Journal of Urban Economics 66 (2009) 90–102

Contents lists available at ScienceDirect

Journal of Urban Economics

www.elsevier.com/locate/jue

What moves housing markets: A variance decomposition of the rent–price ratio ☆

Sean D. Campbell [a], Morris A. Davis [b,*], Joshua Gallin [a], Robert F. Martin [a]

[a] Board of Governors of the Federal Reserve System, Washington, DC, USA
[b] Wisconsin School of Business, University of Wisconsin-Madison, WI, USA

ARTICLE INFO

Article history:
Received 4 March 2009
Revised 19 June 2009
Available online 26 June 2009

JEL classification:
E31
G12
R31

Keywords:
Rent–price ratio
House prices
Housing rents
Interest rates

ABSTRACT

We apply the dynamic Gordon growth model to the housing market in 23 US metropolitan areas, the four Census regions, and the nation from 1975 to 2007. The model allows the rent–price ratio at each date to be split into the expected present discounted values of rent growth, real interest rates, and a housing premium over real rates. We show that housing premia are variable and forecastable and account for a significant fraction of rent–price ratio volatility at the national and local levels, and that covariances among the three components damp fluctuations in rent–price ratios. Thus, explanations of house-price dynamics that focus only on interest rate movements and ignore these covariances can be misleading. These results are similar to those found for stocks and bonds.

1. Introduction

The boom and bust to house prices and housing returns over the past 12 years is likely unprecedented in the United States. According to data from the Bureau of Economic Analysis and MacroMarkets LLC, real house prices in the United States increased by about 6-1/2% per year over the 1997–2006 period. To put this growth in context, over the decade spanning 1987–1996, the same data sources suggest that real house prices in the United States did not increase at all; and, the available evidence suggests that real house prices in the United States increased by less than 2% per year in real terms over the 1950–1996 period (Davis and Heathcote, 2007; Shiller, 2005).

From year-end 2006 through the first quarter of 2009, real house prices have fallen by 34%, and many expect house prices to continue to fall over the next few quarters. Extraordinary events in the financial sector and the macroeconomy as a whole have accompanied this decline of house prices. The fall in house prices triggered a wave of mortgage defaults and home foreclosures, perhaps because some borrowers did not fully understand the terms of their mortgage contract (Bucks and Pence, 2008) or perhaps because a significant portion of homeowners chose to strategically default once their mortgage was sufficiently under water (Haughwout et al., 2008; Foote et al., 2008). The increase in default rates on mortgages lead to a collapse in the price of mortgage-backed securities, which likely contributed to a run on the "shadow" banking system (Gorton, 2009) and sharp devaluation of stock prices. According to data from the *Flow of Funds Accounts of the United States*, the decline in house prices and stock prices reduced household net worth by 20% in nominal terms ($13 trillion) from mid-2007 through year-end 2008. The loss of wealth was associated with a sharp decline in consumer spending via standard "wealth-effect" arguments (Davis and Palumbo, 2001) leading to the contraction of real GDP and the current recession.

With this background in mind, the goal of this paper is to examine time-series fluctuations in house prices and the returns to housing using tools that have proved successful in characterizing the nature of returns in the stock and bond markets. Specifically, we start with the definition of the one-period return to housing. It can be shown that this definition implies that the ratio of housing rents to house prices, the "rent–price ratio," must be equal to the present discounted value of expected future housing service flows and the expected future returns to housing assets. The expected future returns to housing assets can further be split into the sum of expected future risk-free rates of interest and expected

☆ We thank Mike Gibson, Michael Palumbo, David Reifschneider, Tom Tallarini, two anonymous referees and the editor. The views expressed in this paper are those of the authors and should not be attributed to the Board of Governors of the Federal Reserve System or other members of its staff.

* Corresponding author. Address: Wisconsin School of Business, Department of Real Estate and Urban Land Economics, 5261 Grainger Hall, 975 University Avenue, Madison, WI 53706, USA. Fax: +1 608 262 8775.
E-mail address: mdavis@bus.wisc.edu (M.A. Davis).

0094-1190/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jue.2009.06.002

future premia paid to housing over the real risk-free rate. This model is known in the finance literature as the dynamic version of the Gordon growth model (Campbell and Shiller, 1988a,b). The approach is equivalent to assuming that house prices are the discounted sum of housing rents, where the growth rate of housing rents and required return to housing can vary over time. It is precisely this variation over time in expected required returns and expected growth rate of housing rents that yields changes in relative house prices, enabling us to study the factors responsible for time-series changes to housing valuations.

To put the dynamic Gordon growth model to practice, at each point in time we need to measure expectations of the expected present value of risk-free interest rates, housing premia, and rent growth. Our strategy, which is common in the finance literature, is to specify that households form expectations using a VAR with fixed coefficients. We use the VAR to directly compute expected future real risk-free rates and expected housing risk premia and then, given the accounting identity that we document, identify expected future rents as a residual given data on rent–price ratios. This approach accounts for all of the observed variation in rent–price ratios; it also facilitates comparisons of our results to results from other asset markets (Bernanke and Kuttner, 2005; Campbell, 1991; Campbell and Ammer, 1993; Shiller and Beltratti, 1992; Vuolteenaho, 2002). With time-series estimates of the expected real risk-free rate of interest, the expected risk premium to housing, and the expected growth rate of rents in hand, we use variance decompositions to detail how each of these three components contributed to the volatility of rent–price ratios over the 1975–2007 period.

We do not analyze data on the experience of individual housing units. Rather, we perform our analysis on averages for owner-occupied housing in each of 28 housing markets – 23 metropolitan, 4 regional, 1 national – at the semi-annual frequency from 1975 to 2007. As such, our unit of analysis can be thought of as "portfolios" of individual houses. As we show, rent–price ratios were roughly stable in most markets from 1975 to 1996, but declined precipitously after 1997 in almost all of the markets we examine. Shiller (2005) argues that the behavior of house prices since 1997 has no precedent in the twentieth century. With this in mind, we conduct separate variance decompositions of the 1975–1996, or "pre-boom," period, and the 1997–2007, or "boom," period to ensure that our conclusions are not driven exclusively by recent experience.

We have two main findings that are largely robust to time period. First, we find that changes in expected future housing premia are an important source of volatility in rent–price ratios. For example, at the national level, variation in housing premia is the dominant source of variation in rent–price ratios during the 1975–1996 period and an important source of variation during the 1997–2007 period. More generally, we find that time-varying premia are an important feature of housing markets at the national, regional, and metropolitan levels. Second, we find that the covariances between the three components dampen total volatility of rent–price ratios. In particular, we find that expected future premia and rent growth tend to be positively correlated and expected future real risk-free rates and premia tend to be negatively correlated. The latter implies that, historically speaking, house prices have not fully capitalized changes to expected future real risk-free rates.

While many features of housing markets seem to be fundamentally different than those of stock and bond markets – for example, search frictions may play an important role in the liquidity of any given house (Wheaton, 1990) – we find important similarities between returns to housing markets and returns to financial assets that have not been previously recognized. To start, housing returns and returns to financial assets exhibit substantial variation in premia over real rates. In terms of volatility, housing premia contribute to housing valuations in much the same way as stock and bond premia contribute to stock and bond valuations. Further, our finding that expected future rent growth and premia tend to be positively correlated is also consistent with Vuolteenaho's (2002) finding that expected future dividends and premia tend to be positively correlated at the firm-level.

To put our paper in context, we are the first to use the dynamic Gordon growth model to study valuations of owner-occupied real estate across a large number of geographic markets, and the first to document the similarities of valuations in housing markets and those of stock and bond markets. Previous authors (Himmelberg et al., 2005) have used the static version of the Gordon growth model to study rent–price ratios in housing markets. Recently, the dynamic version of the Gordon growth model has been applied to study valuations in commercial real estate (Plazzi et al., 2006) and to examine the linkages of money illusion and house-price inflation in national rent–price ratios (Brunnermeier and Julliard, 2008).

Our results and analysis have some important implications for current analysis and policy. For example, many housing-market analysts have argued that the run-up of house prices from 2002 to 2006 was the result of an unexpectedly low Federal Funds rate (Taylor, 2007) or (related) a sharp decline in mortgage rates in the early 2000s (Himmelberg et al., 2005). Additionally, some have proposed reducing the rate of interest on a 30-year fixed rate mortgage for the purposes of stabilizing the level of house prices (Hubbard and Mayer, 2008). Our finding that the expected net present value of the risk premium for housing and the risk-free rate of interest are negatively correlated implies that the link between the level of house prices and real interest rates is more complex than these interpretations of history suggest. Indeed, our results provide evidence that changes in risk-free interest rates may not have done much to change housing valuations over the 1975–2007 period.

Recently, the Federal Open Market Committee (FOMC) announced in a press release dated March 18, 2009 that it will purchase up to \$1.25 trillion in mortgage backed securities in 2009 to "Provide greater support to mortgage lending and housing markets".[1] While this policy will likely improve the availability of credit to home buyers, our results suggest that the effect of this policy on the level of house prices is less clear.

The rest of this paper proceeds as follows. Section 2 describes our implementation of the dynamic Gordon growth model and Section 3 discusses the data. In Section 4, we outline the VAR model and report estimation results. Section 5 details the results of all of our variance decompositions. In Section 6 we conclude and discuss directions for future research.

2. The Campbell–Shiller decomposition

Consider the one-period gross real return to housing

$$\frac{P_{t+1} + R_{t+1}}{P_t}, \tag{1}$$

where P is the real price of housing and R is housing rents. We can use the method of Campbell and Shiller (1988a,b) to rewrite this gross return using a log-linear approximation that sets the log of the rent–price ratio at date $t, \log(R_t/P_t) \equiv r_t - p_t$, equal to the expected net present value of all future (date $t+1+j$ for $j = 0, \ldots, \infty$) real rates of return to housing and real growth in housing rents,

[1] See the press release dated March 18, 2009 at the Federal Reserve Board web site, available at http://www.federalreserve.gov/newsevents/press/monetary/20090318a.htm.

$$r_t - p_t = k + E_t\left[\sum_{j=0}^{\infty}\rho^j\varphi_{t+1+j} - \sum_{j=0}^{\infty}\rho^j\Delta r_{t+1+j}\right]$$
$$\rho = (1 + e^{\overline{r-p}})^{-1} \qquad (2)$$
$$k = (1-\rho)^{-1}[\ln(\rho) + (1-\rho)\ln(1/\rho - 1)],$$

where φ is the log of the gross real return to housing, r is the log of real housing rents, ρ is a discount factor related to the average of the rent–price ratio (written as $e^{\overline{r-p}}$), and k is a constant of linearization. If we define the return to housing, φ, as the sum of a real risk-free interest rate, i, and the per-period premium over that rate, $\pi = \varphi - i$, then the log rent–price ratio can be rewritten as the sum of three components related to the expected present value of future real risk-free interest rates (hereafter called real interest rates), housing premia, and rent growth,[2]

$$r_t - p_t = k + E_t\sum_{j=0}^{\infty}\rho^j i_{t+1+j} + E_t\sum_{j=0}^{\infty}\rho^j\pi_{t+1+j} - E_t\sum_{j=0}^{\infty}\rho^j\Delta r_{t+1+j}, \qquad (3)$$

or

$$r_t - p_t = k + \mathscr{I}_t + \Pi_t - \mathscr{G}_t, \qquad (4)$$

where $\mathscr{I}_t, \Pi_t$, and $\mathscr{G}_t$ represents households' time-t expectation of the present value of real interest rates, premia and rent growth:

$$\mathscr{I}_t = E_t\sum_{j=0}^{\infty}\rho^j i_{t+1+j}$$
$$\Pi_t = E_t\sum_{j=0}^{\infty}\rho^j\pi_{t+1+j} \qquad (5)$$
$$\mathscr{G}_t = E_t\sum_{j=0}^{\infty}\rho^j\Delta r_{t+1+j}.$$

The representation of the rent–price ratio in Eq. (4) is the dynamic version of the classic Gordon growth model that was first developed by Campbell and Shiller (1988a,b). Campbell and Shiller used this framework to analyze the determinants of variability of the dividend–price ratio in the aggregate stock market.[3] Since then, it has been applied to fixed-income markets by Shiller and Beltratti (1992) and Campbell and Ammer (1993), to firm-level stock returns by Vuolteenaho (2002), and to commercial real estate by Plazzi et al. (2006).

In the field of real-estate finance, a well-studied version to Eq. (4) expresses the level of the rent–price ratio as

$$R_t/P_t = i_t + \pi - g_{t+1}, \qquad (6)$$

where i_t is some current real interest rate, π is an assumed constant housing premium over this rate and g_{t+1} is the expected capital gain or loss to housing over some future horizon. Versions of this expression can be found in Cutts et al. (2005), Gallin (2008), Himmelberg et al. (2005), Verbrugge (2008), and elsewhere. There are three important differences between Eqs. (4) and (6). First, Eq. (4) allows for a time-varying housing premium. Second, although both equations recognize that prices are forward-looking, Eq. (4) explicitly accounts for the dynamics of each component of the rent–price ratio, while Eq. (6) combines all future considerations into expected future capital gains. Third, as long as real interest rates and housing premia are stationary, Eq. (4) ties appreciation of house prices to growth in rents over the long-run, whereas Eq. (6) does not.

2.1. Implementing the dynamic Gordon growth model

Our strategy to implement Eq. (4) for the case of owner-occupied housing is to use publicly available data on house prices, rents, interest rates, and various macroeconomic variables to construct time-series estimates of $\mathscr{I}_t, \Pi_t$, and $\mathscr{G}_t$. Following Campbell (1991) and Campbell and Ammer (1993), we use a vector autoregressive (VAR) approach to construct our estimates of these expectations variables. Define

$$Z_t = (i_t, \pi_t, \Delta r_t, x_t')', \qquad (7)$$

where x_t is a column vector that contains other variables (discussed in more detail later in the paper) that are useful for predicting i_t, π_t, and Δr_t. For reasons discussed later, we specify that households assume that Z_t follows a first-order VAR, i.e.,

$$Z_t = AZ_{t-1} + \varepsilon_t. \qquad (8)$$

Given an estimate of A, denoted $\widehat{A}$, our estimates of the present values $\mathscr{I}_t, \Pi_t$, and $\mathscr{G}_t$ are the first three elements of

$$\widehat{A}(I - \rho\widehat{A})^{-1}Z_t, \qquad (9)$$

where I denotes the identity matrix. We denote the VAR-based estimates of these three present values as $\widehat{\mathscr{I}}_t, \widehat{\Pi}_t$, and $\widehat{\mathscr{G}}_t$, respectively.

The dynamic accounting identity in Eq. (4) implies that $\mathscr{I}_t, \Pi_t$, and $\mathscr{G}_t$ fully describe the rent–price ratio up to a constant of linearization. However, if households do not literally form forecasts according to the process specified by Eq. (8), our estimate of expectations terms given by Eq. (9) will not match actual expectations maintained by households. In this case, our predicted rent–price ratio, defined by

$$\widehat{r_t - p_t} = k + \widehat{\mathscr{I}}_t + \widehat{\Pi}_t - \widehat{\mathscr{G}}_t, \qquad (10)$$

will differ from the actual rent–price ratio at each point in time, which implicitly defines a discrepancy, e_t,

$$\widehat{r_t - p_t} = r_t - p_t + e_t. \qquad (11)$$

In the rest of the paper, we call e_t the "forecast discrepancy," in the sense that our VAR model does not yield expectations that are perfectly aligned with observed movements in the rent–price ratio in accordance with Eq. (4).

Following the literature, we treat the present value of future rent growth as a residual. That is, we relabel the quantity $\widehat{\mathscr{G}}_t + e_t$ as $\mathscr{E}_t$ so that Eq. (11) has the form

$$r_t - p_t = k + \widehat{\mathscr{I}}_t + \widehat{\Pi}_t - \mathscr{E}_t, \qquad (12)$$

Treating the forecast discrepancy this way yields the following decomposition of the time-series variance of the rent–price ratio in any given geographic area:

$$var(r_t - p_t) = var(\widehat{\mathscr{I}}_t) + var(\widehat{\Pi}_t) + var(\mathscr{E}_t) + 2cov(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t)$$
$$- 2cov(\widehat{\mathscr{I}}_t, \mathscr{E}_t) - 2cov(\widehat{\Pi}_t, \mathscr{E}_t) \qquad (13)$$

All elements of this decomposition related to $\mathscr{E}_t$ are described in the paper in terms of the contribution of "rent growth".[4]

We also use variance decompositions to document how local factors that are unrelated to aggregate movements and trends affect variation in rent–price ratios. Define a variable with an asterisk to be the difference between a variable and its national-level value. With this notation, the "relative" rent–price ratio is defined as

[2] In the rest of the paper, we drop the distinction between the rent–price ratio and the log rent–price ratio when the context is clear.

[3] In the classic Gordon growth model, growth rates and rates of return are constant, such that $R/P = i + \pi - \Delta r$.

[4] Of course, we could have used the VAR to construct estimates for G and combined the forecast discrepancy with $\widehat{\mathscr{I}}_t$ or $\widehat{\Pi}_t$, or even leave the discrepancy as its own component. We choose to combine e_t with $\widehat{\mathscr{G}}_t$ so our results are directly comparable with the existing finance literature.

$$r_t^* - p_t^* = r_t - p_t - (r_t^{US} - p_t^{US}) = \widehat{\Pi}_t^* - \mathscr{E}_t^*$$
$$\widehat{\Pi}_t^* = \widehat{\Pi}_t - \widehat{\Pi}_t^{US} \tag{14}$$
$$\mathscr{E}_t^* = \mathscr{E}_t - \mathscr{E}_t^{US}.$$

Note that we have imposed that the expected present value of real interest rates in any geographic market, $\widehat{\mathscr{I}}_t$, is the same as in the aggregate, $\widehat{\mathscr{I}}_t^{US}$.[5] The variance decomposition on the relative rent–price ratio is simply

$$var(r_t^* - p_t^*) = var\left(\widehat{\Pi}_t^*\right) + var(\mathscr{E}_t^*) - 2cov\left(\widehat{\Pi}_t^*, \mathscr{E}_t^*\right). \tag{15}$$

3. Data

3.1. House prices and rents

The complete set of data we use in our study is available for download at http://www.morris.marginalq.com/whatmoves.html.

For prices, we measure changes to the price of owner-occupied housing using the repeat-transactions house price index published by the Federal Housing Finance Agency (FHFA). For rents, like Case and Shiller (1989) and other studies, we assume that the growth rate of rents paid by renters is identical to the growth rate of rents accruing to owner–occupiers. Any trends in the implicit rental price of owner-occupied housing that are not captured by trends in the explicit rental price of tenant-occupied housing will affect our analysis and conclusions, and our results should be interpreted with this caveat in mind. Data on the growth rate of rents paid by renters are reported by the Bureau of Labor Statistics (BLS) for 23 separate metropolitan markets, the four Census regions, and the nation.[6] Note that the region-level rent indexes are not simple averages of the component metropolitan markets in our analysis: the BLS collects rent data for over 80 metro markets and uses those data for regional and national estimates, but only publishes metro-area indexes for selected markets. Because the BLS data are reported at a semi-annual frequency for several markets, we conduct all our analysis at this frequency. Where needed, we average quarterly readings to convert to a semi-annual frequency.

For each market, we convert the nominal growth rate of rents and prices to real growth rates by deflating each local-area BLS and FHFA index using the national CPI excluding shelter. We then use micro data from the 2000 Decennial Census of Housing (DCH) to benchmark the level of the rent–price ratio in 2000, employing a procedure described by Davis et al. (2008). This benchmarking is required to obtain values for ρ and k in Eq. (2) for each MSA. Summarizing our benchmarking procedure: for each MSA we regress gross rents paid by renters on a set of housing characteristics; we use the resulting regression coefficients to predict the rental value of all owned units in the MSA; and, we set the rent–price ratio equal to the average imputed rent of owned units divided by the average value of these units. Note that the results we report are not sensitive to small changes in the benchmarked levels of rent–price ratios.

Given our use of metro-area rental and house price indexes, we wish to make clear that we are not studying the returns to owning any given house in a given geographic area. Rather, given the data we use in this study, we are studying the rent–price ratio, rents, prices and returns for a portfolio of homes in each geographic area. The returns to owning any given house in one of the geographic areas we study is likely to be much more volatile than we report: see, for example, Case and Shiller (1989), Quigley and Van Order (1995), Deng et al. (2000) and Flavin and Yamashita (2002).

Column 1 of Table 1 lists our sample's beginning date, which varies by market according to data availability; the end date of the sample is 2007:H2 for all markets. Columns 2 through 4 display the value of the rent–price ratio, R_t/P_t, at the beginning of the sample, in 1996:H2, and at the end of the sample for each of the 28 markets we study. The rent–price ratio is graphed in Fig. 1 for the nation and the four Census regions.

Table 1
Summary data on rent–price ratios, 1975–2007.

	Date of first observation (1)	Rent–price ratio, Ann. Pct.		
		Initial (2)	1996:H2 (3)	2007:H2 (4)
USA	1975:H1	5.53	4.97	3.75
Midwest	1978:H1	4.95	5.06	3.96
Chicago	1975:H2	5.68	4.63	3.43
Cincinnati	1976:H1	4.99	4.76	4.12
Cleveland	1975:H2	6.03	4.97	4.78
Detroit	1976:H2	6.66	4.98	4.45
Kansas City	1976:H1	5.78	5.85	4.84
Milwaukee	1977:H1	5.01	5.04	3.48
Minneapolis	1976:H1	6.86	6.21	3.92
St. Louis	1975:H2	6.58	5.67	4.11
Northeast	1978:H1	5.95	4.76	3.18
Boston	1978:H1	6.74	4.15	2.68
New York	1976:H1	6.51	3.98	2.35
Philadelphia	1976:H1	5.91	5.39	3.53
Pittsburgh	1976:H2	5.68	5.48	4.56
South	1978:H1	5.21	5.55	4.16
Atlanta	1976:H2	5.81	5.92	4.35
Dallas	1976:H1	5.69	6.45	5.24
Houston	1976:H1	5.69	7.33	5.60
Miami	1978:H1	6.64	5.70	2.78
West	1978:H1	4.34	4.32	2.69
Denver	1976:H2	8.22	6.32	4.74
Honolulu	1977:H1	4.59	3.19	1.98
Los Angeles	1975:H1	6.06	4.23	2.13
Portland	1976:H1	7.17	4.93	2.94
San Diego	1976:H1	5.94	4.75	2.85
San Francisco	1975:H2	6.70	4.20	2.27
Seattle	1975:H2	8.77	5.16	2.94
Metro median		6.03	5.04	3.53

Column (1) lists the date of the initial observation in each housing market. Columns (2)–(4) list the annualized rent–price ratio, reported as a percentage, in each housing market at the initial observation date, 1996:H2, and 2007:H2. The median rent–price ratio across the metropolitan markets is reported in the final row.

3.2. Real interest rates

For the benchmark real interest rate i, we use an estimate of the ex-ante real expected yield on a 10-year US Treasury bond. We use this particular measure because, on average, households move every 9 years and mortgages are usually prepaid after 10 years, and further, the use of this interest rate facilitates comparisons of our results with other recent studies of the returns to housing: for example, see Cutts et al. (2005), Gallin (2008), Himmelberg et al. (2005) and Meese and Wallace (1994). We define the real rate as the nominal 10-year Treasury yield less the median reading of 10-year inflation expectations from professional forecasters. This inflation reading is taken from the Blue Chip Economic Indicator forecast for the 1975:H1 to 1991:H1 period and from the Livingston survey from 1991:H2 to the end of the sample. Our estimates

[5] Although the real interest rate is identical across markets, its expected present value depends on the discount factor ρ, which varies by location. However, the cross-sectional variation in ρ is small enough that the error caused by this approximation is unimportant.

[6] Across our 23 metropolitan areas, the median of the correlation of the real growth rates of the BLS tenant rent index, the rent index used in this study, and the BLS owner equivalent rent index is 0.90.

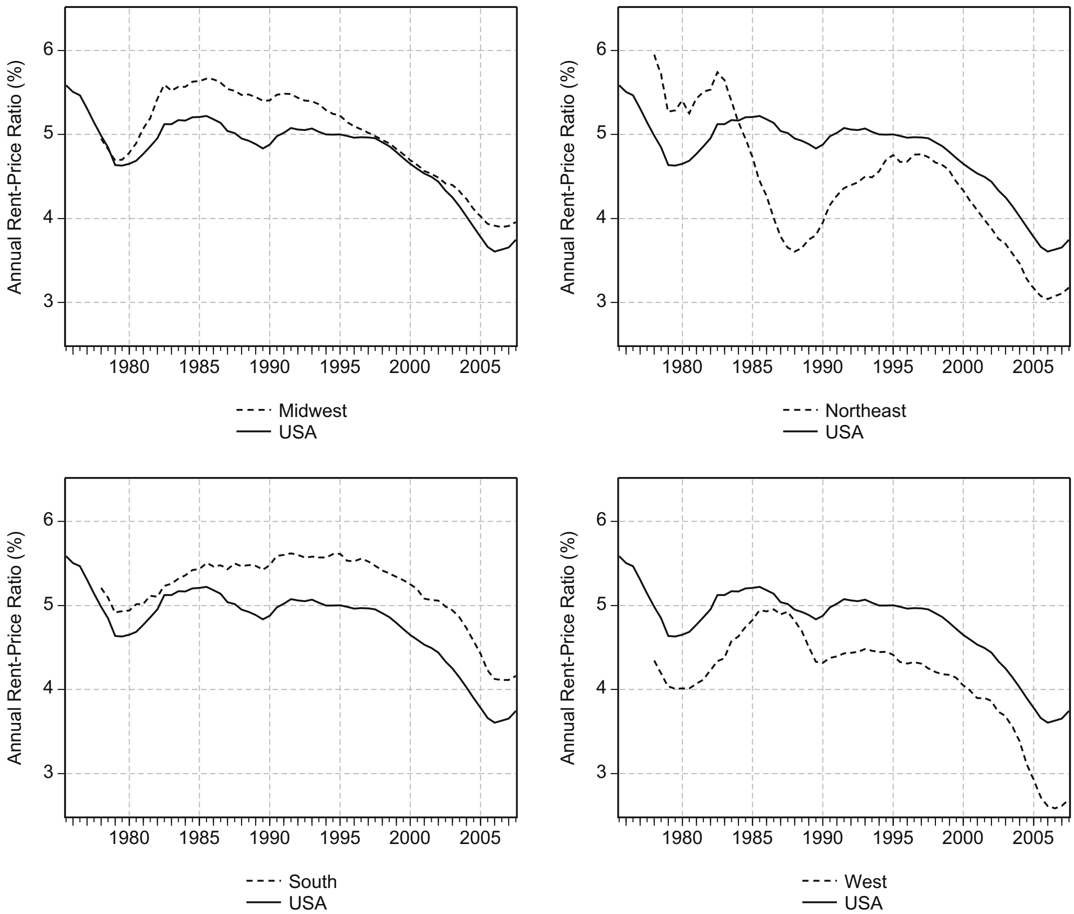

Fig. 1. The rent–price ratio for the Midwest, Northeast, South, and West Census regions are plotted as dashed lines and the rent–price ratio for the USA is plotted as a solid line. The house price data are from FHFA and the rent data are from the BLS. To fix the level of the rent–price ratio in 2000 in each series, the price and rent data are benchmarked to micro data from the 2000 Decennial Census of Housing (DCH). See the paper for details.

of housing premia are therefore relative to the real yield on a 10-year Treasury.

3.3. Housing returns and premia

For each of the 28 markets in our sample, the real return to owner-occupied housing (for a portfolio of homes in each market, as mentioned earlier) is computed as

$$\varphi_t = \frac{R_t + P_t - P_{t-1}}{P_{t-1}}. \tag{16}$$

The excess return or premium to housing paid above the real 10-year Treasury is computed as

$$\pi_t = \varphi_t - i_t. \tag{17}$$

In these equations, R_t is the real (inflation-adjusted) rent accruing to homeowners between $t-1$ and t, P_t is the real price of housing at time t, and i_t is the real yield on the 10-year Treasury bond as of time t. Note that we do not adjust returns for depreciation or for taxes.[7]

In Table 2 we list the sample mean ("Avg."), standard deviation ("SD"), and autocorrelation coefficient (ρ) for the annualized real growth rate of rents (columns 1–3), annualized real housing returns (4–6), and annualized excess returns (7–9) for our 28 housing markets over the entire 1975–2007 sample. The autocorrelation coefficient is computed from a regression of each semi-annual variable on its first lag. Shown in columns 4 and 7, real housing returns in the aggregate averaged 6.5% per year and excess returns averaged about 3% over our full sample. Across our metropolitan areas, average real housing returns range from 5.4% (Cincinnati) to almost 10% (Seattle).

Shown in columns 5 and 8, the standard deviations of real and excess housing returns vary from about 2.5% to 7.0% per year,[8] depending on the market and whether the total or excess return is considered. Excess returns tend to be more variable than total returns everywhere, and the returns of larger and more geographically diverse markets such as the nation and the four Census regions tend to be less volatile than returns of individual metro markets. Real housing returns are typically 1-1/2 to 3 times more volatile than rent growth, computed as column 5 divided by column 2. This ratio is less than half that of US equities; total returns for the S&P 500 were roughly 5-1/2 times more volatile than dividend growth over the 1975–2007 period. In general, the first-order autocorrelation of real rent growth (column 3) and housing returns (columns 6 and 9) are of the same order of magnitude, on average about 0.6. The persistence of real and excess returns is high relative to the returns on other assets such as stocks and bonds, but is in line with the original findings of Case and Shiller (1989) who document similar patterns in the returns to housing in four major metro markets (Atlanta, Chicago, Dallas, and San Francisco) over the 1970–1986 period.

---

[7] Allowing for depreciation expenses would shift the average level of returns and thus the discount factor in each geographic area, but would not affect any of our main results. We do not account for taxes at the household level (that is, we study pre-tax returns) so our results are comparable to those in the finance literature.

[8] We exclude Honolulu, an obvious outlier, from this calculation.

Table 2
Summary data on the annualized real growth rate of rents Δr_t, the real annualized return to housing φ_t, and the excess return to housing π_t – the average, standard deviation, and autocorrelation coefficient (ρ), 1975–2007.

	Δr_t			φ_t			π_t		
	Avg. (1)	SD (2)	ρ (3)	Avg. (4)	SD (5)	ρ (6)	Avg. (7)	SD (8)	ρ (9)
USA	0.42	1.33	0.58	6.47	2.48	0.65	2.99	3.13	0.74
Midwest	−0.08	1.32	0.63	5.73	2.55	0.71	2.08	3.18	0.75
Chicago	0.49	1.64	0.48	6.82	3.72	0.64	3.32	4.20	0.69
Cincinnati	−0.06	1.53	0.34	5.36	2.60	0.70	1.84	3.22	0.76
Cleveland	−0.14	1.72	0.42	5.88	3.38	0.63	2.37	3.93	0.67
Detroit	−0.12	1.68	0.47	6.81	5.07	0.58	3.26	5.44	0.63
Kansas City	-0.02	1.89	0.58	6.23	3.14	0.37	2.71	3.74	0.54
Milwaukee	−0.16	1.79	0.33	6.09	3.32	0.59	2.50	4.00	0.66
Minneapolis	0.11	1.87	0.45	7.95	3.57	0.46	4.43	4.12	0.58
St. Louis	−0.19	1.46	0.51	6.96	4.04	0.10	3.45	4.51	0.24
Northeast	0.80	1.68	0.69	7.50	5.12	0.75	3.85	5.24	0.73
Boston	1.03	2.37	0.71	8.39	6.46	0.79	4.75	6.38	0.76
New York	0.80	1.72	0.68	8.15	5.65	0.74	4.63	5.72	0.73
Philadelphia	0.53	1.85	0.62	7.48	4.20	0.70	3.95	4.48	0.70
Pittsburgh	−0.17	1.86	0.35	6.12	3.16	0.64	2.57	3.69	0.68
South	0.15	1.31	0.55	6.16	2.15	0.53	2.52	2.90	0.69
Atlanta	0.22	2.37	0.63	6.62	2.34	0.48	3.07	2.66	0.58
Dallas	0.03	2.21	0.67	6.08	3.99	0.25	2.55	4.42	0.36
Houston	−0.37	2.80	0.59	6.00	4.10	0.43	2.48	4.68	0.56
Miami	0.50	2.13	0.27	9.04	5.66	0.18	5.39	6.30	0.35
West	0.91	1.52	0.48	6.77	3.26	0.65	3.13	3.94	0.72
Denver	0.26	2.52	0.63	8.45	4.25	0.75	4.90	4.70	0.78
Honolulu	0.69	2.37	0.41	7.66	13.87	−0.32	4.08	13.89	−0.30
Los Angeles	1.32	1.88	0.60	8.66	6.50	0.82	5.17	7.00	0.83
Portland	0.12	1.63	0.57	8.70	5.09	0.40	5.18	5.71	0.50
San Diego	1.36	2.45	0.71	8.34	6.36	0.77	4.82	6.81	0.79
San Francisco	1.37	2.48	0.58	9.05	6.01	0.81	5.55	6.38	0.81
Seattle	0.78	2.03	0.60	9.95	5.46	0.58	6.45	5.96	0.64
Metro median	0.22	1.88	0.58	7.48	4.20	0.59	3.95	4.68	0.66

We report the sample average (Avg.), standard deviation (SD), and first-order autocorrelation coefficient (ρ) of the annualized real growth rate of rents, Δr_t, in columns (1)–(3), the return to housing, φ_t, in columns (4)–(6) and the housing premium, $\pi_t = \varphi_t - i_t$, in columns (7)–(9) for each market over the 1975:H1-2007:H2 period. ρ is computed from a regression of the reported variable on its first lag. The median sample statistic across the metropolitan markets is reported in the final row.

3.4. Macroeconomic conditions

Columns 1-9 of Table 3 list means, standard deviations, and first-order autocorrelation coefficients of real per-capita income growth (ΔY_t), employment growth (ΔL_t), and population growth (ΔN_t). We include these variables in the vector x_t of Eq. (7) to help forecast real rates, housing premia, and rent growth. The Bureau of Economic Analysis (BEA) publishes these data at an annual frequency. We convert each of these variables to a semi-annual frequency by assuming each variable is constant throughout the calendar year. The reported autocorrelation coefficient is computed from a regression of the generated semi-annual variable on its second lag. Comparing Tables 2 and 3, the standard deviation of income and employment growth (columns 2 and 5 of Table 3) are similar to that of rent growth (column 2 of Table 2), and smaller than that of housing premia (column 5 of Table 2). Perhaps not surprisingly, population growth (column 8 of Table 3), is considerably less variable than either income growth or employment growth. The autocorrelation coefficients of these macroeconomic variables (columns 3, 6, and 9 of Table 3) are rather uniform across our 28 markets.

4. Specification and estimates of the VAR

A separate first-order VAR corresponding to Eq. (8) is estimated for each of the 28 housing markets in our sample. Parsimony motivates our choice of a first-order specification since, with semi-annual data spanning 33 years, we have at most 66 observations with which to estimate coefficients.[9] For each market, the forecasting equations for $i_t, \Delta r_t$, and π_t are

$$i_t = \delta_0 + \delta_{\Delta r}\Delta r^{US}_{t-1} + \delta_\pi \pi^{US}_{t-1} + \delta_i i_{t-1} + \delta_{\Delta Y}\Delta Y^{US}_{t-2} + \delta_{\Delta L}\Delta L^{US}_{t-2} + \delta_{\Delta N}\Delta N^{US}_{t-2} + \varepsilon^i_t \quad (18)$$

$$\Delta r_t = \gamma_0 + \gamma_{\Delta r}\Delta r_{t-1} + \gamma_\pi \pi_{t-1} + \gamma_i i_{t-1} + \gamma_{\Delta Y}\Delta Y_{t-2} + \gamma_{\Delta L}\Delta L_{t-2} + \gamma_{\Delta N}\Delta N_{t-2} + \varepsilon^{\Delta r}_t \quad (19)$$

$$\pi_t = \beta_0 + \beta_{\Delta r}\Delta r_{t-1} + \beta_\pi \pi_{t-1} + \beta_i i_{t-1} + \beta_{\Delta Y} Y_{t-2} + \beta_{\Delta L}\Delta L_{t-2} + \beta_{\Delta N}\Delta N_{t-2} + \varepsilon^\pi_t, \quad (20)$$

where the *US* superscript refers to a national-level variable and local variables lack a superscript.

Shown in Eq. (18), our specification assumes that the real interest rate depends on only national-level variables.[10] In contrast, all variables used in forecasting rent growth and premia are measured at the local market level, except the real interest rate.[11]

[9] The AIC and BIC lag length criteria (not shown) also select a first-order model in many of our markets.

[10] In all of the forecasts we use the second lag, rather than the first lag, of the macroeconomic condition variables: $\Delta Y_{t-2}, \Delta L_{t-2}$, and ΔN_{t-2}. As noted earlier, these data are observed at the annual frequency and are assumed to be constant throughout the year. Thus, once these variables are converted to the semi-annual frequency, the second lag is used in the regression to ensure that macroeconomic conditions data from year $y-1$ are always used to forecast variables in year y.

[11] Of course, we could have specified that regional-level or nearby metropolitan-level variables also forecast local premia and rent growth, but allowing for this richness in the model would result in a large reduction in degrees of freedom relative to the sample size.

Table 3
Summary data on real per-capita income growth ΔY_t, employment growth ΔL_t, and population growth ΔN_t – the average, standard deviation, and autocorrelation coefficient (ρ), 1975–2007.

	ΔY_t			ΔL_t			ΔN_t		
	Avg. (1)	SD (2)	ρ (3)	Avg. (4)	SD (5)	ρ (6)	Avg. (7)	SD (8)	ρ (9)
USA	1.77	1.70	0.54	1.82	1.28	0.62	1.05	0.20	0.93
Midwest	1.39	1.92	0.45	1.21	1.43	0.71	0.43	0.43	0.94
Chicago	1.70	2.02	0.39	0.97	1.72	0.61	0.54	0.60	0.94
Cincinnati	1.79	1.86	0.54	1.63	1.90	0.78	0.72	0.38	0.80
Cleveland	1.55	2.00	0.52	0.52	1.87	0.69	−0.12	0.52	0.91
Detroit	1.34	2.96	0.67	0.59	2.94	0.73	0.13	0.77	0.87
Kansas City	1.64	1.94	0.63	1.61	1.85	0.73	0.95	0.44	0.71
Milwaukee	1.56	1.82	0.56	0.96	2.07	0.71	0.34	0.53	0.85
Minneapolis	2.01	2.07	0.56	2.09	1.95	0.72	1.31	0.51	0.82
St. Louis	1.77	1.77	0.43	1.14	1.64	0.59	0.38	0.36	0.83
Northeast	2.03	1.93	0.66	1.20	1.33	0.74	0.41	0.30	0.88
Boston	2.36	2.44	0.71	1.10	2.29	0.80	0.53	0.45	0.84
New York	2.21	2.21	0.63	0.85	1.59	0.77	0.40	0.59	0.93
Philadelphia	2.00	1.65	0.54	0.98	1.42	0.71	0.34	0.39	0.86
Pittsburgh	1.82	1.50	0.60	0.43	1.65	0.74	−0.43	0.52	0.89
South	1.82	1.51	0.40	2.24	1.13	0.64	1.55	0.42	0.90
Atlanta	1.72	2.48	0.73	3.08	2.34	0.80	2.81	0.70	0.82
Dallas	1.84	2.29	0.61	2.93	2.34	0.78	2.57	0.91	0.79
Houston	1.96	2.82	0.54	2.49	3.07	0.73	2.37	1.97	0.85
Miami	1.72	1.98	0.47	2.58	2.00	0.69	2.03	1.25	0.93
West	1.44	1.87	0.61	2.25	1.53	0.74	1.73	0.55	0.93
Denver	1.91	2.09	0.63	2.48	2.23	0.79	2.03	1.28	0.89
Honolulu	1.45	2.05	0.66	1.12	1.66	0.79	0.72	0.90	0.74
Los Angeles	1.36	2.17	0.64	1.80	2.48	0.76	1.60	0.98	0.94
Portland	1.47	2.02	0.61	2.48	2.56	0.79	1.88	1.12	0.87
San Diego	2.00	2.31	0.62	2.65	2.40	0.87	1.91	1.67	0.86
San Francisco	2.26	3.24	0.52	1.65	2.48	0.81	1.12	0.80	0.89
Seattle	2.09	2.35	0.45	2.79	2.41	0.82	1.91	1.18	0.85
Metro median	1.79	2.07	0.60	1.63	2.07	0.76	0.95	0.70	0.86

We report the the sample average (Avg.), standard deviation (SD), and autocorrelation coefficient (ρ) of the annualized growth rate of real income, ΔY_t, in columns (1)–(3), employment growth, ΔL_t, in columns (4)–(6) and population growth, ΔN_t, in columns (7)–(9) for each market over the 1975–2007 period. ρ is computed from a regression of the reported variable on its second lag. The median sample statistic across the metropolitan markets is reported in the final row.

The VAR is closed by specifying a law of motion for national level rent growth and premia as well as local and national income growth, employment growth, and population growth, $(\Delta r_t^{US}, \pi_t^{US}, \Delta Y_t, \Delta L_t, \Delta N_t, \Delta Y_t^{US}, \Delta L_t^{US}, \Delta N_t^{US})$. We specify that each of these variables depends on one lag of the real rate, rent growth, premia, income growth, employment growth and population growth; also, local macroeconomic variables are specified to depend only on lags of local variables and national macroeconomic variables are specified to depend only on lags of national variables.

We estimate each VAR as a SUR system and the estimation results are summarized in the three panels of Table 4. Top to bottom, the panels summarize the results of the forecasting equations for the real interest rate, Eq. (18), real rent growth, Eq. (19), and housing premia, Eq. (20). Rather than report the point estimates for every equation and market, in columns 1–6 we report the 25th percentile, median, and 75th percentile of each point estimate across markets. Column 7 lists the 25th percentile, median, and 75th percentile of the p-values from Wald tests that none of the included variables forecast the corresponding dependent variable. Column 8 summarizes the distribution of p-values from Wald tests that none of the included macroeconomic variables forecast the corresponding dependent variables, and columns 9–11 summarize the distribution of adjusted R^2s (hereafter $\bar{R}^2$) over the full sample, the pre-boom period (1975–1996), and the recent boom period (1997–2007). All $\bar{R}^2$s are computed using parameter estimates from the full sample, implying that sub-sample $\bar{R}^2$s need not be positive. In the bottom row of each panel we report the fraction of times, across the 28 markets, that the parameter estimate or model statistic reported in columns 1–8 is significant at the 5% level. In the case of the coefficient estimates, the significance level is computed using a Wald test that the estimate is zero.

We now turn to the results of the real rate forecasting equation (top panel). The specification of the real rate forecast is identical in each of the 28 markets. Accordingly, differences in estimated parameters and model statistics result either from differences in the available sample period in each market or differences in the manner in which the other included equations of the VAR influence the SUR estimator.[12] As expected, the distribution of parameter estimates and model statistics reported in columns 1–8 is narrow.

The Wald test of no predictability and the $\bar{R}^2$ both strongly indicate that real rates are predictable. In all cases, we estimate a significant degree of persistence in real rates, and the median coefficient estimate on the lagged real rate is 0.99. This evidence is consistent with Rose (1988) and Rapach and Wohar (2004) who also report that real interest rates are quite persistent. Because the distribution of coefficients on other included variables is tightly clustered around zero, our estimated model for forecasting the real rate of interest closely resembles a simple AR(1) specification with a large autoregressive coefficient.

Shown in the middle panel of Table 4, rental growth also appears quite predictable in each of the 28 housing markets in our sample. The Wald test of no predictability (column 7) is rejected at or below the 5% level in almost every market, and the interquartile range of the full sample $\bar{R}^2$ runs between 25% and 45% .

[12] The independent variables appearing in Eqs. (18)–(20) are not identical and thus the SUR estimator is not equivalent to equation-by-equation OLS.

Table 4
Summary of VAR estimation results.

Dependent variable		Coefficient estimates on						p-value, Wald tests		$\overline{R}^2$		
		Δr_{t-1} (1)	π_{t-1} (2)	i_{t-1} (3)	ΔY_{t-1} (4)	ΔL_{t-1} (5)	ΔN_{t-1} (6)	All vars (7)	Macro only (8)	75–07 (9)	75–96 (10)	97–07 (11)
i_t	25th Percentile	−0.13	0.03	0.98	−0.05	−0.02	−0.09	0.00	0.26	0.81	0.75	0.45
	Median	−0.11	0.03	0.99	−0.05	0.00	0.01	0.00	0.38	0.82	0.77	0.54
	75th Percentile	−0.10	0.05	1.02	−0.03	0.02	0.09	0.00	0.54	0.82	0.78	0.56
	Fraction of 28 markets significant at 5%	0.39	0.04	1.00	0.04	0.00	0.00	1.00	0.00	NA	NA	NA
Δr_t	25th Percentile	0.22	0.02	0.30	0.04	−0.09	0.03	0.00	0.02	0.25	0.20	−0.16
	Median	0.28	0.05	0.36	0.10	−0.06	0.29	0.00	0.13	0.39	0.37	−0.04
	75th Percentile	0.41	0.08	0.55	0.16	0.02	0.52	0.00	0.43	0.45	0.50	0.28
	Fraction of 28 markets significant at 5%	0.75	0.14	0.61	0.25	0.11	0.21	0.96	0.39	NA	NA	NA
π_t	25th Percentile	0.00	0.13	−1.47	0.11	−0.17	−0.82	0.00	0.07	0.47	0.35	−0.31
	Median	0.29	0.46	−0.78	0.26	−0.03	0.19	0.00	0.25	0.52	0.46	0.10
	75th Percentile	0.50	0.54	−0.30	0.41	0.13	0.58	0.00	0.37	0.59	0.54	0.34
	Fraction of 28 markets significant at 5%	0.29	0.82	0.43	0.36	0.04	0.07	1.00	0.25	NA	NA	NA

In this table, we summarize the results of the real rate, i_t, real rent growth, Δr_t, and premia, π_t, forecasting models. Columns (1)–(6) summarize the estimated coefficients from the VAR. Columns (7)–(8) summarize the results of the Wald test that none of the included regressors forecast the dependent variable (column 7), and the Wald test that none of the included macroeconomic variables, $(\Delta Y_t, \Delta L_t, \Delta N_t)$, forecast the dependent variable. Columns (9)–(11) summarize the fit of the forecasting model by reporting the in-sample R^2 over the full sample, column (9), the 1975–1996 sample, column (10), and the 1997–2007 sample, column (11).
The results of the real rate forecasting model, $i_t = \delta_0 + \delta_{\Delta r}\Delta r^{US}_{t-1} + \delta_\pi \pi^{US}_{t-1} + \delta_i i_{t-1} + \delta_{\Delta Y}\Delta Y^{US}_{t-2} + \delta_{\Delta L}\Delta L^{US}_{t-2} + \delta_{\Delta N}\Delta N^{US}_{t-2} + \varepsilon^i_t$, are summarized in the top panel. The results of the real rent growth forecasting model, $\Delta r_t = \gamma_0 + \gamma_{\Delta r}\Delta r_{t-1} + \gamma_\pi \pi_{t-1} + \gamma_i i_{t-1} + \gamma_{\Delta Y}\Delta Y_{t-2} + \gamma_{\Delta L}\Delta L_{t-2} + \gamma_{\Delta N}\Delta N_{t-2} + \varepsilon^{\Delta r}_t$, are summarized in the middle panel. The results of the housing premia forecasting model, $\pi_t = \beta_0 + \beta_{\Delta r}\Delta r_{t-1} + \beta_\pi \pi_{t-1} + \beta_i i_{t-1} + \beta_{\Delta Y} Y_{t-2} + \beta_{\Delta L}\Delta L_{t-2} + \beta_{\Delta N}\Delta N_{t-2} + \varepsilon^\pi_t$, are summarized in the bottom panel.
The first three rows of each panel display the 25th,50th and 75th percentile of the empirical distribution of the VAR coefficients and test statistics across the 28 housing markets in our sample. The fourth row of each panel reports the fraction of markets in which the associated coefficient estimate or test statistic is significant at or below the 5% level.

As the sub-sample $\overline{R}^2$ measures indicate, the degree of in-sample fit is different in the pre- and post-boom periods. Shown by the row marked "Fraction of Metro Areas Significant at 5%," lagged real rates and rent growth significantly predict future rent growth. The local macro variables significantly forecast rent growth in about 40% of markets (column 8), suggesting that local economic factors are sometimes important determinants of future rent growth.

Finally, shown in the bottom panel of Table 4, we find that housing premia are predictable in every one of the 28 markets included in our study. Specifically, the Wald test of no predictability is rejected at or below the 5% significance level in every market; the median full-sample $\overline{R}^2$ is roughly 50% and the inter-quartile range of the full sample $\overline{R}^2$ runs from 47% to 59%. The lagged housing premium is a significant predictor of future housing premia in over 80% of the markets in our sample, results that are qualitatively similar to those of Case and Shiller. One important difference between our results and Case and Shiller's is that we find that in 25% of our metro areas, local macroeconomic variables such as lagged rent growth and lagged income growth also appear to predict future premia, but that lagged premia have predictive power even after including these other variables.

The sub-sample $\overline{R}^2$ of the forecasting model for housing premia is much lower during the 1997–2007 period than in the 1975–1996 period, and in many cases is negative. The model typically under-predicts recent housing premia largely because these premia became much more persistent during the boom. In other words, the poor in-sample fit during the 1997–2007 period results from imputing too little predictability to future premia. In what follows, we discuss the consequences of the forecasting model's inability to track premia over the boom period for our variance decompositions.

5. Decomposing the variability of rent–price ratios

The results of the previous section show that real rates, rental growth, and housing premia exhibit significant predictable variation. In this section, we examine how our estimates of changes to expectations of each of these components determines the overall volatility of rent–price ratios.

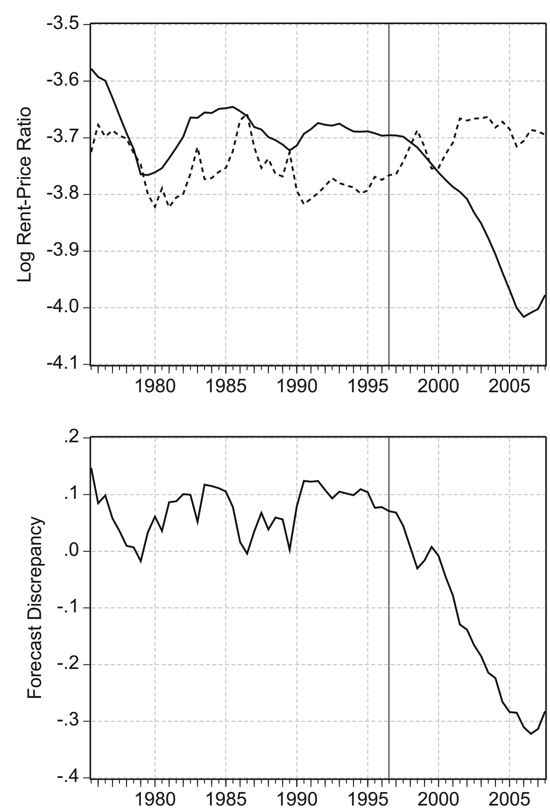

Fig. 2. In the top panel, we graph the actual ($r_t - p_t$, sold line) and forecasted ($\widehat{r_t - p_t}$, dashed line) log rent–price ratio for the USA. In the bottom panel, we show the forecast discrepancy resulting from differencing the actual and forecasted log rent–price ratios that are graphed in the corresponding top panel. A solid line is imposed at 1996:H2 for reference.

Table 5
Variance decomposition of the rent–price ratio, 1975–1996.

	Variance		Variance shares			Covariance shares		
	$r-p$ (1)	$\widehat{r-p}$ (2)	$\widehat{\mathscr{I}}_t$ (3)	$\widehat{\Pi}_t$ (4)	$\mathscr{E}_t$ (5)	$(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t)$ (6)	$(\widehat{\mathscr{I}}_t, \mathscr{E}_t)$ (7)	$(\widehat{\Pi}_t, \mathscr{E}_t)$ (8)
USA	0.04	0.04	0.95	2.51	1.04	−2.27	−0.36	−0.87
Midwest	0.05	0.06	0.81	4.34	2.43	−2.81	1.86	−5.63
Chicago	0.07	0.08	0.45	0.63	0.47	−0.73	0.32	−0.14
Cincinnati	0.06	0.03	0.62	0.92	0.80	−1.26	0.60	−0.67
Cleveland	0.07	0.03	0.46	0.39	0.97	−0.52	−0.12	−0.18
Detroit	0.11	0.11	0.13	1.50	1.78	−0.48	0.07	−1.99
Kansas City	0.08	0.06	0.27	0.70	0.93	−0.82	0.33	−0.41
Milwaukee	0.10	0.08	0.21	0.78	0.87	−0.63	0.40	−0.63
Minneapolis	0.06	0.04	0.40	0.91	1.13	−1.16	0.27	−0.56
St. Louis	0.06	0.03	0.67	1.36	1.15	−1.74	1.08	−1.52
Northeast	0.14	0.18	0.14	1.99	2.37	0.47	−0.58	−3.40
Boston	0.18	0.17	0.05	1.39	1.77	0.28	−0.22	−2.27
New York	0.24	0.15	0.07	0.53	2.37	0.23	−0.43	−1.78
Philadelphia	0.10	0.07	0.35	1.14	2.88	0.47	−1.02	−2.82
Pittsburgh	0.07	0.03	0.36	0.31	0.50	0.03	−0.07	−0.14
South	0.04	0.04	0.84	1.84	0.98	−2.45	0.31	−0.52
Atlanta	0.03	0.03	1.78	3.56	1.57	−4.01	−0.58	−1.32
Dallas	0.09	0.06	0.30	1.70	0.98	−1.33	0.45	−1.10
Houston	0.11	0.07	0.14	1.08	1.11	−0.72	0.39	−0.99
Miami	0.04	0.12	0.84	8.06	5.82	−2.22	0.82	−12.32
West	0.06	0.10	0.79	3.29	1.65	−3.04	1.62	−3.31
Denver	0.06	0.07	0.43	4.76	4.70	−2.05	0.71	−7.57
Honolulu	0.19	0.06	0.19	0.34	1.01	−0.25	0.19	−0.48
Los Angeles	0.15	0.23	0.13	3.32	2.83	−0.57	0.17	−4.89
Portland	0.11	0.05	0.14	0.73	1.06	−0.60	0.33	−0.66
San Diego	0.11	0.16	0.15	2.89	1.43	−0.72	0.19	−2.94
San Francisco	0.16	0.13	0.09	0.80	0.90	−0.39	0.03	−0.43
Seattle	0.14	0.05	0.09	0.79	1.02	−0.50	0.11	−0.51
Metro median	0.10	0.07	0.27	0.92	1.11	−0.63	0.19	−0.99

In this table, we report the results of the variance decomposition defined by Eq. (13), $var(r_t - p_t) = var(\widehat{\mathscr{I}}_t) + var(\widehat{\Pi}_t) + var(\mathscr{E}_t) + 2cov(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t) - 2cov(\widehat{\mathscr{I}}_t, \mathscr{E}_t) - 2cov(\widehat{\Pi}_t, \mathscr{E}_t)$, over the 1975–1996 period.
Column (1) reports the variance of the observed log rent–price ratio, $r_t - p_t$. Column (2) reports the variance of the predicted log rent–price ratio implied by the VAR, $\widehat{r_t - p_t}$. The results of the variance decomposition are reported in columns (3)–(8). Columns (3)–(5) report the share of the variance in the rent–price ratio accounted for by variation in expected future real rates, $var(\widehat{\mathscr{I}}_t)$, expected future premia, $var(\widehat{\Pi}_t)$, and expected future rent growth, $var(\mathscr{E}_t)$. Columns (6)–(8) report the share of the variance in the rent–price ratio accounted for by the covariance between expected future real rates and premia, $2cov(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t)$, expected future real rates and rent growth, $-2cov(\widehat{\mathscr{I}}_t, \mathscr{E}_t)$, and expected future premia and rent growth, $-2cov(\widehat{\Pi}_t, \mathscr{E}_t)$.

Before discussing the details of the variance decompositions we briefly summarize the behavior of the VAR model's ability to track the broad movements in the rent–price ratio over the sample period, 1975–2007. We focus our discussion on the national level results, but the results of the regional and metropolitan markets are qualitatively similar.

In the upper panel of Fig. 2, we plot the log of the rent–price ratio computed from the national VAR, $\widehat{r_t - p_t}$, as a dashed line and the actual log of the national rent–price ratio, $r_t - p_t$, as a solid line over the entire 1975–2007 period. In the lower panel we display the forecast discrepancy, e_t. For clarity, we plot a solid vertical line at 1996:H2 in both the top and bottom panels. Prior to 1996:H2, the VAR-computed rent–price ratio seems to capture low frequency boom-bust movements in the actual rent–price ratio in the sense that the local peaks and troughs of both series are similar. Recall that the VAR is not constructed with the explicit goal of fitting the historical rent–price ratio. Rather the VAR is constructed to fit historical patterns of real rates, risk premia, and rental growth. It just so happens that when we construct a rent–price ratio given the VAR-based estimates (the "VAR-computed rent–price ratio"), it matches low-frequency movements in the actual rent–price ratio.[13] After 1996:H2, however, the two series diverge remarkably. The estimated rent–price ratio rises slightly whereas the actual rent–price ratio declines by a considerable amount.

The fact that the VAR-computed rent–price ratio does not track the decline in the actual rent–price ratio is evident in the plot of the forecast discrepancy in the bottom panel. It is not completely surprising that the VAR model proves inadequate in this respect. One reason for the miss is the unprecedented decline in the ratio. A more fundamental reason involves the covariance of real interest rates and risk-premia. Himmelberg et al. (2005) attribute much of the housing boom to a decline in real interest rates. However, in the next section we will show that prior to 1996, movements in real rates have been accompanied by opposite movements in housing risk-premia such that, on net, changes to real interest rates have not been accompanied by large changes in housing valuations. Explaining the recent housing boom (and the subsequent bust) is beyond the scope of this paper. However, it is worth noting that our results suggest that either expectations about future conditions were not consistent with history or the covariance structure of real rates and housing risk-premia have changed. Such changes could have been brought about by changes in lending standards (Dell'Arriccia et al., 2008; Keys et al., 2009) or by other factors.

5.1. Sources of rent–price ratio variability: 1975–1996

Table 5 displays the results of the variance decompositions for the 28 housing markets in our sample over the 1975–1996 period.

[13] In Table 5, we show that for the median metro area, the variance of the VAR-computed rent–price ratio is equal to 70% of the variance of the actual rent–price ratio over the 1975–1996 period.

Table 6
Variance decomposition of the rent–price ratio, 1997–2007.

	Variance		Variance shares			Covariance shares		
	$r-p$ (1)	$\widehat{r-p}$ (2)	$\widehat{\mathscr{I}}_t$ (3)	$\widehat{\Pi}_t$ (4)	$\mathscr{E}_t$ (5)	$(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t)$ (6)	$(\widehat{\mathscr{I}}_t, \mathscr{E}_t)$ (7)	$(\widehat{\Pi}_t, \mathscr{E}_t)$ (8)
USA	0.11	0.03	0.09	0.22	1.15	−0.26	0.48	−0.67
Midwest	0.09	0.03	0.03	0.34	0.35	−0.11	−0.09	0.48
Chicago	0.12	0.06	0.04	0.13	0.94	−0.11	0.25	−0.25
Cincinnati	0.05	0.02	0.39	0.31	0.32	−0.27	0.36	−0.11
Cleveland	0.03	0.02	1.90	0.83	1.14	0.27	−2.22	−0.93
Detroit	0.07	0.09	0.16	2.00	1.81	0.82	−0.73	−3.06
Kansas City	0.07	0.05	0.38	0.62	0.98	−0.89	1.06	−1.14
Milwaukee	0.13	0.02	0.04	0.05	0.68	−0.05	0.23	0.03
Minneapolis	0.17	0.04	0.04	0.06	0.94	−0.09	0.30	−0.24
St. Louis	0.11	0.02	0.09	0.16	0.88	−0.21	0.46	−0.39
Northeast	0.16	0.07	0.01	0.16	1.81	0.02	−0.10	−0.89
Boston	0.17	0.10	0.01	0.67	0.67	0.13	−0.07	−0.42
New York	0.20	0.06	0.01	0.18	0.61	0.02	−0.11	0.29
Philadelphia	0.17	0.03	0.02	0.05	0.90	0.00	0.08	−0.05
Pittsburgh	0.06	0.02	0.09	0.13	0.52	−0.03	0.14	0.16
South	0.10	0.02	0.02	0.06	1.09	−0.07	−0.01	−0.10
Atlanta	0.11	0.02	0.06	0.07	0.65	0.00	0.21	0.01
Dallas	0.08	0.03	0.13	0.38	1.25	−0.39	0.66	−1.03
Houston	0.08	0.04	0.20	1.11	2.09	−0.90	1.12	−2.63
Miami	0.27	0.09	0.00	0.12	1.68	−0.02	0.03	−0.82
West	0.18	0.04	0.01	0.05	1.33	−0.01	−0.02	−0.36
Denver	0.10	0.03	0.09	0.16	0.55	0.07	0.21	−0.08
Honolulu	0.22	0.03	0.03	0.09	1.74	−0.01	−0.25	−0.60
Los Angeles	0.26	0.14	0.01	0.35	1.01	0.00	0.04	−0.41
Portland	0.18	0.06	0.03	0.19	1.35	−0.14	0.26	−0.68
San Diego	0.22	0.08	0.02	0.22	0.54	0.00	0.15	0.06
San Francisco	0.22	0.09	0.03	0.24	1.03	0.05	0.12	−0.46
Seattle	0.18	0.06	0.03	0.24	1.51	−0.16	0.32	−0.95
Metro median	0.13	0.04	0.04	0.19	0.94	−0.02	0.21	−0.41

In this table, we report the results of the variance decomposition defined by Eq. (13), $var(r_t - p_t) = var(\widehat{\mathscr{I}}_t) + var(\widehat{\Pi}_t) + var(\mathscr{E}_t) + 2cov(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t) - 2cov(\widehat{\mathscr{I}}_t, \mathscr{E}_t) - 2cov(\widehat{\Pi}_t, \mathscr{E}_t)$, over the 1997–2007 period.
Column (1) reports the variance of the observed log rent–price ratio, $r_t - p_t$. Column (2) reports the variance of the predicted log rent–price ratio implied by the VAR, $\widehat{r_t - p_t}$. The results of the variance decomposition are reported in columns (3)–(8). Columns (3)–(5) report the share of the variance in the rent–price ratio accounted for by variation in expected future real rates, $var(\widehat{\mathscr{I}}_t)$, expected future premia, $var(\widehat{\Pi}_t)$, and expected future rent growth, $var(\mathscr{E}_t)$. Columns (6)–(8) report the share of the variance in the rent–price ratio accounted for by the covariance between expected future real rates and premia, $2cov(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t)$, expected future real rates and rent growth, $-2cov(\widehat{\mathscr{I}}_t, \mathscr{E}_t)$, and expected future premia and rent growth, $-2cov(\widehat{\Pi}_t, \mathscr{E}_t)$.

The reported shares in this table correspond to the share attributable to each element of Eq. (13), and thus sum to 1.0. For example, the reported covariance shares of $(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t)$, $(\widehat{\mathscr{I}}_t, \mathscr{E}_t)$, and $(\widehat{\Pi}_t, \mathscr{E}_t)$ in columns 6–8 refer to shares of the variance of $r-p$ attributable to $2cov(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t)$, $-2cov(\widehat{\mathscr{I}}_t, \mathscr{E}_t)$, and $-2cov(\widehat{\Pi}_t, \mathscr{E}_t)$, respectively.[14]

Before examining the variance decompositions, we briefly discuss the overall variability of the actual rent–price ratio, $r-p$, and its estimate, $\widehat{r-p}$. The volatility of the estimated rent–price ratio (column 2) is similar to the variability of the actual log rent–price ratio (column 1) at the national, regional and metropolitan market level. The variance of the estimated and actual rent–price ratio are identical at the national level, and when considering all markets, the variance of the estimated rent–price ratio is typically 70% as large as the variance in the actual rent–price ratio. This suggests that the VAR model for expectations captures important movements in rent–price ratios over this period. Also note that rent–price ratios are about 30% more volatile at the metropolitan market level than at the national level, suggesting that some local factors average out in the aggregate.

[14] Thus, if we report a negative contribution of $(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t)$ to the variance of $r-p$, it means changes to these two variables are negatively correlated. But in the case of in $(\widehat{\mathscr{I}}_t, \mathscr{E}_t)$ and $(\widehat{\Pi}_t, \mathscr{E}_t)$, if we report a negative contribution of the covariance of these variables to the overall variance of $r-p$ or $\widehat{r-p}$, it means that changes to the two variables are positively correlated.

At the national level, variation in expected future premia, $\widehat{\Pi}_t$, is estimated to be the largest source of variability of rent–price ratios; variation in expected future rent growth, $\mathscr{E}_t$ is similar to that of expected future real rates, $\widehat{\mathscr{I}}_t$. The results at the metropolitan level indicate that while variation in expected future premia plays a key role, variation in expected future rent growth is about as important. For example, at the median for the 23 metropolitan markets (reported in the final row of the table), variation in expected future rent growth and premia account for similar shares of the volatility of rent–price ratios.

The covariances among the components serve to dampen fluctuations in rent–price ratios; columns 6–8 provide some details. We find that expected future premia and real rates (column 6) are negatively correlated in all markets except those in the Northeast, and that expected future premia and rent growth (column 8) are positively correlated in every market (which implies a negative contribution). The positive correlation of premia and rent growth reduces the volatility of rent–price ratios by approximately the same extent at the metropolitan level (median of −0.99) as at the national level (−0.87). All of these results also hold for the full 1975–2007 sample of data (not shown).

These results are quite consistent with some key findings of variance decompositions from the stock and bond markets. First, we find that in the aggregate, variation in expected future premia is the largest source of rent–price ratio variation. In particular, we find that variation in future premia is roughly 2-1/2 times as large as variation in future rent growth (columns 4 and 5). Bernanke and Kuttner (2005) find, over a similar time period, that expected future equity premia are roughly three times as variable as expected future dividend growth in the aggregate stock market. Second, we find that expected future premia are positively correlated with expected future rent growth, which reduces the overall volatility of rent–price ratios. In his study of firm-level stock returns, Vuolteenaho (2002) also finds that expected future premia and fundamentals are positively correlated.

Vuolteenaho suggests that the positive correlation of expected premia and fundamentals is consistent with a behavioral story that stock prices under-react to positive changes in expected future fundamentals: at the same time that good news about future fundamentals arrives, raising prices, required returns tend to rise, depressing prices. This behavioral interpretation could also be applied to the housing market. The positive correlation between expected future rent growth and premia that we document could simply indicate that house prices do not increase by "enough" during periods of rising rent growth, which mechanically implies a contemporaneous increase in housing premia. This interpretation may also be applied to our finding that premia and real rates are negatively correlated, in the sense that house prices do not increase by "enough" when real rates decline. Whether or not the negative covariance shares we report are the consequence of behavioral or fully-rational and forward-looking decision making is beyond the scope of this paper. We find it interesting, however, that housing markets and other financial markets exhibit key similarities that suggests the patterns we document may be linked to fundamental characteristics of investor preferences or constraints.

5.2. Sources of rent–price ratio variability: 1997–2007

Table 6 reports the results of the variance decompositions for the housing-boom period of 1997–2007 using the identical layout as in Table 5. As mentioned, much of the change in the rent–price ratio in this period is not explained by changes to the VAR-based estimates of $\mathscr{I}_t$, Π_t, and $\mathscr{G}_t$. As a result, the variance of the predicted rent–price ratio (column 2) is typically less than half as large as the variance of the actual rent price ratio (column 1) over this period.

Our results indicate that variation in expected future rent growth played the dominant role in accounting for variation in rent–price ratios from 1997 to 2007. In particular, at both the national level and for the median of the metro-area markets, variation in rent growth is estimated to be roughly 5 times larger than variation in premia. This is due entirely to the behavior of the forecast discrepancy. If we were to associate this discrepancy with either of the other two components $(\widehat{\mathscr{I}}_t, \widehat{\Pi}_t)$ during this time period, we would find it to be the dominant source of variation. Even so, with the "deck stacked" against finding significant variation in premia, we find that such variation is important over the 1997–2007 period. Across all our markets, variation in premia typically accounts for roughly 20% of the variation of rent–price ratios over this period, which is substantial considering the pronounced decline of these ratios and corresponding increase in their volatility.

The covariance shares indicate that the correlations between the three components of the rent–price ratio dampened total variability over the 1997–2007 period, as they did from 1975 to 1996. Specifically, we find that the sum of the three covariance shares is negative or zero in all but seven markets; future premia and real rates are negatively correlated, also in all but eight markets; and premia and rent growth are positively correlated in all but six markets. Thus, although the contribution of rent growth in the indirect decomposition changes quite significantly between the 1975–1996 and 1997–2007 periods, the pattern in the covariance shares is fairly stable over these two periods.

5.3. Sources of relative rent–price ratio variation: 1975–2007

In this section we document that our results to this point are not driven solely by common trends across all markets. Our approach is to examine the variability of relative rent–price ratios, $r_t^* - p_t^*$ which is defined in Eq. (14) as the difference between a local and the national rent–price ratio. The variance decompositions are performed according to Eq. (15) and the results are reported in Table 7. We report results from the full 1975–2007 sample only because the large decline in the rent–price ratio that occurred starting in 1997 was common to the US and almost every market in our sample: this can be seen by comparing columns 3 and 4 of Table 1.

The first three columns of Table 7 report the variance of the rent price ratio, $r_t - p_t$, the variance of the relative rent–price ratio, $r_t^* - p_t^*$, and the variance of the forecast of the relative rent price ratio, $\widehat{r_t^* - p_t^*}$. The relative rent–price ratio is typically 50% to 75% as variable as the actual rent–price ratio, indicating that local factors play a significant role in generating rent–price ratio variability. Thus, while the traditional maxim that all real estate is local is likely an overstatement, local factors are certainly important. Also, the forecast of the relative rent–price ratio is typically almost as

Table 7
Variance decomposition of the deviations of the rent–price ratio from the US, 1975–2007.

	Variances			Variance shares		Cov. Share
	$r-p$ (1)	r^*-p^* (2)	$\widehat{r^*-p^*}$ (3)	$\widehat{\Pi}_t^*$ (4)	$\mathscr{E}_t^*$ (5)	$\left(\widehat{\Pi}_t^*, \mathscr{E}_t^*\right)$ (6)
Midwest	0.11	0.03	0.07	9.05	6.60	−14.64
Chicago	0.12	0.05	0.05	0.91	2.34	−2.24
Cincinnati	0.09	0.05	0.04	1.18	3.07	−3.25
Cleveland	0.09	0.07	0.05	0.79	2.97	−2.76
Detroit	0.21	0.15	0.12	0.91	1.57	−1.48
Kansas City	0.09	0.07	0.04	0.19	1.15	−0.34
Milwaukee	0.16	0.08	0.06	0.95	1.55	−1.49
Minneapolis	0.18	0.08	0.03	0.18	1.27	−0.45
St. Louis	0.13	0.04	0.03	1.17	2.69	−2.86
Northeast	0.17	0.12	0.15	1.30	1.46	−1.75
Boston	0.24	0.18	0.18	1.56	1.26	−1.82
New York	0.27	0.21	0.13	0.48	1.74	−1.22
Philadelphia	0.16	0.08	0.05	0.97	2.34	−2.30
Pittsburgh	0.10	0.06	0.06	2.02	3.16	−4.17
South	0.09	0.04	0.03	1.99	1.05	−2.04
Atlanta	0.10	0.02	0.02	5.90	6.10	−11.00
Dallas	0.09	0.12	0.05	0.48	0.83	−0.31
Houston	0.11	0.14	0.04	0.39	0.68	−0.07
Miami	0.24	0.14	0.11	0.50	2.31	−1.81
West	0.17	0.07	0.09	1.61	2.45	−3.05
Denver	0.14	0.06	0.04	2.73	3.83	−5.56
Honolulu	0.23	0.17	0.06	0.35	1.36	−0.71
Los Angeles	0.24	0.15	0.18	1.84	2.11	−2.95
Portland	0.24	0.16	0.04	0.13	1.49	−0.62
San Diego	0.22	0.12	0.11	1.11	1.17	−1.28
San Francisco	0.25	0.16	0.11	0.58	1.41	−1.00
Seattle	0.25	0.17	0.04	0.24	1.61	−0.85
Metro median	0.16	0.12	0.05	0.91	1.61	−1.49

In this table, we report the results of the variance decomposition of the relative log rent–price ratio series, $r_t^* - p_t^* \equiv r_t - p_t - (r_t^{US} - p_t^{US})$, defined by Eq. (15), $var(r_t^* - p_t^*) = var\left(\widehat{\Pi}_t^*\right) + var(\mathscr{E}_t^*) - 2cov\left(\widehat{\Pi}_t^*, \mathscr{E}_t^*\right)$, over the 1975–2007 period.
Column (1) reports the variance of the raw log rent–price ratio, $r_t - p_t$. Column (2) reports the variance of the relative log rent–price ratio, $r_t^* - p_t^*$. Column (3) reports the variance of the predicted log relative rent–price ratio, $\widehat{r_t^* - p_t^*}$.
Columns (4)–(6) report the results of the variance decomposition. Column (4) reports the share of relative rent–price ratio variation accounted for by variation in relative expected future premia, $var\left(\widehat{\Pi}_t^*\right)$. Column (5) reports the share of relative rent–price ratio variation accounted for by variation in expected future relative rent growth, $var(\mathscr{E}_t^*)$. Column (6) reports the share of relative rent–price ratio variation accounted for by the covariance in relative premia and rent growth, $-2cov\left(\widehat{\Pi}_t^*, \mathscr{E}_t^*\right)$.

volatile as the relative ratio itself suggesting that the VAR model captures a good deal of the variation in relative rent–price ratios.

The results of the variance decomposition are reported in columns 4–6. Variation in relative premia (column 4) and the relative expected future rent growth (column 5) are both important for overall variation in relative rent–price ratios. Further, the estimated covariance share for relative expected future rent growth and premia is negative in every market. Markets that experience an increase in expected future rent growth, relative to the nation, simultaneously experience an increase in premia, also relative to the nation, that ultimately has a moderating influence on prices. These two findings are consistent with the results reported in Tables 5 and 6.

6. Conclusion

In this paper we have applied the dynamic Gordon growth model to study the fundamental sources of variation in rent–price ratios in 23 metropolitan housing markets, four regional markets, and the national housing market over the 1975–2007 period. We have decomposed the variance of rent–price ratios using an approach common in the finance literature that uses VAR-based forecasts of expected future real interest rates and housing premia, and identifies expected future rent growth residually. We have examined the data over the 1975–1996 and 1997–2007 period separately in light of the unprecedented decline in rent–price ratios that occurred during the latter period. For the regional and local markets, we have also studied the variability of changes relative to the national market.

Using a VAR approach, we show that real interest rates, housing premia, and rent growth all exhibit substantial predictable variation. Each of these three components of the rent–price ratio makes a significant contribution to its variation. The resulting variance decompositions that we perform reveal some important characteristics of housing market behavior.

First, we find that housing premia exhibit significant variation. Over the 1975–1996 period, variation in risk premia was the principal source of variation of rent–price ratios at the national level. Variation in premia is an important source of rent–price ratio variation at the regional and metropolitan market level as well. Even during the boom period, 1997–2007, our approach (which residually attributes most of variation in rent price ratios to changes to expected future rent growth) indicates that variation in premia accounts for about 20% of variation in all of our markets. In our relative decompositions that control for movements in housing premia at the aggregate level, we found that variation in risk-premia remains an important source of variation in rent–price ratios.

Second, we consistently find that the covariances among the three components of the rent–price ratio dampen fluctuations in the rent–price ratio. In particular, we show that the covariance between expected future premia and rents is positive in most markets, a finding that is robust to controlling for movements in aggregate premia and rents. The positive covariance between expected future premia and rents is similar to findings from the stock market that expected future dividends and premia tend to be positively correlated, thus providing an additional similarity between the behavior of housing markets and financial markets.

The application of the dynamic Gordon growth model to housing markets that we have examined in this paper provides insight into the fundamental sources of variability in housing valuations. Aside from providing direct evidence on the nature of fluctuations in rent–price ratios, the framework that we adopt allows for a meaningful comparison of housing and other financial assets. Although housing markets are quite different from traditional financial markets in both form and function, we find a number of interesting similarities. We conclude that understanding the underlying structural links between housing and financial markets is likely to be a fruitful area of future research.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jue.2009.06.002.

References

Bernanke, B., Kuttner, K., 2005. What explains the stock market's reaction to Federal Reserve policy? Journal of Finance 60 (3), 1221–1257.

Brunnermeier, M., Julliard, C., 2008. Money illusion and housing frenzies. Review of Financial Studies 21 (1), 135–180.

Bucks, B., Pence, K., 2008. Do borrowers know their mortgage terms? Journal of Urban Economics 64 (2), 218–233.

Campbell, J., Shiller, R., 1988a. The dividend–price ratio and expectations of future dividends and discount factors. Review of Financial Studies 1 (3), 195–228.

Campbell, J., Shiller, R., 1988b. Stock prices, earnings and expected dividends. Journal of Finance 43 (3), 661–676.

Campbell, J., 1991. A variance decomposition for stock returns. Economic Journal 101 (405), 157–179.

Campbell, J., Ammer, J., 1993. What moves the stock and bond markets? A variance decomposition for long-term asset returns. Journal of Finance 48 (1), 3–37.

Case, K., Shiller, R., 1989. The efficiency of the market for single-family homes. American Economic Review 79 (1), 125–137.

Cutts, A., Green, R., Chang, Y., 2005. Did changing rents explain changing house prices during the 1990s? Working Paper, American Real Estate and Urban Economics Association.

Davis, M., Lehnert, A., Martin, R., 2008. The rent–price ratio for the aggregate stock of owner-occupied housing. Review of Income and Wealth 54 (2), 279–284.

Davis, M., Heathcote, J., 2007. The price and quantity of residential land in the United States. Journal of Monetary Economics 54 (8), 2595–2620.

Davis, M., Palumbo, M., 2001. A primer on the economics and time series econometrics of wealth effects. Federal Reserve Board of Governors: Finance and Economics Discussion Series Working Paper 2001–09.

Dell'Arriccia, G., Igan, D., Laeven, L., 2008. Credit booms and lending standards: evidence from the subprime mortgage market. Working Paper, International Monetary Fund.

Deng, Y., Quigley, J., Van Order, R., 2000. Mortgage terminations, heterogeneity and the exercise of mortgage options. Econometrica 68 (2), 275–308.

Flavin, M., Yamashita, T., 2002. Owner-occupied housing and the composition of the household portfolio. American Economic Review 92 (1), 345–362.

Foote, C., Gerardi, K., Willen, P., 2008. Negative equity and foreclosure: theory and evidence. Journal of Urban Economics 64 (2), 234–245.

Gallin, J., 2008. The long-run relationship between house prices and rents. Real Estate Economics 36 (4), 635–658.

Gorton, G., 2009. Slapped in the face by the invisible hand: banking and the panic of 2007. Working Paper, Yale University.

Himmelberg, C., Mayer, C., Sinai, T., 2005. Assessing high house prices: bubbles, fundamentals, and misperceptions. Journal of Economic Perspectives 19 (4), 67–92.

Haughwout, A., Peach, R., Tracy, J., 2008. Juvenile delinquent mortgages: bad credit or bad economy? Journal of Urban Economics 64 (2), 246–257.

Hubbard, G., Mayer, C., 2008. First, let's stabilize home prices. Wall Street Journal: Opinion section, October 2.

Keys, B., Mukherjee, T., Seru, A., Vig, V., 2009. Financial regulation and securitization: evidence from subprime loans. Journal of Monetary Economics 56 (5).

Meese, R., Wallace, N., 1994. Testing the present value relation for housing prices: should I leave my house in San Francisco. Journal of Urban Economics 35 (3), 245–266.

Plazzi, A., Torous, W., Valkanov, R., 2006. Expected returns and the expected growth in rents of commercial real estate. Working Paper, University of California, Los Angeles.

Quigley, J., Van Order, R., 1995. Explicit tests of contingent claims models of mortgage default. The Journal of Real Estate Finance and Economics 11 (2), 99–117.

Rapach, D., Wohar, M., 2004. The persistence in international real interest rates. International Journal of Finance and Economics 9 (4), 339–346.

Rose, A., 1988. Is the real interest rate stable. Journal of Finance 43 (5), 1095–1112.

Shiller, R., 2005. Irrational Exuberance, second ed. Princeton University Press, Princeton, NJ.
Shiller, R., Beltratti, A., 1992. Stock prices and bond yields: can their comovements be explained in terms of present value models? Journal of Monetary Economics 30 (1), 25–46.
Taylor, J., 2007, Housing and monetary policy, Working Paper, Stanford University.
Verbrugge, R., 2008. The puzzling divergence of rents and user costs, 1980–2004. Review of Income and Wealth 54 (4), 671–699.
Vuolteenaho, T., 2002. What drives firm-level stock returns. Journal of Finance 57 (1), 233–264.
Wheaton, W., 1990. Vacancy, search, and prices in a housing market matching model. Journal of Political Economy 98 (6), 1270–1292.