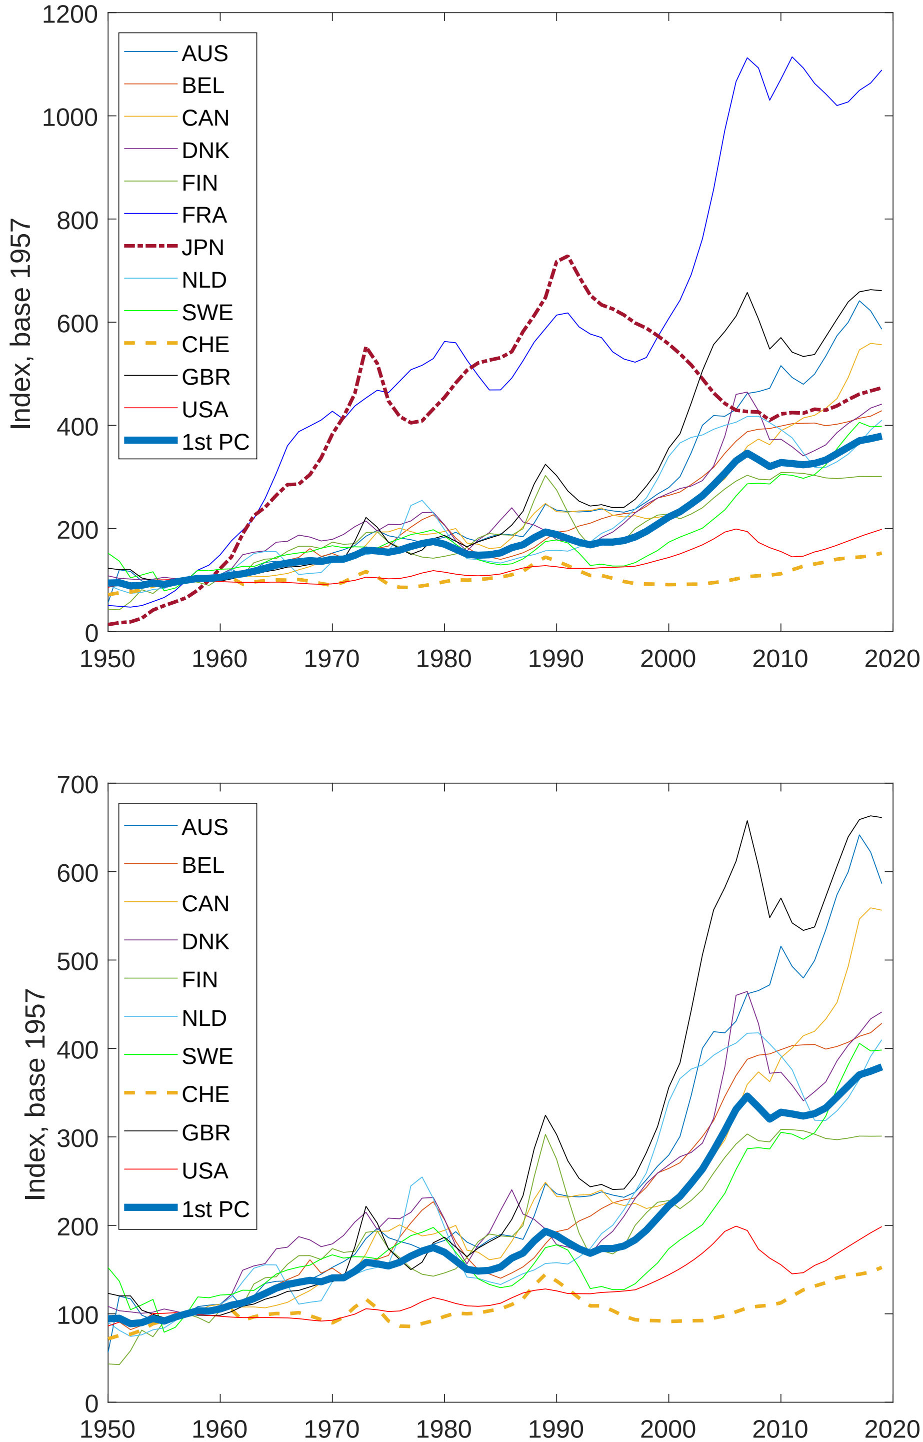

Figure 1: Real house price index, 1957 = 100. The sample is 1950-2019, except CAN, which is from 1957. The thick blue line is the 1st principal component of the countries in the sample, computed using the method of Barigozzi et al. (2021), which extends the principal component analysis to non-stationary data. The 1st principal component is representative of house prices in AUS, BEL, CAN, DNK, FIN, FRA, NLD, SWE, GBR, and USA (the 'G10 countries'). The bottom chart is a zoom-in of the upper chart by removing FRA and JPN.

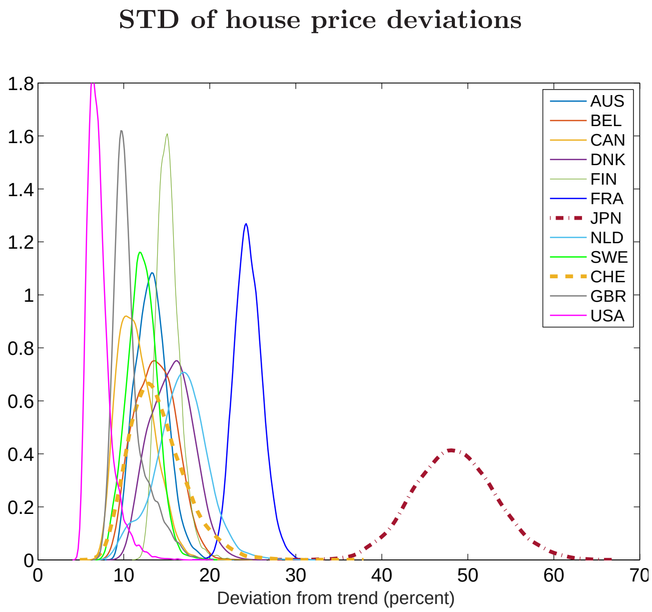

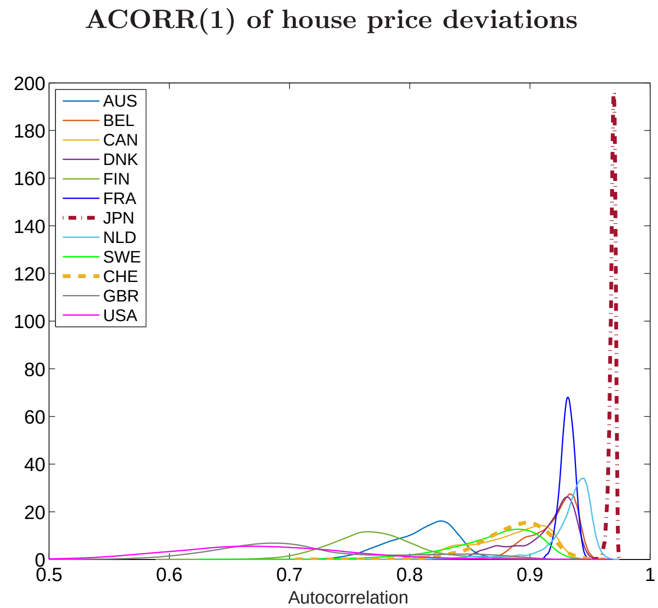

Figure 2: Properties of the model. The standard deviation and the first-order autocorrelation of the percentage deviations of house prices from trend. The distributions of STD and ACORR(1) are based on the posterior distributions of the parameter estimates.

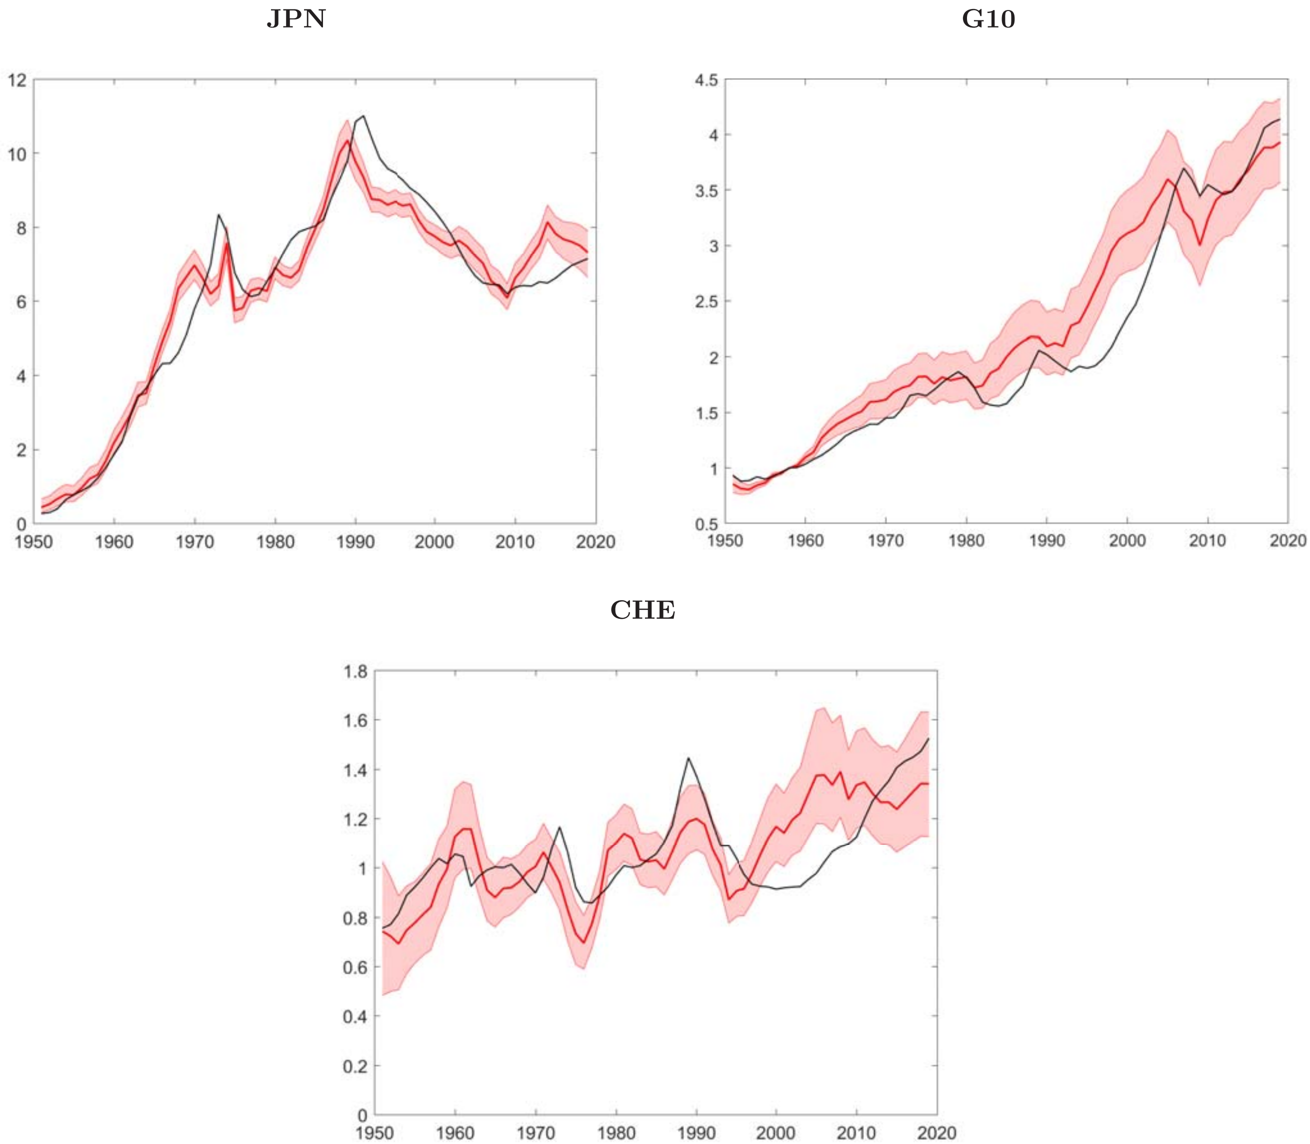

Figure 3: Model vs. data. The thick red line is the median and the shaded areas are the 90% error bands obtained from the posterior distribution of the parameter estimates. The black line is the data. The data are a real house price index, 1957 = 1. For the G10, both the data and the model are based on the 1st principal component of the ten countries.

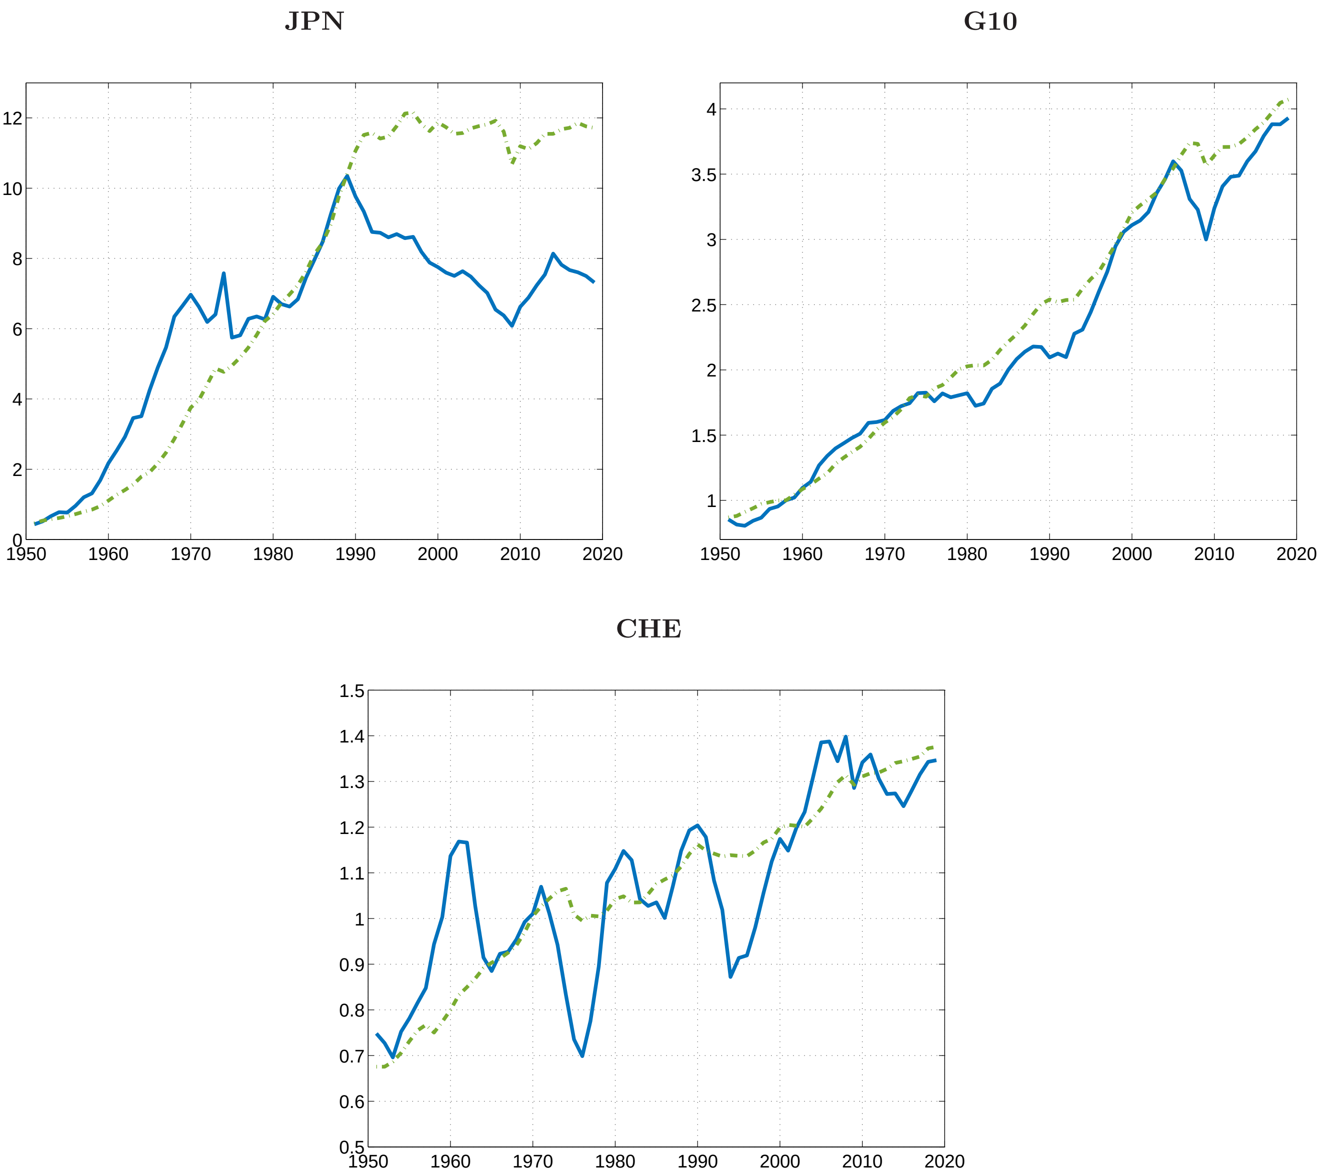

Figure 4: The role of expectations and the interest rate. The median of the model house prices (solid line) and the median of the model stochastic trend (dash-dotted line). The stochastic trend does not depend on expectations and the interest rate. For the G10 countries, the chart plots the 1st principal component of the medians.

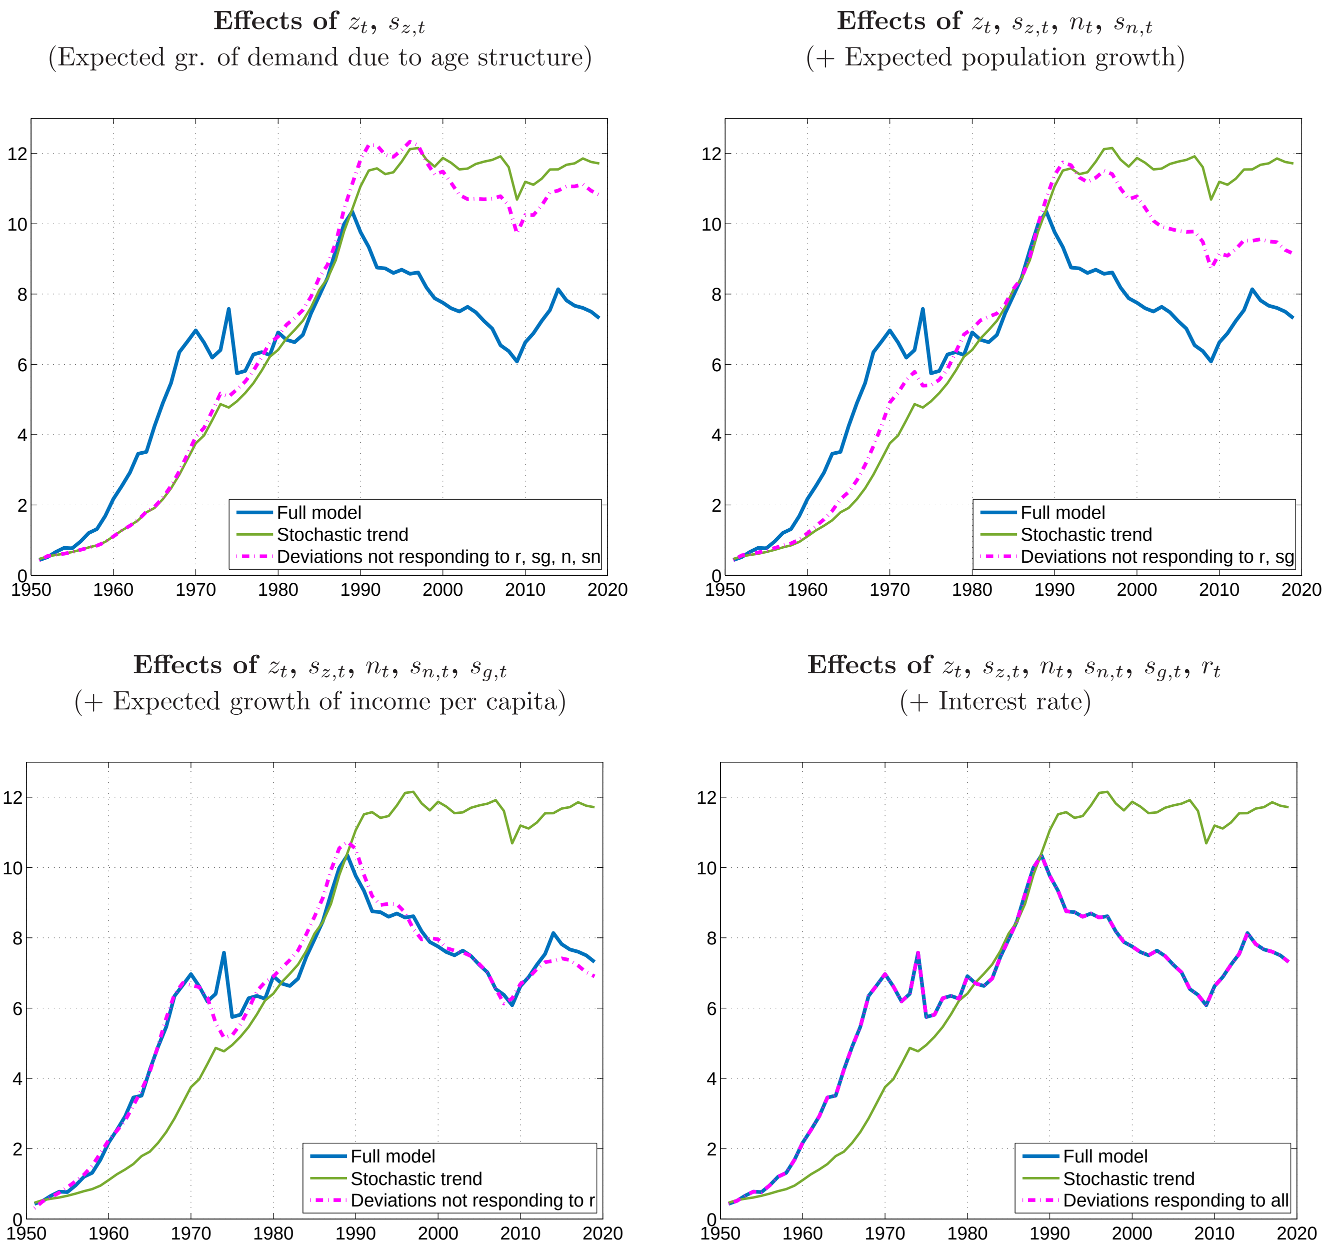

Figure 5: Japan—the marginal contribution of the state variables to the deviations of house prices from the stochastic trend. The solid thick line = the full model; the solid thin line = stochastic trend only; the dash-dotted line = the model with only the state variables in the respective chart title affecting the deviations from trend.

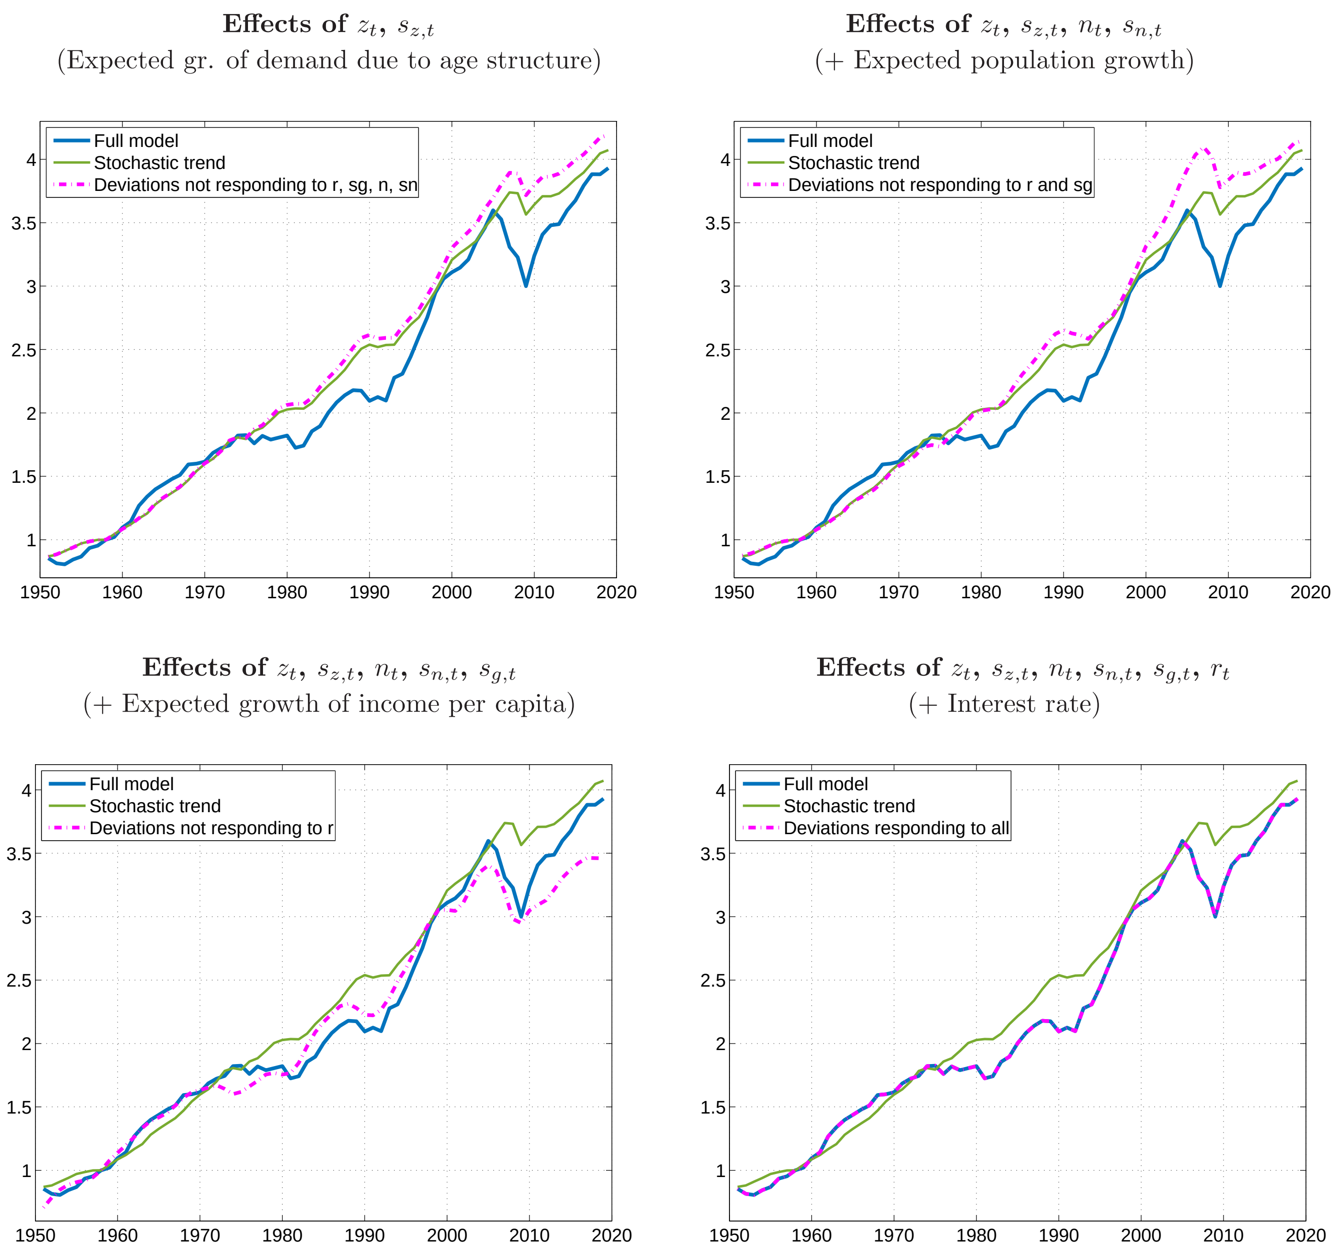

Figure 6: G10 countries—the marginal contribution of the state variables to the deviations of house prices from the stochastic trend. The solid thick line = the full model; the solid thin line = stochastic trend only; the dash-dotted line = the model with only the state variables in the respective chart title affecting the deviations from trend. Based on the 1st principal component of the medians.

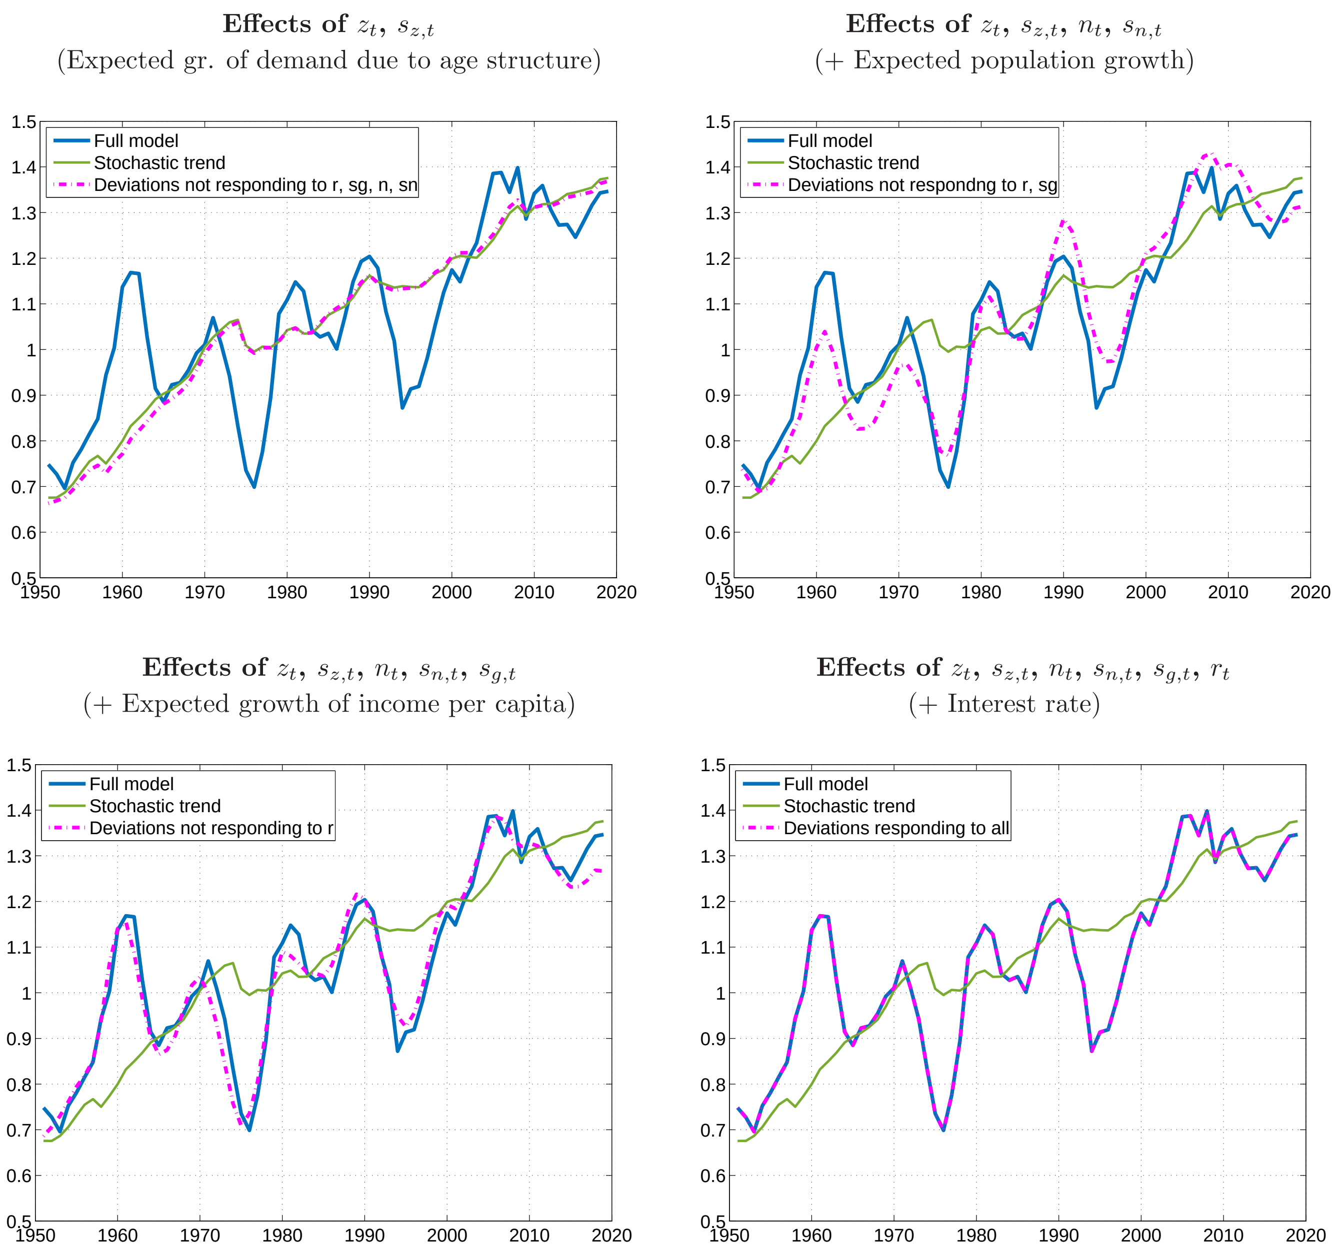

Figure 7: Switzerland—the marginal contribution of the state variables to the deviations of house prices from the stochastic trend. The solid thick line = the full model; the solid thin line = stochastic trend only; the dash-dotted line = the model with only the state variables in the respective chart title affecting the deviations from trend.

Figure 8: Realizations of n_t and $s_{g,t}$. For $s_{g,t}$ the plots are based on the medians of the posterior distributions from the panel estimation. For the G10 countries, the charts plot the cross-country average +/- one standard deviation of the country-specific paths.