What Drives Housing Prices?

James A. Kahn[1]

Wharton School, University of Pennsylvania, and New York University

February 2009

[1]Please direct all correspondence to the author at: jakahn@wharton.upenn.edu or Finance Department, Wharton School, University of Pennsylvania, 2300 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104-6367. The author acknowledges use of computer routines described in Kim and Nelson (1998), as well as helpful comments by Mark Bils, and by seminar participants at Columbia, NYU, CUNY-Baruch, University of North Carolina, and the Federal Reserve Banks of New York, Atlanta, and St. Louis. JEL Codes: E22,E32,O41,O51. Keywords: Housing Prices, Residential Investment, Productivity Growth.

Abstract

This paper develops a growth model with land and housing services that explains much of the amplitude and timing of medium frequency house price fluctuations over the last forty years. House prices are predicted to have a "bubbly" appearance, with housing wealth rising faster than income for an extended period before collapsing and experiencing an extended decline. The analysis suggests that the current downturn in the housing sector was triggered by a productivity slowdown that began in 2004. A more general implication is that policies geared toward artificially boosting house prices through credit market interventions are unlikely to be effective.

From 1996 to its peak in early 2007, the real quality-adjusted price of new houses in the United States appreciated by 33 percent in real terms, about the same as the increase in real GDP over the same period. Prices of existing homes grew even faster, by more than 60 percent. Residential investment as a share of GDP peaked at 6.3 percent in the fourth quarter of 2005, the highest since 1950. The real estate boom, as well as the recent dramatic downturn, has drawn increased attention to the causes and effects of fluctuations in housing prices and investment. Of course this was not the first time housing boom and bust: a similar episode took place beginning in the late 1960s, when prices boomed for a decade before slowing and even declining for the subsequent 15 years. Markets in particular geographical areas have exhibited still greater volatility (cf. Himmelberg et al, 2005). More generally, it has been observed (e.g. Kahn, 2008, Ortalo-Magné and Rady, 2006) that the amplitude of house price variation tends to exceed that of income variation. These phenomona have been variously attributed to speculative bubbles, at least in specific localities (Case and Shiller, 2003), expansions and contractions in monetary policy (Iacoviello and Neri, 2006), or to changing credit market frictions—innovations such as new types of mortgage instruments, or breakdowns such as the recent subprime mortgage crisis. Such explanations suggest that market imperfections, credit market frictions, or irrationality may have played an important role in housing sector fluctuations.

This paper argues that changes in trend productivity growth are the key driver of these medium-term movements in housing prices. It develops a stochastic growth model in which land, capital, and labor are inputs to production, and housing and non-housing consumption provide utility to consumers. Technical progress in output other than housing services is presumed to have a Markov regime-switching component as in Kahn and Rich (2007, hereafter KR). The calibrated model, together with a plausible learning process for trend productivity, can explain the qualitative—and much of the quantitative—behavior of housing prices since the 1960s, including the recent slowdown. In particular, it shows that the regime-switching behavior of productivity growth, which KR finds to be an accurate depiction of postwar data, can give housing prices a “bubbly” appearance in which housing wealth rises faster than income for an extended period, then collapses and experiences an extended decline. The model also provides a partial rationale for the beliefs of investors (and mortgage issuers) in the housing boom in the early part of this decade, as it suggests that the bust that occurred was a low-probability event viewed from the perspective of the early part of the decade.

A key parameter turns out to be the elasticity of substitution between housing and non-housing consumption. This parameter has been featured in many studies related to housing (e.g. Li et al, 2008; Piazzesi and Schneider, 2007; Flavin and Nakagawa, 2004.). At the same time, many studies of housing have assumed, presumably for convenience, a value of one for this elasticity (e.g. Iacoviello and Neri, 2006). We provide evidence based on both aggregate and microeconomic data that this elasticity is considerably less than one. This low elasticity plays a crucial role in the model's ability to explain both qualitatively and quantitatively the magnitude of housing price fluctuations. In particular it helps to account for a “multiplier” effect that changes in housing price growth rates tends to magnify changes in underlying economic growth. Kahn (2008) finds this multiplier effect in international panel data as well.

Beyond the recent boom and bust, there is a more general view that credit market frictions play an important role in house price fluctuations. This paper does not directly refute that view, as it has no frictions in the model; rather, it can be viewed as complementary in providing an additional mechanism to explain the volatility of house prices. But much of that literature is premised on the difficulty of accounting for the magnitude of house price volatility. The results in this paper demonstrate that even a relatively simple growth model without financial frictions can account for much of the medium-term variation in house prices, and therefore sharply limits the role of such frictions—arguably to transitory price volatility, fluctuations in ownership rates, transactions rates, and the like.

1 Background

Figure 1 depicts the behavior of inflation-adjusted housing prices over the last few decades, according to several popular series. The one available going back the farthest (to 1963) is the quality-adjusted price index for new homes, published by the Census Bureau. An index of existing home prices based on repeat sales[1] is available quarterly back to 1975, while the Case-Shiller index, also based on repeat sales, goes back to 1987. We can see that where they overlap, the series behave similarly except for having different trends. Both the Census and OFHEO indexes peak in 1979 or so, then decline until the mid-1980s. All three series have another small peak at around 1990, followed by flat or declining real prices until the second half of the 1990s. Then all three series take off and increase dramatically through 2006 before turning down. The real value of housing wealth, as measured by flow of funds data, has grown an average of 4.6 percent since 1952. This compares with 3.4 percent growth of private net worth excluding real estate, and 3.5 percent growth of personal consumption expenditures over the same time period. Figure 2 plots the ratio of nominal housing wealth to nominal consumption expenditure. This ratio has increase by more than 50 percent since 1952. Figure 3 plots the much more volatile ratio of housing wealth to total net worth. While the enormous volatility of non-real estate wealth (mostly the stock market) hinders precise inferences about relative trends, the upward drift of this ratio is apparent, and not just the result of the runup in real estate wealth over the last decade . The bottom line is that real estate has gone from 27 percent of net worth in 1952 to 39 percent by 2008.[2]

One possible explanation for the relative increase in housing prices and asset values is a simple income effect, or non-homogeneity in preferences. As people get wealthier, they may prefer to have more of their consumption coming from housing services, the price of which will tend to rise because of its being relatively intensive in land, a fixed factor. The (nominal)

[1] This index is published by the Office of Federal Housing Enterprise Oversight (OFHEO).

[2] Davis and Heathcote (2005), however, argue that there are problems with the Flow of Funds data, particularly over long periods of time, and construct their own measures of housing wealth (though only going back to 1975) that exhibit a less clearcut trend.

share of housing services in GDP has gone from 7.5 percent in 1952 to over 10 percent in 2005. The share of housing services in consumer expenditures has gone from 12.2 percent to 14.6 percent over the same period, but in fact has been without any meaningful trend since 1960 or so. In any case, evidence against this proposition will be presented below. Another driver of housing prices could be differing technical progress trends in housing services versus other goods, as in Baumol (1967). The relatively large share of land and structures, two inputs usually thought to be less amenable to embodied technical progress, in the value of housing makes this story plausible. This paper will argue that the timing of low-frequency changes in both housing prices and productivity suggests that this mechanism is important.

There is some evidence that in fact the increase in housing wealth does not stem from an increase in the value of houses per se, but rather from the increase in the value of the land upon which they are built. First, a price index that include the value of land, the Conventional Mortgage Home Price Index, has increased approximately 0.75% faster than indexes that do not, such as the Census's Composite Construction Cost index, on an annual basis. Davis and Heathcote (2004) compute a land price index based on this type of differential and find that land values have increased at an average annual rate of approximately 3.5% (inflation-adjusted) over the period 1975-2005. That price index may, however have an upward bias from not adequately adjusting for quality changes. Land price series available from the Bureau of Labor Statistics (see Figure 4) suggest behavior that is closer to the behavior of new home prices in Figure 1. We focus on the Census Bureau's quality-adjusted price of new homes in part because of concern over this bias, but also because it goes back to 1963.

Finally, Figure 5 depicts the behavior of the HP-trend component of productivity growth (relative to a linear trend) over the postwar period. While its pattern is similar to low-frequency movements in land and housing prices, the downturn in productivity clearly precedes the downturn in housing and land prices by several years. KR provide more detailed econometric estimates of a regime-switching model of the sort incorporated below into this paper, and find significant regime changes corresponding to the shifts depicted in the figure.

2 Related Literature

Research on aggregate housing prices has emphasized demographics, income trends, and government policy as fundamental drivers. In one well-known study, Mankiw and Weil (1989) argued that population demographics were the prime determinant, and predicted that prices would fall in the subsequent two decades with the maturation of the baby boom generation and resulting decline in the growth rate of the prime home-owning age group. While their prediction proved inaccurate, Martin (2005) renews the argument for an important role for demographics. Glaeser et al (2005) argues that price increases since 1970 largely reflect artificial supply restrictions. Gyourko et al (2006) also cite inelastically supplied land as a key driver of the phenomenon they call "superstar cities." Van Nieuwerbergh and Weill (2006), however, argue that so long as there are regional markets in which such restrictions are not present, the aggregate impact of restrictions in some local markets is likely to be modest—in other words, they primarily affect the cross-sectional distribution of housing prices as opposed to the aggregate. Iacoviello and Neri (2006) examine the role of monetary policy with credit market frictions.

While Ortalo-Magné and Rady (2006) argue that credit market frictions play an important role in house price volatility, their model has a fixed supply of housing and no rental market alternative. Kiyotaki et al (2007) find that credit market frictions primarily affect own vs. rent decisions as opposed to prices. Piskorski and Tchistyi (2008) examine optimal mortgage lending in a setting where housing prices obey essentially the same type of regime-switching behavior assumed here, and find that "many features of subprime lending observed in practice are consistent with economic efficiency and rationality of both borrowers and lenders," though, as they point out, there may be negative externalities associated with massive defaults in a downturn.

Case and Shiller (2003) and Himmelberg et al (2005) investigate the bubble hypothesis, looking across a large number of cities, and both suggest that the phenomenon is limited to a few localities. As with the research above on inelastic land supplies, these papers emphasize the cross-sectional variation of house prices across metropolitan areas rather than aggregate time series variation. Consistent with the approach in this paper, Attanasio et al (2005) find that "common causality" is an important driver of the comovement of house prices and consumption, as opposed to wealth or the collateral channels.

One important innovation in this paper is to allow for unbalanced sectoral growth. General equilibrium models with production have generally either assumed Cobb-Douglas preferences (e.g. Davis and Heathcote, 2005, Kiyotaki et al., 2007, Iacoviello and Neri, 2006) or have abstracted from longer-term growth issues (e.g. Van Nieuwerburgh and Weill, 2007). This is the first housing model (to my knowledge) with production that features (approximately) balanced aggregate growth and systematically varying sectoral shares due to non-unit elastic preferences. The importance of this is that it is consistent with aggregate growth facts as well as with the evidence on substitution elasticities found by numerous authors (see the discussion below), and also enables the model to match the volatility of housing prices in a plausible and disciplined way. The modeling approach is based on recent work of Ngai and Pissarides (2007).

3 A Simple Static Model

There are two goods, c and h, each produced by combining capital and land:

$$c = AK_c^{\alpha}L_c^{1-\alpha} \tag{1}$$

$$h = K_h^{\beta}L_h^{1-\beta} \tag{2}$$

where A is relative TFP in the c sector, we assume $\alpha > \beta$ (h is more land-intensive) and

$$K_c + K_h = K \quad (3)$$

$$L_c + L_h = L. \quad (4)$$

A representative consumer has a utility function

$$u(c,h) = \begin{cases} \left[\omega c^{(\epsilon-1)/\epsilon} + (1-\omega)\, h^{(\epsilon-1)/\epsilon}\right]^{\epsilon/(\epsilon-1)} & \epsilon \neq 1, \epsilon > 0 \\ \omega \ln c + (1-\omega) \ln h & \epsilon = 1 \end{cases}. \quad (5)$$

We can examine the impact of A on equilibrium prices and quantities.

It is straightforward to show that the relative price of h in units of c, denoted by p satisfies

$$p = A \left(\frac{1-\alpha}{1-\beta}\right)^{1-\alpha} \left(\frac{\alpha}{\beta}\right)^{\alpha} \left(\frac{K_h}{L_h}\right)^{\alpha-\beta}. \quad (6)$$

More specifically, think of the measure of "houses" as corresponding to the measure of K_h, each with L_h/K_h units of land. Now consider the impact of a once and for all increase in A. If $\epsilon = 1$, then p moves one for one with A, while the factor proportions remain fixed in each sector, as the income and substitution effects on h exactly offset. By contrast, if $\epsilon < 1$, factors shift toward the h sector, both sectors become less land-intensive, and p goes up disproportionately to the increase in A. Aggregate consumption (in units of c) $c + ph$ increases by in proportion to A.

Table 1 gives some results for p and K_h/L_h for a doubling of A from 1 to 2, for various values of ϵ. The other parameters are as follows: $\omega = 0.8$, $\alpha = 0.7$, $\beta = 0.2$. For $\epsilon < 1$, houses become less land intensive (implicitly, land becomes more expensive relative to capital), and p increases by more than one-for-one with A.

4 The Dynamic Model

This section presents a general equilibrium growth model that is capable of capturing the important stylized facts about housing and the economy. The model has two sectors, a "manufacturing" sector that produces non-housing related goods and services, as well as the capital (structures and durable goods) that go into housing services. A second sector uses capital, labor, and land to produce a flow of housing services. The model exhibits approximately balanced aggregate growth, but with unequal growth across sectors. We then consider the behavior of the model under a regime-switching specification for productivity growth in the manufacturing sector.

4.1 Firms and Consumers

Competitive final goods firms produce two types of goods: A "manufactured" good Y_m, and housing services Y_h. Under perfect competition the final goods firms make zero profits and have perfectly elastic supplies of Y_m and Y_h at the above prices. The production functions for the two types of goods are

$$Y_j = A_j K_j^{\alpha} L_j^{\beta_j} N_j^{1-\alpha-\beta_j}$$

for $j = m, h$, where K_j is capital allocated to j, L_j is land, and N_j is labor input. The goods producers rent inputs in competitive markets. In particular, capital is rented from final goods producers of Y_m.In the j sector, the representative firm's nominal profit in period t is given by

$$P_{jt}Y_{jt} - W_t N_{mt} - R_{\ell t} L_{mt} - R_{kt} K_{mt} \tag{7}$$

where R_ℓ and R_k are nominal rental rates for land and capital respectively, and W_t is the nominal wage.

There are N_t representative agents at time t supplying N_t labor, where N is exogenous, growing exponentially at constant rate ν.[3] Let C denote the aggregate non-housing consumption good, and H aggregate housing services. We let $c \equiv C/N$ and $h \equiv H/N$ denote per capita quantities. The representative consumer then cares about c and h, and dislikes working. He solves the problem

$$\max\ U = E_0 \left\{ \sum_{t=0}^{\infty} (1+\rho)^{-t} \ln \left(\left[\omega_c c_t^{(\epsilon-1)/\epsilon} + \omega_h h_t^{(\epsilon-1)/\epsilon} \right]^{\epsilon/(\epsilon-1)} \right) \right\} \tag{8}$$

subject to

$$P_{m,t+s}\left(c_{t+s} + \iota_{t+s}\right) + P_{h,t+s} h_{t+s} + V_{t+s}\left[(1+\nu)\,\ell_{t+s} - \ell_{t+s-1}\right] + b_{t+s}/\left(1+R_{t+s}\right)$$

$$\begin{aligned} &\leq \; b_{t+s-1} + W_{t+s} + (1+\nu)\,R_{k,t+s} P_{m,t+s-1} k_{t+s-1} \\ &\qquad + R_{\ell,t+s} V_{t+s-1} \ell_{t+s-1} + d_{m,t+s} + d_{h,t+s} \end{aligned} \tag{9}$$

$$(1+\nu)\,k_{t+s} = (1-\delta)\,k_{t+s-1} + z\left(\iota_{t+s-1}/k_{t+s-1}\right) k_{t+s-1} \qquad (j = m, h) \tag{10}$$

where ι_t denotes total capital investment at date t, d_{jt} nominal dividends (for simplicity assumed to be distributed in a lump-sum fashion) from the profits of intermediate goods producers in sector j, b_t nominal one-period discount bonds, V_t the price of land at date t, and k_t and ℓ_t per capita capital and land holdings at date t. The constraints reflect the fact that population is growing, so that per capita stocks get deflated at rate ν. Both k_t and ℓ_t denote the sum of capital and land in both sectors. The function $z(x)$ reflects adjustment costs, which will be discussed in more detail below.

4.2 Equilibrium Growth

We assume (necessary for balanced growth) that capital's share α and depreciation δ are the same in both sectors, but labor's share is higher in manufacturing (implying that land's

[3] It is straightforward to endogenize work effort, but this has little impact on any results.

share is higher in the housing sector, i.e. $\beta_h > \beta_m$). Let c and h denote per capita quantities of C and H, while k, ℓ, k_i, ℓ_h refer to per worker quantities in sector i (e.g. $k_{ht} \equiv K_{ht}/N_{ht}$, $k_t \equiv K_t/N_{t+1}$, i.e. no subscript refers to aggregates), while $n_{it} \equiv N_{it}/N_t$, $(i = m, h)$.[4] Given the assumption of perfect competition, the equilibrium is the result of maximizing U from (8) subject to aggregate resource constraints.

$$
\begin{aligned}
c_t + i_t &= A_{mt} k_{mt}^{\alpha} \ell_{mt}^{\beta_m} n_{mt} && (11) \\
(1+\nu) k_t - (1-\delta) k_{t-1} &= z\left(i_t / k_{t-1}\right) k_{t-1} && (12) \\
h_t &= A_{ht} k_{ht}^{\alpha} \ell_{ht}^{\beta_h} n_{ht} && (13) \\
k_{mt} n_{mt} + k_{ht} n_{ht} &= k_{t-1} && (14) \\
\ell_{mt} n_{mt} + \ell_{ht} n_{ht} &= \ell_t && (15) \\
n_{mt} + n_{ht} &= 1. && (16)
\end{aligned}
$$

Total land $\bar{L}$ is assumed fixed, so $\ell_t/\ell_{t-1} = (1+\nu)^{-1}$. Average technological progress in sector i, i.e. the average growth rate of A_i, is denoted γ_i $(i = m, h)$. Note that the timing assumptions in (14) and (15) are such that while aggregate capital k is chosen one period ahead of time, and total land and labor are exogenous, for simplicity the sectoral allocations are determined contemporaneously.

Note that technical progress in the h sector is unrelated to technological progress in construction. (In fact, home construction occurs in the m sector in this model.) Rather, it refers to an increase in the housing services from given stocks of K_h, L_h, and labor inputs N_h. What this means in practice depends on exactly what the term "housing services" encompasses, and on how one measures K_h. In the model it is assumed for simplicity to be indistinguishable from K_m other than by its allocation to the h sector. In particular, it is assumed to have the same price as K_m and C. In principle it would include both residential

[4] The derivations here draw on Ngai and Pissarides (2007), albeit in discrete time, and adding a fixed factor with heterogeneous technology.

structures and housing service-related consumer durables (home appliances). N_h would include both non-market and market labor involved in the production of housing services such housekeeping and maintenance.

The model abstracts from a number of potentially important factors. First and foremost, the housing and construction sectors are heavily affected by government intervention, both via distortionary taxation and regulations. In particular, much land in the U.S. (and in most other countries as well) is neither residential nor commercial, and is either owned or heavily restricted in its use by the government. Second, there is tremendous heterogeneity in land and housing values. Land near navigable bodies of water, or ports, or along coastlines is much more valuable than land that does not have these features. Obviously this model will have nothing directly to say about the cross-sectional distribution of land values or housing prices (though many of the factors that affect them over time undoubtedly come into play in the cross-section as well). Nonetheless if all of these factors remain relatively constant over time, then ignoring them in a model such as this should not be too great a sin.

The static first-order conditions for the problem (8) can be shown to imply that p_t, the relative price of housing services in terms of manufactured goods, satisfies

$$p_t = \frac{A_{mt}}{A_{ht}} \left(\frac{\beta_m}{\beta_h} \right)^{\beta_h} \left(\frac{1-\alpha-\beta_h}{1-\alpha-\beta_m} \right)^{\alpha+\beta_h-1} \ell_{mt}^{-(\beta_h-\beta_m)} \tag{17}$$

Thus as in the static example, growth in the price of housing services reflects both relative productivity growth in manufacturing and the increasing scarcity of land.

4.3 Aggregate Growth under Certainty

Let total expenditure $c + ph$ be denoted by x. It also turns out that $\mu_m = x^{-1}$ (see the proof in the Appendix), hence $\mu_{mt}/\mu_{mt-1} = x_{t-1}/x_t$. We will define aggregate balanced growth under certainty as an equilibrium path in which x and k both grow at a constant rate, and in which the interest rate (i.e. the marginal product of capital) is also constant.

We will also assume that $z(i/k) = i/k$ and $z'(i/k) = 1$ at the steady state value of i/k, so that adjustment costs are zero on the balanced growth path. Balanced growth clearly requires that $A_{mt}k_{mt}^{\alpha-1}\ell_{mt}^{\beta_m}$ be constant, which amounts to a linear restriction on the growth rates in the m sector of the capital-labor ratio, technological progress, and the land-labor ratio. Therefore, let

$$Z_t \equiv \left[A_{mt}\ell_{mt}^{\beta_m}\right]^{1/(1-\alpha)} \tag{18}$$

and define variables with "~" over them to be deflated by Z_t, e.g. $\tilde{k}_{mt} \equiv k_{mt}/Z_t$. We then have

$$x_t/k_{t-1} + i_t/k_{t-1} = \tilde{k}_{mt}^{\alpha-1} \tag{19}$$

$$(1+\nu)\,k_t/k_{t-1} = z\left(i_t/k_{t-1}\right) + 1 - \delta \tag{20}$$

$$(x_{t+1}/x_t)(1+\nu)(1+\rho)\,q_t = \alpha\tilde{k}_{mt+1}^{\alpha-1} + q_{t+1}\left[z\left(i_{t+1}/k_t\right) + 1 - \delta\right] - \left(i_{t+1}/k_t\right) \tag{21}$$

With $\tilde{k}_m$ constant under balanced growth, k and x both grow at the same constant rate.

From the resource constraints and first-order conditions for maximization, we can show that

$$k_{mt}\left[\frac{1-\alpha-\beta_m}{1-\alpha-\beta_h}n_{ht} + n_{mt}\right] = k_{t-1}. \tag{22}$$

Now let

$$Q_t \equiv 1 + \tau n_{ht} \tag{23}$$

where

$$\tau \equiv \frac{\beta_h - \beta_m}{1-\alpha-\beta_h}.$$

Then we have $k_{mt} = k_{t-1}/Q_t$, and we can define $\hat{k}_t \equiv k_t/\left(Z_tQ_t\right)$. This gives a normalization of k_t that is constant on the balanced growth path. Note that if $\beta_m = \beta_h$, then $Q = 1$ and we would have $k_{mt} = k_{ht} = k_{t-1}$. But with $\beta_m > \beta_h$, $Q > 1$ and n_h and n_m are changing over time (unless $\epsilon = 1$). In particular, if $\epsilon < 1$ and $\gamma_m \geq \gamma_h$, then n_h (and hence Q) grows over time.

Thus, strictly speaking, balanced growth requires one of these knife-edge conditions: $\epsilon = 1$, $\beta_m = \beta_h$, or

$$(1 + \gamma_m)(1 + \nu)^{\beta_h - \beta_m} = 1 + \gamma_h \quad (24)$$

None of these is very palatable: Below I provide evidence that ϵ is substantially less than one, and cite other studies with similar findings. I also show that ϵ has important implications for housing price dynamics, so assuming $\epsilon = 1$ for convenience is not innocuous. Similar comments apply to the assumption $\beta_m = \beta_h$. To assume (24) is less problematic, as it is tantamount to assuming that p_t does not grow over time. It does imply that $\gamma_h > \gamma_m$, which is hard to believe but also hard to refute directly since γ_h is difficult to measure. If (24) fails to hold (say if $\gamma_m \geq \gamma_h$, so that p drifts higher over time), the dynamic response of the model will be a function of the level of p_t—in particular the expected growth rate varies over time and is only asymptotically constant.

It turns out, however, that the consequences of assuming this when it is false are in fact innocuous: Both the variation in the growth rate over time and the differences in dynamics are tiny. Thus when (24) does not hold, the model exhibits near-balanced aggregate growth and unbalanced sectoral dynamics, as in Ngai and Pissarides (2007).[5] Alternatively, we will see later that over a wide range of parameters, growth is so close to balanced even when p is growing over time that it is reasonable to treat it as balanced for computational purposes. So the bottom line is that at least for reasonable parameters the model implies that growth is (approximately) balanced even when sectoral growth is unbalanced because p grows over time.

[5] In the stochastic growth case considered below, however, TFP will have a unit root, so it will be necessary to take account of the imbalances away from the balanced growth path. For example, an impulse response to the growth rate of A_{mt} results in a permanent level effect, and the new balanced growth path (assuming (24) holds again) has slightly different dynamics than the original path. Again, the differences are turn out to be tiny, but this has to be verified.

On the balanced aggregate growth path, ZQ grows at a constant rate. In fact it is straightforward to show that its growth rate g satisfies

$$\left[(1+\gamma_m)(1+\nu)^{-\beta_m}\right]^{1/(1-\alpha)} \equiv 1+g^* \equiv G^* \tag{25}$$

We then have $\tilde{k}_{mt} = k_{mt}/Z_t = k_{t-1}/(Q_t Z_t) = \hat{k}_{t-1}Q_{t-1}Z_{t-1}/(Q_t Z_t) \equiv \hat{k}_{t-1}/G_t$. Aggregate output per capita (in terms of manufactured goods), which we denote y_t, is $A_{mt}k_t^{\alpha}\ell_{mt}^{\beta_m}n_{mt} + p_t A_{ht}k_{ht}^{\alpha}\ell_{ht}^{\beta_h}n_{ht}$, or (after substituting for p_t and simplifying as before):

$$y_t = A_{mt}k_{mt}^{\alpha}\ell_{mt}^{\beta_m}Q_t = \tilde{k}_{mt}^{\alpha}Z_t Q_t, \tag{26}$$

so we can also define $\hat{y}_t = y_t/(Z_t Q_t) = \tilde{k}_{mt}^{\alpha} = \left[\hat{k}_{t-1}/G_t\right]^{\alpha}$ and $\hat{x}_t = x_t/(Z_t Q_t)$.

We can now characterize the dynamics in terms of stationary variables:

$$\hat{x}_t + \hat{\imath}_t = \left[\hat{k}_{t-1}/G_t\right]^{\alpha} \tag{27}$$

$$G_t(1+\nu)\hat{k}_t = z\left(G_t\hat{\imath}_t/\hat{k}_{t-1}\right)\hat{k}_{t-1} + (1-\delta)\hat{k}_{t-1} \tag{28}$$

$$\begin{aligned}(\hat{x}_{t+1}/\hat{x}_t)G_t(1+\nu)(1+\rho)q_t &= \alpha\left[\hat{k}_t/G_t\right]^{\alpha-1} - G_t\hat{\imath}_{t+1}/\hat{k}_t \\ &\quad + q_{t+1}\left[z\left(G_t\hat{\imath}_{t+1}/\hat{k}_t\right) + 1 - \delta\right]\end{aligned} \tag{29}$$

$$q_t = z'\left(G_t\hat{\imath}_t/\hat{k}_{t-1}\right)^{-1} \tag{30}$$

Since along the balanced growth path G_t is constant, $z(x) = x$, and $q = 1$, the aggregate economy behaves exactly as the standard neoclassical growth model. The innovation in this paper is to simultaneously characterize the behavior of sectoral variables, and in particular housing prices and investment, within the aggregate steady state.

The sectoral variables can be solved for directly as functions of the aggregates. Only in the knife-edge cases of $\epsilon = 1$ or (24) will these variables exhibit balanced growth in the sense of either being constant or growing at the same rate as the aggregate economy. Details are provided in the Appendix.

Although land is not explicitly priced in the model, we can compute its shadow rental price ζ_t in terms of manufactured goods:

$$\zeta_t = \beta_m A_{mt} k_{mt}^{\alpha} \ell_{mt}^{\beta_m - 1} \tag{31}$$

To a first approximation we can say that the land rental price grows at rate $g + \nu$ on the balanced growth path—exactly $g + \nu$ if $\epsilon = 1$, a bit faster if $\epsilon < 1$ and p is growing.

4.4 Stochastic Growth

We suppose that the growth rate of A_h is fixed at γ_h, but that of A_m follows a Markov regime-switching process:

$$A_{mt}/A_{mt-1} = (1 + \tilde{\gamma}_{mt})\, \eta_t/\eta_{t-1} \tag{32}$$

where

$$\tilde{\gamma}_{mt} = \begin{cases} \gamma_m^1 & \text{if } \xi_t = 1 \\ \gamma_m^0 & \quad \xi_t = 0 \end{cases} \tag{33}$$

η_t is a transitory disturbance, and ξ_t is a state variable with Markov transition matrix Θ, where $\Theta[i,j] = \Pr\left(\xi_t = j | \xi_{t-1} = i\right)$. Since the columns of Θ must sum to one, we write it as

$$\Theta = \begin{bmatrix} \theta_1 & 1 - \theta_0 \\ 1 - \theta_1 & \theta_0 \end{bmatrix}. \tag{34}$$

If the diagonal elements of Θ are close to one, the growth states will be highly persistent, and a shift from one state to the other will carry with it a sizeable adjustment in the long-term level of A_m.

The log deviation version of G_t can be written as

$$G_t = \frac{1}{1-\alpha}\left[\frac{\tilde{\gamma}_{mt} - \bar{\gamma}_m}{1 + \bar{\gamma}_m} + \Delta\eta_t + \frac{\tau n_h}{1 + \tau n_h}\Delta n_{ht} - \beta_m \Delta\ell_{mt}\right]$$

We suppose that $\eta_t = \phi_1 \eta_{t-1} + \phi_2 \eta_{t-1} + \upsilon_t$, where υ_t is i.i.d. with a zero mean. In what follows, we will first assume that economic agents observe both z_t and η_t before making their period t decisions. Later we will consider the possibility that they only observe G_t and must estimate z_t and η_t given the history of G_t.

4.5 Asset Prices

Thus far we have only described the behavior of the price of housing services and rental prices for land. The term "housing prices" generally refers to asset values of homes, both the structures and the land. In this model we can calculate the value of what might be called "real estate wealth," which would be the total value of capital and land allocated to the housing services sector. The value of the capital is just $K_h = k_h n_h$. The asset value of the land $L_h = \ell_h n_h$ requires some computation, as described below. Given a land price, which we will denote by V_t (expressed in terms of m sector output), valuing a representative house requires constructing an index, because the composition of the representative house changes over time due to changes in the price of land. Given a path $\{K_{ht}, L_{ht}\}$ we will define a "constant-quality" house price index P_{ht} as a Laspeyres index by choosing a base year, say $t = 0$, and setting $P_{ht} = 100 \left(V_t L_{h0} + K_{h0}\right) / \left(V_0 L_{h0} + K_{h0}\right)$.

We know that V_t is the present discounted value of the stream of rents $\{\upsilon_t\}$:

$$V_t = \zeta_t + E_t \left\{\Phi_{t,1} V_{t+1}\right\} = E_t \left\{\sum_{\tau=0}^{\infty} \Phi_{t,\tau} \zeta_{t+\tau}\right\} \tag{35}$$

where

$$\Phi_{t,\tau} = \frac{\mu_{m,t+\tau}}{\mu_{mt} \left(1+\nu\right)^{\tau} \left(1+\rho\right)^{\tau}} = \frac{x_t}{x_{t+\tau} \left(1+\nu\right)^{\tau} \left(1+\rho\right)^{\tau}} \tag{36}$$

is the stochastic discount factor. On the balanced growth path we have $\Phi^{-1} = (1+g)(1+\rho)(1+\nu)$, and $\hat{\upsilon}_t$, as mentioned previously, is (for plausible parameters) almost constant but technically a function of A_{mt}/A_{ht} and N_t (for $\epsilon < 1$ it is increasing in both arguments). Hence while the capital stock and the aggregate output grow at $g + \nu$, the price of land, and hence the price of "houses" (capital plus land in the h sector) grows at a rate (slightly) faster than $g + \nu$.

4.6 Calibration

Most of the parameters take on standard values for quarterly data: $\alpha = 0.33$, $\nu = 0.0025$, $\delta = 0.02$.[6] The parameters β_h and β_m should reflect the shares of land in the cost of housing services and non-housing output respectively. The work of Davis and Heathcote (2004) suggests a value for β_h of 0.5. Data from the Bureau of Labor Statistics capital tables suggest a value for β_m of 0.05.[7] Since housing services represent about 20 percent of overall consumer expenditures, we set $\omega_h = 0.2$, $\omega_c = 0.8$, though since expenditure shares vary it is necessary to choose the starting level of A_m/A_h appropriately. We set the time preference rate ρ equal to 0.01. Finally, we choose the parameters of the regime-switching process for productivity to correspond roughly to the results in KR: $(\gamma_m^1 - \beta_m \nu) / (1 - \alpha) = 0.029$, $\gamma_m^0 = 0.013$, $\theta_1 = 0.99$, $\theta_0 = 0.983$. Thus high growth regimes are slightly more persistent than low-growth, and implied the overall mean growth rate of A_m, $\bar{\gamma}_m$, is 2.31 percent.

4.7 The Elasticity of Substitution between Housing and other Consumption

The first-order conditions of the model imply a relationship between the expenditure ratio for housing and non-housing consumption and the relative price.

$$\frac{\omega_c^\epsilon}{\omega_h^\epsilon} \frac{p_t h_t}{c_t} = p_t^{1-\epsilon} \tag{37}$$

[6] The choice for δ represents a compromise between conventional values for structures and equipment for the sake of tractability. Allowing δ to differ for K_h and K_m complicates the math without altering the predictions of the model in any significant way.

[7] Of course there is great variability across subsectors: Manufacturing has a land share of about 0.02, as do most service sectors, while wholesale and retail trade is about 0.07, and agriculture about 0.3. But weighting the various sectors by their share in GDP yields a figure of 0.05.

The long-term behavior of aggregate expenditures on housing services suggests a unit income elasticity for such expenditures, but price inelastic (i.e. $\epsilon < 1$). This is because the ratio expenditures on housing services to non-housing consumption expenditures has no long-run trend, but is positively correlated with the relative price of housing services, at least as measured by NIPA. Figure 6 presents annual data going back to 1929 of the two series, which show a positive relationship for most of the sample, though recently (since roughly 1990) they have diverged. The magnitude of the elasticity, however, is difficult to infer from time series data, given that both the ratio and the price are endogenous variables. In addition, whereas the nominal expenditures on h and c may be measured accurately, there may be substantial error in measuring the true relative price. Whereas the Boskin Commission had estimated an upward bias in CPI rents, Gordon and van Goethem (2005) argue for a downward bias averaging roughly 0.5 percent annually, but varying over time.

As a consequence, we instead examine evidence from micro data. Specifically, we examine data from the Consumer Expenditure Survey (CEX) to gauge the extent of housing service expenditure share variability as a function the relative price of housing services. To do so we construct rent (or owner's equivalent rent) relative to other expenditures, and match this up with data on housing prices by region, total expenditures, and demographic controls. As the above condition suggests, we can obtain an estimate of ϵ by a suitable regression of the nominal expenditure ratio on relative price.

Looking at micro data solves several problems. First, it is reasonable to take the price, which is based on regional CPI measures of owner's equivalent rent relative to the CPI excluding shelter, as exogenous to individual households. Second, the price measurement issue alluded to above is arguably mitigated by reliance on relative prices across regions. Third, we can include specific demographic controls to account for variation in, for example, ω_c/ω_h.

Consider the following model for individual i at date t in region j:

$$\ln\left[p_{jt}h_{it}/\left(x_{it}-p_{jt}h_{it}\right)\right]=a_j+b\ln x_{it}+(1-\epsilon)\ln p_{jt}+z_{it}'\theta+u_{itj}.$$

Here a_j reflects some constant region-specific factor that might affect expenditure shares. For example, if living in a region provides some amenities such as inexpensive public transportation or moderate weather that substitute for other expenditures (automobiles, heating oil), a_j might be positive. Note that we have to assume that a_j is constant over time or else we would not be able to identify the price effect. The coefficient b would reflect a wealth effect to the extent it differs from zero, and the coefficient on p_{jt} has the interpretation indicated: if $\epsilon < 1$, relative expenditures on housing services increase with their cost. The model also includes a set of demographic variables z_{it}, which would include things like family size and number of wage-earners. Note that we would expect a negative coefficient on an indicator of two adults working, as this would typically result in less household production and more expenditures on non-housing goods and services. Finally, the error term u_{ijt} represents idiosyncratic variation in preferences for housing services (in the model represented by ω_h/ω_c), measurement error of the dependent variable, and other omitted variables.

If the CEX had a true panel structure we could difference out the a_j. We could also allow for individual fixed effects. Unfortunately each individual household observation is present for at most four quarterly observations, so the panel aspect is probably useless (the other explanatory variables are not likely to change in a meaningful way over the course of a household's participation). Consequently we choose to pool the sample in levels and use regional dummies to capture the region effects. We also just use one (the first) observation per household, to avoid giving more weight to households simply because they remain in the sample.

Another consideration is adjustment costs. Moving is costly, so many households may be passively accepting changes in their expenditure shares for a while until they decide to move.

Consider an (s, S) framework in which households have fixed costs of moving. They remain where they are and simply absorb any price variation so long as the price p is between the lower and upper bounds s and S; otherwise they move and obtain a new level of expenditure $ph \in (s, S)$. If this is the case, we want to gauge the potential bias in estimating ϵ.

While one might think that the tendency for households to move infrequently would bias the estimate of ϵ toward zero, the results of Caplin and Spulber (1987) suggests that if p is stationary and not very persistent, and households experience idiosyncratic shocks in addition to changes in p, then households will be uniformly distributed in the interval (s, S). And if, on the other hand, p has a unit root or is trending over time and $\epsilon \neq 1$, (s, S) will drift along with p, but at any point in time households will still be distributed uniformly throughout the interval. Consequently observed expenditure ratio likely equal the "desired" or long-run ratios plus some noise. The covariance of this error with p should be zero; that is, households are just as likely to to be above as below their desired ratio. Thus there is no reason to think that the estimate of ϵ should be biased in one direction or another.

The total expenditure variable x_{it} is likely to be measured with error. Such error would tend to produce a potentially large downward bias in the estimate of b because x_{it} enters the denominator of the dependent variable. Fortunately we have candidates for instruments: demographic variables such as race and education that are plausibly unrelated to the u_{itj}.

For nominal expenditures on housing services we use rent for renters and owner's equivalent rent (OER) on the primary residence for owners.[8] Consequently we subtract mortgage and home equity interest from total expenditures, as well as property taxes and expenditures on various categories of home repairs and maintenance that would normally be included in rent. We assume that OER is not intended to include utilities, so for owners utility expenditures are included in total but not housing expenditures. For renters, reported utility

[8] We do not include rent or OER on vacation or second homes. Though it might be desirable to do so, especially for second homes, it is difficult to distinguish, for example, someone who owns a vacation home but rents it out for 48 weeks of the year and uses it for four weeks, on the one hand, from someone who uses it every weekend. The former should be treated as equivalent to someone who rents a vacation home for four weeks, which we regard as something qualitatively different from housing services.

expenditures are also included only in total expenditures, but a dummy variable is included for renters that report zero utilities expenditures (on the assumption that they are included in rent).

The results of this estimation exercise are shown in Table 2. Because of an unexplained break in the level of the dependent variable that occurs in 1993 (presumably because of some change in variable definitions), all of our estimates included a dummy variable for post-1993 data. We present results with and without constraining $b = 0$, and with and without instrumenting for x. Without instrumenting for x, the estimate of b is an implausibly large negative number, and the estimate of ϵ is only 0.134. Instrumenting brings $\hat{b}$ down to a more reasonable -0.254, and increases $\hat{\epsilon}$ to around 0.2. Imposing $b = 0$ results in $\hat{\epsilon} = 0.284$.

The fact that the estimates of ϵ are relatively insensitive to the treatment of other variables in the equation provides some reassurance that it is well below one. This is consistent with a long history of studies that find housing demand to have a low price elasticity (e.g. Hanushek and Quigley, 1980; Polinsky and Ellwood, 1979), or other estimates of ϵ based on micro data. For example, Flavin and Nakagawa (2004) use PSID data and estimate ϵ to be 0.13. Using a different methodology based on asset prices, Lustig and Van Nieuwerburgh (2007) set $\epsilon = 0.05$ in their calibration. The significant negative estimate of b is problematic for the assumption of homotheticity, but is hard to square with the relatively flat expenditure share in Figure 6. If true it would tend to mute the effects described in the simulation, as growth in the demand for housing services would not be as responsive to expected productivity growth. For this reason we will be conservative in our choice of ϵ and set it equal to 0.3 for our benchmark case.

A final comment on adjustment costs: Households only infrequently adjust their consumption of housing services. Since these costs are primarily of the non-convex variety (search, financial and legal transactions, and moving costs) it is likely that when they do adjust, they do so by a discrete amount. While a fully worked out argument is beyond the scope of this paper, results such as Caplin and Spulber (1987) suggest that it is reasonable to assume that households will be randomly distributed around their ideal level. For example, if house prices and income are drifting up over time, households that move infrequently will purchase a house that anticipates where they would like to be several years in the future. So some households will be above and some below their ideal.

4.8 Model Simulations

The unit root assumption on productivity growth necessitates gauging the evolution of the model economy when (24) does not hold. Specifically, suppose $(1+\gamma_m)(1+\nu)^{\beta_h-\beta_m} > 1+\gamma_h$ so that p grows over time. We can simulate the model under constant TFP growth around fixed values of $\hat{k}$ and $\hat{x}$ and find the value for g that matches the actual average growth rate of QZ (which is the growth rate of per capita capital and output) in the simulation. We then examine how much the growth rate of QZ varies over time, and what that implies for variation in $\hat{k}$ and $\hat{x}$. If it turns out to be miniscule, then balanced growth is a good approximation and we can linearize around fixed values of $\hat{k}$ and $\hat{x}$ in the usual way.

Table 3 provides the results of this exercise for various values of γ_h implying different average growth rates of p. The simulations were over 50 years (200 quarterly observations), using the parameters described above. The results show that even for the case in which p grows at one percent annually for 50 years (the growth rates in the table are quarterly percentages), the change in the growth rate is tiny and economically inconsequential in terms of its impact on the level of k and x.

Given these findings, we can separate the model conveniently into its dynamic aggregate component, which is approximately the neoclassical growth model, and the sectoral variables, which do not have a simple steady state representation, but are static functions of the aggregate state variables. The connection between the two components is in the growth rate G_t, which depends on the changes in sectoral variables from $t-1$ to t. Thus we use standard methods (e.g. Uhlig, 1997—see the Appendix for more details) to obtain a solution for the linearized aggregate model. It should be noted that parameters related to the sectoral dimension of the model do not enter into this part of the problem. These include β_h, ϵ, γ_h, ω_c, and ω_h. For given realizations of the exogenous disturbances, time paths for the aggregates can be computed, they can be converted back to levels and become inputs to solving the for the sectoral variables (e.g. p, n_m, n_h, etc.) period-by-period.

The key to doing interesting simulations is to take the peculiar error structure of the disturbance process into account. Even though the conditional expectation of the errors in the z_t process (the regime states) is zero, actual realizations of zero are not possible, and in fact given the values of q_1 and q_0, a small error (of absolute value $1-q_0$ or $1-q_1$) that leaves z unchanged is highly likely in any given time period. So rather than consider a one-time shock to v, it makes sense to consider a single large shock (a regime-switch) followed by a sequence of identical small shocks that leave the regime unchanged for an extended period of time. Such a path is more like a modal outcome rather than the improbable mean.

Figure 7 gives an example of this type of simulation. The economy is in the low growth regime in periods 1 to 11, and then switches to the high growth regime, where it remains. The figure plots the behavior of the asset price of a house (a fixed-weight combination of capital and land—see below) against per capita income, for $\epsilon = 0.3$ and 0.9. House prices are clearly much more responsive, both at impact and during the regimes, in the $\epsilon = 0.3$ case. When the regime shift occurs, the price jumps about 5 percent if $\epsilon = 0.3$ versus around 2 percent if $\epsilon = 0.9$. During the high-growth regime the growth rate of the price is about 0.5 percent (annualized) faster if $\epsilon = 0.3$ versus $\epsilon = 0.9$. Prices actually accelerate as long as the economy remains in the high-growth regime, the more so the lower is ϵ.

4.9 Regime Uncertainty

The model solution and simulations above assume perfect knowledge about the growth regime. This is unlikely to be a realistic assumption. Fortunately, the framework in KR provides a natural mechanism for extracting what economic agents know about the growth regime from the behavior of other economic variables. It is a dynamic factor model with

Markov regime-switching in the stochastic growth component (see KR for details).

Figure 8 provides an updated estimate of the so-called smoothed (incorporating all available data through 2007:Q3) and zero-lag (incorporating data for each observation only up to that date) estimates of regime probabilities, incorporating information about the growth regime from data on productivity, labor compensation, aggregate consumption, and aggregate hours of work. The zero-lag estimates provide a one basis for what economic agents might have thought at the time. Note in particular the recent signs of a shift back to the low-growth regime, roughly coincident with the sudden end of the housing boom.

As an extreme benchmark, we can simulate the model under the assumption of perfect information, i.e. that agents actually observe the regime shifts and level shocks η in real time. Figure 9 illustrates the behavior of housing prices according to the model under this extreme assumption, where the simulation assumes, based on Figure 8, that regime shifts occurred in 1973, 1996, and 2004. As suggested by the discussion in the previous section, a regime shift triggers both a level and growth rate change in housing prices. Transitory shocks have relatively little impact since they are not confused with permanent shocks. Note that for these and subsequent simulations, because the linear trend depends on the unobservable γ_h, we just assume a value for γ_h such that the trends line up exactly, and judge the model by its ability to match deviations from the linear trend. By that standard, the perfect information assumption clearly misses on both timing and amplitude.

There are, however, three potential sources of incomplete or imperfect information. First, agents may not observe η_t, the persistent but transitory disturbance to A_{mt}, separately from v_t, the regime-dependent error. This is a standard inference problem addressed by Hamilton (1994) and others. As new data arrive, agents observe A_{mt} and update their assessment of the current regime. It turns out that although this aspect of imperfect information smooths out the implied price series somewhat, the impact is limited by the fact that with full knowledge of the underlying parameters, agents can infer regime switches relatively quickly—typically within two years of when the full sample indicates the switches occurred.

Consequently the resulting price series exhibits most of the flaws of the series under perfect information in Figure 9, attenuated only slightly.

This sort of rational expectations updating is implausible in anything other than a stationary environment in which agents know (or have had sufficient time to learn) the underlying parameters of the stochastic processes. As of the early 1970s there had been no experience in the United States with a sustained productivity slowdown since the Great Depression. To believe that the low productivity growth beginning in 1973 was a change in the trend would have required vivid imagination or exceptional insight. Other than considering the experience of other countries, there would have been no realistic way to form estimates of either the alternative mean growth rate or the transition probabilities.

Thus, a second source of imperfect information is the result of agents not knowing the parameters of the productivity process. One way to approach this is to suppose that expectations are formed by estimating the regime-switching dynamic factor model in "real time"—that is, based on data only available at each point in time.[9] This is similar to least squares learning, except that there is no simple recursion to update the parameter estimates. While estimation of a regime-switching model when there is no regime switch (yet) in the data is potentially problematic, as some parameters are not identified, the point in real time at which the estimation begins to detect a regime switch provides a plausible mechanism for dating when economic agents might have done so (apart from the fact that the econometric techniques had not been developed yet!) It turns out that rolling estimation of the KR model annually through the 1970s confirms that it fails even to detect a second regime until around 1978, and only puts high probability on it in 1979 (see Figure 10a). After that the estimation converges quickly to close to the full sample estimates. It is worth noting that 1979 is also the year when a number of studies about "the productivity slowdown" began to appear (e.g. Denison, 1979; Norsworthy et al., 1979), though even these were primarily

[9] Edge et al (2007) consider a similar question in a linear context using a Kalman filter. Kahn and Rich (2006) apply this paper's approach to the 1990's and find that the model detects the productivity boom relatively quickly—within a year or two of when, with hindsight, it apparently began. This is because by the 1990s the parameter estimates have converged close to the full sample estimates.

retrospective and did not take clear stands on how likely the slowdown was to persist. Edge et al (2007) also report that official forecasts of long-run productivity growth drifted down slowly in the 1970s, but remained above 2 percent through 1978, and then plummeted in 1979 to 1.5 percent.

Yet a third complication is illustrated by Figure 10b. Whereas the KR model estimated with data available as of the end of July 2007—recall Figure 8—shows the current productivity slowdown beginning (with high probability) in 2004, data available prior to July 2007 had not indicated a slowdown with any substantial probability. Even as of June 2007, Figure 10b indicates that the low-growth regime probabilities were in the vicinity of 0.1. Then the benchmark revisions that came out in July 2007 revised productivity growth downward over the previous three years. Thus, in addition to the difficulties presented by learning the parameters of the process, and by—given the parameters—assessing the regime probabilities in real time as the data arrive, there is the third difficulty that the data are regularly revised. While it is possible that the downward revisions did not come as a complete surprise, it is reasonable to assume that there was significant delay in recognizing the magnitude of the productivity slowdown. Hence the real-time estimates from the KR model would appear to be more realistic than the full sample estimates underlying Figure 8.

The general equilibrium model in this paper does not have the complexity that would allow the methods of KR to be applied or simulated directly; in particular, there is only one stochastic variable, so there would be no point to dynamic factor analysis. On the other hand, using actual productivity data would result in unrealistically poor assessments of when the regime switches occurred (see KR). Instead, we construct an artificial productivity process as in $(32)-(34)$, imposing regime switches that correspond to those found in the data, i.e. in 1973, 1996, and 2004. For the η_t process we use the estimated common transitory process from the KR model, scaled so as to obtain similar regime inferences to those obtained by KR.[10] The resulting "productivity" series is then used to estimate the parameters of the

[10] Similar results were obtained using Monte Carlo simulations for the η process, but of course the business cycle peaks and troughs did not generally line up with those in the data.

process along with regime probabilities as of each date from 1963:Q1 forward. This results in a set of parameters associated with each date, which then are used to simulate the model. The only calibrated (fixed) parameters are the discount rate ρ, the depreciation rate δ, and the share parameters β_h and β_m.

More formally, the idea here is to assume that agents observe $A_{mt} = A_{mt-1}(1 + \tilde{\gamma}_{mt})\eta_t/\eta_{t-1}$, but whereas the true process, and accordingly the values of $\tilde{\gamma}_{mt}$ and η_t, are as described earlier, agents instead have time-varying estimates of the parameters $\hat{\gamma}^1_{mt}$, $\hat{\gamma}^0_{mt}$, $\hat{\theta}_{1t}$, $\hat{\theta}_{0t}$, probabilities $\hat{\pi}_{t|t} = \Pr(\xi_t = 1|\Omega_t)$, and transitory terms $\hat{\eta}_{s|t} = E(\eta_s|\Omega_t)$ where Ω_t represents data observed through period t. (We assume only that they know the production function parameters α, β_h, and β_m, and the parameters of the utility function.) They obtain estimates of the parameters by using maximum likelihood (or, rather, an approximation[11] based on Kim, 1994). At each date they update their estimates upon observing the latest realization of A_{mt}, and revise their expectations of future exogenous and endogenous variables accordingly. In other words, by construction agents' estimates of the various components of g_t add up to the actual g_t, which is observed, but have different implications for expectations compared to the complete information case where the parameters are known and the shocks observed. (See the Appendix for additional details)

Figure 11 provides the results of the model simulation of housing prices under the real-time learning assumption, in comparison with the Census and OFHEO indexes. All series are detrended, since the model does not have any prediction about the price trend. While the model misses on the full amplitude of the price fluctuations, it gets most of it: The amplitude of the model's predicted price path is roughly plus or minus 10 percent relative to trend, while both the new and existing home price series get to around 15 percent above trend, and the existing home price series plunges to more than 15 percent below trend in the mid-nineties. With regard to timing, the model tends to anticipate the actual price

[11] The approximation arises because the inferences about η_t and η_{t-1} should involve expectations conditional on not only the current regime, but on the whole history of regimes. Instead these are collapsed into single unconditional values.

increases, especially in the 1960s, but gets the timing of the peaks and troughs remarkably well.

Another check on the model is its predictions about investment. Clearly the lack of adjustment costs at the sectoral level will mean that sectoral investment will be excessively volatile at high frequencies. Aggregate adjustment costs appear to have little impact on this volatility: With relative price changes, the model makes it too easy to flip capital from one sector to the other. Nonetheless we can still examine the model's predictions and focus on the lower frequencies of interest. Figure 12 displays the behavior of residential investment (defined as the change in residential capital), along with actual series. Both series are detrended, the model's by the steady state growth trend, and the data by a Hodrick-Prescott trend for real GDP. Just as the model overpredicts prices in the early part of the sample, it overpredicts residential investment as well. Overall the model hits the medium frequency movements reasonably well, albeit predicting somewhat greater medium-frequency amplitude than is in the data. The only substantive miss after 1975 is the model's prediction of a mini-cycle in the second half of the 1990s that never occurred. This is the mirror image of the model's slightly premature prediction of a price runup beginning at the same time. It appears that the artificial learning that takes place in the model is a little too rapid, at least in response to favorable news.

The finding that the model slightly underpredicts the amplitude of price fluctuations and overpredicts the amplitude of investment fluctuations is not surprising given the lack of any adjustment costs or other real frictions. In addition, the model omits demographics, which may have played some role, particularly in the early part of the sample in the wake of the baby boom: There was an explosion of family formation and residential investment following World War II, and the offspring in those families had for the most part not yet come of age. This may have held down both prices and investment in the 1960s and early 1970s, and boosted them in the late 1970s and 1980s and again in recent years during the "echoes" of the baby boom.

Overall, given its parsimony and simplicity, and the fact that it has only one driving force, the model explains a substantial part of both price and quantity fluctuations. This is noteworthy considering that the key parameters—the elasticity of substitution, the land share parameters in the production functions, and the productivity process, are calibrated either from micro data or, in the case of the productivity process, from empirical exercises that do not take the housing sector into account. Certainly there are parameters that would result in an even better fit, but the goal here was not to explain everything, only to quantify the role of changes in trend productivity.

5 Conclusions

This paper has developed a growth model with land, housing services, and other goods and shown that it is capable both qualitatively and quantitatively of explaining a substantial portion of the movements in housing prices and residential investment over the past 40 years, including the recent downturn. The paper also uses micro data to calibrate a key cross-elasticity parameter that governs the relationship between productivity growth and home price appreciation. The matching of the model to the data relies not on fitting the overall trend (which depends on an unobservable), but on the deviations from that trend as a function of productivity growth. The calibrated model under rational expectations can explain some of the acceleration in housing prices that occurred both in the 1960s and since the mid-1990s, and also suggests a contributing explanation for the recent downturn, but fails to get the full amplitude and timing of the fluctuations. When a realistic model of learning is added in lieu of rational expectations, however, the model does much better on both. In particular, the continued boom in housing prices in the 1970s is largely explained by the time it took for agents to figure out that the productivity slowdown was quasi-permanent. Finally, the paper also has some success in matching the low frequency behavior of housing investment, in particular the boom that began in the late 1990s. Even so, future work will incorporate adjustment costs, and other sources of shocks to capture more accurately the high frequency behavior of both prices and quantities.

Another implication of this analysis is that it sheds new light on recent subprime mortgage crisis and the ripple effects on financial markets. If we accept the idea that the housing downturn was triggered by a productivity slowdown, and that this was a low probability event in light of what was known prior to 2007 about productivity growth, it may be that subprime mortgage lending was, ex ante, economically rational. The probability of a productivity slowdown—and therefore of a reversal in housing prices—such as the one that occurred in 1973 could reasonably have been viewed as small over the relevant horizon for mortgage lenders.. Indeed, a recent paper by Piskorski and Tchistyi (2008) makes precisely this point. They examine optimal mortgage lending in a setting where housing prices obey essentially the same type of regime-switching behavior, and find that "many features of subprime lending observed in practice are consistent with economic efficiency and rationality of both borrowers and lenders," though, as they point out, there may be negative externalities associated with massive defaults in a downturn.

Moreover, if indeed house prices are driven to a great extent by long-term fundamentals like productivity growth, and not so much by credit market frictions, this has important implications not just for a historical understanding of the causes of the housing bust that began in 2005, but also for policies aimed at shoring up house prices by providing foreclosure relief, subsidized mortgages, and the like. The results in this paper suggest that such policies are likely to have limited impact, since they do not affect the economy's long-term growth potential. More generally, the results suggest that policies aimed at stimulating the economy by supporting house prices are getting it backwards: The best way to support house prices is to undertake policies that support long-term economic growth.

6 Appendix

6.1 Solving the Closed Economy Model under Complete Information

Proof that $\mu_m = x^{-1}$: $\mu_m = u_c = \phi_c/\phi$:

Since ϕ is homogeneous of degree one, $\phi = \phi_c c + \phi_h h$. Also $p = \phi_h/\phi_c$, so $\phi = \phi_c (c + ph) = \phi_c x$. Hence $\phi_c/\phi = x^{-1}$.

Modeling the regime shift process:

Since the stationary distribution of ξ is $\xi^* \equiv \begin{bmatrix} \frac{1-\theta_0}{2-\theta_1-\theta_0} & \frac{1-\theta_1}{2-\theta_1-\theta_0} \end{bmatrix}'$, the average growth rate of A_m is

$$\bar{\gamma}_m = \frac{1-\theta_0}{2-\theta_1-\theta_0}\gamma_m^1 + \frac{1-\theta_1}{2-\theta_1-\theta_0}\gamma_m^0. \tag{38}$$

For concreteness we will call $\xi = 1$ the "high-growth" regime, and $\xi = 0$ the "low-growth" regime, i.e. we assume $\gamma_m^1 > \gamma_m^0$.

Elaborating on Hamilton (1994), we can describe μ as an AR(1) process. We have

$$\xi_t = 1 - \theta_0 + (\theta_1 + \theta_0 - 1)\,\xi_{t-1} + v_t \tag{39}$$

where $E_{t-1}(v_t) = 0$, and is distributed as follows:

$$v_t = \begin{cases} 1-\theta_1 & \text{Prob } \theta_1 \\ -\theta_1 & \text{Prob } 1-\theta_1 \end{cases} \text{if } \xi_{t-1} = 1 \quad \begin{cases} -(1-\theta_0) & \text{Prob } \theta_0 \\ \theta_0 & \text{Prob } 1-\theta_0 \end{cases} \text{if } \xi_{t-1} = 0 \;. \tag{40}$$

Note that while $E\left(v_t|\xi_{t-1}\right) = 0$, v_t is not identically distributed over time, as the conditional distribution depends on ξ_{t-1}.

The first-order conditions for the planner are as follows: Letting $\phi(c,h) \equiv \left[\omega_c c^{(\epsilon-1)/\epsilon} + \omega_h h^{(\epsilon-1)/\epsilon}\right]^{\epsilon/(\epsilon-1)}$

we have:

$$\begin{aligned}
\omega_c \phi^{-(\epsilon-1)/\epsilon} c_t^{-1/\epsilon} &= \mu_{mt} &(41)\\
\omega_h \phi^{-(\epsilon-1)/\epsilon} h_t^{-1/\epsilon} &= \mu_{ht} &(42)\\
\mu_{mt}\beta_m A_{mt} k_{mt}^{\alpha} \ell_{mt}^{\beta_m-1} &= \mu_{ht}\beta_h A_{ht} k_{ht}^{\alpha} \ell_{ht}^{\beta_h-1} &(43)\\
\mu_{mt} A_{mt} k_{mt}^{\alpha-1} \ell_{mt}^{\beta_m} &= \mu_{ht} A_{ht} k_{ht}^{\alpha-1} \ell_{h}^{\beta_h} &(44)\\
\mu_{mt} A_{mt} k_{mt}^{\alpha} \ell_{mt}^{\beta_m} &= \mu_{ht} A_{ht} \times &(45)\\
& \left[\alpha k_{ht}^{\alpha-1} \ell_{ht}^{\beta_h} k_{mt} + \beta_h k_{ht}^{\alpha} \ell_{ht}^{\beta_h-1} \ell_{mt} + (1-\alpha-\beta_h) k_{ht}^{\alpha} \ell_{ht}^{\beta_h}\right] &(46)\\
\lambda_t \; z'(i_t/k_{t-1}) &= \mu_{mt} &(47)\\
\lambda_t (1+\nu)(1+\rho) &= E_t \Big\{ \mu_{mt+1} A_{mt+1} \alpha k_{mt+1}^{\alpha-1} \ell_{mt+1}^{\beta_m} + &(48)\\
& \lambda_{t+1} \left[z(i_{t+1}/k_t) - (i_{t+1}/k_t) z'(i_{t+1}/k_t) + 1 - \delta \right] \}
\end{aligned}$$

μ_{mt}, μ_{ht}, and λ_t are shadow prices on the resource constraints (11), (13), and (12). Note that in the absence of adjustment costs, i.e. when $z(x) = x$, we have $\lambda_t = \mu_{mt}$, and (48) becomes

$$\mu_{mt}(1+\nu)(1+\rho) = E_t \left\{ \mu_{mt+1} \left[A_{mt+1} \alpha k_{mt+1}^{\alpha-1} \ell_{mt+1}^{\beta_m} + 1 - \delta \right] \right\} \tag{49}$$

which is just the familiar condition that the intertemporal marginal rate of substitution equals the marginal product of capital.

To solve the model, first we linearize the system

$$\begin{aligned}
\hat{x}_t + \hat{\imath}_t &= \left[\hat{k}_{t-1}/G_t\right]^{\alpha}\\
G_t(1+\nu)\hat{k}_t &= z\left(G_t \hat{\imath}_t/\hat{k}_{t-1}\right)\hat{k}_{t-1} + (1-\delta)\hat{k}_{t-1}\\
(1+\nu)(1+\rho) q_t &= E_t \Big\{ (\hat{x}_t/\hat{x}_{t+1}) G_t^{-1} \left[\alpha \left[\hat{k}_t/G_t\right]^{\alpha-1} + \right.\\
& \left. q_{t+1} \left[z\left(G_t \hat{\imath}_{t+1}/\hat{k}_t\right) + 1 - \delta \right] - G_t \hat{\imath}_{t+1}/\hat{k}_t \right] \Big\}\\
q_t &= z'\left(G_t \hat{\imath}_t/\hat{k}_{t-1}\right)^{-1}
\end{aligned}$$

around the quasi-steady state values $\hat{k}$, $\hat{x}$, and G.. After some rearranging, and letting $R \equiv \alpha \left[\hat{k}/G\right]^{\alpha-1} + 1 - \delta = (1+\rho)(1+\nu)G$, the linearized versions of the four equations can be expressed as

$$
\begin{aligned}
G(1+\nu)\hat{k}_t &= R\left(\hat{k}_{t-1} - G_t\right) - \hat{x}_t\left[G\hat{x}/\hat{k}\right] \\
\hat{x}_t\left[G\hat{x}/\hat{k}\right] + \hat{\imath}_t\left[G\hat{\imath}/\hat{k}\right] &= \left[R - (1-\delta)\right]\left(\hat{k}_{t-1} - G_t\right) \\
Rq_t &= E_t\Big\{\left[\hat{x}_t - \hat{x}_{t+1} - G_t\right]R \\
&\quad -\alpha(1-\alpha)\left(\frac{\hat{k}}{G}\right)^{\alpha-1}\left(\hat{k}_t - G_t\right) + q_{t+1}\left(G\hat{\imath}/\hat{k} + 1 - \delta\right)\Big\} \\
0 &= q_t + z''\left(G\hat{\imath}/\hat{k}\right)\left(G_t + \hat{\imath}_t - \hat{k}_{t-1}\right)
\end{aligned}
$$

Note that

$$
\begin{aligned}
\alpha\left(\frac{\hat{k}}{G}\right)^{\alpha-1} + 1 - \delta &= (1+\rho)(1+\nu)G \\
G\hat{x}/\hat{k} &= \left(\frac{\hat{k}}{G}\right)^{\alpha-1} + 1 - \delta - (1+\nu)G.
\end{aligned}
$$

The log deviation version of G_t can be written as

$$
G_t = \frac{1}{1-\alpha}\left[\frac{\tilde{\gamma}_{mt} - \bar{\gamma}_m}{1+\bar{\gamma}_m} + \Delta\eta_t + \frac{\tau n_h}{1+\tau n_h}\Delta n_{ht} - \beta_m \Delta\ell_{mt}\right] \tag{50}
$$

where Δn_{ht} and $\Delta\ell_{mt}$ are deviations from their local means given growth at G. Thus G_t reflects the endogenous resource shift toward or away from the h sector in response to the shock.[12]

The near-balanced growth behavior of the system means that we can appoximate the

[12] Note that the unit root in productivity means that, strictly speaking, the dynamics of the model are not independent of the initial position in levels, as n_h will will not be stationary. It turns out, however, that this variation is miniscule, because n_h evolves slowly and the coefficient on $(n_{ht} - n_{ht-1})$ above is sufficiently small and insensitive to n_h so that this can be neglected.

growth rates of the sectoral variables by linearizing the system

$$
\begin{aligned}
p_t h_t &= A_{mt} k_{mt}^{\alpha} \ell_{mt}^{\beta_m} (1+\tau) n_{ht} / (1+\tau n_{ht}) && (51)\\
1 &= \omega_c \phi (c_t, h_t)^{-(\epsilon-1)/\epsilon} c_t^{-1/\epsilon} x_t && (52)\\
h_t / c_t &= p_t^{-\epsilon} (\omega_h/\omega_c)^{\epsilon} && (53)\\
\ell_{mt} &= \frac{\bar{L}}{N_t} \frac{\beta_m/\beta_h}{n_{ht}(1+\tau-\beta_m/\beta_h)+\beta_m/\beta_h} && (54)\\
p_t &= \frac{A_{mt}}{A_{ht}} \left(\frac{\beta_m}{\beta_h}\right)^{\beta_h} (1+\tau)^{1-\alpha-\beta_h} \ell_{mt}^{-(\beta_h-\beta_m)}. && (55)\\
k_{mt} &= k_{t-1} / (1+\tau\, n_{ht}) && (56)
\end{aligned}
$$

in logarithmic first differences around the values that obtain under aggregate growth of G and population growth ν. Here we can assume that k_m and x grow at G, while h, c, n_h, and ℓ_m each grow at locally constant rates, which as we have seen is a good approximation over periods of many decades. This results in a system of the form:

$$\Phi \Delta s_t = \Omega \Delta m_t \tag{57}$$

where $s_t \equiv [p_t\ c_t\ h_t\ \ell_{mt}\ n_{ht}\ k_{mt}]'$ and $\Delta m_t \equiv \begin{bmatrix} \gamma_{mt} & \Delta x_t & \Delta k_{t-1} & 0 & 0 \end{bmatrix}'$, all in deviations. Φ is 6×6 and Ω is 6×5.

The local constant growth rates can be found from

$$\Phi^{-1} \Omega \Delta \bar{m} \tag{58}$$

where $\Delta \bar{m} = \begin{bmatrix} \gamma_{mt} & g & g & \nu & \gamma_h \end{bmatrix}'$. We then have

$$
\begin{aligned}
\Delta s_t &= \Phi^{-1} \Omega \Delta m_t && (59)\\
G_t &= \zeta' \Delta m_t = \pi' \Phi^{-1} \Omega \Delta m_t + \frac{1}{1-\alpha} \gamma_{mt} && (60)
\end{aligned}
$$

wher ζ is 5×1 and $\pi = \left[\begin{array}{ccccccc} 0 & 0 & 0 & \frac{\beta_m}{1-\alpha} & \tau n_h / (1 + \tau n_h) & 0 & \frac{1-\alpha-\beta_m}{1-\alpha} \end{array} \right]'$ from (50) above. Thus the growth rate can be expressed in terms of the first differences of the aggregate exogenous and endogenous variables.

The result is a system is of the form (if we neglect adjustment costs for the sake of exposition):

$$
\begin{aligned}
0 &= \mathbf{A} \begin{bmatrix} \hat{k}_t \\ \hat{k}_{t-1} \\ \hat{x}_t \\ \hat{x}_{t-1} \end{bmatrix} + \mathbf{B} \begin{bmatrix} \hat{k}_{t-1} \\ \hat{k}_{t-2} \\ \hat{x}_{t-1} \\ \hat{x}_{t-2} \end{bmatrix} + \mathbf{C} G_t + \mathbf{D} \Lambda_t \\
0 &= E_t \left\{ \mathbf{F} \begin{bmatrix} \hat{k}_{t+1} \\ \hat{k}_t \\ \hat{x}_{t+1} \\ \hat{x}_t \end{bmatrix} + \mathbf{G} \begin{bmatrix} \hat{k}_t \\ \hat{k}_{t-1} \\ \hat{x}_t \\ \hat{x}_{t-1} \end{bmatrix} + \mathbf{H} \begin{bmatrix} \hat{k}_{t-1} \\ \hat{k}_{t-2} \\ \hat{x}_{t-1} \\ \hat{x}_{t-2} \end{bmatrix} + \mathbf{J} G_{t+1} + \mathbf{K} G_t + \mathbf{L} \Lambda_{t+1} + \mathbf{M} \Lambda_t \right\} \\
\Lambda_{t+1} &= \mathbf{N} \Lambda_t + \Xi_{t+1}
\end{aligned}
$$

where

$$
\Lambda_t = \left[\begin{array}{ccc} \hat{\xi}_t & \eta_t & \eta_{t-1} \end{array} \right]',
$$

$\hat{\xi}_t \equiv \xi_t - \bar{\xi}$, and

$$\mathbf{A} = \begin{bmatrix} (1+\nu)\,G & 0 & G\hat{x}/\hat{k} & 0 \\ 0 & -\zeta_1 & -\zeta_2 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$\mathbf{B} = \begin{bmatrix} -R & 0 & 0 & 0 \\ 0 & \zeta_1 & \zeta_2 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \end{bmatrix}, \ \mathbf{C} = \begin{bmatrix} R \\ 1 \\ 0 \\ 0 \end{bmatrix}, \ \mathbf{D} = \begin{bmatrix} 0 & 0 & 0 \\ -\zeta_1 - \frac{1}{1-\alpha}\frac{\gamma_m^1 - \gamma_m^0}{1+\bar{\gamma}_m} & -1 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\mathbf{F} = \begin{bmatrix} 0 & 0 & -R & 0 \end{bmatrix}, \ \mathbf{G} = \begin{bmatrix} -(1-\alpha)\left[R-(1-\delta)\right] & 0 & R & 0 \end{bmatrix}$$

$$\mathbf{H} = \begin{bmatrix} 0 & 0 & G\hat{\imath}/\hat{k} + 1 - \delta & 0 \end{bmatrix}$$

$$\mathbf{J} = (1-\alpha)\left[R-(1-\delta)\right] - R$$

$$\mathbf{K} = 0, \ \mathbf{L} = 0, \ \mathbf{M} = 0$$

$$\mathbf{N} = \begin{bmatrix} \theta_1 + \theta_0 - 1 & 0 & 0 \\ 0 & \phi_1 & \phi_2 \\ 0 & 1 & 0 \end{bmatrix}$$

and

$$\Xi_t = \begin{bmatrix} v_t \\ u_t \\ 0 \end{bmatrix}$$

where v_{1t} and v_{2t} are as defined earlier. We can then use the method of undetermined coefficients outlined by Uhlig (1997) to find the solution of the model in the form

$$\begin{aligned}
\begin{bmatrix} \hat{k}_t \\ \hat{k}_{t-1} \\ \hat{x}_t \\ \hat{x}_{t-1} \end{bmatrix} &= \mathbf{P} \begin{bmatrix} \hat{k}_{t-1} \\ \hat{k}_{t-2} \\ \hat{x}_{t-1} \\ \hat{x}_{t-2} \end{bmatrix} + \mathbf{Q}\Lambda_t, \; G_t = \mathbf{R} \begin{bmatrix} \hat{k}_{t-1} \\ \hat{k}_{t-2} \\ \hat{x}_{t-1} \\ \hat{x}_{t-2} \end{bmatrix} + \mathbf{S}\Lambda_t \\
\Lambda_t &= N\Lambda_{t-1} + \Xi_t.
\end{aligned}$$

where in this we know that the $(2,1)$ and $(4,3)$ elements of $\mathbf{P}$ are 1 and the remaining elements of the two rows are zeros. Given paths for $\hat{k}_t$, $\hat{x}_t$, and A_{mt} we can compute the path of s_t, i.e. the levels of the sectoral variables and work effort, using the system $(51)-(56)$ The rates of change of these variables will be very similar to those computed from (59), the linearized system.

6.2 Learning

Now the system describing the equilibrium is a time-dependent version of what we had under complete information:

$$\begin{aligned}
0 &= \mathbf{A}_t \begin{bmatrix} \hat{k}_t \\ \hat{k}_{t-1} \\ \hat{x}_t \\ \hat{x}_{t-1} \end{bmatrix} + \mathbf{B}_t \begin{bmatrix} \hat{k}_{t-1} \\ \hat{k}_{t-2} \\ \hat{x}_{t-1} \\ \hat{x}_{t-2} \end{bmatrix} + \mathbf{C}_t G_t + \mathbf{D}_t \Lambda_t \\
0 &= E_t \left\{ \mathbf{F}_t \begin{bmatrix} \hat{k}_{t+1} \\ \hat{k}_t \\ \hat{x}_{t+1} \\ \hat{x}_t \end{bmatrix} + \mathbf{G}_t \begin{bmatrix} \hat{k}_t \\ \hat{k}_{t-1} \\ \hat{x}_t \\ \hat{x}_{t-1} \end{bmatrix} + \mathbf{H}_t \begin{bmatrix} \hat{k}_{t-1} \\ \hat{k}_{t-2} \\ \hat{x}_{t-1} \\ \hat{x}_{t-2} \end{bmatrix} + \mathbf{J}_t G_{t+1} + \mathbf{K}_t G_t + \mathbf{L}_t \Lambda_{t+1} + \mathbf{M}_t \Lambda_t \right\}.
\end{aligned}$$

The exogenous state vector Λ_t is reassessed at each date according to the procedure described in the text. The solution will be similarly time-dependent:

$$\begin{bmatrix} \hat{k}_t \\ \hat{k}_{t-1} \\ \hat{x}_t \\ \hat{x}_{t-1} \end{bmatrix} = \mathbf{P}_t \begin{bmatrix} \hat{k}_{t-1} \\ \hat{k}_{t-2} \\ \hat{x}_{t-1} \\ \hat{x}_{t-2} \end{bmatrix} + \mathbf{Q}_t \Lambda_t$$

$$G_t = \mathbf{R}_t \begin{bmatrix} \hat{k}_{t-1} \\ \hat{k}_{t-2} \\ \hat{x}_{t-1} \\ \hat{x}_{t-2} \end{bmatrix} + \mathbf{S}_t \Lambda_t.$$

As with many learning models, agents do not take into account the fact that there is parameter uncertainty and that their beliefs about the parameter values and shocks will change over time

References

[1] Baumol, William J., "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *American Economic Review* 57(3), June 1967, 415–426.

[2] Attanasio, O., L. Blow, R. Hamilton, A. Leicester, "Consumption, House Prices and Expectations," Bank of England Working Paper Series no. 271 (2005).

[3] Benhabib, J., R. Rogerson, and R. Wright, "Homework in Macroeconomics: Aggregate Fluctuations," *Journal of Political Economy* 99(6), 1166-1187.

[4] Caplin, A., and D. Spulber, "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics* 102 (November 1987), 703-726.

[5] Case, K.E., and Shiller, R.J., "Is There a Bubble in the Housing Market?" *Brookings Papers on Economic Activity*, No. 2, 2003

[6] Davis, M., and J. Heathcote, "The Price and Quantity of Residential Land in the United States," manuscript, 2004.

[7] Davis, Morris A., and Jonathan Heathcote, "Housing and the Business Cycle," *International Economic Review* 46 (2005) , 751–784

[8] Denison, Edward, *Accounting for Slower Economic Growth: The United States in the 1970s* (Brookings Institution, 1979).

[9] Edge, R., T. Laubach, and J. Williams, "Learning and Shifts in Long-run Productivity Growth," *Journal of Monetary Economics* 54(6), November 2007, 2421-2438.

[10] Flavin, M., and S. Nakagawa, "A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence," *American Economic Review* 98, March 2008 , 474-495.

[11] Glaeser, E., and J. Gyourko, "Why Have House Prices Gone Up?" American Economic Review Papers and Proceedings 97(2), May 2005, 329-333.

[12] Gordon, R. and vanGoethem, T. "A Century of Downward Bias in the Biggest CPI Component: The Case of Rental Shelter, 1914-2003," in E. Berndt and C. Hulten, eds., Hard-to- Measure Goods and Services: Essays in Honor of Zvi Griliches, Conference on Research in Income and Wealth (Chicago: University of Chicago Press for NBER, 2006).

[13] Greenwood, J., and Z. Hercowitz, "The Allocation of Capital and Time over the Business Cycle" *The Journal of Political Economy* 99(6), December 1991, 1188-1214.

[14] Hamilton, James, *Time Series Econometrics*, Princeton: Princeton University Press, 1994.

[15] Hanushek, Eric A., and John M. Quigley, "What is the Price Elasticity of Housing Demand?" *The Review of Economics and Statistics* 62(3), August 1980, 449-454

[16] Himmelberg, C., C. Mayer, and T. Sinai, "Assessing High House Prices: Bubbles: Fundamentals, and Misperceptions" Federal Reserve Bank of New York Staff Report no. 218, September 2005.

[17] Iacoviello, M., and Neri, S., "The Role of Housing Collateral in an Estimated Two-Sector Model of the US Economy," manuscript, 2006.

[18] Kahn, James A., "Understanding the World Housing Boom," manuscript, 2008.

[19] Kahn, J.A., and R. Rich, "Tracking Productivity Growth in Real Time" *Current Issues in Economics and Finance*, Federal Reserve Bank of New Yori, 2006.

[20] Kahn, J.A., and R. Rich, "Tracking the New Economy: Using Growth Theory to Detect Changes in Trend Productivity," *Journal of Monetary Economics* 54(6), September 2007, 1670-170.

[21] .Kim, C.-J., and C. R. Nelson, “State-Space Models with Regime Switching: Classical and Gibbs-sampling Approaches with Applications,” MIT Press, 1998.

[22] Kiyotaki, N., A. Michaelides, and K. Nikolov, “Winners and Losers in Housing Markets,” manuscript, 2007.

[23] Martin, R.F., “The Baby Boom: Predictability in House Prices and Interest Rates,” International Finance Discussion Papers No. 847, Federal Reserve Board, November 2005.

[24] Norsworthy, J.R., M.J. Harper, and K. Kunze, “The Slowdown in Productivity Growth: Analysis of Some Contributing Factors,” *Brookings Papers on Economic Activity*, 1979(2), 387-421.

[25] Li, W., H.Liu, R. Yao, “Housing and Non-Housing Consumption over the Life-Cycle: A Structural Estimation,” manuscript, 2008.

[26] Lustig, H., and S. Van Niewerburgh,. “Can Housing Collateral Explain Long-Run Swings in Asset Returns?” manuscript, 2007.

[27] Mankiw, N.G., and D.N. Weil, “The Baby Boom, the Baby Bust, and the Housing Market,” *Regional Science and Urban Economics* 19(2), May 1989, 235-258.

[28] Ngai, L. Rachel, C.A. Pissarides, “Structural Change in a Multi-Sector Model of Growth,” *American Economic Review* 97(1), March 2007, 429-443.

[29] Ortalo-Magné, F., and S. Rady, “Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints,” *Review of Economic Studies* 73, April 2006, 459-485.

[30] Piazzesi, M., and M. Schneider, “Housing, Consumption, and Asset Pricing,” *Journal of Financial Economics* 83, March 2007, 531-569.

[31] Piskorski, T., and A.Tchistyi, “Stochastic House Appreciation and Optimal Mortgage Lending” manuscript, 2008.

[32] Polinsky, A. Mitchell and David T. Ellwood,“An Empirical Reconciliation of Micro and Grouped Estimates of the Demand for Housing,” *The Review of Economics and Statistics* 61(2), May, 1979, 199-205

[33] Van Nieuwerburgh, S., and P.-O. Weill, “Why Has House Price Dispersion Gone Up,” NBER Working Paper No. W12538, September 2006.

[34] Uhlig, Harald, “A Toolkit for Analyzing Nonlinear Dynamic Stochastic Models Easily.” Manuscript, 1997

Table 1: Results from Static Model

	$\epsilon = 1$		$\epsilon = .5$		$\epsilon = .2$	
	p	K_h/L_h	p	K_h/L_h	p	K_h/L_h
$A = 1$	0.722	0.168	0.809	0.210	0.893	0.257
$A = 2$	1.444	0.168	1.769	0.252	2.126	0.363

Table 2: Parameter Estimates from CEX Household Data

Parameter			
$\hat{\epsilon}$	0.134	0.195	0.284
	(0.042)	(0.046)	(0.052)
$\hat{b}$	-0.743	-0.254	—
	(0.003)	(0.009)	
Instruments for x	N	Y	—
R^2	0.575	0.464	0.317

Table 3: Deviations from Balanced Growth

	γ_h	0.0050	0.0041	0.0032
	avg. p growth*	0	0.126	0.253
QZ growth*	mean	0.5587	0.5644	0.5724
	min	0.5587	0.5629	0.5654
	max	0.5587	0.5661	0.5810
$\hat{k}$	min	25.143	25.102	24.928
	max	25.143	25.071	25.078
$\hat{x}$	min	2.215	2.2136	2.2128
	max	2.215	2.2126	2.2080

*quarterly percentage growth rate (not annualized)

Figure 1: Alternative Home Price Indexes (Inflation-Adjusted)

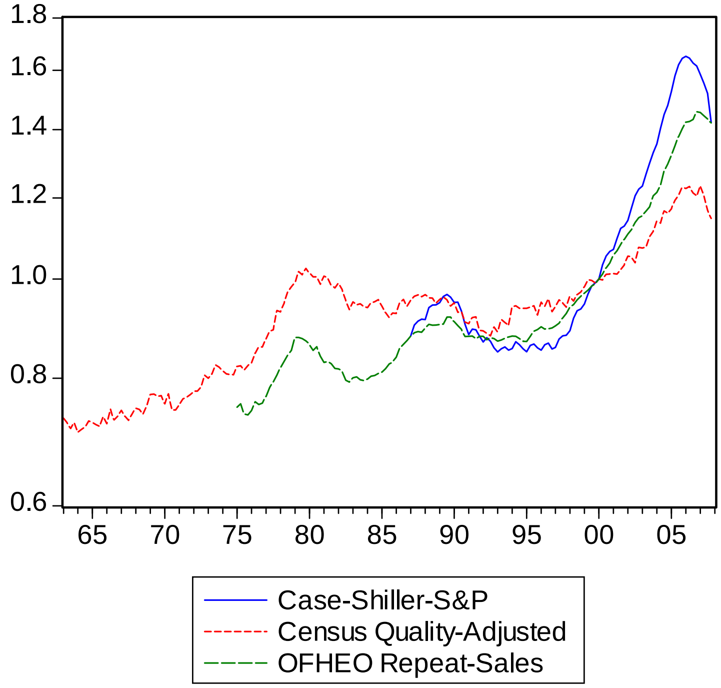

Note: Logarithmic scale, 2000:Q1 = 1.00

Figure 2: Ratio of Housing Wealth to Consumption

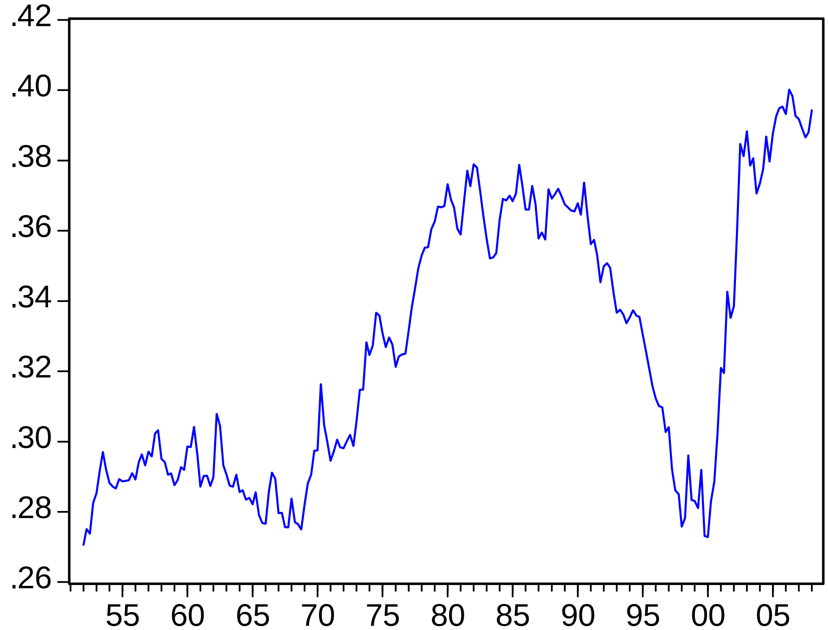

Figure 3: Ratio of Housing Wealth to Total Net Worth

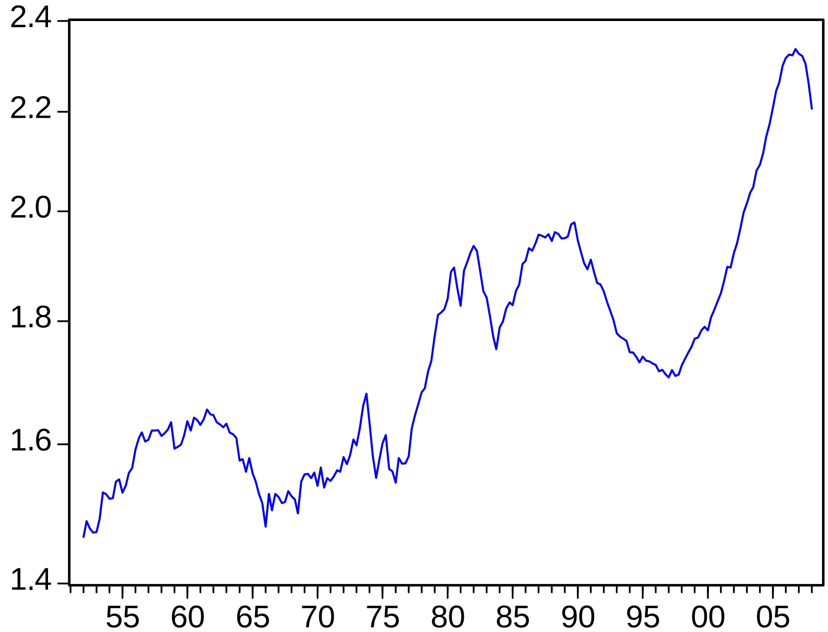

Source: Flow of Funds

Figure 4: Inflation-Adjusted Land Prices
(relative to linear trend)

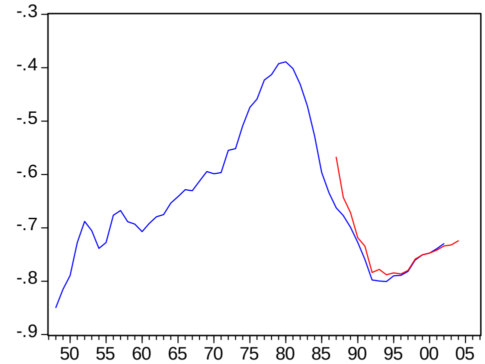

Source: BLS

Figure 5: Non-Farm Productivity
(HP-smoothed, relative to linear trend)

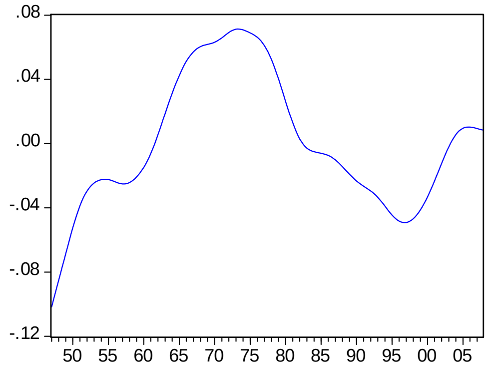

Source: BLS

Note: Both series are in logarithms. The land series are from different vintages of BLS data

Figure 6: Housing Services: Expenditures and Prices

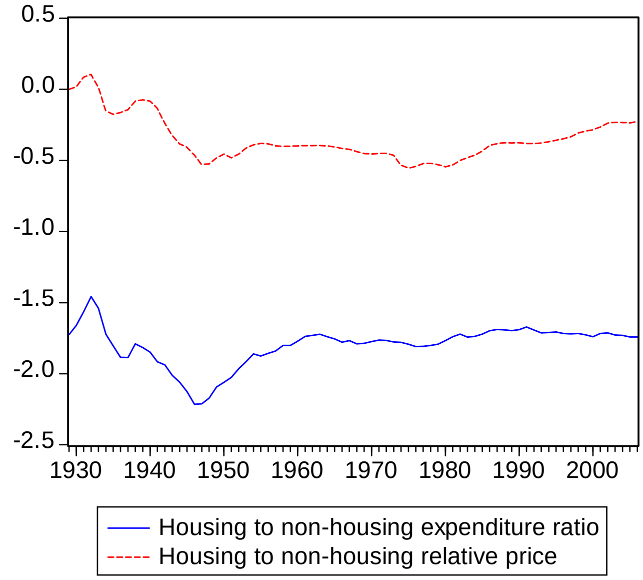

Note: The chart depicts the logarithm of the ratios

Figure 7: Housing Price Response to
Low-to-High Growth Regime Switch

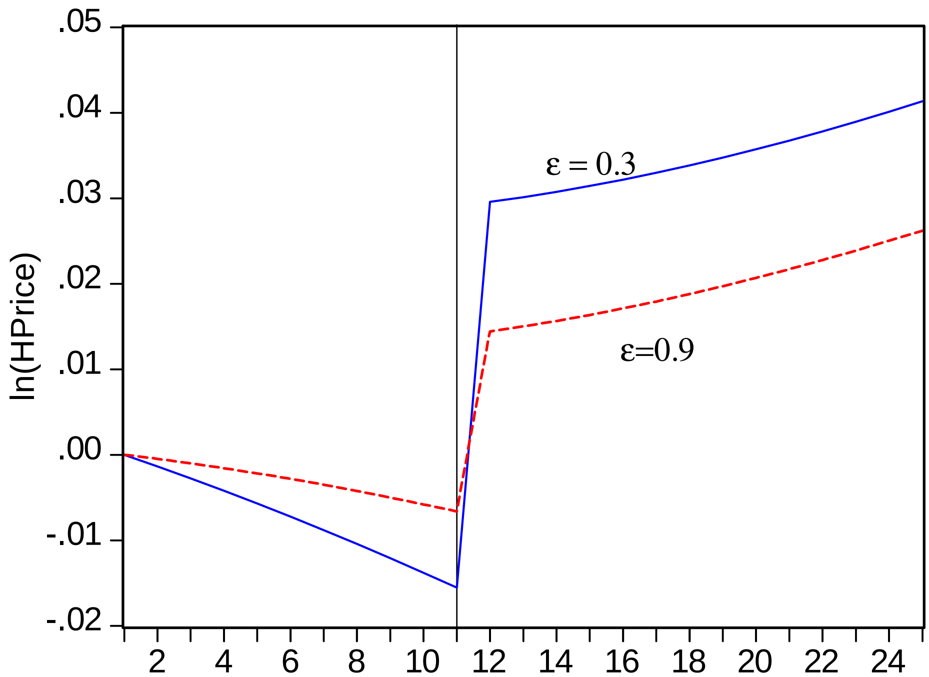

Figure 8: High-Growth Regime Probabilities

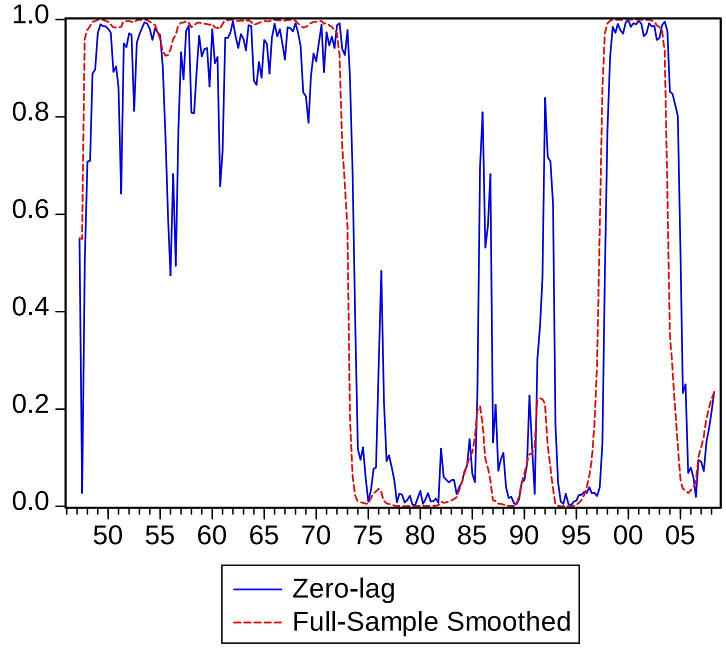

Calculations based on Kahn and Rich (2007)

Figure 9: Model Simulation of Housing Prices

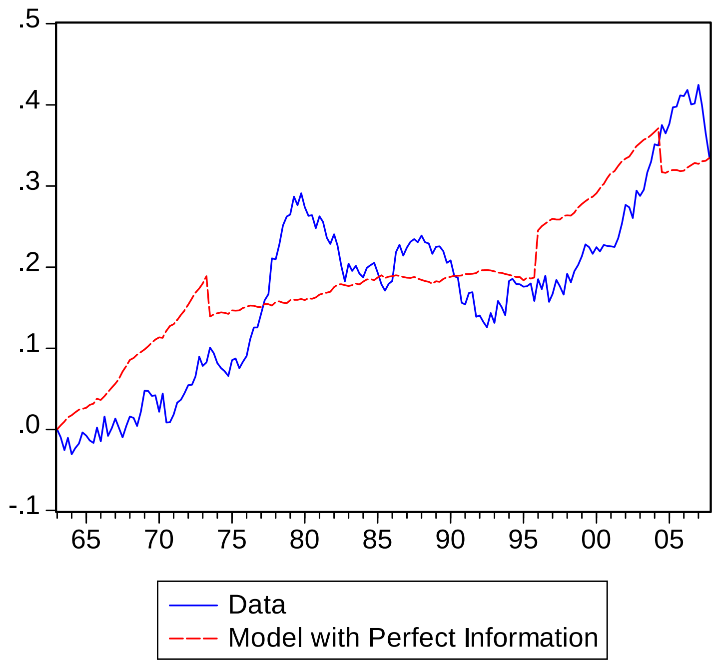

Figure 10a: Real-Time Probabilities of Low-Growth Regimes

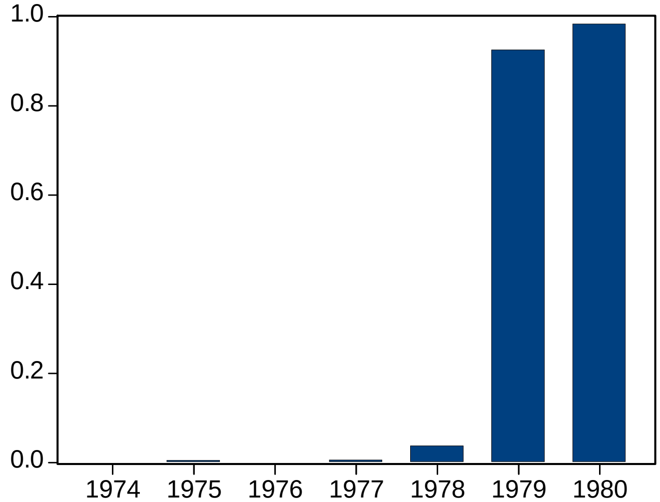

Figure 10b: Real-Time Low-Growth Regime Probabilities since 2005

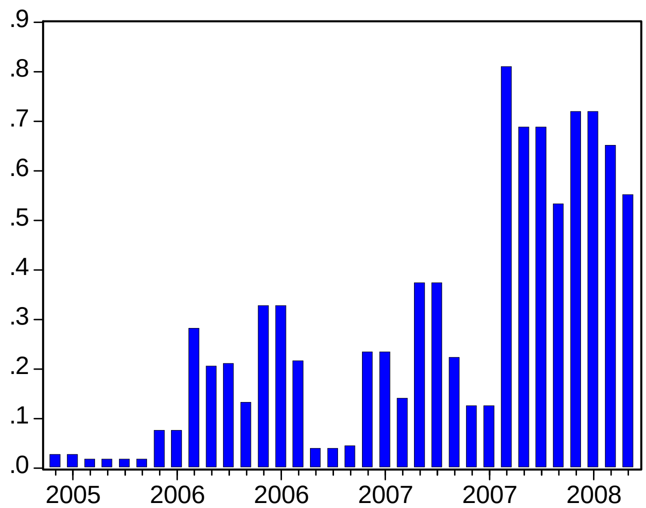

Figure 11: Model vs. Actual Housing Prices (detrended)

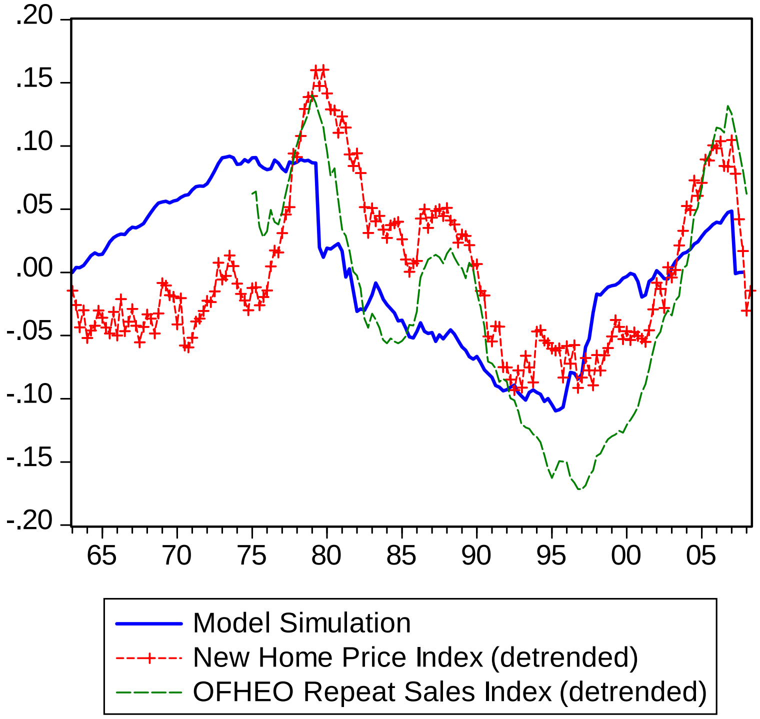

Figure 12: Model vs. Actual Residential Investment (detrended)