The Econometrics Journal

Econometrics Journal (2011), volume **14**, pp. C1–C24.
doi: 10.1111/j.1368-423X.2010.00319.x

A hierarchical factor analysis of U.S. housing market dynamics

EMANUEL MOENCH[†] AND SERENA NG[‡]

[†]*Federal Reserve Bank of New York, 33 Liberty St., New York, NY 10045, USA.*
E-mail: emanuel.moench@ny.frb.org

[‡]*Columbia University, 420 W. 118 St., MC 3308, New York, NY 10025, USA.*
E-mail: serena.ng@columbia.edu

First version received: July 2009; final version accepted: March 2010

Summary This paper studies the linkages between housing and consumption in the United States taking into account regional variation. We estimate national and regional housing factors from a comprehensive set of U.S. price and quantity data available at mixed frequencies and over different time spans. Our housing factors pick up the common components in the data and are less affected by the idiosyncratic noise in individual series. This allows us to get more reliable estimates of the consumption effects of housing market shocks. We find that shocks at the national level have large cumulative effects on retail sales in all regions. Though the effects of regional shocks are smaller, they are also significant. We analyse the driving forces of housing market activity by means of factor-augmented vector autoregressions. Our results show that lowering mortgage rates has a larger effect than a similar reduction of the federal funds rate. Moreover, lower consumer confidence and stock prices can slow the recovery in the housing market.

Keywords: *FAVAR*, *Hierarchical factor models*, *Mixed sampling frequency*, *Missing values*.

1. INTRODUCTION

This paper provides a quantitative assessment of the dynamic effects of housing shocks on retail sales. The econometric exercise consists of estimating national and regional housing factors from large non-balanced panels of data. The economic analysis consists of studying the dynamic response of national as well as regional retail sales to housing market shocks, and assessing the sensitivity of the housing factor to economic conditions and stimulus. As a by-product, we obtain estimates of 'house price factors' that summarize the common information in the different published house price indicators.

Our approach has three features. First, we make a distinction between 'national' and 'regional' housing markets in recognition of the fact that not all variations in housing are pervasive. Second, we consider shocks to the housing 'market' as opposed to shocks to home 'prices' only. The analysis thus makes extensive use of housing data rather than relying on a single measure of house prices. Third, we use diverse measures of price and volume to determine the (latent) state of the housing market. The non-balanced panel of data covers series sampled at different frequencies and that are available over different time spans.

Our analysis consists of two steps. We first use a dynamic hierarchical (multi-level) factor model to disentangle information on the housing market into national, regional and series-specific components. For each region, we embed the estimated national and regional housing factors along with other variables that control for the effects of regional business cycles into factor-augmented vector autoregressions (FAVAR). An analysis of the impulse responses then allows us to study the propagation mechanism of regional and national housing shocks.

Several considerations motivate our analysis. First, over the last few years, the U.S. housing market has experienced an extended period of expansion followed by an abrupt and pronounced downturn. Newspaper articles and the media often suggest that a housing boom stimulates, while a housing bust slows non-housing activity, in particular consumption. For example, reporting on a speech made by Federal Reserve Chairman Greenspan, the *Los Angeles Times* (March 9, 1999) wrote, 'Capital generated by a booming housing market have probably spurred consumer spending and given a strong boost to the U.S. Economy'. In its October 16, 2006, issue, *Newsweek* magazine wrote that, 'if home prices drop too much, the damage to consumer confidence and spending won't be easily offset'.

Although two-thirds of U.S. households are homeowners, evidence on the macroeconomic effects of housing shocks have been limited to point estimates from micro data for certain demographic groups or short panels, neither of which are ideal for studying the dynamic response to housing shocks at the aggregate level. There are few VARs estimated for consumption and housing at the national level and even fewer at the regional level because it is difficult to find housing and consumption data that are available for a long enough period to make estimation of a VAR suitable. We circumvent this problem by not restricting ourselves to house price data alone. We also do not restrict ourselves to data sampled at the same frequency. Instead, we extract common housing factors at the national and regional level from data on house prices and housing market quantities. We then use FAVARs to trace out the aggregate effects of housing shocks without fully specifying the structure of the housing market.

Second, the notion of a 'U.S. housing market' disregards the fact that there is substantial variation in housing activity across markets, a point also raised by Calomiris et al. (2008). Indeed, of the four regions defined by the Census Bureau, the Northeast (including New York and Massachusetts) and the West (including California, Arizona and Nevada) have historically had more active housing markets than the South (including Texas, Florida and Virginia) and the Midwest regions (including Illinois, Ohio and Michigan). Consumers in regions not used to large variations in the housing market might respond differently from those that are accustomed to housing cycles. Consumption responses might also depend on regional business cycle conditions. It is very much an empirical matter whether regional differences in the consumption response to housing shocks exist.

Third, there exist numerous measures of house prices, including the well-known Case–Shiller price index, the indices published by the Federal Housing Finance Administration (FHFA, formerly the Office of Federal Housing Enterprise Oversight or OFHEO), and indices published by the National Association of Realtors (NAR). Some series are available for time spans longer than 30 years, while others are available for a little over a decade. Calomiris et al. (2008) note that many issues surrounding empirical estimates of the wealth effect of housing relate to the definition of the house price. We address this problem by estimating a house price factor that extracts the common variations underlying all indicators of house prices, thereby filtering out idiosyncratic noise.

While price is a key indicator of the state of the housing market, data on the volume of transactions are also available. Leamer (2007) argues that the 'volume cycle' rather than the 'price cycle' is what makes housing important in U.S. business cycles. Some regional markets may be more prone to high volatility in house prices and to high volume of transactions than others. In order to assess how the activities in the housing 'market' (as opposed to house prices) affect consumption, we construct broad-based measures of the state of the housing markets at both the regional and the national level. In our analysis, this is handled using a hierarchical factor model framework.

Our results can be summarized as follows. First, we find that the national and regional factors are of comparable order of importance in three of the four regions, but regional variations are much more important in the West. Second, there is a marked difference between the house price and the housing market factor in the Midwest and the South. Third, retail consumption in all regions responds positively to a national housing shock with the peak occurring about 15 months after the shock. A two-standard-deviation shock in the housing factor can reduce aggregate consumption by as much as 8%, all else equal. However, the consumption responses are largely driven by responses to house price shocks. Shocks to volume have significantly smaller effects. A FAVAR in variables that might affect the housing factor indicates that interest rate cuts will stimulate housing activity, with reductions in mortgage rates potentially having a larger impact than similar cuts in the fed funds rate. Moreover, consumer confidence and stock prices both have a significant effect on housing market activity.

2. THE DATA

We assemble a data set that covers the four census regions Northeast, Midwest, West and South. We also have aggregate measures of housing activity for the United States, which we will refer to as 'national' data. We combine information from various sources, including federal agencies and private institutions. Instead of focusing on house price indicators alone, our data set consists of both price and volume information. This allows us to capture the different dimensions of the housing market. The data are further sampled at different frequencies: some series are monthly, and some are quarterly. Moreover, some series start as early as 1963, while some are not available until 1990. We thus face an unbalanced panel. We take January 1973 to be the beginning of our sample, and the last data entry is May 2008. The key series are listed in Tables 1 and 2. To motivate the analysis to follow, it helps to have a quick review of the data used.

As house price is the key indicator of housing wealth, most studies quite naturally use some measure of house prices to study the effect of housing market shocks on real economic activity. However, there exist several indicators of house prices published by various data providers, each employing different data sources and aggregation methods. While these house price series are correlated, each has some distinctive features. Ideally, one would therefore like to use a genuine measure of price movement. In this paper, this is taken to be variations that are common to all observed price measures.

In general terms, the literature distinguishes between three main ways to measure house prices. As discussed in Rappaport (2007), the simplest method computes the average or median of house prices observed in a period. This 'simple approach' can be volatile due to a changing composition of high and low priced units. The 'repeat sales method' focuses on houses that have been sold more than once. It provides a price index and does not measure the price level itself. Furthermore, the number of repeat transactions can be small relative to total transactions, and it is subject to continual revisions. The third is the 'hedonic approach' which uses statistical methods to control for differences in quality. As a general matter, price measures based on repeat

Table 1. Regional data.

Series	Source	Frequency	First obs.	TFCode
		Price data		
Median Sales Price of Single-Family Existing Homes	NAR	mly	Jan1968	2
Single-Family Median Home Sales Price	CENSUS	qly	Q1-1968	2
Average Existing Home Prices	NAR	qly	Q1-1970	2
Average New Home Prices	NAR	qly	Q1-1970	2
CMHPI	FHLMC	qly	Q1-1970	2
FHFA Purchase-Only Index	FHFA	mly	Jan1986	2
FHFA Home Prices	FHFA	qly	Q1-1970	2
		Volume data		
New One-Family Houses Sold	CENSUS	mly	Jan1968	1
New One-Family Houses For Sale	CENSUS	mly	Jan1968	1
Single-Family Housing Units under Construction	CENSUS	mly	Jan1980	1
Multi-Family Units under Construction	CENSUS	mly	Jan1980	1
Homeownership Rate	CENSUS	qly	Q1-1968	1
Homeowner Vacancy Rate	CENSUS	qly	Q1-1968	1
Rental Vacancy Rate	CENSUS	qly	Q1-1968	1

Note: This table reports the regional housing data used in the estimation. All variables are available for each of the four Census regions ‘Northeast’, ‘MidWest’, ‘South’ and ‘West’. The column ‘TFCode’ documents the transformations applied to the raw series prior to estimation. TFCode = 1 refers to annual growth rates, TFCode = 2 refers to annual differences of annual growth rates.

transactions are often thought to give more precise estimates of house price appreciation. Prices subject to compositional effects are believed to be better at measuring the amount required to purchase housing than at estimating the rate of house price changes.

The Federal Housing Finance Administration (FHFA) indices include only homes with mortgages that conform to Freddie Mac and Fannie Mae guidelines. Data are available at the national, regional and state levels, as well as for the major metropolitan areas. They are based on transactions and appraisals, and are then adjusted for appraisal bias. The FHFA also publishes a purchase-only index that excludes refinancing. These indices equally weight prices regardless of the value of the house. The coverage of the indices is broad because Freddie Mac and Fannie Mae provide loans throughout the country. However, the so-called jumbo loans over $417,000 are not included.

The S&P/Case–Shiller home price indices, published by Fiserv Inc., are based on information from county assessor and recorder offices. The index started with data from 10 cities in 1987 but was extended to cover 20 cities in 2000. The Case–Shiller indices do not use data from 13 states and have incomplete coverage for 29 states. Compared to the FHFA, the Case–Shiller indices thus have a narrower geographical coverage. However, homes purchased with subprime and other unconventional loans are included in the indices. As they cover defaults, foreclosures and forced sales, these indices show more volatility than the FHFA indices. Note also that the Case–Shiller indices are value weighted and hence give more weight to higher priced homes.

Table 2. National data.

Series	Source	Frequency	First obs.	TFCode
		Price data		
Median Sales Price of Single-Family Existing Homes	NAR	mly	Jan1968	2
Median Sales Price of Single-Family New Homes	CENSUS	mly	Jan1968	2
Single-Family Median Home Sales Price	CENSUS	qly	Q1-1968	2
Average Existing Home Prices	NAR	mly	Jan1994	2
Average New Home Prices	CENSUS	mly	Jan1970	2
S&P/Case–Shiller Home Price Index	S&P	qly	Q1-1982	2
CMHPI	FHLMC	qly	Q1-1968	2
FHFA Purchase-Only Index	FHFA	mly	Jan1986	2
FHFA Home Prices	FHFA	qly	Q1-1970	2
		Volume data		
New One-Family Houses For Sale	CENSUS	mly	Jan1968	1
Housing Units Authorized by Permit: One-Unit	CENSUS	mly	Jan1968	1
Multi-Family Units under Construction	CENSUS	mly	Jan1968	1
Multi-Family Permits United States	CENSUS	mly	Jan1968	1
Multi-Family Starts United States	CENSUS	mly	Jan1968	1
Multi-Family Completions	CENSUS	mly	Jan1968	1
Homeowner Vacancy Rate	CENSUS	qly	Q1-1968	1
Homeownership Rate	CENSUS	qly	Q1-1968	1
Rental Vacancy Rate	CENSUS	qly	Q1-1968	1

Note: This table reports the national housing data used in the estimation. The column 'TFCode' documents the transformations applied to the raw series prior to estimation. TFCode = 1 refers to annual growth rates, TFCode = 2 refers to annual differences of annual growth rates.

The FHFA and the Case–Shiller indices are both based on repeat sales. In contrast, the NAR report the mean/median purchase prices of homes directly. The NAR represents real estate professionals and has close to 2000 local associations and boards offering multiple listing services. The NAR surveys a fixed subset of its associations. Based on reported transactions from the sample, the NAR calculates a median price for each of the four Census Bureau regions. The national price is then taken as a weighted average of the regional medians. The NAR price indices can be volatile due to compositional changes. An increase in the difference between high priced relative to low priced units will increase the regional and hence the national median. The NAR indices are, however, available for each region on a monthly basis over a long time period.

The Bureau of Census publishes several house price series. A monthly national series is available since 1963, but the regional data are available only quarterly. The Census also provides an average price of new homes of constant quality from 1977 onwards on a quarterly basis, both for the United States and for the four regions. The indices are based on a monthly survey of residential construction activity for single-family homes. These indices are also subject to compositional effects that might arise from the sales sample rather than any true changes in price. The Census Bureau also publishes an index of one-family homes sold based on the hedonic approach.

The Conventional Mortgage Home Price Index (CMHPI) is provided by Freddie Mac. It is calculated on a quarterly basis at both the national and regional level from 1975 onwards. The index is based on conventional conforming mortgages for single-unit residential houses that were purchased or securitized by Freddie Mac or Fannie Mae. The CMHPI overlaps with the FHFA series to some extent. We are primarily interested in the common variations that underlie these series.

In addition to prices, volume data on transactions and turnovers are also informative about the level of housing activity. Dieleman et al. (2000) found that demographic changes are largely responsible for turnovers in the housing market, three-quarters of which are generated by renters. In contrast, house prices are mainly affected by household income. In a frictionless world, prices adjust and sales occur instantly after a shock. But frictions in the housing market might prevent house prices from adjusting, which slows turnover. Stein (1995) observed a positive contemporaneous correlation between changes in house prices and sales and that there is more intense trading activity in rising markets than in falling markets. He suggests that downpayment and other borrowing constraints might be responsible for market frictions. Case and Shiller (1989) argue that the rational response in a falling market is for a homeowner to hold on to his/her investment in anticipation of higher future returns. Berkovec and Goodman (1996) suggest that transactions might act as a forward indicator of price changes. As Stein (1995) suggests, if an initial shock knocks prices down, the loss on existing homes could undermine the ability of would-be movers to make downpayments on new homes. This lack of demand could further depress prices.

The transactions data used in our factor analysis include new and existing single-family homes under construction, sold and for sale, as well as data on employment in the construction sector. Additionally, the Census bureau publishes data on homeowner vacancy rates, home-ownership rates, as well as the rental vacancy rate. These latter indicators are informative about the tightness of the prevailing housing markets.

It is also useful to discuss data that are available but are not used in our factor analysis. The BLS publishes data on housing starts, permits as well as rent. Housing starts and permits are informative about the future as opposed to the present market conditions. We do not use these variables in the factor analysis as our model only allows for variables to load on contemporaneous and lagged factor observations. The rent data are based on the Consumer Expenditure Survey of which two-thirds of the sample are homeowners. To the extent that rent captures the capitalized value of housing, it provides a measure of the fundamental instead of the market value of houses. During periods of speculative housing booms, rents and house prices can diverge quite substantially. However, rent is regulated in many areas. We also have data on prices of mobile homes. These data are not used in our analysis, which focuses on single-unit and multi-family housing.

Finally, we note that all price variables are deflated by the (all items) CPI to control for increases in the overall price level. The NAR data are not seasonally adjusted. We use the X11 seasonal filter in Eviews to adjust these series. In all cases, we first annualize the data by taking year-to-year differences of the log level of the series. This means that for monthly data, we take the log difference of a series over a 12-month period. For a quarterly series, we take the log difference over four quarters. Since many series remain non-stationary, we transform the series into annual differences of the annual growth rates before estimating the factors. The effective sample is thus January 1975 to May 2008.

3. ECONOMETRIC METHODOLOGY

Our econometric framework is set up with three issues in mind. First, shocks to the housing sector need not be the same as shocks to house prices. Second, there are substantial variations in housing market conditions across regions. Third, we have non-balanced panels of data.

To deal with the aforementioned issues, we use an extension of the 'Dynamic Hierarchical Factor Model' framework developed in Moench et al. (2009). In such a model, variations in an economic time series can be idiosyncratic, common to the series within a block, or common across blocks. Here, we treat a block (identified as b) as one of the four major geographical regions, the Northeast (NE), Midwest (M), South (S) and West (W). More precisely, we posit that for each $b =$ NE, M, S and W, we observe N_b housing indicators which have zero mean and unit variance. The data, stacked in the vector X_{bt}, have a factor representation given by

$$X_{bt} = \Lambda_{Gb}(L)G_{bt} + e_{Xbt}, \tag{3.1}$$

where $\Lambda_{Gb}(L)$ is an $N_b \times k_{Gb}$ matrix polynomial in L of order s_{Gb}. According to the model, the housing indicators in a block are driven by a set of k_{Gb} regional factors denoted G_{bt}, and idiosyncratic components e_{Xbt}. Stacking up G_{bt} across regions yields the $K_G \times 1$ vector $G_t = (G_{1t}\ G_{2t}\ \ldots G_{Bt})'$. Observed indicators of the national housing market are stacked into a $K_Y \times 1$ vector Y_t. At the national level, we assume that

$$\begin{pmatrix} G_t \\ Y_t \end{pmatrix} = \Lambda_F(L)F_t + \begin{pmatrix} e_{Gt} \\ e_{Yt} \end{pmatrix}, \tag{3.2}$$

where K_F factors, collected into the vector F_t, capture the comovement common to all regional factors, and where $\Lambda_F(L)$ is a $(K_G + K_Y) \times K_F$ matrix polynomial of order s_F.

In the housing application under consideration, AR(1) dynamics are assumed throughout. Thus,

$$F_t = \Psi_F F_{t-1} + \epsilon_{Ft}, \tag{3.3}$$

$$e_{Gbt} = \Psi_{G.b} e_{Gbt-1} + \epsilon_{Gbt}, \tag{3.4}$$

$$e_{Yjt} = \Psi_{Y.j} e_{Yjt-1} + \epsilon_{Yjt}, \tag{3.5}$$

$$e_{Xbt} = \Psi_{X.b} e_{Xbt-1} + \epsilon_{Xbt}, \tag{3.6}$$

where Ψ_F is a diagonal $K_F \times K_F$ matrix, $\Psi_{G.b}$ is a diagonal $k_{Gb} \times k_{Gb}$ matrix, $\Psi_{Y.j}$ is a diagonal $K_Y \times K_Y$ matrix and $\Psi_{X.b}$ is a diagonal $N_b \times N_b$ matrix. Furthermore,

$$\begin{aligned}
\epsilon_{Fkt} &\sim N\big(0, \sigma^2_{Fk}\big), & k &= 1, \ldots, K_F, \\
\epsilon_{Gbjt} &\sim N\big(0, \sigma^2_{Gbj}\big), & j &= 1, \ldots, k_{Gb}, \\
\epsilon_{Yjt} &\sim N\big(0, \sigma^2_{Yj}\big), & j &= 1, \ldots, K_Y, \\
\epsilon_{Xbit} &\sim N\big(0, \sigma^2_{Xbi}\big) & i &= 1, \ldots, N_b.
\end{aligned}$$

Equations (3.1) and (3.2) constitute a three-level factor model—the level-one variations are due to ϵ_{Xbit}, the level-two variations are due to ϵ_{Gbt} and the level-three variations are due to ϵ_{Ft}. To identify the sign of the factors and loadings separately, we set the upper-left element of $\Lambda_F(0)$ to one, and for $b = 1, \ldots, B$, we also set the upper-left element of $\Lambda_G(0)$ equal to one.[1]

Our model can be seen as having two sub-models, each with a state space representation. Specifically, if G_{bt} was observed, (3.2) and (3.3) is a standard dynamic linear model where the latent factor is F_t. Then equations (3.1), (3.4) and (3.6) constitute the second dynamic linear model where the latent vector is G_t. It is in principle possible to use variables that are only available at the national level to estimate F_t. However, to the extent that G_{bt} are correlated across regions, they also convey information about F_t which we exploit in the estimation.

The three-level model implies that

$$X_{bt} = \Lambda_{Gb}(L)\Lambda_{Fb}(L)F_t + \Lambda_{Gb}(L)e_{Gbt} + e_{Xbt},$$

where Λ_{Fb} is the sub-block of Λ_F corresponding to block b. This is in contrast to a two-level factor model that consists of only a common and an idiosyncratic component. Omitting variations at the block level amounts to lumping $\Lambda_{Gb}(L)e_{Gbt}$ with e_{Xbt} in the estimation of F. This can result in an imprecise estimation of the common factor space if variations in e_{Gbt} are large. Moreover, explicitly specifying the block structure facilitates interpretation of the factors.

In our setup, the regional factors contain information about the state of the housing sector at the national level. A different formulation of regional effects is a model specified as

$$X_{bit} = b_i F_t + c_{bi} e_{Gbt} + e_{bit}.$$

Fu (2007) uses such a model to decompose house prices in 62 U.S. metropolitan areas into national, regional and metro-specific idiosyncratic factors using quarterly FHFA price data from 1980 to 2005. A similar model was also used by Del Negro and Otrok (2005) to estimate the common component of quarterly FHFA price data from 1986 to 2005. Stock and Watson (2008) use a variant of this model to analyse national and regional factors in housing construction. Kose et al. (2008) use it to study international business cycle comovements. Our model is more restrictive in that the responses of shocks to F_t for all variables in block b can only differ to the extent that their exposure to the block-level factors differs. However, the additional structure we impose makes the model more parsimonious, and it is easy to accommodate observed aggregate indices Y_t in estimating F_t.

Numerous methods are available to estimate two-level factor models. For models with a few numbers of series, maximum likelihood is widely used. For large dimensional factor models, the method of principal components is popular. The factor model used in the present study and introduced in Moench et al. (2009) is a multi-level extension of the simple two-level factor model considered in various previous contributions. We use Markov chain Monte Carlo (MCMC) methods (specifically a Gibbs sampling algorithm) to estimate the posterior distribution of the parameters of interest and the latent factors. Unlike the method of principal components, estimation via Gibbs sampling requires parametric specification of the innovation processes. However, one practical advantage of MCMC is that credible regions can be conveniently computed. In contrast, there exists no inferential theory for multi-level models estimated by the method of principal components.

[1] In the more general case with multiple factors at both levels, the k_{Gb} regional factors could, for example, be identified by requiring that for each b, $\Lambda_{Gb}(0)$ is a lower triangular matrix with diagonal elements of unity. Similarly, the K_F factors could be identified by requiring that $\Lambda_F(0)$ is lower triangular, again with ones on the diagonal. In such a case of multiple factors, the ordering of the variables might potentially have an impact on the factor estimates.

The problem considered here is non-standard because not every series on the housing market is available on a monthly basis. Aruoba et al. (2008) also consider estimation of latent factors when the data are sampled at mixed frequencies. However, we have the additional problem that not all our data series are available over the same time span. For example, house price data are available from NAR since 1970, from the FHFA at the regional level since 1975, and the Case–Shiller index is available since 1987.

The first problem that not all data are available on a monthly basis is easily handled in the Bayesian framework using data augmentation techniques. Suppose for now that we have data over the entire sample for X_{bt}, but it is only available on a quarterly instead of a monthly basis. The monthly value X_{bt} when t does not correspond to a month during which new data are released has conditional mean

$$X_{bt|t-1} = \Lambda_{Gb}(L)G_{bt} + e_{Xbt|t-1},$$

where $e_{Xbt|t-1} = \Psi_{X.b}e_{Xbt-1}$. A monthly observation of X_{bt} with conditional mean $X_{bt|t-1}$ and variance σ_{Xb} is obtained by taking a draw from the normal distribution with this property. Similarly, if the mth aggregate indicator is available on a quarterly basis, we make use of the fact that

$$Y_{t,m} = \Lambda'_{Fm}(L)F_t + e_{Yt,m}.$$

Conditional on F_t, Λ_{Fm} and $\Psi_{Y.m}$, the monthly value of $Y_{t,m}$ when data are not observed at time t can be drawn from the normal distribution with mean $\Lambda'_{Fm}(L)F_t + e_{Yt,m|t-1}$ and variance σ^2_{Ym}, where $e_{Yt,m}$ is an autoregressive process with parameters Ψ_{Ym}.

As for the second problem that some data are missing for the early part of the sample, assume for the sake of discussion that the data (when they are available) come on a monthly basis. For the sub-sample over which data are not available, we fill the data with the value 'NaN'. In a state space framework, these values contain no new information and contribute zero to the Kalman gain. To implement this, let X^+_{bt} be the subset of X_{bt} for which data are available at time t. If X_{bt} is $N_b \times 1$, and X^+_{bt} is $N^+_{bt} \times 1$, we work with the measurement equation:

$$X^+_{bt} = \Lambda^+_{Gb}(L)G_{bt} + e^+_{Xbt},$$

where var(e^+_{Xbt}) is an $N^+_{bt} \times N^+_{bt}$ matrix. Equivalently, let W_t be an $N^+_{bt} \times N_b$ selection matrix so that $X^+_{bt} = W_t X_{bt}$ is the $N^+_{bt} \times 1$ vector of variables that contain new information at time t. The measurement equation in terms of X_{bt} (which contains missing values) is

$$W_t X_{bt} = W_t \Lambda_{Gb}(L)G_{bt} + W_t e_{Xbt}. \tag{3.7}$$

This is equivalent to using the entire $T \times N_b$ matrix X_b, which is padded with zeros when missing values are encountered, and then setting the Kalman gain to zero.

The third problem which makes our MCMC algorithm non-standard relates to the fact that G_{bt} conveys information about F_t. More precisely,

$$\Psi_{Gb}(L)G_{bt} = \alpha_{F.bt} + \epsilon_{Gbt}, \tag{3.8}$$

where $\alpha_{F.bt} = \Psi_{G.b}(L)\Lambda_F(L)F_t$ depends on t. Given a draw of F_t, this can be interpreted as a time-varying intercept that is known for all t. By conditioning on F_t, our updating and smoothing equations for G_t explicitly take into account the information carried by F_t.

Summarizing, when a data series is unavailable for part of the sample, they are 'zeroed out' in the measurement equation. When a series is quarterly instead of monthly, then over the sample for which the data are available, we use data augmentation techniques to draw the monthly values. In Moench et al. (2009), we show that a simple extension of the algorithm in Carter and Kohn (1994) allows estimation of three-level models that takes into account the dependence of G_b on F. In this paper, we further modify the algorithm to accommodate the first two problems. Precisely, denote the observed national indicators Y_t and the observed regional indicators X_{bt}, $b = 1, \ldots, B$. Let

$$\begin{aligned}
\Sigma_{Xb} &= \operatorname{diag}\left(\sigma^2_{Xb1}, \ldots, \sigma^2_{XbN_b}\right),\\
\Sigma_{Gb} &= \operatorname{diag}\left(\sigma^2_{Gb1}, \ldots, \sigma^2_{Gbk_b}\right),\\
\Sigma_{Y} &= \operatorname{diag}\left(\sigma^2_{Y1}, \ldots, \sigma^2_{YK_Y}\right),\\
\Sigma_{F} &= \operatorname{diag}\left(\sigma^2_{F1}, \ldots, \sigma^2_{Fk_F}\right).
\end{aligned}$$

These matrices are of dimension $N_b \times N_b$, $k_{Gb} \times k_{Gb}$ and $K_F \times K_F$, respectively. Collect $\{\Lambda_{G1}, \ldots, \Lambda_{GB}\}$ and Λ_F into $\mathbf{\Lambda}$, $\{\Sigma_{X1}, \ldots, \Sigma_{XB}\}$, $\{\Sigma_{G1}, \ldots, \Sigma_{GB}\}$, Σ_Y and Σ_F into $\mathbf{\Sigma}$, and $\{\Psi_{X1}, \ldots, \Psi_{XB}\}$, $\{\Psi_{G1}, \ldots, \Psi_{GB}\}$, Ψ_Y, and Ψ_F into $\mathbf{\Psi}$. We first use data available for the entire sample to construct principal components. These are used to initialize $\{G_{bt}\}$ and $\{F_t\}$. Based on these estimates of the factors, initial values for Λ_{Gb}, Ψ_b, Σ_{Gb} and Σ_F are obtained. Each iteration of the sampler then consists of the following steps:

(1) Conditional on $\mathbf{\Lambda}$, $\mathbf{\Psi}$, $\mathbf{\Sigma}$, $\{G_t\}$ and $\{Y_t\}$, draw $\{F_t\}$.
(2) Conditional on $\{F_t\}$, draw Ψ_F, Σ_F and Λ_F.
(3) For each b, conditional on $\mathbf{\Lambda}$, $\mathbf{\Psi}$, $\mathbf{\Sigma}$ and $\{F_t\}$, draw $\{G_{bt}\}$ taking into account time-varying intercepts.
(4) For each b, conditional on $\{G_{bt}\}$ and $\{Y_t\}$, draw Ψ_{Gb} and Σ_{Gb}.
(5) For each b, conditional on $\{G_{bt}\}$, draw Λ_{Gbi}. Also draw Ψ_{Xbi} and σ^2_{Xbi}.
(6) Data augmentation:
  (i) For each b and conditional on $\{G_{bt}\}$ and the parameters of the model, sample monthly values for elements of $\{X_{bt}\}$ that are observed at lower frequencies.
  (ii) Conditional on $\{F_t\}$ and the parameters of the model, sample monthly values for those $\{Y_t\}$ that are observed at lower frequencies.
  (iii) Draw Ψ_Y using the augmented data vector for $\{Y_t\}$.

We assume normal priors centred around zero and with precision equal to 1 for elements of $\mathbf{\Lambda}$ and $\mathbf{\Psi}$, and inverse gamma priors with parameters 4 and 0.01 for elements of $\mathbf{\Sigma}$.[2] Given conjugacy, Λ_{Gb}, Λ_F, Ψ_{Xbi}, Ψ_{Gb} and Ψ_F in steps 4 and 5 are simply draws from the normal distributions whose posterior means and variances are straightforward to compute. Similarly, σ^2_{Gb} and σ^2_{Xbi} are draws from the inverse chi-square distribution. Notice that the model for (G_{bt}, Y_t) is linear in F_t and it is Gaussian. We can therefore run the Kalman filter forward to obtain the conditional mean of F_t at time T and the corresponding conditional variance. We then draw F_T from its conditional distribution, which is normal, and proceed backwards to generate draws $F_{t|T}$ for $t = T - 1, \ldots, 1$ using the algorithm suggested by Carter and Kohn (1994) and detailed in Kim and Nelson (2000). Draws of $\{G_{bt}\}$ can be obtained in a similar manner, as the model for X_{bt} is linear in G_{bt} and is Gaussian. This basic algorithm is modified to deal with a time-varying intercept in the transition equation for G_{bt} and missing values as discussed earlier.

[2] We use the equivalence of the inverse gamma and scale-inverse χ^2 distribution in our procedure and effectively sample variance parameters based on the χ^2 distribution.

3.1. *Estimates of F_t and G_{bt}*

Our sample starts in 1975:01 and ends in 2008:05. The base case price factors are estimated from seven series for each of the four regions, plus nine national price series. The base case housing market factors are estimated from 14 price and volume series for each of the four regions, plus 17 national series.

The model has a number of parameters that we need to specify. The final model assumes $k_{Gb} = 1$ for all b and $K_F = 1$.[3] As discussed earlier, we set the upper-left element of the factor loading matrices $\Lambda_F(0)$ and the four $\Lambda_{Gb}(0)$ equal to one in order to separately identify the signs of the factors and factor loadings. We order the NAR's 'Median Sales Price of Single Family Existing Homes' first both at the national and at the regional level. This implies that our estimated factors will have a positive correlation with the corresponding NAR price indices, respectively. As stated earlier, we assume e_{Xbi}, e_{Gb}, e_{Yj} and e_F to be AR(1) processes. Moreover, we let $s_{Gb} = s_F = 2$ so that the factors at both the regional and the national level are allowed to have a lagged impact of up to two periods on the respective observed variables. This allows us to accommodate lead–lag relationships between the housing cycles across the different regions and the nation as a whole. We begin with 20,000 burn in draws. We then save every 50th of the remaining 50,000 draws. These 1000 draws are used to compute posterior means and standard deviations of the factors and parameters.

Table 3. Estimates of Ψ and σ_F: house price model.

F_p	var(F_p)	Ψ_F	s.e.	σ_F^2	s.e.
	0.639	0.942	0.054	0.020	0.011
$G_{p,b}$	var($G_{p,b}$)	Ψ_{Gb}	s.e.	σ_{Gb}^2	s.e.
NE	1.012	0.633	0.239	0.079	0.050
W	2.066	0.744	1.058	0.126	1.300
MW	0.524	0.187	0.107	0.280	0.229
South	0.074	−0.018	0.002	0.079	0.000

Decomposition of variance:

	ϵ_F	s.e.	ϵ_{Gb}	s.e.	ϵ_{Xb}	s.e.
NE	0.340	0.087	0.239	0.053	0.421	0.051
W	0.247	0.083	0.271	0.060	0.482	0.056
MW	0.341	0.080	0.159	0.058	0.500	0.075
South	0.241	0.070	0.083	0.021	0.677	0.061

Table 3 reports estimates of the dynamic parameters and the variance of shocks to the common, regional and series-specific components. The unconditional variance of F_p and $G_{p,b}$ are denoted var(F_p) and var($G_{p.b}$), while the variances of ϵ_{F_p} and $\epsilon_{G.b}$ are denoted $\sigma^2_{F_p}$ and $\sigma^2_{G_p.b}$, respectively. The common price factor, F_p, is more persistent but has a smaller variance than the regional factors, $G_{p,b}$. Of the four regional factors, the West is the most persistent. Even though

[3] We considered $k_{Gb} = K_F = 2$, but the additional factors tend to have little variability and were subsequently dropped.

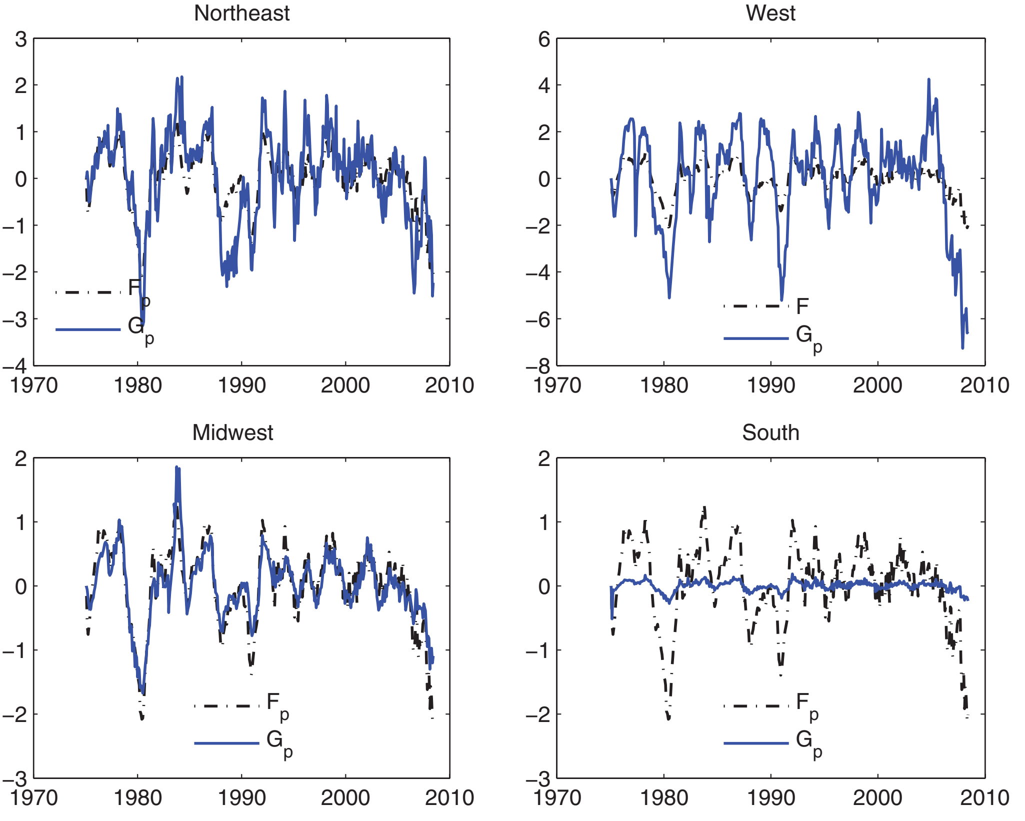

Figure 1. National and regional house price factors.

house price shocks in the Midwest are larger, the house price factor in the West has a larger unconditional variance once persistence is taken into account.

Figure 1 presents the estimated posterior mean of the national (F_p) versus the regional house price factors (G_p) for each of the four Census regions. The sample period is 1975:01–2008:05. The standard deviation of F_p is 0.636. The national factor (solid line) is notably smoother than the regional factors (dash-dotted line). The Northeast experienced housing busts in the early 1980s and the late 1980s that were much more pronounced than the national market. However, throughout the 1990s, the Northeast market is stronger than the national market. The West experienced a sharp decline in house prices in the mid-1970s and again in the early 1990s. These variations are larger than what was recorded for other periods, or in any of the other regions. Because of these two episodes, the G_p for the West has a standard deviation of 2.066, much larger than the 1.012 observed for the Northeast. The Midwest and the South have more tranquil housing markets. The standard deviation of $G_{p,b}$ are 0.515 and 0.078, respectively.

Table 3 also reports a decomposition of variance in the series used to estimate the factors. We find that shocks to the national house price factor, ϵ_{Fp}, account for 34%, 24.7%, 34.1% and 24.1% of house price variations in the four regions, respectively. Shocks to the regional factor ϵ_{Gp} have a share of 23.9% in the Northeast, 27.1% in the West, 15.9% in the Midwest and 8.3%

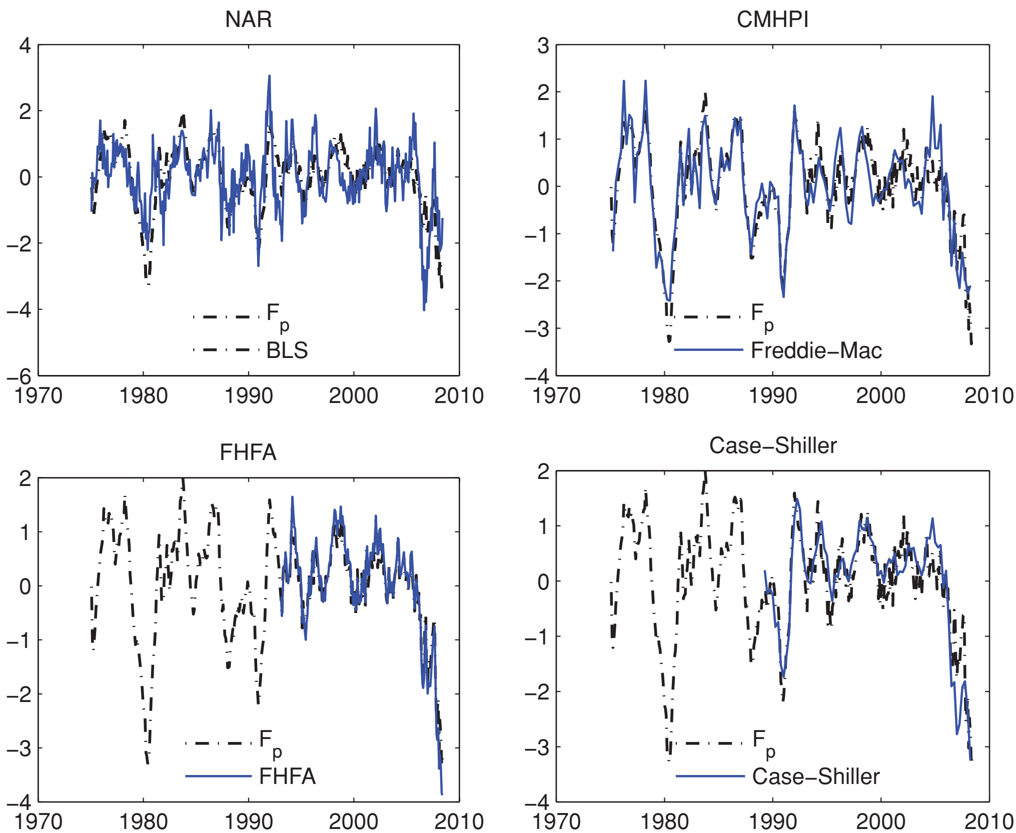

Figure 2. National house price factor and leading house price indices.

in the South. Series-specific shocks account for the remaining 42.1%, 48.2%, 50% and 67.7% of the variation in the regional house prices series. Notably, the factor structure is strongest in the Northeast and is weakest in the South.

Figure 2 plots the estimated posterior mean of the national house price factor (F_p) versus four leading national house price indices. 'NAR' is the Median Sales Price of Single Family Existing Homes from the NAR; 'CMHPI' is the CMHPI from Freddie Mac; 'FHFA' is the Purchase-Only House Price Index from the Federal Housing Finance Administration (formerly OFHEO); 'Case–Shiller' is the S&P/Case–Shiller Home Price Index published by Fiserv Inc. While the NAR series is available at the monthly frequency, the latter three indices are only available quarterly. Our estimated house price factor is a monthly time series. The sample period is 1975:01–2008:05.

As expected, F_p is somewhat smoother than the individual price series because F_p is essentially a weighted average of all price indices. The Census monthly price index has a correlation of 0.67 with F_p, while the Conventional Mortgage Home (quarterly) price index has a correlation with F_p of 0.93. Notably, both series are more volatile than F_p. The monthly FHFA series that is available since 1986 is also highly correlated with F_p: computed over the sample for which the two series overlap, the correlation coefficient is 0.99. The correlation between the

Table 4. Estimates of Ψ and σ_F: housing market model.

F_{pq}	$\text{var}(F_{pq})$	Ψ_F	s.e.	σ_F^2	s.e.
	0.554	0.896	0.069	0.028	0.020
$G_{pq,b}$	$\text{var}(G_{pq,b})$	Ψ_G	s.e.	σ_{Gb}^2	s.e.
NE	1.025	0.530	0.300	0.145	0.109
W	2.019	0.766	1.636	0.139	2.341
MW	0.127	0.206	0.009	0.101	0.002
South	0.591	0.504	0.159	0.080	0.169

Decomposition of variance:

	ϵ_F	s.e.	ϵ_{Gb}	s.e.	ϵ_{Xb}	s.e.
NE	0.164	0.036	0.147	0.028	0.689	0.028
W	0.055	0.023	0.247	0.048	0.698	0.044
MW	0.114	0.029	0.128	0.019	0.758	0.016
South	0.130	0.032	0.154	0.027	0.717	0.026

Case–Shiller index and F_p is 0.91. Notice, however, that there are significant differences between F_p and the indicators in recent years. The four price indices seem to show sharper declines in house prices than the house price factor which incorporates information from various series.

One of our objectives is to investigate whether the house price cycle differs from the housing cycle, where the latter is defined based on data on prices as well as volume. Table 4 reports the posterior mean of the dynamic parameters and the variance of the shocks for the housing model. To distinguish them from the house price factor, F_p, we denote the national housing factor by F_{pq} and the regional housing factors by G_{pq}. The common factor is still highly persistent. While a regional factor is not evident in house prices in the South, the data on volume help to isolate this factor.

A decomposition of variance of the housing market model reveals that national and regional shocks are equally important in the Northeast, the Midwest and the South, while regional shocks in the West are more important than the national shocks. However, the result that stands out is that idiosyncratic variation in the housing market data in all four regions are relatively more important than the common shocks and dominate the total variations in the data.

Figure 3 shows the estimated posterior mean of the national housing factor, F_{pq}, using both price and volume data, and the estimated posterior mean of the national house price factor, F_p, exclusively based on home price data. The sample period is 1975:01–2008:05.

At the national level, the correlation between the house price factor F_p and the housing market factor F_{pq} is 0.85. In spite of this strong correlation and as seen from Figure 3, there is a notable difference between the two series during peaks and troughs. The discrepancy has been especially pronounced since 2007. While F_p is -2.006 at the end of our sample in May 2008, almost four standard deviations below the mean, F_{pq} is -0.517, roughly one standard deviation below the mean. The drop in housing activity as indicated by F_{pq}, estimated using both house price and quantity information, is thus less severe.

This section has focused on different measures of housing market activity, and several conclusions can be drawn. First, there is substantial regional variation in housing market activity with the regional component playing the largest role in the West. Second, our house price factor

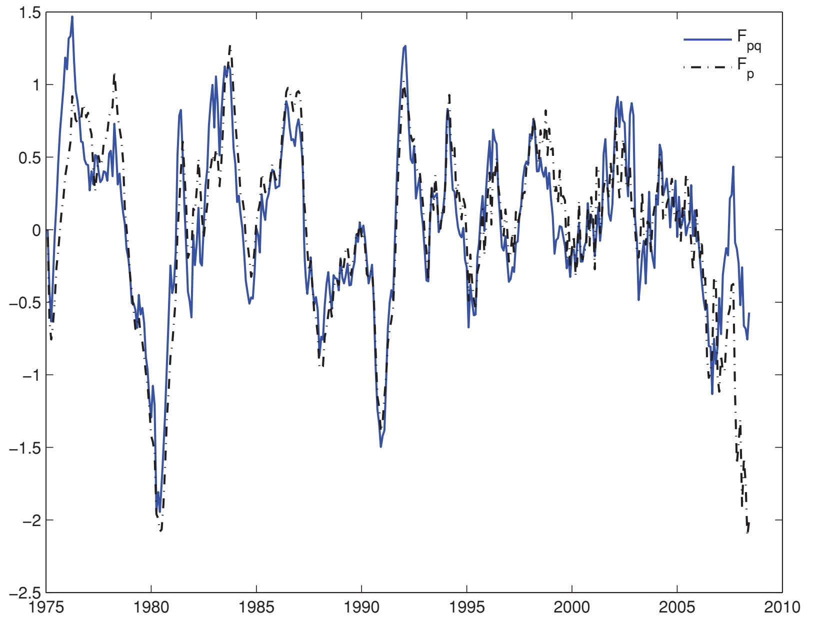

Figure 3. National housing factor and national house price factor.

is highly correlated with each of the four widely used house price indices with the important difference that our F_p is smoother. Third, G_{pq} is generally similar to G_p except in the South. At the national level, F_p and F_{pq} are well synchronized, but the decline of F_{pq} since 2007 is much less pronounced than F_p. These observations suggest that there are important idiosyncratic movements in observed housing data. A particular house price series will not, in general, be representative of the true level of activity underlying the housing market. The more data we use to estimate the factors, the better we are able to 'wash out' the idiosyncratic noise. However, all indicators point to a sharp decline in housing market activity since 2007. This decline is pervasive and occurs at both the regional and national levels. We next investigate whether shocks to the housing market affect consumption.

4. HOUSING AND CONSUMPTION

There exists little work on the regional aspect of housing variations. Using a dynamic Gordon model, Ng and Schaller (1998) find that regional housing bubbles have predictive power for future consumption. Campbell et al. (2008) find that housing premia are variable and forecastable and account for a significant fraction of the variation in the rent–price ratio at the national and regional levels. However, they do not assess the consumption effects of housing. One reason why there are so few estimates of the regional effects of housing is data limitation. Not only is it difficult to find regional housing data over a long time period, government statistical agencies do not publish consumption data at the regional level. We use regional retail sales data provided by the Census Bureau until 1997 and continued by the Bank of Tokyo-Mitsubishi (BTM) since then. These series are available monthly from 1970 onwards, both for each of the four Census regions, and also for the United States as a whole. This retail sales (which we will simply refer to as consumption) series is not seasonally adjusted. We run it through the X11 filter in Eviews, and deflate by the all-items CPI. We then analyse the log annual difference of this seasonally adjusted, real retail sales series.

4.1. Estimates from FAVAR

We are interested in quantifying the response of retail sales consumption to changes in regional and national housing market conditions. Retail sales, while representing only a sub-category of total consumption, have the advantage of being available for the main Census regions. We can therefore study the effects of housing market shocks on consumption both at the regional and the national level. The foregoing discussion suggests that a strong housing market will increase consumption of some but decrease the consumption of others. As those affected may have different propensities to consume, the aggregate effect of changes in housing market conditions on consumption is an empirical matter.

Our analysis is based on factor-augmented vector autoregressions (FAVAR), a tool for analysing macroeconomic data popularized by Bernanke et al. (2005). While a conventional VAR is an autoregressive model for a vector of observed time series, a FAVAR augments the observed vector of variables by a small set of latent factors often estimated by the method of principal components. Bai and Ng (2006) showed that if $\sqrt{T}/N \to 0$ as $N, T \to \infty$, the estimated factors that enter the FAVAR can be treated as though they are observed. The method of principal components is not, however, well suited for the present analysis for two reasons. First, the number of series available for analysis is much smaller than the typical large dimensional analysis in which principal components is applied. Second, we have a non-balanced panel with data sampled at mixed frequencies. Both problems are more easily handled by Bayesian estimation. Accordingly, our FAVAR is based on Bayesian estimates of the factors.

Our first set of FAVARs consist of five variables, respectively, F_{pq}, $G_{pq,b}$, U, U_b and RS_b, where U is the national unemployment rate, U_b is the unemployment rate for region b, and RS_b is the linearly detrended logarithm of real retail sales for region b. The variables U_b and U allow us to control for regional and aggregate business cycle conditions. We use the housing factors instead of the house price factors as these provide more comprehensive measures of the housing market. The regional unemployment rates are available only from 1976 onwards. Thus, for this exercise, the sample is 1976:01–2008:05. The standard deviations of the estimated factors used in the FAVAR are given in Table 4.

We identify shocks to the housing market factor using a simple recursive identification scheme where the variables are ordered as they appear earlier. This identification implies that the national housing factor does not react to regional housing market shocks, regional and national unemployment shocks as well as regional consumption shocks within the same month. This assumption appears reasonable given that it usually takes at least a few weeks from the time a decision is made to purchase, sell or construct a home before an actual transaction is being made. Our identification also implies that regional retail sales can respond within the same month to both national and regional housing shocks as well as national and regional labour market shocks, as would be the case if households can adjust their consumption decisions rather quickly in response to various kinds of economic shocks. As for the particular ordering between the national and regional housing market factors, it also appears plausible to suppose that national housing market shocks may have an immediate impact on regional housing market dynamics, whereas the reverse does not hold.

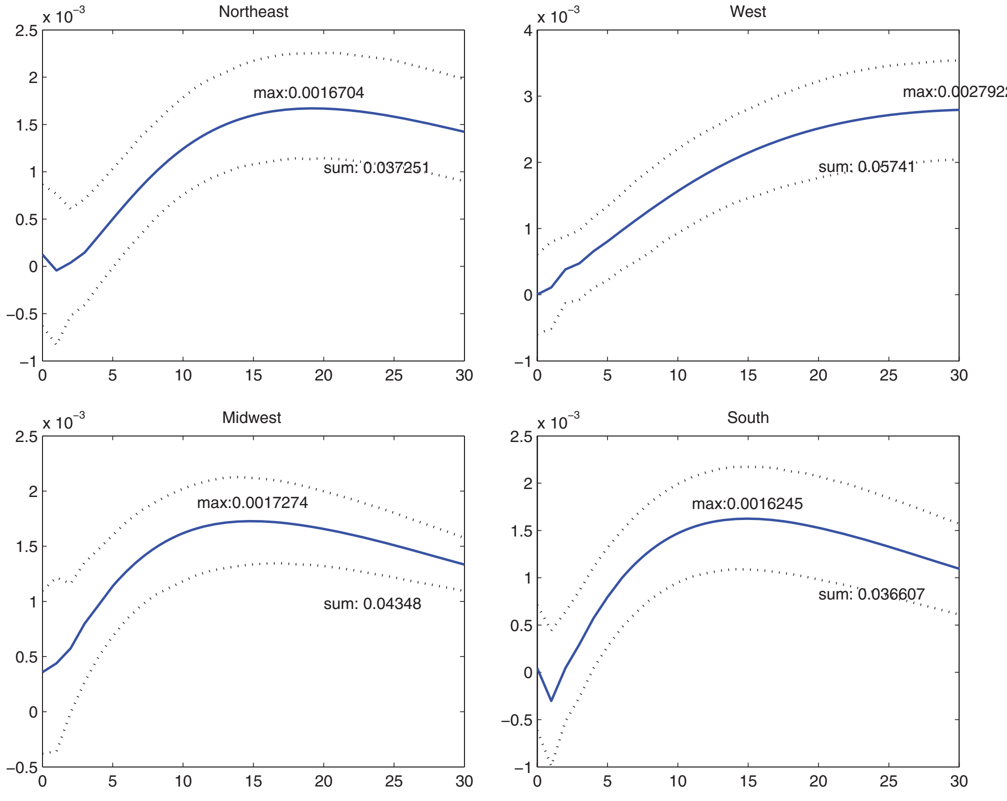

Figure 4. Impulse responses of regional retail sales to national housing shocks.

The impulse response functions are obtained as follows. First recall that we saved 1000 of the 50,000 draws of $G_{pq,b}$ and F_{pq} from the Gibbs sampler. For each draw of the regional and national housing factors, we estimate a five-variable FAVAR with two lags for each of the four regions. The estimated FAVARs are then used to obtain impulse response of regional retail sales to shocks in F_{pq} and $G_{pq,b}$. We also estimate a three-variable VAR in F_{pq}, U and RS to study the impulse response of aggregate retail sales to housing market factor shocks. Repeating this for each of the 1000 draws of F_{pq} and $G_{pq,b}$ gives a set of impulse responses from which we can compute the posterior means and percentiles of the posterior distribution.

Figure 4 reports the estimated posterior mean and 90% probability interval of impulse responses to a one-standard-deviation shock in F from the FAVARs discussed in Section 4.1. For each region b, these contain the following five variables: F_{pq}, $G_{pq,b}$, U, U_b and RS_b, where F_{pq} and $G_{pq,b}$ denote the national and regional housing factor, U and U_b the national and regional unemployment rate, and RS_b the linearly detrended logarithm of real retail sales for region b. We identify shocks using a recursive identification scheme of the five variables ordered in the way they appear earlier. The FAVAR has two lags. The sample period is 1976:01–2008:05.

The effect of a one-standard-deviation shock to F is positive in all four regions. The response of retail sales is hump shaped. The shock triggers a permanent increase in the level of retail sales. The effects, similar across regions, peak about 10 months after the shock. The cumulative effects

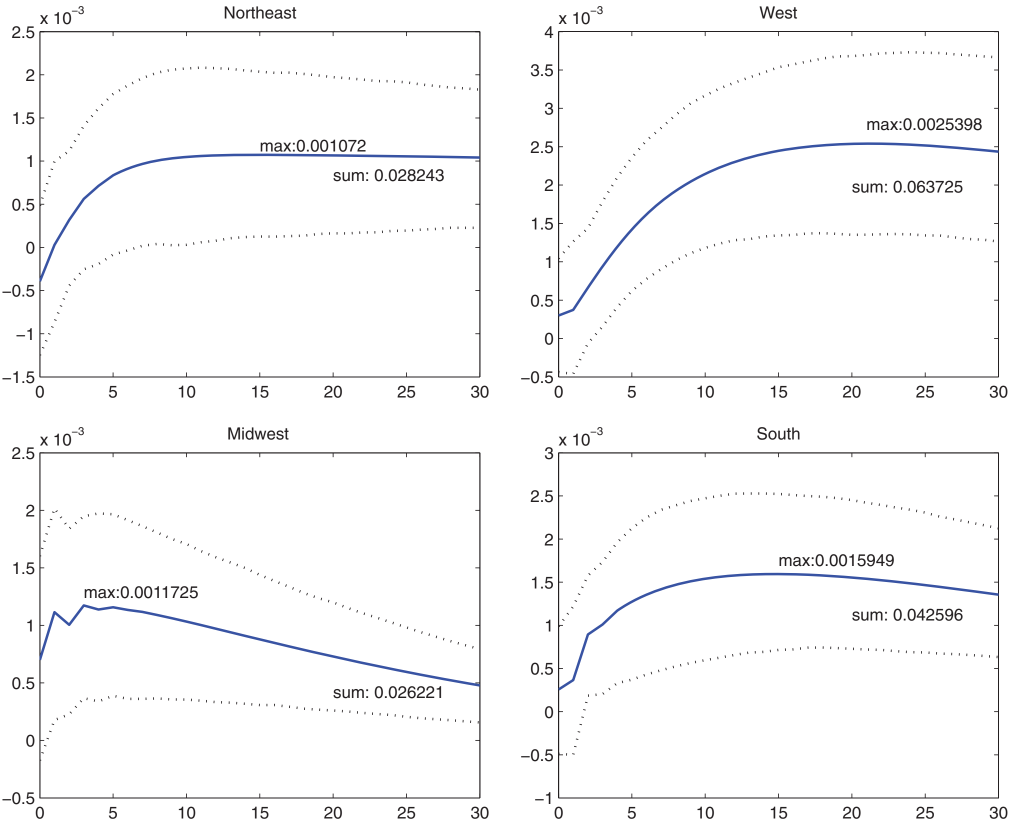

Figure 5. Impulse responses of regional retail sales to regional housing shocks.

in the four regions over two and a half years are 0.037, 0.057, 0.043 and 0.037, respectively. Hence, according to our estimates, a shock to the national housing market may result in a long-run increase of regional consumption between 3.7% and 5.7% above its trend level, all else being equal. While the positive consumption effect would be consistent with the idea that homeowners take advantage of a hot housing market, trading down their homes to enjoy realized capital gains, a more likely interpretation is that the positive response is due to the collateral effect brought about by increased home equity.

Figure 5 is constructed similarly to Figure 4 and graphs the impulse response to a standard deviation shock in G_{pq}. Because F_{pq} is also in the FAVAR and is ordered first, shocks to F_{pq} can trigger a contemporaneous response in the regional housing market factor, while the reverse is not true. The results can thus be interpreted as shocks to regional housing market activity that are orthogonal to national housing market shocks. As indicated by wider probability bands, the effects of shocks to G_{pq} are statistically less well determined than those to F_{pq}. While the effects at the peak are about the same as the response to a national housing factor shock, the cumulative effects of regional housing shocks are smaller and differ substantially across regions. The long-run effects are 0.028, 0.064, 0.026 and 0.043, respectively, implying an increase of regional consumption due to regional housing market shocks between 2.6% and 6.4% above its trend level, all else equal. Notably, the effects are largest in the West where variations in G_{pq} are also relatively more important. Even in this region, we find the effects of regional housing shocks on consumption to be less pronounced than the national shocks.

In unreported results, we find that shocks to volume tend to reduce retail sales for a few months after the shock, and have no substantial long-run effects. Furthermore, the consumption response to G_{pq} also seems to be largely due to the response to G_p. Thus, the consumption responses we observe in Figures 4 and 5 are largely a consequence of shocks to house prices rather than housing volume.

Given the heterogeneity in response across regions, what is the consumption response at the national level? To assess this question, we estimate a FAVAR in three variables: aggregate retail sales, *RS*, aggregate unemployment, U, and the national housing market factor, F_{pq}. We again identify shocks to the housing market factor using a recursive identification scheme where the ordering is as the variables appear earlier. The economic reasoning behind this approach follows the discussion earlier.

Figure 6 displays the estimated posterior mean and 90% probability interval of impulse responses from the three-variable FAVAR discussed in Section 4.1. This contains national retail sales, *RS*, the national unemployment rate, U, and the national housing market factor, F_{pq}. We identify shocks using a recursive identification scheme of the three variables in the order they appear earlier. The FAVAR has two lags. The sample period is 1976:01–2008:05.

The top panel in Figure 6 shows that the response of aggregate retail sales to a one-standard-deviation shock in F_{pq} is positive with a maximum effect occurring about 15 months after the shock, and a cumulative effect of 0.04 after 30 months. This implies that a one-standard-deviation shock may push aggregate retail sales 4% above their trend level. At the same time, unemployment falls by over 10 basis points as housing market activity increases. Figure 6 also shows how F_{pq} responds to its own shock. The response is gradual, and the half-life of the shock is about 10 months.

We have presented results for counter-factual increases in housing market activity. A policy question of interest is the quantitative consumption effect as the national housing market contracts. Figure 4 implies that consumption is expected to fall immediately in all regions as a result of a housing market shock, all else equal. As noted earlier, the housing factor in May 2008 was -0.571, about one standard deviation below the mean. A one-standard-deviation F_{pq} shock in the FAVAR is about 0.25. A two-standard-deviation shock to F_{pq} can thus have a cumulative consumption effect on the West of 2×0.057, or about 11%, and about 8% in the other three regions.

The consumption effect is much larger if we look at the house price factor, recalling that F_p is estimated to be -2.006 in May 2008, almost four standard deviations below the mean. Our results then suggest that at its worst (about 15 months after the shock), consumption can fall by 1.2% with an even higher cumulative effect than a shock to housing activity. The consumption effect thus depends on whether we think house prices alone reflect the state of the housing market, or whether volume information should be taken into account.

4.2. *What affects the housing factor?*

The slumping U.S. housing market has been a deep concern for private citizens and policy makers alike. While as of May 2008, the last data point in our sample, our housing factor was only one standard deviation below average, housing market activities have further decelerated since. This raises the question of what might stimulate housing activity.

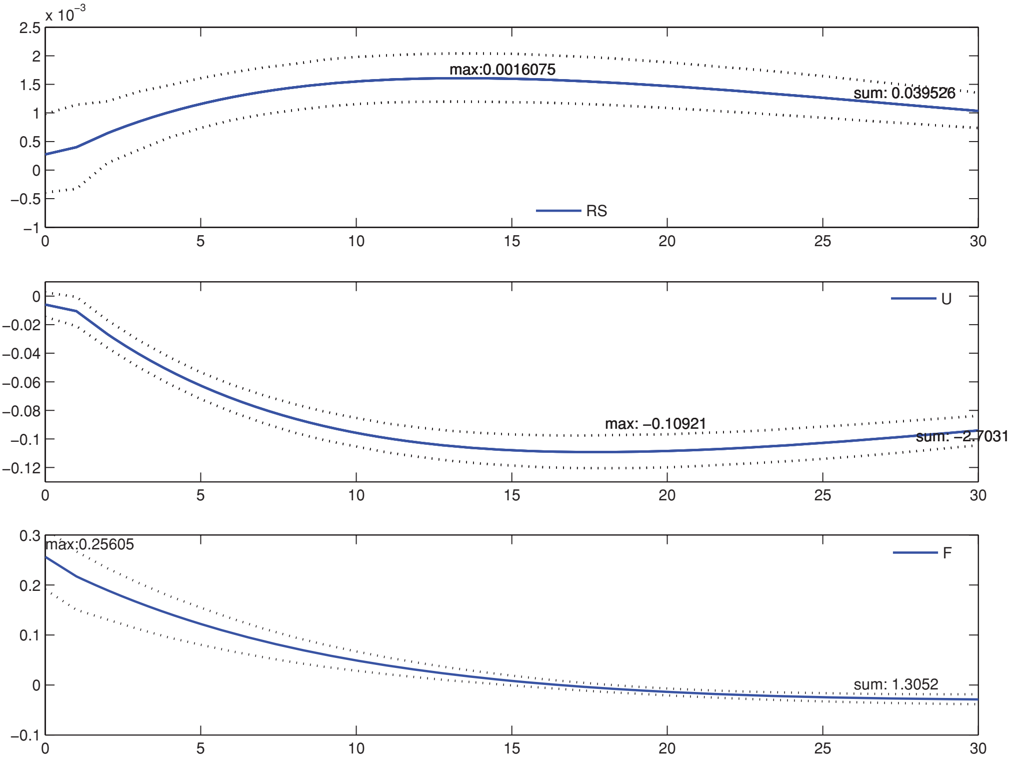

Figure 6. Impulse responses of national retail sales to national housing shocks.

To address this question, we consider a FAVAR with six variables and four lags at the national level. These variables in the order they enter the FAVAR are the unemployment rate (U), the fed funds rate (FF), a 30-year effective mortgage rate (MR), our housing factor (F_{pq}), the University of Michigan's survey of consumer sentiment ($MICH$), and the log of the S&P 500 index (SP). The unemployment rate captures aggregate business cycle dynamics. The fed funds rate and the effective mortgage rate measure tightness in the money and loans market. The Michigan survey measures confidence for the economy, and the S&P 500 index is a proxy for changes in financial wealth. Arguably, each of these variables can be thought of having an effect on the level of housing activity.

We identify shocks in the FAVAR using a recursive ordering of the variables as they appear earlier. This ordering implies that the unemployment rate does not respond within the month to any of the shocks but its own. The federal funds rate is ordered second and hence responds on impact only to unemployment shocks and monetary policy shocks. The 30-year effective mortgage rate is ordered third which implies that it is assumed to respond on impact to unemployment and monetary policy shocks, but with a one-month lag to shocks to the housing market as measured by our estimated housing factor, consumer confidence and the S&P 500 index. We put the housing factor in fourth position, which implements the assumption that housing market activity cannot respond within the month to shocks to consumer confidence and stock prices. Finally, the fact that consumer confidence and the S&P 500 are ordered last implies

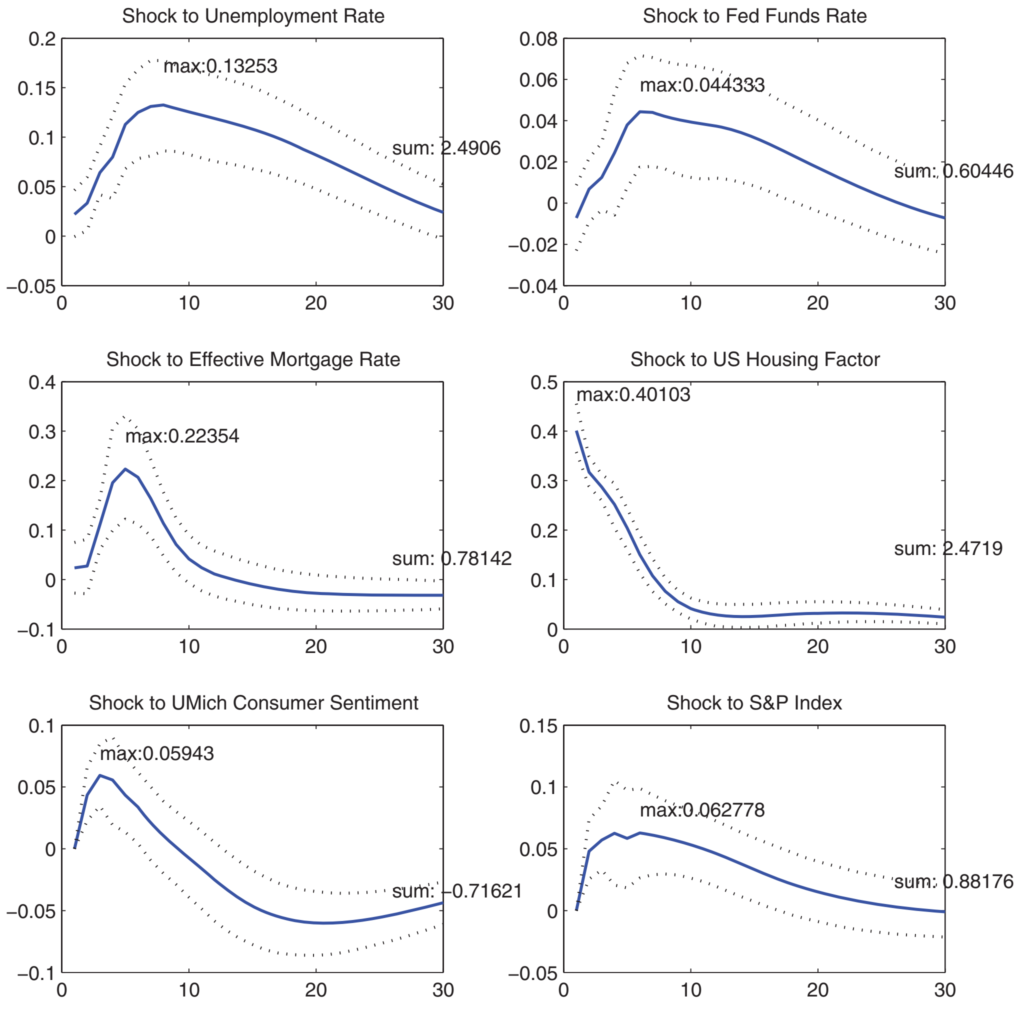

Figure 7. Impulse responses of national housing factor to different shocks.

that these two variables can respond within the month to unemployment, interest rate and housing market shocks.

Figure 7 shows the estimated posterior mean and 90% probability interval of impulse responses of F_{pq} to shocks to each of the six variaties from the FAVAR discussed in Section 4.2.[4] This contains the national unemployment rate, U, the fed funds rate, FF, a 30-year effective mortgage rate from Freddie Mac, MR, our national housing factor, F_{pq}, the University of Michigan's survey of consumer sentiment, $MICH$, and the log of the S&P 500 index, SP. We identify shocks using a recursive identification scheme of the variables in the order they appear earlier. The FAVAR has four lags. The sample period is 1976:01–2008:05. We normalize the

[4] Unreported results show that changes to the ordering of the six variables do not qualitatively affect our conclusions.

shocks to the federal funds rate and the effective mortgage rate to have a contemporaneous -25 basis point impact on itself, respectively. All other shocks are one-standard-deviation shocks. Reductions in both interest rates lead to an increase in housing market activity. The maximum effect of a 25 basis points cut in the fed funds rate on the housing factor is 0.026 and the cumulative effect is 0.35. By contrast, the maximum effect of a 25 basis points reduction in the effective mortgage rate is 0.13 and the cumulative effect after 30 months equals 0.46. Hence, reducing mortgage rates has a larger maximum and cumulative effect on the housing factor than an equivalent cut of the fed funds rate. This result suggests that direct policy interventions in the mortgage market may represent an effective way to revive the housing market.

Interestingly, we find that a positive one-standard-deviation shock to the unemployment rate boosts housing activity. This effect is due to a strong reduction of the federal funds rate following a negative shock to real activity. We further find that a one-standard-deviation increase in stock prices has a positive effect on housing activity both in the short run and in the long run. The maximum effect on F_{pq} is 0.033, recorded three periods after the shock, and the cumulative effect is 0.51 which is about equal to two standard deviations of the housing factor. A one-standard-deviation increase in consumer confidence boosts the housing factor in the short term, but interestingly has a negative cumulative effect.

An overview of the results suggests that all else equal, housing market activity can be stimulated by a cut in the federal funds rate and a reduction of mortgage rates, the latter potentially having larger effects. Increases in stock prices will positively affect housing market activity in the short run and in the long run. Increased consumer confidence is found to stimulate housing market activity in the short run, but may have a negative long-run effect on housing.

We also re-estimate the previous FAVAR using three different national house price series available for our full data sample as well as our estimated national house price factor, each standardized to have the same unconditional variance. We do not report these results here, but restrict ourselves to noting that the various indicators imply quite different reactions of house prices to the six shocks. The particular choice of house price measure thus potentially has a large impact on the conclusions one may reach from a quantitative analysis such as the one carried out earlier.

5. CONCLUSION

This paper provides three new perspectives on the effects of housing shocks on consumption. First, we distinguish between house price shocks and shocks to general activity in the housing market. Second, we analyse regional as well as national data. Third, our housing shock is not tied to a specific house price series. Instead, we extract house price and housing market factors from a large number of housing indicators. Our results indicate that in spite of large idiosyncratic variations, there is a national and a regional housing component in each of the regions, though the regional component is more important than the national component in the West. The aggregate response of consumption to national housing shocks is hump shaped. According to our estimates, the drop in housing market activity that began in 2006 can lead to a significant decline of consumption, all else being equal. Interest rate cuts can stimulate housing activity, and directly targeting lower mortgage rates may be an effective way to revive the housing market. However, without a boost in consumer confidence and the stock market, the housing market can remain depressed for a prolonged period of time. Our econometric framework permits a block structure and can handle data of mixed frequencies as well as missing data. Latent factors can also co-exist with observed factors. The methodology can be useful in other applications.

ACKNOWLEDGMENTS

This paper was presented at The Econometrics Journal Special Session of the 2009 RES Conference in Surrey, and the 2009 Summer Meeting of the Econometrics Society in Boston. The authors are grateful to Chris Otrok, Dan Copper, seminar participants at Columbia University and the Bundesbank for helpful comments and discussions. Evan LeFlore provided valuable research assistance. The second author would like to acknowledge financial support from the National Science Foundation under grant SES 0549978. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York or the Federal Reserve System.

REFERENCES

Aruoba, B., F. Diebold and C. Scotti (2008). Real-time measurement of business conditions. Working paper, University of Maryland.

Bai, J. and S. Ng (2006). Confidence intervals for diffusion index forecasts and inference with factor-augmented regressions. *Econometrica 74*, 1133–50.

Berkovec, J. A. and J. L. Goodman (1996). Turnover as a measure of demand for existing homes. *Real Estate Economics 24*, 421–40.

Bernanke, B., J. Boivin and P. Eliasz (2005). Factor augmented vector autoregressions (FVARs) and the analysis of monetary policy. *Quarterly Journal of Economics 120*, 387–422.

Calomiris, C., S. Longhofer and W. Miles (2008). The foreclosure–house price nexus: lessons from the 2007–2008 housing turmoil. NBER Working Paper No. 14294, National Bureau of Economic Research.

Campbell, S., M. Davis, J. Gallin and R. Martin (2008). What moves housing markets? A variance decomposition of the rent–price ratio. Working paper, Board of Governors of the Federal Research System.

Carter, C. K. and R. Kohn (1994). On Gibbs sampling for state space models. *Biometrika 81*, 541–33.

Case, K. and R. Shiller (1989). The efficiency of the market for single family homes. *American Economic Review 79*, 125–37.

Del Negro, M. and C. Otrok (2005). Monetary policy and the house price boom across U.S. states. Working Paper No. 2005-24, Federal Reserve Bank of Atlanta.

Dieleman, F., W. Clark and M. Deurloo (2000). The geography of residential turnover in twenty-seven large US metropolitan housing markets, 1985–1995. *Urban Studies 37*, 223–45.

Fu, D. (2007). National, regional and metro-specific factors of the U.S. housing market. Working Paper No. 0707, Federal Reserve Bank of Dallas.

Kim, C. and C. Nelson (2000). *State Space Models with Regime Switching*. Cambridge, MA: MIT Press.

Kose, A., C. Otrok and C. Whiteman (2008). Understanding the evolution of world business cycles. *International Economic Review 75*, 110–30.

Leamer, E. (2007). Housing is the business cycle. NBER Working Paper No. 13248, National Bureau of Economic Research.

Moench, E., S. Ng and S. Potter (2009). Dynamic hierarchical factor models. Working paper, Columbia University.

Ng, S. and H. Schaller (1998). Do housing bubbles affect consumption? Working paper, Columbia University.

Rappaport, J. (2007). A guide to aggregate house price measures. *Economic Review, Federal Reserve Bank of Kansas City 2*, 41–71.

Stock, J. H. and M. W. Watson (2008). The evolution of national and regional factors in U.S. housing construction. Working paper, Princeton University.