Hot and Cold Housing Markets: International Evidence

Jose Ceron
Omega Capital

Javier Suarez*
CEMFI, CEPR and ECGI

September 2006

Abstract

This paper examines the experience of fourteen developed countries for which there are about thirty years of quarterly inflation-adjusted housing price data. Price dynamics is modeled as a combination of a country-specific component and a cyclical component. The cyclical component is a two-state Markov switching process with parameters common to all countries. We find that the latent state variable captures previously undocumented changes in the volatility of real housing price increases. These volatility phases are quite persistent (about six years, on average) and occur with about the same unconditional frequency over time. In line with previous studies, the mean of real housing price increases can be predicted to be larger when lagged values of those increases are large, real GDP growth is high, unemployment falls, and interest rates are low or have declined. Our findings have important implications for risk management in regard to residential property markets.

Key words: housing prices, cycles, volatility, Markov switching.

JEL codes: E32, G15, R31

*We thank Robert Korajczyk, Claudio Michelacci, Pedro Mira, Francisco Peñaranda, Gabriel Perez-Quiros, Rafael Repullo, Enrique Sentana, and the seminar audiences at the Banco de España, European Central Bank, and the Gutmann Center Symposium on Real Assets and Portfolio Management for useful comments. We also thank the Bank for International Settlements for kindly providing us with the country series of inflation-adjusted residential property prices. Address for correspondence: Javier Suarez, CEMFI, Casado del Alisal 5, 28014 Madrid, Spain. Tel.:+34-914290551. Fax: +34-914291056. E-mail: suarez@cemfi.es

1 Introduction

Many expert and media accounts of residential property markets describe their dynamics as the recurrence of booms and busts.[1] In a typical boom or hot phase, transactions are abundant, average selling times are short, and prices tend to grow fast. In a bust or cold phase, there are fewer transactions, average selling times are longer, and price growth moderates or becomes negative. The empirical literature on housing markets recognizes the importance of swings in housing prices and their relationship to changes in the liquidity of the market, but it is somewhat less assertive about the cyclical pattern.[2] For instance, Case and Shiller (1989) document the autocorrelation in housing price increases, which is suggestive of imperfections that make residential property markets informationally inefficient but does not imply the existence of cycles. In contrast, Muellbauer and Murphy (1997) and Herring and Wachter (1999) explicitly refer to booms and busts in housing prices and admit the non-linearities that they imply.[3] The studies by Englund and Ioannides (1997), Capozza et al. (2002), Tsatsaronis and Zhu (2003), and Borio and Mcguire (2004), among others, report significant correlations between real housing price growth and variables such as real GDP, unemployment, interest rates, and inflation, which suggests that property prices might feature a cyclical pattern, if only because of the convolution of the cyclicality of the other variables. Finally, swings in housing prices have been shown to be positively correlated with the volume of transactions (Stein, 1995) and negatively correlated with average selling times (Krainer, 2001).

Theoretical work on housing markets focuses on providing explanations for the persistence of price changes, as well as the correlations between price growth, volumes, and selling times. The most recent theories typically attribute the fluctuations to shocks to the demand for housing services or to buyers' income, and put the emphasis

[1] The Special Report on "The global housing boom: In come the waves" published in *The Economist* on 16th June 2005 provides a good illustration of these views.

[2] See Cho (1996), for an early survey of the literature.

[3] Muellbauer and Murphy (1997) argue that introducing a cubic transformation of lagged price changes helps capturing the strong inertia of the peaks and troughs in price growth.

on the role played by either search frictions (Wheaton, 1990; Williams, 1995; Krainer, 2001; Novy-Marx, 2005) or financial imperfections (Stein, 1995; Ortalo-Magne and Rady, 2005). For instance, Krainer (2001) shows that positive (negative) shocks to the value of the housing service flow can produce a positive relationship between the liquidity of the market and its prices by increasing (decreasing) the opportunity cost of failing to complete a transaction. From a different perspective, Ortalo-Magne and Rady (2005) develop a dynamic model in which shocks to the income of the financially-constrained first-time home buyers cause overreactions in prices and volumes. Finally, there is a less structured line of thought, recently surveyed by Case and Shiller (2003), that imputes part of the runups in housing prices to expectations of subsequent price increases, like in the usual theories of bubbles.[4]

To the extent that search frictions, financial constraints, and expectational feedback might have varying importance across the cyclical phases of the housing market, one might expect price dynamics to differ across the hot and cold phases. The explicit empirical modeling of cyclical regime switches in price dynamics, which is the focus of this paper, confirms that this is indeed the case. Quite remarkably, though, our analysis documents the existence of sizable and significant fluctuations in the volatility of price growth, a phenomenun not explicitly predicted by any of the theoretical studies. Hence, our findings can be seen more as a challenge for future theoretical work than as evidence in support of current theories.

We consider a panel of inflation-adjusted residential property price indices that covers about thirty years of quarterly data from fourteen OECD countries. Real housing price growth is modeled as the combination of a country-specific component and a cyclical component. The country-specific component is intended to capture unobservable cross-country heterogeneity, while the cyclical component is a two-state Markov switching process a la Hamilton (1989) with parameters common to all coun-

[4]It is argued, however, that, in contrast to stock prices, housing prices rarely crash since many sellers resist to cut down their prices or sell at a loss—see Genesove and Mayer (2001) for evidence on this.

tries.[5] Besides the genuine interest of international evidence, with the multi-country approach we aim to improve the reliability of the estimates of the cyclical component, especially in a context where we fear that the available time series might be short relative to the average length of the underlying cycles.[6] Our formalization, inspired in standard panel data techniques, allows us to focus on the cyclical component by simply standardizing each country's price increases and then fitting a common Markov switching model to the standardized time series.[7]

The estimation yields three set of results. First, the dynamics of real housing prices is characterized by two rather persistent states that mostly differ in the volatility of price increases. Specifically, the variance of the unpredictable part of quarterly price increases in the high volatility state is almost four times as big as in the low volatility state. The low volatility state is associated with phases of higher growth, occurs with an unconditional probability of 47% and has an expected duration of 23 quarters. The high volatility state is associated with phases of lower price growth, occurs with a frequency of 53% and lasts, on average, about 26 quarters. Second, in addition to the latent state variable, a number of lagged macroeconomic variables have significant predictive power for the expected growth rate of real housing prices. Specifically, the prediction of quarterly growth rates depends positively on the lagged quarterly rate of real GDP growth, negatively on the lagged one-year variation in the unemployment rate and also negatively on the lagged long-term nominal interest rate. We find no evidence of the effect of these variables to be state dependent. Third, even after controlling for the effect of the latent state variable and the explanatory variables, the quarterly growth rate of real housing prices exhibits significant positive autocorrelation.

[5] Because of its reference to an explicit probabilistic model, Hamilton's approach is more suitable for forecasting purposes than the approaches that focus on the dating of peaks and troughs according to some a priori definition of a cycle—see, for example, Harding and Pagan (2001).

[6] Anecdotal evidence (confirmed by our results below) suggests that a typical housing cycle phase is much longer (about six years) than a typical business cycle phase (two to three years).

[7] Importantly, the parameters of the Markov switching model are the same for all the standardized series, but each country's latent state variable is treated as an independent realization of the latent state process.

The most novel results are obviously those related to the Markov switching structure. The second and third sets of findings are consistent with previous evidence and imply that allowing for a Markov switching structure in the mean of price increases does not deprive the usual macroeconomic variables and lagged price increases of their predictive power. The picture that the Markov switching structure draws does not correspond to a simple story of rise and decline. Rather than large differences in the expected growth of real housing prices (which seem better captured by the observable explanatory variables), the latent cyclical variable captures striking changes in volatility. A housing cycle features long phases (almost six years on average) of less volatile growth in prices followed by even longer phases (six and a half years on average) of high volatility. In the high volatility phases, expected real housing price growth is not necessarily negative but price declines happen to be much more likely. In fact, when we depict the filtered probabilities of staying in the high volatility state (obtained as a by-product of the estimation) together with the time series of each country's growth in real housing prices, we find that the switches to the high volatility state typically precede or coincide with busts in the usual sense of the word (that is, declining inflation-adjusted prices).

One possible interpretation of our results is that the larger volatility (or lower predictability) of price rises reflects the lower liquidity of the housing market (lower volume of transactions, longer average selling times, less clear expectations of price increase) during cold phases. The fact that larger volatility, rather than lower average price increases, is what best detects a cold phase in price data is consistent with the standard argument that sellers are resistant to cut down prices, especially at the beginning of a cold phase, and that inertia implies only a gradual build-up of confidence and high price growth once the cold phase gets to an end.[8] Similarly, the finding that hot phases are characterized by more predictable price rises is consistent with the arguments emphasizing the importance of price-rise expectations for price

[8] The difference in liquidity across phases could be due to a variation of the so-called *disposition effect* (the tendency of investors to ride losses and realize gains), a pattern of behavior documented for stockholders by Odean (1998), among others, and for homeowners by Genesove and Mayer (2001).

runups.

From a practical perspective, a virtue of our approach is that it allows to diagnose the cyclical position of national housing markets by looking at their inflation-adjusted price indices, which are more readily available and comparable across countries than the data on volumes or selling times. Our results on the cyclical evolution of volatility in housing markets are also important for risk management. In a section below we describe several possible applications of our model in both dimensions, including the construction of an indicator of risk in the housing market that uses a metric based on Value-at-Risk techniques. We illustrate the applications with computations referred to the immediate post-sample quarters of each of the fourteen countries in our data.

The rest of the paper is organized as follows. Section 2 elaborates on the econometric strategy and presents the models to estimate. Section 3 describes the data. In Section 4 we report and comment on the estimation results. In Section 5 we explore the implications of the results for cyclical diagnosis and risk management. Section 6 contains our concluding remarks.

2 Econometric Strategy

Besides the inherent interest of international evidence, a reason for our multi-country approach is the fear that single country experiences might not contain sufficient regime switches so as to reliably estimate the parameters of a model for each country. The idea is to profit from the cross-sectional variation in the data in order to estimate the common features of the cyclical pattern of the various national real housing price indices. Of course, the key latent assumption is that those common features exist.

A problem with the multi-country approach is heterogeneity. Geographical, historical and institutional factors may make residential property prices evolve differently in different countries. Perhaps even simple methodological differences in the construction of each country's price indices may make them show systematic differences in mean or variance. Thus, without properly controlling for the underlying heterogeneity, a regime-switching model estimated with pooled multi-country data, instead of capturing the common structure of the cyclical pattern, might end up associating the latent states with some rather cross-sectional partition of the data.[9]

Our approach to the issue is inspired in standard panel data techniques. Specifically, it is inspired in the idea of coping with unobserved heterogeneity such as, say, country fixed effects by transforming the data in a convenient way. Of course, this approach is only possible if the unobserved heterogeneity and the cyclical component enter the data generating process (DGP) in a suitably separable way. Such separability is part of the identifying assumptions necessary for the estimation of the parameters of interest.

In the rest of this section we first formally explain our approach to the identification problem using a simplified model. Then we describe the two classes of specifications actually used in the empirical part, which are straightforward extensions of the simplified model.

2.1 A Markov model with country heterogeneity

To describe our approach to the problem of identifying the common cyclical pattern of various national real housing price indices in a context with country-level heterogeneity, let us first consider a simple DGP in which the distribution of the variable of interest in country $i = 1, 2, ...N$ and quarter $t = 1, 2, ...T$ is a function of just a latent dichotomous state variable $s_{it} = 1, 2$.[10] Specifically, let y_{it} be the first quarterly difference in quarter t of the log of the real housing price index in country i and consider the following DGP:

$$y_{it} = \omega_i(s_{it}) + \sigma_i(s_{it})\varepsilon_{it}, \tag{1}$$

[9] For instance, each state might identify a group of countries whose indices share similar means and variances rather than a phase of the cycle within each country.

[10] The length of the time series need not be a common T for all the countries in the sample but, for notational simplicity, we will describe the model as if this were the case.

where ε_{it} is iid $N(0,1)$ and s_{it} follows a first-order Markov chain, independent across countries, with transition probabilities:

$$p_i \equiv Pr(s_{it} = 1 \mid s_{it-1} = 1)$$

and

$$q_i \equiv Pr(s_{it}{=}2 \mid s_{it-1}{=}2).$$

This formulation allows for a wide range of country-level heterogeneity: in each country and each state, the variable y_{it} is characterized by potentially different means, $\omega_i(s_{it})$, standard deviations, $\sigma_i(s_{it})$, and state-transition matrices:[11]

$$\left[\begin{array}{cc} p_i & 1-q_i \\ 1-p_i & q_i \end{array} \right].$$

This specification would involve estimating six parameters per country, that is, a total of $6 \times N$ parameters, which is a large number for a Markov switching model.[12] Yet the main problem with this specification is that, by allowing all parameters to vary across countries, their estimation would entirely rely on within-country variability. But, as we have already mentioned, if the time series dimension of the panel is not large enough, the number of regime switches within each country may be too small, limiting the reliability of the estimates of parameters such as p_i and q_i. This problem would be mitigated if one could impose the restriction that parameters such as p_i and q_i are the same in several or all countries and estimate them accordingly.

Within this logic, we propose a parameterization of the country-specific and cyclical components of the DGP described in (1) that will allow us to exploit cross-country variability in the estimation of the cyclical components. In particular, we propose to decompose each country's state-dependent mean and variance in two parts: a part

[11] Extending the model to cases in which $\omega_i(s_{it})$ or $\sigma_i(s_{it})$ are additionally functions of some vector of observable variables x_{it} would be immediate.

[12] In this first specification, however, the estimation could be done country-by-country, decomposing the problem in N estimations of six parameters each. Such simplification is not possible once some commonality or cross-country restrictions are introduced.

that is country-specific but exhibits no state-dependency and another that is state-dependent but has the same parameters in all countries. In particular, our proposal is to constrain (1) by imposing:

$$\omega_i(s_{it}) = \alpha_i + \mu(s_{it}) \cdot \sigma_i, \tag{2}$$

$$\sigma_i(s_{it}) = \sigma_i \cdot \sigma(s_{it}), \tag{3}$$

and

$$p_i = p \text{ and } q_i = q,$$

for all i. We will identify α_i and σ_i^2 as the unconditional mean and variance of y_{it} in each country i by assuming

$$E[\mu(s_{it})] = 0 \text{ and } E[\sigma(s_{it})] = 1 \tag{4}$$

for all i.

Now, if the country-specific moments α_i and σ_i were known, it would be possible to define the standardized transformation, z_{it}, of the original variable of interest, y_{it}:

$$z_{it} = \frac{y_{it} - \alpha_i}{\sigma_i}, \tag{5}$$

and to rewrite the DGP as follows:

$$z_{it} = \mu(s_{it}) + \sigma(s_{it})\varepsilon_{it}, \tag{6}$$

where ε_{it} is iid $N(0,1)$ and s_{it} follows a first-order Markov chain with a state-transition matrix:

$$\begin{bmatrix} p & 1-q \\ 1-p & q \end{bmatrix},$$

common to all countries. Under this formulation, the DGP of the standardized quarterly growth rate of real housing prices, z_{it}, has just six parameters, while the standardization of the original variable y_{it} additionally requires knowing α_i and σ_i for $i = 1, 2, ...N$ (that is, $2 \times N$ country-specific parameters).

In practice, the country-specific moments α_i and σ_i are not known but, for a sufficiently large T, they could be accurately estimated through the corresponding sample moments, which converge in probability to the distributional moments by virtue of standard asymptotics. In the presence of serial correlation, it is harder to ensure that the available T is large enough for the sample moments to be good substitutes for the true values of α_i and σ_i, but we will proceed as if they were. That is, we will construct a sample-based standardized series z_{it} for each country i and estimate the parameters of (6) using the maximum likelihood methods applied to standard Markov switching models.[13] In our inference about (6), we will not take into account the sampling variability in our estimates of α_i and σ_i, but, partly correcting for this, we will not impose the constraints (4) in the estimation.[14]

In the following two subsections we describe the models that we will actually estimate, which imply straightforward generalizations of the term $\mu(s_{it})$ in (6). The other components of the DGP will remain exactly as explained above.

2.2 Autoregressive models

Following Hamilton (1989), we allow for autocorrelation in the evolution of the standardized quarterly growth rate of real housing prices by considering the process:

$$z_{it} = \mu(s_{it}) + \phi\left[z_{it-1} - \mu(s_{it-1})\right] + \sigma(s_{it})\varepsilon_{it} \tag{7}$$

where ε_{it} and s_{it}, are specified as in (6). The autocorrelation parameter ϕ measures the contribution of the deviation of z_{it} with respect to its (state-contingent) mean to the prediction of the corresponding deviation one period ahead. As in all specifications mentioned so far, we will treat ε_{it} and s_{it} as independently distributed across countries. Allowing for country interdependencies is an interesting extension left for future work.

[13] With a slight abuse of notation, we will denote by z_{it} both the true and the sample-based standardizations of y_{it}.

[14] With this approach, the standard maximum likelihood algorithm a la Hamilton (1989) only requires minimal adaptations to account for the fact that our panel involves N independent realizations (i.e., N country time series) of the latent state process.

When estimating (7), we can test whether the means $\mu(s_{it})$ and the variances $\sigma^2(s_{it})$ significantly differ across states, as well as whether the autoregressive parameter ϕ is significantly different from zero.

2.3 Augmented models

We can expand the autoregressive process described in (7) by allowing its mean to depend not only on the latent state variable s_{it} but also on a vector of predetermined explanatory variables, x_{it-1}. In particular, we can replace $\mu(s_{it})$ by

$$\mu(s_{it}, x_{it-1}) = c(s_{it}) + \beta(s_{it})x_{it-1},$$

so that the DGP becomes:

$$z_{it} = c(s_{it}) + \beta(s_{it})x_{it-1} + \phi\left[z_{it-1} - c(s_{it-1}) - \beta(s_{it-1})x_{it-2}\right] + \sigma(s_{it})\varepsilon_{it}, \tag{8}$$

where the specification of ε_{it} and s_{it} remains unchanged.

The vector x_{t-1} may include different lags of predictors of the growth of real housing prices, such as the rates of GDP growth, unemployment, interest, and inflation, whose impact on the expectation of z_{it} is, in principle, allowed to be different across states. The precise definition of the explanatory variables and the way in which they enter the various estimated models is further explained below.

3 Data

For the estimation of the autoregressive specifications, we rely exclusively on country series of Inflation-adjusted Residential Property Prices that come from calculations made by the Bank for International Settlements (BIS) based on national data. The BIS kindly provided us with quarterly series covering a period of about thirty years (1970-2003) for the fourteen developed countries listed on Table 1. The table also contains the periods covered by the available time series in each country and some descriptive statistics of the variable of interest: the quarterly growth rate of real housing prices, y_{it}. This variable is computed as the non-annualized first quarterly difference of the log of the indices provided by the BIS —the units are percentage points as we have multiplied the transformation by 100 to facilitate its interpretation.

Table 1. Descriptive statistics of the variable of interest
Quarterly growth rate of real housing prices, y_{it}
(quarterly percentage rates)

	# Obs.	Sample period	Mean	S.D.	Max	Min
Australia	133	1970.2-2003.2	0.71	2.26	8.30	-5.53
Belgium	90	1981.2-2003.3	0.54	2.66	5.62	-8.59
Canada	134	1970.2-2003.3	0.48	3.03	8.66	-10.39
Denmark	133	1970.2-2003.2	0.30	2.91	10.81	-7.35
Finland	102	1978.2-2003.3	0.45	3.27	10.20	-8.77
Ireland	109	1976.2-2003.2	0.92	3.06	9.87	-6.75
Netherlands	133	1970.2-2003.2	0.72	3.23	11.44	-9.50
New Zealand	55	1990.1-2003.3	0.33	0.95	3.05	-1.50
Norway	134	1970.2-2003.3	0.39	2.88	9.07	-5.69
Spain	66	1987.2-2003.3	1.37	2.35	6.07	-6.55
Sweden	134	1970.2-2003.3	0.05	2.62	8.32	-8.97
Switzerland	134	1970.2-2003.3	0.05	2.22	6.31	-6.29
United Kingdom	133	1970.2-2003.2	0.95	3.28	12.51	-5.85
United States	134	1970.2-2003.3	0.43	0.99	3.09	-2.50

Note: The original indices of Inflation-adjusted Residential Property Prices are BIS calculations based on national data. Growth rates are computed as log differences. # Obs.: Number of observations. S.D.: Standard deviation.

Table 1 reveals a large cross-country variation in the pattern of growth of real housing prices. The mean quarterly growth rate is as low as 0.05% in Sweden and Switzerland, while it is close to or above 1% in Ireland, Spain and the UK; seven countries have their mean rate in the intermediate 0.30-0.60% range. The dispersion of quarterly growth rates around their country means is also quite heterogeneous, ranging form slightly below 1% (in New Zealand and the US) to over 3% (in Canada, Finland, Ireland, the Netherlands, and the UK); in eight countries the standard deviation of y_{it} is in the 2-3% range. A maximum quarterly growth rate as high as 12.51% appears in the UK time series, and a minimum as low as -10.39% appears in the Canada time series. Interestingly, many countries' maximum and/or minimum differ from the country's mean by more than three standard deviations, which, given the length of the available series, is more than what one would expect if the unconditional distribution of y_{it} in each country were normal. However, the regime switching models that we estimate below imply that the unconditional distributions of y_{it} are mixtures of normals, which may capture the fat tails in the data.

We use the country means and standard deviations reported in Table 1 in order to transform the original series y_{it} into the country-standardized series z_{it}. By construction, z_{it} exhibits a sample mean of zero and a variance of one in each country's time series and in the overall sample.

As mentioned above, in the augmented models (8), the prediction equation for z_{it} will include a vector of lags of other variables, x_{it-1}. Specifically, we will use lags of the quarterly growth rate of real GDP, gdp_{it}, the yearly variation in the unemployment rate, Δu_{it}, and the long-term nominal interest rate, r_{it}. The data for the construction of these variables come from the quarterly country series of the OECD Economic Outlook.[15] Table 2 reports the descriptive statistics of these variables for each country, computed over each country's sample period.

4 Results

In this section we first report on the estimation of the autoregressive and augmented models for the standardized quarterly growth rate of real housing prices, z_{it}, that we have described in subsections 2.2 and 2.3, respectively. We also comment on the various possible specifications, the criteria used in order to select among them, and the parameter estimates of our preferred specifications. Secondly, we further dissect the results by describing their implications in terms of the cyclicality of the *original* quarterly growth rate of real housing prices, y_{it}, in each country i. Finally, we focus on two particular countries, the UK and the US, in order to illustrate the applicability of

[15] This source was also used to construct the lags of the inflation rate, the short-term nominal interest rate, and the real interest rates included in some non-reported intermediate specifications. An exception is Ireland for which the long-term nominal interest rate of the first quarters in the sample is the lending rate reported in the International Financial Statistics of the International Monetary Fund.

the estimates of the probability of being in one state or another (which is a by-product of the estimation) for the identification and dating of the cyclical phases registered in their housing markets.

Table 2. Descriptive statistics of the explanatory variables
Variables in the vector x_{it-1}
(quarterly percentage rates)

	gdp		Δu		r	
	Mean	S.D.	Mean	S.D.	Mean	S.D.
Australia	0.81	1.07	0.10	1.15	9.71	3.06
Belgium	0.50	0.69	0.05	0.95	8.30	2.79
Canada	0.80	0.88	0.07	1.17	8.84	2.49
Denmark	0.46	1.04	0.10	0.96	11.52	4.74
Finland	0.65	1.04	0.13	1.86	8.97	2.62
Ireland	1.28	1.23	-0.09	1.69	10.09	3.76
Netherlands	0.65	1.07	0.04	0.96	7.48	1.70
New Zealand	0.66	1.25	-0.08	1.17	7.90	1.98
Norway	0.84	1.96	0.05	0.67	8.72	2.73
Spain	0.78	0.87	-0.25	1.69	9.26	3.55
Sweden	0.51	1.18	0.05	0.92	9.42	2.68
Switzerland	0.35	0.85	0.08	0.51	4.65	1.12
United Kingdom	0.57	1.00	0.07	1.24	9.80	2.92
United States	0.74	0.87	0.05	1.08	7.94	2.35

Note: The original data come from the OECD Economic Outlook. gdp is the quarterly growth rate of real GDP, Δu is the yearly variation in the unemployment rate, and r is the long-term nominal interest rate. S.D.: Standard deviation.

4.1 Estimating the models for z_{it}

Our estimation results are summarized in Table 3. Model 1 is actually a linear, first-order autoregressive model included as a benchmark (arbitrarily, we denote its single state by s=1). Given that z_{it} is a standardized variable and this model allows for just one constant term and one variance, the only parameter of interest in the estimation is the autoregressive coefficient ϕ, which happens to be significantly different from zero, with a point estimate 0.37.

Table 3: Estimation results
(standard errors in parenthesis)

	Model 1	Model 2	Model 3	Model 4	Model 5
$c(1)$	0.003	0.382	0.116	0.391	0.430
	(0.023)	(0.059)	(0.068)	(0.097)	(0.092)
$c(2)$	—	-0.711	-0.055	—	0.599
		(0.095)	(0.063)		(0.121)
gdp_{-1}	—	—	—	0.071	0.062
				(0.019)	(0.019)
$\Delta\mu_{-1}$	—	—	—	-0.153	-0.128
				(0.029)	(0.028)
r_{-1}	—	—	—	-0.110	-0.110
				(0.024)	(0.022)
r_{-4}	—	—	—	0.060	0.046
				(0.025)	(0.023)
ϕ	0.373	0.193	0.394	0.292	0.314
	(0.023)	(0.028)	(0.026)	(0.024)	(0.026)
$\sigma^2(1)$	0.859	0.703	0.319	0.814	0.339
	(0.030)	(0.028)	(0.038)	(0.029)	(0.031)
$\sigma^2(2)$	—	=	1.194	—	1.212
			(0.087)		(0.082)
p	1	0.962	0.955	1	0.957
		(0.010)	(0.013)		(0.012)
q	—	0.923	0.970	—	0.962
		(0.017)	(0.011)		(0.012)
Log-likelihood	-682.25	-645.47	-615.62	-639.12	-571.62

Parameters are defined as in equations (7) and (8). Explanatory variables are defined in Table 2. =: The same as in state 1.

Models 2 and 3 are autoregressive, Markov switching models with two states $s = 1, 2$ of the class described in (7). Model 2 imposes the variance $\sigma^2(s_{it})$ to be equal in both states. Its estimation unveils two regimes with different means asymmetrically positioned around zero: a high (low) growth state in which the estimated mean has a magnitude of about 38% (-71%) of the unit unconditional standard deviation of z_{it}.[16] The high (low) growth state happens with unconditional probability

[16]Given that z_{it} is the result of country-by-country standardization (that is, has zero mean and unit variance), its units of measurement are "standard deviations" of the original country variable

of 67% (33%) and has an expected duration of about 26 (13) quarters. As expected, allowing for regime switching in the mean reduces the quantitative importance of the autoregressive coefficient ϕ relative to Model 1.

Model 3 extends Model 2 by allowing the variance of the process of z_{it} to also differ across states. It turns out that the estimated difference in variances between states is sizable and very significant, while the difference in means between states is substantially smaller than when a single variance is imposed. Actually, the likelihood ratio (LR) test (run using the fact that Model 2 is nested in Model 3) allows us to reject the hypothesis that the variance is equal across states. The picture that Model 3 describes involves a low volatility state and a high volatility state. The high volatility state has a variance (1.19) that almost quadruples that of the low volatility state (0.32), is on average more frequent (60% vs. 40%) and endurable (33 vs. 22 quarters) than the low volatility state, and is associated with a (slightly) lower expected growth in prices. In contrast to Model 2, the point estimate of the autoregressive coefficient ϕ in (7) is as high as it was in Model 1, possibly because the difference in means across states is now smaller and, thus, state persistence plays a smaller role in explaining the persistence of z_{it}.

In Models 4 and 5, we incorporate a vector of explanatory variables x_{t-1} with lags of the quarterly growth rate of real GDP, gdp_{it}, the yearly variation in the unemployment rate, Δu_{it}, and the long-term nominal interest rate, r_{it}. In fact, Model 5 corresponds to the best of the augmented specifications that we tried, whereas Model 4 is just the single-state version of such specification, which we include in order to assess whether the predictive role of the variables in x_{t-1} changes much when moving from a standard linear approach to a Markov switching approach.

The set of regressors in Model 5 was determined after considering both wider and narrower combinations of regressors formed by lags of gdp_{it}, r_{it}, the unemployment rate, and other variables (such as the inflation rate or the short-term interest rate)

y_{it}. Thus, the coefficients $c(1)$ and $c(2)$ measure by how many of these standard deviations y_{it} tends to be above (if positive) or below (if negative) its unconditional mean in each state. The same units of measurement would apply to the standard deviations, $\sigma(1)$ and $\sigma(2)$, attributed to each state.

that turn out to add no significant predictive power when included together with the current regressors. Model 5 includes the yearly variation in the unemployment rate, Δu_{it}, after checking that, in alternative models, the coefficients of the first and fifth lags of the unemployment rate got point estimates with very similar absolute value and opposite signs, and after passing a LR test on the constraint imposed by including its yearly difference. We also allowed the coefficients of the explanatory variables to be state-dependent, but the restriction of them being equal across states was not rejected by the corresponding LR test.[17]

Qualitatively, the Markov switching ingredients of the augmented Model 5 are very similar to those of the autoregressive Model 3. Again, the two identified states differ mainly in their variance, which in the high volatility state almost quadruples that of the low volatility state. Now, however, the high and the low volatility states are closer in terms of unconditional probability of occurrence (53% vs. 47%, respectively) and expected duration (26 vs. 23 quarters). Somewhat surprisingly, the point estimate of the state-contingent intercept $c(s)$ is higher for the high volatility state (s=2) than for the low volatility state (s=1). However, this does not contradict the fact that the high volatility state tends to be associated with a lower expected growth in prices (as found in Model 3), since the conditional expectation of z_{it} is now a function of the regressors, such as gdp_{it-1}, Δu_{it-1}, r_{it-1}, and r_{it-4}, which are probably correlated with the state variable s_{it}.

As for the explanatory variables, the sign and significance of their coefficients are in line with previous empirical studies, despite they did not use a Markov switching approach. This is not surprising since, eventually, the detected Markov structure does more in explaining the variance of z_{it} than its mean. In fact, the estimation of Model 4 evidences that neglecting the underlying Markov switching structure has only marginal effects on the point estimates of the autoregressive parameter ϕ and the parameters associated with the other regressors. However, for crisis forecasting and

[17] We also tried specifications based on country-standardized transformations of the explanatory variables, reaching similar results as in Model 5, both in terms of the characterization of the Markov switching structure and in the sign and significance of the coefficients of the explanatory variables.

risk management in the residential property sector, the Markov switching structure uncovered by Model 5 (which happens to affect mainly to the variance of z_{it}) is absolutely crucial, since both activities are more concerned about the fatness of the tails of the distribution of future values of z_{it} than with their point forecasts (see Section 5 below).

4.2 Implications of the results for country-level cyclicality

Before starting, notice that the two-state Markov switching approach has led to identify two phases in the dynamics of the standardized growth of real housing prices that differ more in volatility than in mean. So, rather than a simple story of booms (high price growth) and busts (low or even negative price growth), the analysis identifies a cyclical pattern characterized by the alternation of phases of less volatile (more predictable) growth and phases of more volatile (less predictable) growth. The latent state variable s seems to capture the differences between the "hot" phases of the housing cycles (possibly characterized by more confidence, more speculative demand, a higher volume of transactions, and shorter average selling times) and the "cold" phases (possibly characterized by more uncertainty, less speculative demand, a lower volume of transactions, and longer average selling times). In fact, if price growth features inertia (as it is the case), the housing market is likely to switch between these phases before price growth visibly changes its trend, so hot and cold phases of housing cycles do not need to be equivalent to booms and busts, in the usual sense.

Table 4 summarizes the country-level implications of the results. In the columns devoted to Model 3, we express the consequences of this model in terms of the state-contingent means, $\omega_i(s_{it})$, and standard deviations, $\sigma_i(s_{it})$, of the *original* quarterly growth rate of real housing prices, y_{it}, in each country i. Recall from (2) and (3) that estimates of $\omega_i(s_{it})$ and $\sigma_i(s_{it})$ can be recovered as a convolution of the country-specific unconditional means and standard deviations reported in Table 1 and the country-invariant parameter estimates in Table 3.

Table 4. Country-level implications of the results
Conditional means and S.D. of y_{it}
(quarterly percentage rates)

	Model 3				Model 5						
	Mean		S.D.		Mean					S.D.	
	s=1	s=2	s=1	s=2	*Shift*	$gdp^{\dagger}$	$\Delta u^{\dagger}$	$r^{\dagger}$	$r^{\ddagger}$	s=1	s=2
Australia	0.97	0.59	1.28	2.47	-0.38	0.15	-0.33	-0.76	-0.44	1.32	2.49
Belgium	0.84	0.39	1.50	2.90	-0.45	0.11	-0.32	-0.82	-0.47	1.55	2.93
Canada	0.83	0.32	1.71	3.31	-0.51	0.17	-0.45	-0.83	-0.48	1.76	3.34
Denmark	0.63	0.13	1.64	3.18	-0.49	0.19	-0.36	-1.52	-0.88	1.69	3.20
Finland	0.82	0.26	1.84	3.57	-0.55	0.21	-0.78	-0.94	-0.55	1.90	3.60
Ireland	1.28	0.76	1.73	3.43	-0.52	0.23	-0.66	-1.27	-0.74	1.78	3.37
Netherlands	1.09	0.54	1.82	3.53	-0.55	0.21	-0.40	-0.60	-0.35	1.88	3.56
New Zealand	0.44	0.28	0.54	1.04	-0.16	0.07	-0.14	-0.21	-0.12	0.55	1.05
Norway	0.72	0.23	1.62	3.14	-0.49	0.35	-0.25	-0.86	-0.50	1.68	3.17
Spain	1.64	1.24	1.32	2.56	-0.40	0.13	-0.51	-0.92	-0.53	1.37	2.59
Sweden	0.35	-0.10	1.48	2.87	-0.44	0.19	-0.31	-0.77	-0.45	1.53	2.88
Switzerland	0.31	-0.07	1.25	2.42	-0.38	0.12	-0.14	-0.27	-0.16	1.29	2.44
UK	1.33	0.77	1.85	3.58	-0.55	0.20	-0.52	-1.05	-0.61	1.91	3.61
US	0.54	0.37	0.56	1.08	-0.17	0.05	-0.14	-0.26	-0.15	0.58	1.09

Note: Computations based in Tables 1 and 2, and the parameter estimates of Models 3 and 5 in Table 3. The variables are defined in Tables 1 and 2. S.D.: Standard deviation. *Shift*: Difference between the estimates of the constant terms in state 1 and state 2, c(1)-c(2). †: Effect of a one-S.D. increase in the variable after one quarter. ‡: Effect of one-S.D. permanent increase in r after four quarters.

The results are self-explanatory. Depending on country specificities, state variability may imply reaching more extreme or more moderate values in the conditional mean and variance of real housing price increases. According to Model 3, the hot phase ($s = 1$) implies expected quarterly growth rates close to or above 1% in Australia, Ireland, the Netherlands, New Zealand, Norway, Spain, and the UK, while the cold phase ($s = 2$) implies expected quarterly growth rates close to zero or negative in Denmark, Sweden, and Switzerland. Perhaps more importantly, the combination of low state-contingent means and high state-contingent standard deviations in the cold phase implies that negative quarterly growth rates are very likely in such a phase in essentially all countries. In the hot phase, such a risk is smaller but still significant—we will readdress this point in Section 5.[18]

[18] The models have been estimated under normality assumptions, recall (1). Under these assump-

The columns devoted to Model 5 are elaborated in a similar manner. The main difference with respect to the columns on Model 3 is that Model 5 specifies the mean of y_{it} as a function of the state variable s and lags of observable variables such as the quarterly growth rate of real GDP, *gdp*, the yearly variation in the unemployment rate, Δu, and the long-term nominal interest rate, r. The columns referred to these variables report the effects on the mean of y_{it} in each country of a one-standard-deviation shock to the corresponding country variable. Notice that there are two columns devoted to the long-term interest rate: the first from the left reports the impact of a one-standard-deviation shock after one quarter; the second contains the net accumulated impact after four quarters, which is smaller in absolute value, reflecting that slightly less than half of the initial effect of a permanent change in nominal interest rates on the growth of real housing prices is transitory, disappearing within one year. In spite of this and with slight variations across countries, it seems that the quantitatively most important effects are those associated with the level of nominal long-term interest rates, followed by those associated with the yearly variation in unemployment, and finally by those associated with GDP growth.

4.3 Two country cases: the UK and the US

The maximum likelihood estimation of a Markov switching model generates, as a by-product, an estimate of the probability with which the latent state variable takes each of its possible values in each observation. These filtered probabilities can be very useful for cyclical diagnosis, that is, the identification and dating of the cyclical phases registered in the analyzed time series. In this section we will illustrate this use by focusing on two of the countries in our sample: the UK and the US. These are good examples first because their residential property markets have been extensively studied before and, so, their booms and busts, in the common sense of these words,

tions, it is immediate to compute the state-contingent quantiles of y_{it} using the means and standard deviations in Table 4 and the cdf of a standard normal random variable.

are well known.[19] A second reason to look at these two countries is that, over the last few decades, the means and variances of the growth of their real housing prices have been very different (see Table 1), so that comparing them may help us to assess the success of our "standardization approach" in dealing with the heterogeneity problem.

Figure 1 refers to the UK experience. As reflected by the dashed lines in the lower panel, Model 5 detects likely cold phases in years 1970-1977, 1980-1983, 1989-1993, and from 2002 to the end of the sample. These intervals can be compared with the periods of falling real housing prices (or busts) that one could mention in view of the solid line: 1973-1977, 1980-1982, and 1990-1995. The comparison suggests that both types of phases largely overlap, although the start of cold phases tends to precede the arrival of a bust, while the start of hot phases sometimes leads and sometimes lags the end of periods of real price declines.

Aggregate fluctuations in housing prices are far more moderate in the US than in the UK: the real price changes shown in Figure 2 have much smaller ranges of variation than their Figure 1 counterparts. This fact is related to the predominantly regional character of the fluctuations registered in the US housing market during the period of analysis.[20] Yet both the original series of real housing price changes and the filtered probabilities of the high volatility state obtained with Models 3 and 5 suggest the existence of distinct phases in national-level dynamics. Specifically, the solid lines show housing busts in the year intervals 1973-1976, 1979-1983, and 1990-1995, while the inference from Model 5 suggests a virtually uninterrupted long cold phase covering years 1970-1983 and another much shorter cold phase in 1990-1992. As just described for the UK, the start of cold phases tends to coincide or lead that of bust periods, while cold phases and busts tend to end around similar dates.

[19] See Bordo and Jeanne (2002), and Helbling (2005) for a description of housing booms and busts in these and other countries in our sample using standard business cycle dating techniques.

[20] See, for example, Abraham and Hendershott (1996).

United Kingdom

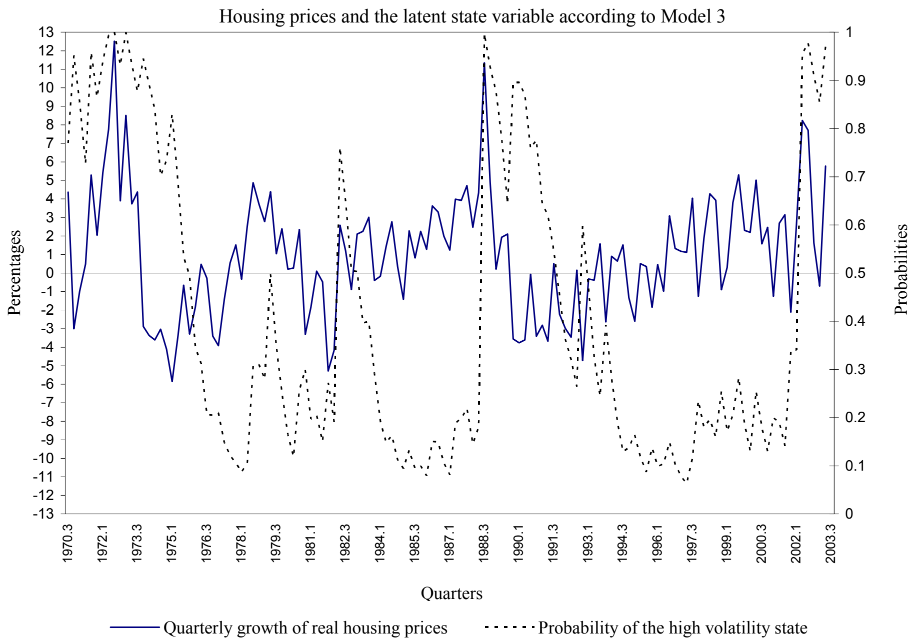

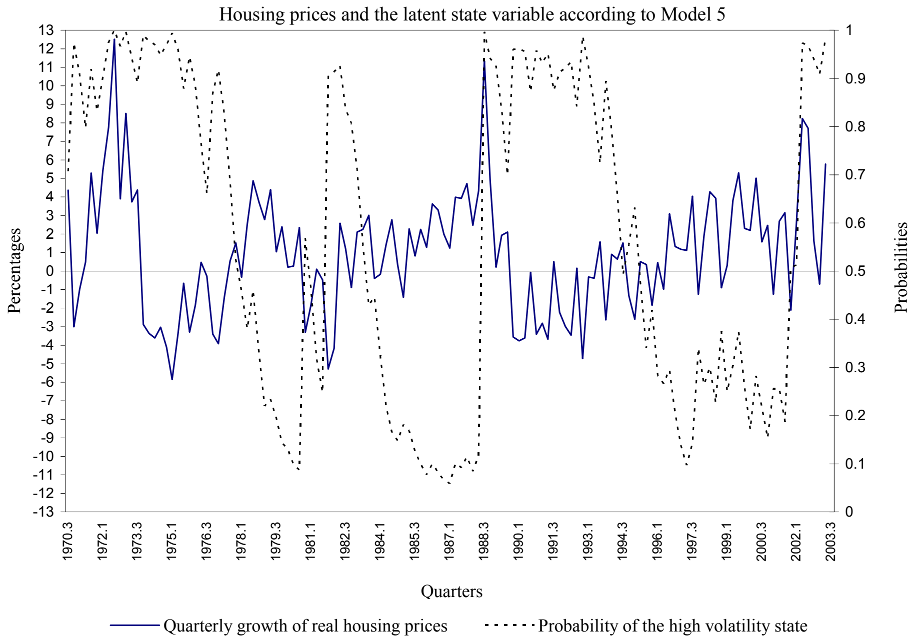

Figure 1. Identification of cold phases in the UK housing market

United States

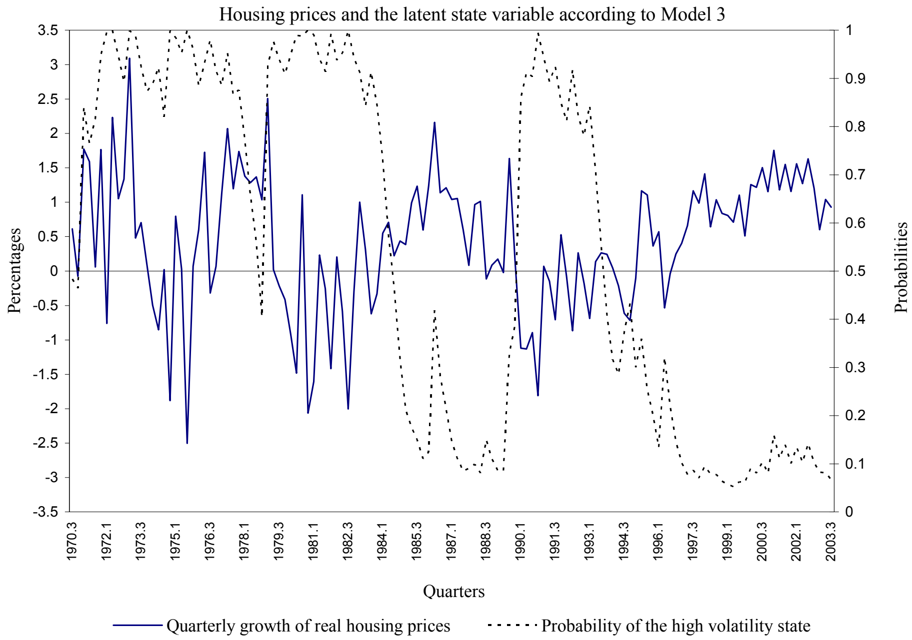

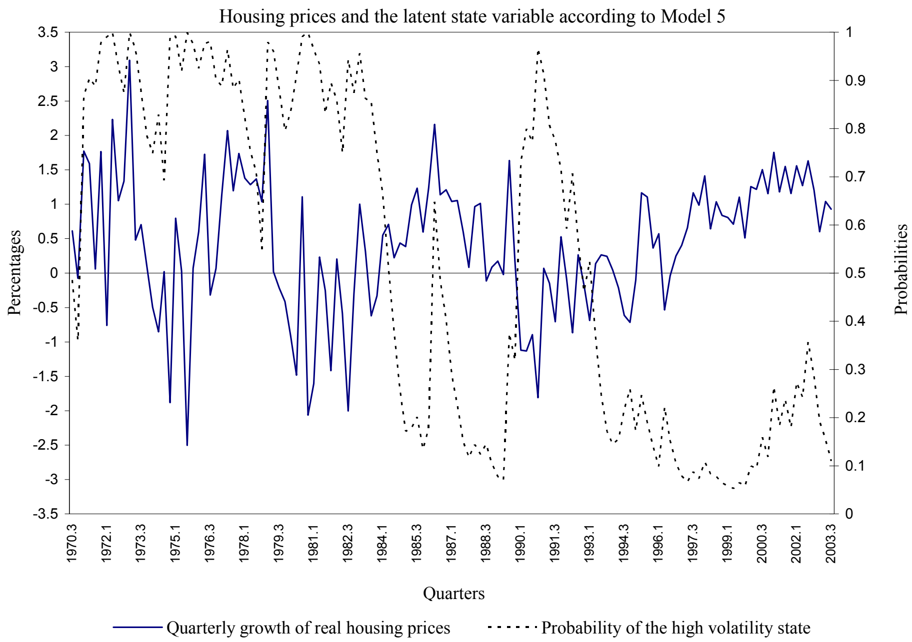

Figure 2. Identification of cold phases in the US housing market

5 Risk Assessment Applications of the Results

The applications described in this section are based on Model 5 and the inferences about its parameters and the values of the latent state variable that emerge from its estimation. Following the standard risk-management practice, we abstract from estimation error and model uncertainty —that is, we will take the point estimates of the parameters and the probabilities of the state variable taking one value or another as if they were the true ones.[21]

In these applications, we use the maximum likelihood estimates of the probability of having reached the high-volatility state in the last sample quarter T in each country i, together with the estimated state transition probabilities, in order to evaluate the probability of the staying in a cold phase (s=2) in the first post-sample quarter $T+1$ conditional on the information available in quarter T, $Pr(s_{iT+1}{=}2 \mid \Omega_T)$. We also use an estimate of the cumulative density function (cdf) of the growth rate of real property prices in country i and quarter $T+1$ conditional on the information available in quarter T, $F(y_{iT+1} \mid \Omega_T)$. In the Appendix, we describe in detail the computation of $Pr(s_{iT+1}{=}2 \mid \Omega_T)$ and $F(y_{iT+1} \mid \Omega_T)$, which is complicated by the non-linearity associated with the dynamics of the latent state variable and its interaction with the autoregressive component of the model—recall equation (8). We show how $F(y_{iT+1} \mid \Omega_T)$ can be written as a mixture of four normal cdfs in which the mixing probabilities depend on the inferred probabilities of staying in each possible phase of the cycle in T and $T-1$ and the state-transition probabilities

The first column of results in Table 5 illustrates the possibility of using the model in order to forecast the cyclical position of each housing market. It reports our estimate of the probability of staying in a cold phase at $T+1$ conditional on the information available at T, $Pr(s_{iT+1}{=}2 \mid \Omega_T)$, where T is the last quarter for which we had data on real housing price increases in each country i (see Table 1). The results imply that in mid-to-end 2003, ten out of the fourteen national housing markets under analysis

[21] Extending risk-management techniques to account for estimation error and model uncertainty is an interesting and active area of research, but lies beyond the scope of this paper.

are predicted to be in a hot phase of their cycles. The exceptions are Australia and the UK, where the cold phase is strongly predicted, as well as Belgium and Norway, with a somewhat weaker phase assignment.

Table 5. Risk assessment applications

Predictions for quarter $T+1$ based on information available in quarter T

(quarterly percentage rates)

	$Pr(s_{iT+1}=2 \mid \Omega_T)$	$Pr(y_{iT+1}<0 \mid \Omega_T)$	$\text{VaR}_{99\%}(y_{iT+1})$
Australia	90.24	22.38	3.89
Belgium	66.38	31.20	5.08
Canada	26.34	36.52	5.00
Denmark	11.04	34.66	3.87
Finland	10.60	32.06	4.18
Ireland	43.11	26.72	5.06
Netherlands	13.49	36.04	4.53
New Zealand	35.30	12.67	1.14
Norway	63.49	40.18	6.04
Spain	16.15	15.06	2.40
Sweden	26.85	44.11	4.69
Switzerland	9.92	37.55	3.00
United Kingdom	93.51	28.49	6.31
United States	14.44	15.34	0.99

Note: These predictions are based on the estimates of Model 5, in Table 3. $Pr(s_{iT+1}=2 \mid \Omega_T)$ is the probability of the high-volatility state at $T+1$; $Pr(y_{iT+1}<0 \mid \Omega_T)$ is the probability of a real price decline at $T+1$; $\text{VaR}_{99\%}(y_{iT+1})$ is the 99%-quantile of $-y_{iT+1}$. All these estimates are conditional on the information available at T.

The second column of results reports our country-level estimates of the probability of suffering a housing market bust in the first post-sample quarter. By a bust we simply mean a decline in the corresponding real housing price index, so we report $Pr(y_{iT+1}<0 \mid \Omega_T)$, which can be obtained by evaluating $F(y_{iT+1} \mid \Omega_T)$ at $y_{iT+1}=0$. At first sight, the numbers may seem surprisingly large, as half of them exceed 30% (although none exceeds 50%). However quarterly declines in the real housing price index are quite frequent in the sample and, if these realizations are just close to zero and not persistent, they need not imply a severe risk. In other words, the conventional notion of a housing bust may not be that useful after all.

The last column in Table 5 assesses the risk of each national housing market in the first post-sample quarter using a more sophisticate metric, based on Value-at-Risk (VaR) techniques. Specifically, it reports an estimate of the critical value w that solves

$$Pr[-y_{iT+1} \leq w \mid \Omega_T] = 0.99.$$

Following risk-management standards, such critical value can be interpreted as the upper bound of a 99%-confidence interval of the form $[-\infty, w]$ for the quarterly percentage *loss* that one might suffer on an investment indexed to real housing prices in country i that were hold from quarter T to quarter $T+1$. Recall that, according to our model, the prediction error of the quarterly variation in each index of real housing prices is generated by a mixture of four normal distributions.[22] So the underlying loss distribution is much richer than in the traditional JP Morgan RiskMetrics approach (which simply assumes normality) and may accommodate asymmetry and kurtosis, making both of them implicitly dependent on the dynamics of the latent state variable.

Table 5 shows that, for the first post-sample quarter, countries such as Belgium, Canada, Ireland, Norway, and the UK could foresee maximum real housing price declines, with a 99% confidence level, in the range of 5-6%. The numbers for Australia, Denmark, Finland, the Netherlands and Sweden are in an intermediate 3-5% range, while for New Zealand, Spain, Switzerland and the US, they are in the most moderate 1-3% range. These results draw a picture that is consistent with the pictures offered in the previous two columns in the table, but not exactly equivalent. For instance, the slight differences between Australia and the UK in terms of the first two indicators, becomes much more visible in terms of the VaR-type indicator. Also, the VaR-type indicator combines complementary ingredients of the other two. For instance, in Belgium, the Netherlands, Norway or Switzerland the risk of a decline in the real

[22] w is simply the 1% quantile of the conditional distribution $F(y_{iT+1} \mid \Omega_T)$, whose derivation is described in the Appendix.

housing price index by the end of 2003 was higher than in the UK, but in terms of the VaR indicator their risks were more modest than in the UK (the riskiest housing market at that point in time). One reason for these differences is that Belgium, the Netherlands, Norway, and Switzerland were less likely to be in (or enter) a high volatility phase than the UK.

6 Conclusions

We have examined the experience of fourteen developed countries for which there are about thirty years of quarterly inflation-adjusted housing price data. Price dynamics has been modeled as a combination of a country-specific component and a cyclical component. For the cyclical component we have postulated a two-state Markov switching process with parameters common to all countries. Our main finding is that the latent cyclical state variable captures previously undocumented changes in the volatility of real housing price increases. Housing cycles feature high and low volatility phases that are quite persistent (about six years, on average) and occur with about the same unconditional frequency over time. These findings have important implications both as formal evidence on the cyclical pattern of housing markets and for cyclical diagnosis and risk management in regard to these markets.

From a technical perspective, a novelty in our regime-switching modeling proposal is to have adopted a multi-country approach. Looking at several countries is a way to learn about the common cyclical patterns of their housing markets in a context where the infrequency of regime switches in each market would impare the reliability of country-by-country estimations. One can think of many robustness checks and extensions with which to enrich our basic analysis. Part of this work, however, will have to wait until longer time series are available for each country and/or for the development of better panel data techniques for regime-switching models.

The cyclical phases identified in this paper correspond widely with the cold and hot phases referred in popular and academic descriptions of the dynamics of the housing market. However, the existing literature is not explicit about the important cyclical changes in price volatility that we document: explaining them is a challenge for future theoretical research. We conjecture that, if price growth volatility reflects market liquidity, then it might be possible to extend some of the current explanations for the correlation between price growth and measures of liquidity such as the volume of transactions or average selling times in order to yield predictions on price growth volatility.

Appendix

The distribution of s_{iT+1} and y_{iT+1} conditional on Ω_T

This Appendix describes how the standard output of the estimation of a Markov switching model using the algorithm developed by Hamilton (1989) can be utilized in the applications contained in Section 5. Readers familiar with Markov switching models will find this material redundant and may skip it.

We are interested in the implications of a DGP such as that implied by Model 5 for the distribution of the state variable in country i and quarter $T+1$, s_{iT+1}, conditional on the information available in quarter T, Ω_T, as well as in the distribution of the growth rate of real property prices in country i and quarter $T+1$, y_{iT+1}, conditional on Ω_T. Given our interest in country-level objects, we can w.l.o.g. simplify the presentation by dropping the country subscript in all expressions below. In addition, we can use (5) and (8) in order to write the DGP in terms of the original growth rate of real property prices:

$$y_t = \omega(s_t, x_{t-1}) + \phi\left[y_{t-1} - \omega(s_{t-1}, x_{t-2})\right] + \sigma \cdot \sigma(s_t)\varepsilon_t, \tag{9}$$

where

$$\omega(s, x) = \alpha + \sigma[c(s) + \beta(s)x],$$

α and σ are the unconditional country-specific mean and standard deviation of y_t, $c(s)$ and $\sigma(s)$ are the state-contingent intercepts and standard deviations of the country-standardized variable z_t, x_{t-1} is the vector of explanatory variables, and $\beta(s)$ is the vector of their possibly state-dependent coefficients.[23]

First we will obtain the distribution of s_{T+1} conditional on Ω_T. As one can clearly see from (9), the autoregressive component in brackets implies that y_t is affected by the pair (s_{t-1}, s_t), as well as the vectors x_{t-2} and x_{t-1}. Thus, conditional on x_{t-2} and x_{t-1}, the DGP in (9) can be seen as a first-order Markov process referenced to a *new*, four-valued state variable $S_t = 1, 2, 3, 4$ that identifies the position of the pair (s_{t-1}, s_t) in the list of its possible realizations $\{(1,1), (1,2), (2,1), (2,2)\}$.

In fact, if we consider a vector of state probabilities $\Gamma_t = (\Gamma_{st})_{s=1,2,3,4}$, with

$$\Gamma_{st} = \Pr(S_t = s \mid \Omega_t),$$

[23] We will proceed as if all parameters were known although, for the computations included in Section 5, their values will be replaced by the corresponding maximum likelihood estimates, which appear in Tables 1 and 3. Actually, Model 5 in Table 3 imposes the constraint $\beta(1) = \beta(2)$, which was not rejected by the data, but in this presentation, and for the sake of generality, we will allow for state contingency in $\beta(s)$.

as the vector of filtered probabilities obtained for $t = 1, 2, ...T$ with Hamilton's algorithm for the estimation of Model 5, then the law of movement of Γ_t is a first-order Markov chain with transition matrix

$$A = \begin{pmatrix} p & 0 & p & 0 \\ 1-p & 0 & 1-p & 0 \\ 0 & 1-q & 0 & 1-q \\ 0 & q & 0 & q \end{pmatrix},$$

where p and q are the transition probabilities that describe the dynamics of the *original* state variable s_t. Since estimating the model yields estimates of A and Γ_T, the vector of state probabilities for the first after-sample quarter, Γ_{T+1}, can be recursively estimated as $A\Gamma_T$ and the estimates of $\Pr(s_{T+1} = 1 \mid \Omega_T)$ (or its complement $\Pr(s_{T+1} = 2 \mid \Omega_T)$) can be obtained by adding up Γ_{1T+1} and Γ_{3T+1} (or Γ_{2T+1} and Γ_{4T+1}). This is how we got the values that appear in the corresponding column of Table 5.

Next we will obtain the distribution of y_{iT+1} conditional on Ω_T. As a first step, let $F_{T+1}(y) = (F_{sT+1}(y))_{s=1,2,3,4}$ denote the vector of cdfs of the variable y_{iT+1} conditional on *both* Ω_T and $S_{T+1} = s$ and let (i, j) denote the pair (s_T, s_{T+1}) identified by $S_{T+1} = s$. Since conditional on *both* Ω_T and $(s_T, s_{T+1}) = (i, j)$, the only random term in (9) is ε_{T+1}, which is $N(0, 1)$, the components of the vector $F_{T+1}(y)$ can be expressed as

$$F_{sT+1}(y) = \Phi\left(\frac{y - \omega(j, x_T) - \phi\left[y_T - \omega(i, x_{T-1})\right]}{\sigma \cdot \sigma(j)}\right), \tag{10}$$

where $\Phi(\cdot)$ is the cdf of a $N(0, 1)$.

Obviously, the state variable S_{T+1} is not incorporated into Ω_T (and actually no S_t is, since the state variable is unobservable). However, the Total Probability Theorem allows us to write the cdf of y_{iT+1} conditional on *just* Ω_T as:

$$F(y \mid \Omega_T) = \Gamma_{T+1} \cdot F_{T+1}(y), \tag{11}$$

where $\cdot$ denotes the inner product. In words, (10) and (11) together imply that $F(y \mid \Omega_T)$ is mixture of four standard normal random variables. Evaluating (11) simply requires evaluating the vector $F_{T+1}(y)$ using x_{T-1} and x_T, and evaluating Γ_{T+1} with the recursion $A\Gamma_T$.

References

Abraham, J., and P. Hendershott (1996). "Bubbles in metropolitan housing markets," *Journal of Housing Research* 7, pp. 191-207.

Bordo, M., and O. Jeanne (2002). "Monetary policy and asset prices: Does 'benign neglect' make sense?," *International Finance* 5, pp. 139-164.

Borio, C., and P. Mcguire (2004). "Twin peaks in equity and housing prices?," *BIS Quarterly Review,* March, pp. 79-93.

Capozza, D., P. Hendershott, C. Mack, and C. Mayer (2002). "Determinants of real house price dynamics," NBER Working Paper No. 9262.

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

Case, K., and R. Shiller (2003). "Is there a bubble in the housing market?," *Brookings Papers on Economic Activity* 2, pp. 299-362.

Cho, M. (1996). "House price dynamics: A survey of theoretical and empirical issues," *Journal of Housing Research* 7, pp. 145-172.

Englund, P., and Y. Ioannides (1997). "House price dynamics: An international empirical perspective," *Journal of Housing Economics* 6, pp. 119-136.

Genesove, D., and C. Mayer (2001). "Loss aversion and seller behavior: Evidence from the housing market," *Quarterly Journal of Economics* 116, pp. 1233-1260.

Hamilton, J. (1989). "A new approach to the economic analysis of nonstationary time series and the business cycle," *Econometrica* 57, pp. 357-384.

Harding, D., and A. Pagan (2002). "Dissecting the cycle: a methodological investigation," *Journal of Monetary Economics* 49, pp. 365-381.

Helbling, T. (2005). "Housing price bubbles – a tale based on housing price booms and busts," *BIS Papers* 21, pp. 30-41.

Herring, R., and S. Wachter (1999). "Real estate booms and banking busts: An international perspective," The Wharton Financial Institutions Center, WP 99-27.

Krainer, J. (2001). “A theory of liquidity in residential real estate markets,” *Journal of Urban Economics* 49, pp. 32-53.

Muellbauer, J., and A. Murphy (1997). “Booms and busts in the UK housing market,” *Economic Journal* 107, pp. 1701-1727.

Novy-Marx, R. (2005). “The microfoundations of hot and cold markets,” mimeo, University of Chicago.

Odean, T. (1998). “Are investors reluctant to realize their losses?,” *Journal of Finance* 53, pp. 1775-1798.

Ortalo-Magne, F., and S. Rady (2005). “Housing market dynamics: On the contribution of income shocks and credit constraints,” *Review of Economic Studies*, forthcoming.

Stein, J. (1995). “Prices and trading volume in the housing market: A model with downpayment effects,” *Quarterly Journal of Economics* 110, pp. 379-406.

Tsatsaronis, K., and H. Zhu (2003). “What drives housing price dymanics: cross-country evidence”, *BIS Quartely Review*, March, pp. 65-78.

Wheaton, W. (1990). “Vacancy, search, and prices in a housing market matching model,” *Journal of Political Economy* 98, pp. 1270-1292.

Williams, J. (1995). “Pricing real assets with costly search,” *Review of Financial Studies* 8, pp. 55-90.