Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

Housing Price Forecastability

A Factor Analysis

Bork, Lasse; Møller, Stig Vinther

Published in:
Real Estate Economics

DOI (link to publication from Publisher):
[10.1111/1540-6229.12185](https://doi.org/10.1111/1540-6229.12185)

Publication date:
2018

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Bork, L., & Møller, S. V. (2018). Housing Price Forecastability: A Factor Analysis. *Real Estate Economics*, 46(3), 582-611. <https://doi.org/10.1111/1540-6229.12185>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Housing price forecastability: A factor analysis*

Lasse Bork[†]

Stig V. Møller[‡]

May 2012

Abstract

We examine US housing price forecastability using a common factor approach based on a large panel of 122 economic time series. We find that a simple three-factor model generates an explanatory power of about 50% in one-quarter ahead in-sample forecasting regressions. The predictive power of the model stays high at longer horizons. The estimated factors are strongly statistically significant according to a bootstrap resampling method which takes into account that the factors are estimated regressors. The simple three-factor model also contains substantial out-of-sample predictive power and performs remarkably well compared to both autoregressive benchmarks and computational intensive forecast combination models.

Keywords: House prices; Forecasting; Factor model; Principal components; Macroeconomic factors; Factor forecast combination; Bootstrap

JEL codes: C53; E3; G1

*We gratefully thank Graham Elliott, Carsten Tanggaard, participants at the 2012 Arne Ryde Workshop in Financial Economics, and seminar participants at CREATES, Aalborg University, and Aarhus University for comments and suggestions. Møller acknowledges support from CREATES, which is funded by the Danish National Research Foundation.

[†]Department of Business and Management, Aalborg University, Denmark. E-mail: bork@business.aau.dk.

[‡]Department of Economics and Business, Aarhus University and CREATES. Address: Fuglesangs Allé 4, 8210 Aarhus, Denmark. E-mail: svm@asb.dk.

1 Introduction

The recent boom and burst of the U.S. housing market strongly emphasizes that movements in housing prices play a significant role for consumer spendings, financial markets, and the macroeconomy as a whole. It follows that building an adequate forecasting model could provide useful information to central banks, financial supervision authorities as well as to other economic agents.

The existing literature on house price forecastability is in some respects still limited. In particular, the existing evidence tends to focus on long-run trends in house prices. The focus also tends to be on only a single or a few selected house price indicators at a time such that a very narrow information set is used to generate forecasts of house prices. Valuation ratios (e.g., price-to-rent or price-to-income ratios) are among the most commonly used predictors of future house prices. These ratios work well as predictive variables at long horizons, but they are not necessarily useful at shorter horizons.¹ Using only a single or a few selected variables at a time to some extent appears inefficient when predicting future house prices because movements in house prices may reflect many different sources of information. Therefore, it may be possible to obtain more accurate house price forecasts by conditioning on a rich information set instead of only a few variables.

Motivated by the above, this paper examines the ability to forecast real house prices at short and long horizons using a common factor approach in which we employ information from a very large number of macroeconomic and financial time series. The basic idea is to

¹See, for example, Campbell et al. (2009) and Gallin (2008) for in-depth analyses of the predictive power of the price-rent ratio.

summarize a large amount of information in a relatively small number of estimated factors and, as the next step, we use these estimated factors to forecast housing price fluctuations. In this way, we are able to make use of a much richer information set in comparison to previous studies on housing price forecastability. Essentially, the methodology that we apply in this paper enables us to condition the house price forecasts on a large information set involving more than 100 macroeconomic and financial variables. This is in sharp contrast to the typical predictive regression where only up to about a handful of observed variables is included in the predictor set. The key to embedding such a large information set in the regressions derives from a factor analysis of the panel where a few latent common factors are shown to describe the majority of the variation of the series in the panel. The factor analysis thus makes it possible to effectively reduce the dimension of the predictor set while still being able to summarize and use the underlying information in the panel. Furthermore, Stock and Watson (2002a, 2008) find that the common factor approach is robust to structural instability that often plagues predictive regressions because of the many different sources of information that shape each factor.

Based on a large panel of 122 economic time series, we show that the estimated common factors contain substantial information about future movements in real house prices. In particular, a three-factor model is able to explain as much as around 50% of the variation in one-quarter ahead growth rates in real house prices. The predictive power also stays high at longer forecasting horizons. The three predictive factors can be interpreted as an economic activity, an inflation, and an interest rate factor. Following Gospodinov and Ng (2011), we conduct inference using a bootstrap procedure to address potential statistical issues originating from our use of estimated regressors and to take into

account possible time series dependencies in the predictive regressions. The bootstrap resampling method suggests that the three factors are generally strongly statistically significant.

Our results are robust using both in-sample and out-of-sample forecasting regressions. In the out-of-sample forecasting, we compare with richly parameterized autoregressive models, and the simple three-factor model consistently beats the autoregressive benchmark across all forecasting horizons. This holds true when taking into account announcement delays of the macroeconomic series. We also conduct out-of-sample forecasting using a factor forecast combination approach in which optimal weights are recursively chosen based on criteria such as the past forecasting performance. Again, the results show that utilizing information from the large panel of economic series leads to smaller forecast errors than autoregressive benchmarks. In addition, we find that the simple three-factor model works remarkably well in comparison to the computational intensive forecast combination models.

Focusing on the very volatile period around the recent boom and burst of the housing market, we find that the three-factor model contains useful information in the sense that it predicts with the right sign. Admittedly, the model does not fully capture the very large growth rates around the peak of the house price boom, but it does predict positive growth rates, i.e., the sign of the forecast is correct. Likewise, the model does not fully capture the sharp decline in house prices when the crash occurred, but it does predict negative growth rates. Thus, also when the crash occurred, the model gets the direction right.

Related literature. Case and Shiller (1989, 1990) were the first to document that housing prices do not follow a random walk. They show that housing returns exhibit positive autocorrelation and that various information variables predict future housing returns. Many recent studies examine predictability in the housing market using the price-rent ratio, which is analogous to the price-dividend ratio often used to forecast the stock market. Gallin (2008) examines the long run relationship between house prices and rents and finds that they are cointegrated and that the price-rent ratio contains useful information for predicting housing returns at long horizons. Favilukis et al. (2012) develop a general equilibrium model of the housing market, and their model implies that a high price-rent ratio predicts low future housing returns. They provide empirical evidence consistent with this implication. Campbell et al. (2009) and Cochrane (2011) also give empirical evidence that a high price-rent ratio is a signal of low future housing returns. Our focus is to forecast the growth in house prices conditioning on a large and more general information set. We find that using the large panel of economic time series gives much better out-of-sample predictive power relative to using the price-rent ratio. Actually, we find that the price-rent ratio performs worse than the historical mean in out-of-sample regressions, i.e., just like Goyal and Welch (2003, 2008) document that the price-dividend ratio is not able to beat the historical mean when forecasting stock returns out-of-sample.

As an exception, Rapach and Strauss (2009) examine housing price forecastability using more than just a single or a few selected predictors. They analyze differences in housing price forecastability across the 20 largest U.S. states in terms of population. One of their findings is that the degree of predictability is lower in states with high average

house price growth (i.e., coastal states) than in states with low average price growth (i.e., interior states). As the focus of our paper is to forecast national house prices and not to examine the cross-sectional variation in forecasting power across regions, we only include national variables in our panel, and this also allows us to base our forecasts on a much larger set of economic variables. Specifically, we use 122 national economic series, while Rapach and Strauss (2009) only use 16. Exploiting a richer information set when forecasting house prices could potentially be very important as it intuitively should lead to more accurate forecasts, and it makes it possible to more fully examine the degree of predictability in housing prices. Another important difference is that our sample period includes the 2007 crash and the subsequent volatile period in the housing market.

The empirical methodology that we apply in this paper to predict housing returns has also been used to predict stock returns (Ludvigson and Ng 2007) and to predict bond returns (Ludvigson and Ng 2009, 2011). The innovative methodological feature of our paper is that we also consider factor forecast combination and use a bootstrap resampling method to take into account that the factors are estimated regressors.

The rest of the paper is organized as follows. Section 2 describes our empirical methodology, section 3 provides empirical results, and section 4 concludes.

2 Empirical methodology

In this section, we first describe how we estimate the factors using asymptotic principal component analysis. Then we describe how we run predictive regressions in which the predictor set includes time series of common factors from the factor analysis.

2.1 Factor model and the estimation of factors

In recent years, factor models have become a standard tool in applied macroeconomics and finance. Essentially, when the number of random sources of variation is less than the number of dependent variables, then a factor model enables the researcher to reduce the dimension of the number of explanatory variables to a few latent factors. Since the first generation of (exact) factor models by Geweke (1977) and Sargent and Sims (1977), a considerable amount of research has been devoted to the econometric theory and empirical analysis of large dimensional dynamic factor models. In particular, building on the approximate factor model of Chamberlain and Rothschild (1983), the large dimensional approximate dynamic factor model is introduced by Forni, Hallin, Lippi and Reichlin (2000, 2004, 2005) in the frequency domain and by Stock and Watson (2002a, 2002b) in the time domain.² They estimate the large dimensional dynamic factor model non-parametrically by dynamic and static principal component methods, respectively, but recently these models have also been estimated by Bayesian methods (Otrok and Whiteman 1998; Kim and Nelson 1999) as well as by maximum likelihood methods (Doz, Giannone and Reichlin 2011a, 2011b; Jungbacker and Koopman 2008).

In our forecast analysis and similar to Ludvigson and Ng (2007, 2009, 2011), we implement the static principal component method of Stock and Watson because this method is fast, easy to implement, and given the sample size we consider, this method also performs similarly well compared to dynamic principal components (Boivin and Ng

²By 'large' we mean large in the cross-section, i.e., large in the number of time series (N), and large in the number of observations (T) of the time series; for instance $N = 100^+$ and $T = 100^+$ depending on the frequency of the data. By 'approximate' we refer to the relaxation of the iid error term assumption in the exact factor model such that the error terms are allowed to be weakly (locally) correlated, cf. Chamberlain and Rothschild (1983).

2005) as well as maximum likelihood methods (Doz, Giannone and Reichlin 2011a). We now briefly describe the static principal component method, and for this purpose we present the dynamic factor model.

2.1.1 Dynamic factor model

Consider a panel of observable economic variables $X_{i,t}$, where i denotes the cross-section unit, $i = 1, \dots, N$, while t refers to the time index, $t = 1, \dots, T$. The panel of observed economic variables is transformed into stationary variables with zero mean and unit variance and denoted by $x_{i,t}$. The key implication of the dynamic approximate factor model is that the variation of each of the N observed variables can be decomposed into a common component, χ_t , that captures the cross-sectional comovement and an idiosyncratic component, ξ_t . Furthermore, the cross-sectional comovement of the variables is entirely driven by $r \ll N$ common factors denoted F_t through series specific factor loadings, Λ_i . For the i th variable we write:

$$\begin{aligned} x_{i,t} &= \chi_{i,t} + \xi_{i,t} \\ &= \Lambda_i^\top F_t + \xi_{i,t} \end{aligned} \tag{1}$$

where F_t is an $r \times 1$ vector, Λ_i is an $r \times 1$ vector of factor loadings for the i th observed variable, and where the idiosyncratic component $\xi_{i,t}$ may have a limited amount of cross-sectional correlation.³ The model in (1) is often labeled as the static form of the dynamic

³Formally, this limited cross-sectional correlation is ensured by imposing the condition that the contribution of the covariance of the idiosyncratic terms to the total covariance of x as N gets large is bounded (by a constant M):

$$\sum_{i=1}^N \sum_{j=1}^N |E[\xi_{i,t} \xi_{j,t}]| \leq M$$

See Stock and Watson (2006) or Ludvigson and Ng (2009).

factor model because the (static) factors F_t only enter contemporaneously, but this is, however, merely a notational artifact.⁴

The dynamic factor model in (1) is estimated by static principal components method which can be seen as a solution to the least squares problem:

$$\min_{F^{(k)}, \Lambda^{(k)}} V(F^{(k)}, \Lambda^{(k)}) \text{ with } V(F^{(k)}, \Lambda^{(k)}) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \left(x_{i,t} - \Lambda_i^{(k)\top} F_t^{(k)} \right)^2 \quad (2)$$

where k refers to number of factors involved in the minimization problem such that $F^{(k)}$ becomes a $T \times k$ matrix of estimated factors and $\Lambda^{(k)} = (\Lambda_1, \Lambda_2, \dots, \Lambda_k)'$ is an $N \times k$ matrix of estimated loadings. As the model in (1) is not econometrically identified, a total of k^2 restrictions are imposed; in this case the standard $k(k+1)/2$ normalization restrictions given by $F^{(k)\top} F^{(k)} = I_k$, and the requirement that $\Lambda^{(k)\top} \Lambda^{(k)}$ is diagonal, which imposes an additional $k(k-1)/2$ restrictions on the symmetric matrix $\Lambda^{(k)\top} \Lambda^{(k)}$. The solution to this least squares problem is $\hat{F}^{(k)} = \sqrt{T} \hat{P}^{(k)}$ and $\hat{\Lambda}^{(k)} = x^\top \hat{F}^{(k)} / T$, where $\hat{P}^{(k)}$ is the $T \times k$ matrix of eigenvectors corresponding to the k largest eigenvalues of $xx^\top / NT$; see, e.g., Stock and Watson (1998). Consistency of the principal component estimator of F_t is shown by Connor and Korajczyk (1986), Stock and Watson (2002a), and Bai and Ng (2006). In particular, Bai and Ng (2006) provide improved rates under which the estimated factors $\hat{F}_t$ can be treated as observed, and hence inference about the parameters obtained in our second-stage predictions are not necessarily affected by the fact that the factors are estimated. However, as a precautionary step we base our

⁴Notice, we could specify a model with the common component given by $\chi_{i,t} = \lambda_i^T(L) \mathbf{f}_t$, where $\lambda_i(L)$ is a $q \times 1$ lag polynomial of finite order s representing the dynamic loadings, and $\mathbf{f}_t$ is $q \times 1$ dimensional vector of dynamic factors. Stacking the loadings and the dynamic factors into the $q(s+1) = r$ vectors Λ_i and F_t , respectively, the *static* representation of the dynamic approximate factor model of Stock and Watson (2002b) follows.

inference on a bootstrap resampling procedure, which we detail below in Section 2.4.

We need to determine the number of factors k involved in the principal component analysis above. Econometric theory for the determination of the number of factors has recently been developed for both the dynamic factor framework (Hallin and Liska 2007; Stock and Watson 2005; Bai and Ng 2007) as well as for the static factor framework (Bai and Ng 2002). We apply the information criterion IC_2 of Bai and Ng (2002) and as detailed in the next section, we find $r = 11$ factors. Accordingly, the principal component analysis in (2) proceeds by setting $k = r = 11$.

2.2 Predictive regressions

Our purpose is to forecast the log real house price growth, $y_{t+h} = 100 \times \ln(P_{t+h}/P_t)$, where P_t denotes the real house price at time t , and h is the forecasting horizon. To do this, we use predictive regressions of the form:

$$y_{t+h} = \alpha + \beta(L)' \hat{f}_t + \gamma(L) y_t + \varepsilon_{t+h} \quad (3)$$

where the vector $\hat{f}_t \subset \hat{F}_t$ contains estimated common factors that are relevant for forecasting h -period ahead real house price growth rates. $\beta(L)$ and $\gamma(L)$ are lag polynomials.

We carry out both in-sample and out-of-sample forecasting analyses. The advantage of in-sample regressions is that all information is exploited, and therefore in-sample forecasting regressions is the most useful when it comes to examining the true relationship between the set of predictors and future house price growth rates. The disadvantage of in-sample forecasting regressions is that it does not tell us whether the forecasting model

would have been useful to an economic agent who operates in real time.

We conduct out-of-sample forecasting based on a recursive scheme using all available data at the time of the forecast. We divide the full sample of $T - h$ observations into an initial estimation period of $T_1 - h$ observations and an out-of-sample period of T_2 observations. Thus, using a procedure where we recursively estimate the common factors as well as the parameters of the model, we generate a series of in total T_2 forecasts of y_{t+h} . If we let the out-of-sample window depend on h rather than the initial estimation period, our results do not change to any noteworthy extent. We generate the out-of-sample forecasts using the unrestricted model given in Eq. (3) and compare with an autoregressive benchmark model in which we set $\beta(L) = 0$.⁵ Let $\hat{\varepsilon}_{U,t+h}^h = y_{t+h} - \left(\hat{\alpha}_{U,t}^h + \hat{\beta}_{U,t}^h(L)' \hat{f}_t + \hat{\gamma}_{U,t}^h(L) y_t \right)$ denote the forecast error of the unrestricted forecasting model, and let $\hat{\varepsilon}_{R,t+h}^h = y_{t+h} - \left(\hat{\alpha}_{R,t}^h + \hat{\gamma}_{R,t}^h(L) y_t \right)$ denote the forecast error of the restricted forecasting model. The out-of-sample statistic that we use is then given by:

$$\text{MSFE-ratio} = \frac{\sum_{t=T_1-h}^{T-h} \left(\hat{\varepsilon}_{U,t+h}^h \right)^2}{\sum_{t=T_1-h}^{T-h} \left(\hat{\varepsilon}_{R,t+h}^h \right)^2} \quad (4)$$

where a MSFE-ratio (Mean-Squared-Forecast-Error) of less than 1 implies that the unrestricted model produces a smaller mean squared forecast error than that of the restricted model.

⁵We use the autoregressive model as our main benchmark, but also show results using a number of alternative benchmarks, which we detail further below.

2.3 Factor forecast combination

Common factors effectively reduce the dimension of the predictor set and alleviate the instability issues arising from, e.g., structural shifts in the regressors. However, it is still a challenging task to specify a good forecasting model as a large number of candidate forecast models could arise out of the eleven factors and their potential lags combined with lagged house price growth rates. To overcome this challenge, we also consider a forecast combination approach; see Stock and Watson (2001), Timmermann (2006) and Aiolfi et al. (2010). In forecast combinations, forecasts from a large number of individual forecast models are combined to produce a weighted forecast, $\hat{y}_{t+h|t}^{(c)}$. This approach is often found to offer good empirical performance over individual model forecasts.

In our application, we combine a large number of forecasts from different univariate and multivariate candidate forecast models involving factors and lagged house price growth rates. In particular, the forecast from the i^{th} candidate model takes the form:

$$\hat{y}_{i,t+h|t} = \hat{\alpha}_i + \hat{\beta}_i(L)' \hat{f}_t + \hat{\gamma}_i(L) y_t \quad (5)$$

where at most four factors of any combination of the eleven factors enters a particular specification. $\hat{\beta}_i(L)$ is at most of order three ($1 + \hat{\beta}_{i1}L + \hat{\beta}_{i2}L^2 + \hat{\beta}_{i3}L^3$), and in case the lagged dependent variable enters, $\hat{\gamma}_i(L)$ is at most of order five ($1 + \hat{\gamma}_{i1}L + \hat{\gamma}_{i2}L^2 + \hat{\gamma}_{i3}L^3 + \hat{\gamma}_{i4}L^4 + \hat{\gamma}_{i5}L^5$). For a given forecast horizon h , a total of $M = 13,488$ models were estimated recursively for each of the T_2 forecast dates. However, once the i^{th} candidate forecast model has been estimated for all possible lag structures, we use the Schwartz information criterion (SIC) to choose the best lag specification. As an example, suppose

that the i^{th} model involves four regressors: $\hat{f}_{1,t}$, $\hat{f}_{2,t}$, $\hat{f}_{3,t}$, and y_t . Given the restrictions on the lag polynomials, a total of 4×6 SIC values are then estimated for this particular model at time t , and for this particular model we then choose the best fitting model. This procedure effectively reduces the number of candidate models to $m = 562$.

The combined forecast is then calculated as:

$$\hat{y}_{t+h|t}^{(c)} = \sum_{i=1}^m \hat{\omega}_{i,t+h|t} \hat{y}_{i,t+h|t} \quad (6)$$

where the time t weight assigned to the forecast from the i^{th} candidate model takes a number of forms in the literature. We consider three popular weighting schemes. The first weighting scheme is based on the Schwartz information criterion (SIC) which can be viewed as Bayesian model averaging weights. The second scheme is based on the past MSFE performance, and the third is based on past discounted MSFE performance. Specifically, we apply the following three weighting schemes:

$$\hat{\omega}_{i,t+h|t} = \left\{ \begin{array}{l} \exp \{ -\Delta SIC_{i,t|t-h}/2 \} / \sum_{i=1}^m \exp \{ -\Delta SIC_{i,t|t-h}/2 \} \\ MSFE_{i,t|t-h}^{-1} / \sum_{i=1}^m MSFE_{i,t|t-h}^{-1} \\ DMSFE_{i,t|t-h}^{-1} / \sum_{i=1}^m DMSFE_{i,t|t-h}^{-1} \end{array} \right\} \quad (7)$$

where $\Delta SIC_{i,t|t-h}$ refers to the difference between the SIC criterion for the i^{th} model at time t minus the time t best-fitting model. The $MSFE_{i,t|t-h}$ is calculated over a window of the previous v periods:

$$MSFE_{i,t|t-h} = \frac{1}{v} \sum_{\tau=t-v}^t (y_{i,\tau} - \hat{y}_{i,\tau|t-h})^2 \quad (8)$$

and $DMSFE_{i,t|t-h}$ refers to the discounted MSFE:⁶

$$DMSFE_{i,t|t-h} = \frac{1}{v} \sum_{\tau=t-v}^t \theta^{t-\tau} (y_{i,\tau} - \hat{y}_{i,\tau|t-h})^2 \quad (9)$$

In the empirical section we provide out-of-sample results from this forecast combination approach.

2.4 Bootstrap

We address potential small sample distortions in the inference about predictive regressions by a non-parametric moving block bootstrap that preserves time-series dependencies. On top of this, the moving block bootstrap also addresses the issue of generated regressors (the factors) in the predictive regressions, although Bai and Ng (2006) have shown that for a large cross-sectional dimension relative to the time-series dimension, we can ignore uncertainty in the factor estimates.

The non-parametric moving block bootstrap method resamples the data in blocks of consecutive observations across all variables in order to reproduce possible time series dependencies due to, e.g., serial correlation and heteroskedasticity and to preserve any cross-sectional dependencies in the data. In our application, the resampling procedure thus needs to draw blocks simultaneously from the dependent variable y and the panel x , and from the resampled panel the factors are subsequently estimated. Then for each bootstrap sample the predictive coefficients in Eq. (3) are computed, and based on a large number of bootstraps the bootstrap confidence intervals can be computed; see the details below.

⁶We use the typical value of $\theta = 0.9$ and set $v = 12$ quarters.

Consider stacking the dependent variable in $\mathcal{Y} = (y_{h+1}, y_{h+2}, \dots, y_T)'$, the panel in $\mathcal{X} = (x_1, x_2, \dots, x_{T-h})'$, and in a similar way K lags of the dependent variable in $\mathcal{W}$. We then represent the matrices $\mathcal{Y}$, $\mathcal{X}$, and $\mathcal{W}$ in a single 'parent' matrix $\mathcal{B}$ of dimension $(T - h) \times (1 + N + K)$:

$$\mathcal{B} = (\mathcal{Y}, \mathcal{X}, \mathcal{W}) \quad (10)$$

which is subsequently block-resampled with replacement yielding a particular bootstrap sample $\mathcal{B}^*$. The factors are re-estimated using the corresponding $\mathcal{X}^*$ resulting in a set of factors F^* from which a subset $f^* \subset F^*$ is used along with K lags of the dependent variable from $\mathcal{W}^*$, as described by the predictive regression in Eq. (3). Before detailing the resampling procedure, it can be noted that for a simple predictive regression in the form of $y_{t+h} = \alpha + \beta' f_t + \gamma y_t + \varepsilon_{t+h}$, the $\mathcal{B}$ matrix is particularly simple:

$$\mathcal{B} = \begin{bmatrix} y_{h+1} & x_{1,1} & \cdots & x_{N,1} & y_1 \\ y_{h+2} & x_{1,2} & \cdots & x_{N,2} & y_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ y_T & x_{1,T-h} & \cdots & x_{N,T-h} & y_{T-h} \end{bmatrix}$$

A given bootstrap sample $\mathcal{B}^*$ is essentially composed of a number of randomly selected blocks of size $w \times (1 + N + K)$ that are stacked upon each other so that the size of $\mathcal{B}^*$ is the same size as $\mathcal{B}$. The number of blocks, b , is the integer of T/w , and the length w of each block can be computed using automatic block-length procedures as in, e.g., Patton et al. (2009)⁷. Specifically, a particular bootstrap sample can be generated

⁷We are grateful to Andrew Patton for providing MATLAB code on his homepage to compute the automatic block-length. We find $w = 20$ and varying this number does not change our inference significantly. We are also grateful to Serena Ng for providing MATLAB code for bootstrapping factor models, which we modified slightly to fit our application.

by drawing with replacement (on a so-called circle) b *iid* uniform random variables, $\{u_i\}_{i=1}^b \in \{1, 2, \dots, (T - h)\}$, where u_i essentially determines at which row in $\mathcal{B}$ we select the i th block $\mathcal{B}_i^*$ with w rows $(u_i, u_i + 1, \dots, u_i + w - 1)$ and $(1 + N + K)$ columns. Then a particular bootstrap sample $\mathcal{B}^*$ can be written as:

$$\mathcal{B}^* = (\mathcal{B}_1^*, \mathcal{B}_2^*, \dots, \mathcal{B}_b^*)'$$

For each of the $j = 1, \dots, 5000$ bootstraps, we run predictive regressions like Eq. (3) but here conditioning on the j^{th} bootstrap data $\{\mathcal{B}_j^*, F_j^*\}$:

$$y_{j,t+h}^* = \alpha_j^* + \beta_j^*(L)' \hat{f}_{j,t}^* + \gamma_j^*(L) y_{j,t}^* + \varepsilon_{j,t+h}^* \quad (11)$$

We collect the estimated coefficients from Eq. (11) in a vector $\hat{\theta}_j^*$ with *s.e.* $(\hat{\theta}_j^*)$ denoting the corresponding HAC standard errors.⁸ Then we construct the following quantity for the ℓ^{th} element of $\hat{\theta}_j^*$:

$$t_{p\ell,j} = \frac{(\hat{\theta}_{\ell,j}^* - \hat{\theta}_\ell)}{s.e.(\hat{\theta}_{\ell,j}^*)} \quad (12)$$

where $\hat{\theta}_\ell$ refers to the estimate from Eq. (3) conditioning on the observed sample. Following Gospodinov and Ng (2011), we let v_p^* and $v_{(1-p)}^*$ denote the p^{th} and $(1 - p)^{th}$ element of the sorted sequence of $t_{p\ell,j}$'s, and obtain the $100(1 - p)\%$ percentile bootstrap confidence interval for $\hat{\theta}_\ell$ as:

$$\left[\hat{\theta}_\ell - s.e.(\hat{\theta}_\ell) v_{(1-p/2)}^*, \hat{\theta}_\ell + s.e.(\hat{\theta}_\ell) v_{p/2}^* \right] \quad (13)$$

⁸The lag truncation in the Newey-West HAC standard errors are set to $h+3$ to accommodate potential increasing residual autocorrelation problems in long-run regressions with overlapping observations. Due to the sign indeterminacy of the principal components, we ensure that the j^{th} bootstrap factor stays positively correlated with the j^{th} parent factor; possibly multiplying $\hat{f}_{j,t}^*$ by (-1) .

These bootstrap confidence intervals are applied to every in-sample regression and are designed to address the issue of generated regressors as well as the issue of time series dependencies including residual autocorrelation due to long-run regressions with overlapping observations.

3 Empirical results

Our sample is quarterly and runs from 1975:1 to 2011:2. We measure nominal house prices based on the all-transactions house price index available from the Federal Housing Finance Administration (FHFA).⁹ We obtain a real house price index by dividing the nominal house price index with the personal consumption deflator available from the Bureau of Economic Analysis. In Figure 1, we depict the time series of real house price growth over our sample period. The figure clearly illustrates the boom in the U.S. housing market from the beginning of the mid-1990s up to around 2006 as well as the crash in house prices in 2007. Furthermore, the figure illustrates an increase in volatility during the recent period of economic crisis.

We use a panel of 122 economic series to estimate the factors. The list of series is provided in the Appendix. The series represent the following categories of macroeconomic time series: output and income; employment, hours and earnings; housing; consumption, orders and inventories; money and credit; bond and exchange rates; consumer, producer and commodity prices; and stock prices.

When choosing how many principal components to retain, we apply the panel infor-

⁹Formerly Office of Federal Housing Enterprise Oversight (OFHEO).

mation criteria developed by Bai and Ng (2002). In particular, when we use their IC_2 criterion, we find that the appropriate number of components is eleven. As we show in Table 1, these eleven factors are able to account for a large part of the total variance in the panel: Almost 80% of the total variance is attributed to these eleven factors. The table also shows that about 50% of the total variability in the panel is accounted for by the first three factors.

In Table 1 we also report 1st order autocorrelation coefficients of the estimated factors. All factors (except the 9th factor) have positive persistence, and the degree of persistence varies somewhat across the factors. The second factor is the one with the highest degree of persistence with an AR(1) coefficient of 0.79, which is far from the unit root.

3.1 In-sample results

We first assess the in-sample predictive power of the estimated common factors. We do so by estimating the forecasting model in Eq. (3) using the full sample. We are only interested in the factors containing useful information about future growth rates in real house prices. In general, across forecast horizons, statistical significance and information criteria suggest that it is sufficient to include the first three factors out of the eleven estimated factors. We also adopt a parsimonious lag structure of order one as higher-order lags of $\hat{f}_t$ and y_t tend not to contribute with much predictive power. In the in-sample analysis, we therefore restrict the attention to the following simple three-factor

model:¹⁰

$$y_{t+h} = \alpha + \beta_1 \hat{f}_{1,t} + \beta_2 \hat{f}_{2,t} + \beta_3 \hat{f}_{3,t} + \gamma y_t + \varepsilon_{t+h} \quad (14)$$

For each forecast horizon that we consider ($h = 1, 2, 4$, and 8), Table 2 reports OLS estimates of the slope coefficients, 90% bootstrap confidence intervals of the slope coefficients, the adjusted R^2 -statistic, and 90% bootstrap confidence intervals of the adjusted R^2 -statistic.

Table 2 shows that the three estimated predictive factors are in almost all cases statistically significant according to the bootstrap resampling method. The only exception is when using $h = 8$, where the slope coefficient on the first factor is inside the bootstrap confidence interval. Likewise, the lagged house price growth rate is also in general statistically significant across horizons (only exception is when using $h = 2$). Looking at the adjusted R^2 statistic, the model is able to explain as much as 53.7% of the variation in one-quarter ahead house growth rates. The predictive power of the model stays high when increasing the forecast horizon: 53.2% at the 1-year horizon and 41.0% at the 2-year horizon. Thus, the model works well at various forecasting horizons. We also find that the bootstrap confidence interval for the adjusted R^2 statistic is well above 0 for all forecast horizons.

If we drop y_t from the model and only consider the information from the panel of economic series, the model continues to do well in predicting future house prices. Using the three predictive factors alone and excluding y_t from the set of predictors, the adjusted R^2 s are given by 51.7%, 48.9%, 48.1%, and 36.8% at horizon $h = 1, 2, 4$, and 8 .¹¹

¹⁰In the out-of-sample analysis, we experiment with forecast combination techniques where the lag length and the factor combination are recursively chosen based on Schwartz information criterion weights as well as inverse MSFE-weights (see section 3.2).

¹¹A table with results is available upon request.

Comparing with the adjusted R^2 s from Table 2, we thus observe a relatively small fall in predictive power when leaving out y_t from the model. The high predictive power of the model is therefore not driven by the information contained in past house price growth rates.

The three-factor model is able to forecast about 50% of the variation in the one-quarter ahead housing return. This high degree of forecastability suggests that our results are not only statistically significant, but also economically significant. Further, the strong predictive power of the factor model that utilizes information from a very large number of economic variables also suggests that it may be insufficient and misleading to base house price forecasts on only a single or a few selected predictors. In unreported results, we find that the price-rent ratio – the most commonly used house return predictor – does not contain much predictive power in comparison to the three-factor model and is statistically insignificant when adding it to the model. This is the case at both short and long horizons. Detailed results are available upon request.

To illustrate how well our model works in-sample, Figure 2 plots the actual times series of real house price growth together with the forecasted values from the augmented three-factor model. The model convincingly captures large swings as well as peaks and troughs in housing returns. Even in the recent period with rapid house price fluctuations the model works reasonably well in predicting house prices. We also stress that our model does not fully capture the house price boom in the period from around 2004 to 2006 where the house price growth rates were very large. In this period the model does predict positive growth rates in real house prices, but it tends to underestimate the level of the growth rates. Similarly, the model does predict negative growth rates in 2007 when

the crash occurred, but it does not fully capture the sharp decline in house prices. Still, the sign of the forecast is correct and the model seems to contain important predictive ability even in periods with very volatile growth rates.

3.1.1 Factor interpretation

In the following, we give an economic interpretation of our estimated common factors. We do this by looking at the R^2 s from regressing each of the series in the panel on each of the factors one at a time. Figure 3 illustrates that $\hat{f}_{1,t}$ is an economic activity factor. It loads heavily on industrial production and employment data. From Figure 4, we see that $\hat{f}_{2,t}$ may be interpreted as an inflation factor since it loads most heavily on the various price indices that we transformed to standardized quarterly inflation rates. The factor also loads heavily on interest rates spreads, but whereas the various inflation rates are positively correlated with $\hat{f}_{2,t}$, the interest rate spreads are negatively correlated with $\hat{f}_{2,t}$.¹² We find this inverse relationship between inflation and interest rate spreads quite intuitive as high inflation rates would force the Federal Reserve to increase the short-term monetary policy rate relative to the long-term interest rates in order to bring down inflation. Figure 5 indicates that $\hat{f}_{3,t}$ loads heavily on notably the longer-term interest rate related series (first differences) as well as housing variables. There is an inverse relationship between changes in interest rates and housing variables like housing starts, whereas the relationship is positive for houses for sale relative to houses sold. Accordingly, rising long-term interest rates makes it more difficult to finance houses and increases the selling period. Unreported results show that the absolute correlations between interest rates

¹²Notice also that to enhance interpretation we can rotate the factors and loadings by multiplying by (-1) . This only changes the sign of the regression coefficients, not the magnitude.

changes and $\hat{f}_{3,t}$ are more than twice as high as the absolute correlations with housing variables until summer 2009. However, for the last two years of the sample, housing starts and house selling periods begin to dominate $\hat{f}_{3,t}$. On this background, we interpret the third factor as primarily an interest rate factor. Consistent with this interpretation, Thom (1985) provides evidence that housing starts are significantly influenced by interest rates and not the other way around.

We now relate the factor interpretations to the sign of the slope coefficients in the predictive regressions (see Table 2). The economic activity factor ($\hat{f}_{1,t}$) has a positive slope coefficient for every forecast horizon h in the predictive regressions. It implies that expected housing returns move procyclical, i.e., expected house price changes are high when economic conditions are good and low when economic conditions are bad. This is in contrast to return predictability evidence on the stock market (see, e.g., the survey of Cochrane 2007), the bond market (see, e.g., Ludvigson and Ng 2009), and the currency market (see, e.g., Lustig et al. 2010). In these asset markets, expected returns move countercyclical because investors require a higher expected return in times of bad economic conditions. One reason why we should not expect to see the same predictability patterns in the housing market is that housing is both a consumption good and an investment good. Further, due to various market frictions on the housing market (transaction costs, search costs, transaction time, financing constraints, short sale constraints, etc.), it is very likely that new information is only reflected fully in house prices with a time lag. This time lag also implies that when the affordability of households improves in times of good economic conditions, we should expect to see a positive impact on future house price growth rates. Our findings of procyclicality in expected house price

changes is consistent with Case and Shiller (1990) who find that employment and income variables have a positive relation with future housing returns.

The inflation factor ($\hat{f}_{2,t}$) has a negative slope coefficient for every forecast horizon h in the predictive regressions, so low inflation rates predict high real house price changes. Fama and Schwert (1977), Fama (1981), among others, provide evidence of a negative relation between inflation and real returns on the stock market. Moreover, Brunnermeier and Julliard (2008) document a negative relation between inflation and housing returns due to money illusion. In times when the inflation is low, investors that suffer from money illusion tend to underestimate the real interest rate and, hence, tend to underestimate real mortgage payments. In turn, this drives house prices up.

The interest rate factor ($\hat{f}_{3,t}$) has a negative slope coefficient for every forecasting horizon h . This is a very intuitive result: Low interest rates lead to cheaper mortgage loans, which again lead to larger mortgages and higher expected house prices. Because housing starts variables load negatively on $\hat{f}_{3,t}$, increasing housing starts, due to easy house financing, coincides with higher expected house prices.

3.2 Out-of-sample results

So far we have focused on in-sample forecasting regressions. To check the robustness of our in-sample results, we now turn to the out-of-sample evidence. In-sample forecasting makes use of full-sample estimates, while out-of-sample forecasting tests how sensitive the model is towards unstable predictive coefficients. Also, in the in-sample forecasting regressions, we generate the predictive factors once using the full sample of information, while in the out-of-sample regressions, we estimate the predictive factors recursively. In

the out-of-sample analysis, we thereby address the potential concern of look-ahead bias. Given the fact that many of the time series in our panel are subject to data revisions, another potential concern is that the series that are available today are different from the series available in real time. We only have access to vintage data for a limited part of our panel of time series, and we do not have vintage data for the house price index for a sufficiently long period of time.¹³ We therefore follow Rapach and Strauss (2009), Ludvigson and Ng (2007, 2009, 2011), among many others, and conduct out-of-sample forecasting using the today-available time series. We do, on the other hand, take into account that many of the series in our panel are announced with a delay of up to one quarter. Thus, in the out-of-sample analysis where the goal is to mimic the situation of a real-time forecaster, we lag our predictive variables an additional quarter.

In our out-of-sample forecasting exercise, we use an initial estimation period of 20 years from 1975:1 to 1994:4, and the out-of-sample period thus runs from 1995:1 to 2011:2 (16.5 years). This period covers the recent boom and burst of the U.S. housing market. We conduct out-of-sample forecasting using two approaches. In the first approach the real time forecaster chooses the factors to be included in the model based on the initial estimation period and sticks with that model throughout the out-of-sample period. We recursively estimate the factors and the model parameters, but the choice of factors stays the same. In the second approach the real time forecaster utilizes factor forecast combination models. The out-of-sample forecasting power of the two approaches is compared with that of an autoregressive model in which the optimal number of lags is recursively chosen based on the Schwartz information criterion. We denote this model

¹³In fact, the first vintage for the national house price index is from August 2010 in St. Louis Fed's ALFRED real-time database.

by AR(SIC). To execute our program code within a reasonable time frame, we impose restrictions on the lag polynomials; see section 2.3. Specifically, the maximum number of lags allowed is six. However, the SIC criterion usually selects less than six lags.

Three-factor model. The Bai and Ng (2002) panel information criterion suggests that the panel of economic series is well described by eleven factors. However, it is not necessarily the case that all eleven factors are useful in predicting house price growth rates. Statistical significance and various information criteria suggest only to include the first three out of the eleven factors. This is the case when estimating the model in Eq. (3) on the full sample, but also when estimating the model on the initial estimation period. To control for autocorrelation we augment the factor model with lagged house price growth rates. As with the autoregressive benchmark, we choose the optimal number of lags in the factor model recursively based on the SIC.

The first row in Table 3 reports the out-of-sample results for augmented three-factor model. Across all forecasting horizons $h = 1, 2, 4, 8$, the three-factor model yields a lower mean-squared-forecast-error than the AR(SIC) model. By consistently beating the AR(SIC) model, the out-of-sample evidence confirms the in-sample evidence that the three-factor model contains useful information for predicting future house price growth rates.

Forecast combination. We now consider the results from the factor forecast combination approach which combines forecasts from a large number of individual models. The individual models are based on various combinations of the eleven common factors as well as lags of the house price growth rate as detailed in section 2.3. We choose the

optimal lag structure of the individual models recursively using the SIC. We then calculate forecasts of the house price growth rates by weighting the forecasts of the individual models. We employ three different weighting schemes. The first scheme is based on the SIC criterion, the second scheme is based on the past MSFE performance, and the third scheme is based on past discounted MSFE performance.

Table 3 reports the MSFE-ratio of the forecast combination models relative to the AR(SIC) model. It can be seen that the forecast combination models generally beat the autoregressive benchmark, especially at the long horizons. We also observe that the SIC weighting scheme tends to perform equally well as the two MSFE weighting schemes, and that the discounting procedure in DMSFE does not seem to matter for the forecast performance.

Interestingly, the simple three-factor model performs very well compared to both the AR(SIC) model as well as to the much more computational intensive factor forecast combination models. In fact, the three-factor model performs better in MSFE ratio terms than the AR(SIC) model and the three variations of factor forecast combination methods for all forecast horizons up to one year, while for a two-year forecast horizon the forecast combination methods perform slightly better than the simple three-factor model.

One of the explanations for the success of the forecast combination approach in empirical research is the robustness of this method towards structural shifts in one or more of the variables in the predictor set. Because the combination forecast is a weighted average of many candidate forecasting models, the fact that a few of the models become unstable does not change the forecast significantly. In this perspective, the moderate forecasting performance of the factor forecast combination in comparison to the three-factor model

is probably not surprising as the three-factor model is already robust towards structural shifts. But still, it is striking that among all these combinations, the factor combination approach could not find a particular weighting of various factor configurations that outperforms the simple three-factor model. Thus, the first three factors seem to be the relevant macroeconomic factors for forecasting the growth in real house prices.

Alternative benchmarks. We have also made comparisons with more simple benchmarks than the AR(SIC) model, such as the sample mean growth rate in house prices and autoregressive models with a fixed lag structure through time. Given the strong focus on the price-rent ratio, see, e.g., Gallin (2008), Campbell et al. (2010), Cochrane (2011), Favilukas et al. (2012), we have also compared with the price-rent ratio. As we show in Table 4, the three-factor model strongly outperforms these alternative benchmarks. Interestingly, we find that the price-rent ratio generates worse out-of-sample forecasts than the historical mean. This finding relates to Goyal and Welch (2003, 2008) who provide evidence that the price-dividend ratio and a number of other predictors have worse out-of-sample performance on the stock market than do the historical mean stock return. Thus, the price-rent ratio shares the same lack of ability to predict returns out-of-sample as do the price-dividend ratio.

4 Conclusions

This paper examines the ability to forecast real house price changes using a common factor approach in which we exploit information from 122 economic time series. Using three common factors that together account for about 50% of the variability in the panel,

we are able to explain more than 50% of the variation in one-quarter ahead growth rates in real house prices. The forecasting power of the three-factor model also stays high at longer horizons. The estimated predictive factors are generally strongly statistically significant according to a bootstrap resampling method, which we design to address statistical issues such as time series dependencies and the use of estimated regressors in the forecasting regressions.

The strong degree of predictability that we document using our panel approach suggests that it is insufficient and misleading to form house price forecasts based on a limited set of economic time series. As an illustration of this point, we find that the price-rent ratio — one of the most widely used house price indicators — performs worse than the historical mean in out-of-sample regressions. By contrast, the predictive power of the three-factor model is robust in out-of-sample regressions. The model strongly beats the historical mean, but also performs remarkably well compared to both an autoregressive benchmark with a rich lag structure as well as to much more computational intensive factor forecast combination models.

Appendix

This appendix presents the series in the panel. The first column of the table contains transformation codes where "lvl" indicates an untransformed series, say $X_{i,t}$. " Δ lvl" means $X_{i,t} - X_{i,t-1}$, "ln" means $\ln X_{i,t}$, and " Δ ln" means $\ln X_{i,t} - \ln X_{i,t-1}$. The second column contains longer descriptions of the variables. All series were downloaded from St. Louis Fed's FRED database.

Output and income

Δ ln	Personal Income (Chained 2005 Dollars, SA)
Δ ln	Disposable Personal Income (Chained 2005 Dollars, SA)
Δ ln	Personal Income Excluding Current Transfer Receipts (Chained 2005 Dollars, SA)
Δ ln	Gross Domestic Product (Chained 2005 Dollars, SA)
Δ ln	Industrial Production Index - Total Index (SA)
Δ ln	Industrial Production - Durable Manufacturing (SA)
Δ ln	Industrial Production Index - Final Products (SA)
Δ ln	Industrial Production Index - Consumer Goods (SA)
Δ ln	Industrial Production Index - Durable Consumer Goods (SA)
Δ ln	Industrial Production Index - Nondurable Consumer Goods (SA)
Δ ln	Industrial Production Index - Business Equipment (SA)
Δ ln	Industrial Production Index - Materials (SA)
Δ ln	Industrial Production Index - Durable Goods Materials (SA)
Δ ln	Industrial Production Index - Nondurable Goods Materials (SA)
Δ ln	Industrial Production Index - Manufacturing (SA)
lvl	Napm Production Index (SA)
lvl	Capacity Utilization (SA)

Employment, hours and earnings

Δ ln	Civilian Labor Force (Thous., SA)
Δ ln	Civilian Employment (Thous., SA)
Δ lvl	Unemployed (Thous., SA)
Δ lvl	Average Duration of Unemployment (Weeks, SA)
Δ ln	Civilians Unemployed - Less Than 5 Weeks (Thous., SA)
Δ ln	Civilians Unemployed for 5-14 Weeks (Thous., SA)
Δ ln	Civilians Unemployed - 15 Weeks & Over (Thous., SA)
Δ ln	Civilians Unemployed for 15-26 Weeks (Thous., SA)
Δ ln	Civilians Unemployed for 27 Weeks and Over (Thous., SA)
Δ ln	All Employees: Total Nonfarm (Thous., SA)
Δ ln	All Employees: Total Private Industries (Thous., SA)
Δ ln	All Employees: Goods-Producing Industries (Thous., SA)
Δ ln	All Employees: Mining and Logging (Thous., SA)
Δ ln	All Employees: Construction (Thous., SA)
Δ ln	All Employees: Manufacturing (Thous., SA)
Δ ln	All Employees: Durable Goods (Thous., SA)
Δ ln	All Employees: Nondurable Goods (Thous., SA)
Δ ln	All Employees: Service-Providing Industries (Thous., SA)
Δ ln	All Employees: Trade, Transportation & Utilities (Thous., SA)
Δ ln	All Employees: Wholesale Trade (Thous., SA)
Δ ln	All Employees: Retail Trade (Thous., SA)
Δ ln	All Employees: Financial Activities (Thous., SA)
Δ ln	All Employees: Government (Thous., SA)
Δ ln	All Employees: Information Services (Thous., SA)
Δ ln	All Employees: Professional & Business Services (Thous., SA)
lvl	Average Weekly Hours of Production and Nonsupervisory Employees: Goods
Δ lvl	Average Weekly Hours of Production and Nonsupervisory Employees: Construction
lvl	Napm Employment Index
Δ ln	Average Hourly Earnings of Production and Nonsupervisory Employees: Goods
Δ ln	Average Hourly Earnings of Production and Nonsupervisory Employees: Construction
Δ ln	Average Hourly Earnings of Production and Nonsupervisory Employees: Manufacturing
Δ ln	Average Hourly Earnings of Production and Nonsupervisory Employees: Total Private

Housing

ln	Housing Starts Total: New Privately Owned Housing Units Started (Thous., SA)
ln	Housing Starts in Midwest Census Region (Thous., SA)
ln	Housing Starts in Northeast Census Region (Thous., SA)
ln	Housing Starts in South Census Region (Thous., SA)
ln	Housing Starts in West Census Region (Thous., SA)
ln	New One Family Houses Sold: United States (Thous., SA)
ln	New Private Housing Units Authorized by Building Permits (Thous., SA)
ln	New Privately-Owned Housing Units Under Construction: Total (Thous., SA)
ln	New Homes Sold in the United States (Thous)
lvl	Median Number of Months on Sales Market
lvl	Ratio of Houses for Sale to Houses Sold (SA)

Consumption, orders and inventories

lvl	Purchasing Managers' Index
lvl	Napm New Orders Index
lvl	Napm Supplier Deliveries Index
lvl	Napm Inventories Index
Δ ln	Personal Consumption Expenditures (Chained 2005 Dollars)

Money and credit

Δ ln	Commercial and Industrial Loans at All Commercial Banks (SA)
Δ ln	Consumer (Individual) Loans at All Commercial Banks (SA)
Δ ln	Currency Component of M1 (SA)
Δ ln	M1 Money Stock (SA)
Δ ln	M2 Money Stock (SA)
Δ ln	Real Estate Loans at All Commercial Banks (SA)
lvl	Personal Saving Rate (%)
Δ ln	Total Consumer Credit Outstanding (SA)
Δ ln	Home Mortgages - Liabilities - Balance Sheet of Households and Nonprofit Organizations
Δ ln	Household Sector: Liabilities: Household Credit Market Debt Outstanding
Δ ln	Debt Outstanding Domestic Nonfinancial Sectors - Household, Consumer Credit Sector
Δ ln	Debt Outstanding Domestic Nonfinancial Sectors - Household, Home Mortgage Sector
Δ ln	Owners' Equity in Household Real Estate - Net Worth - Balance Sheet of Households and Nonprofit Organizations
Δ ln	Real Estate - Assets - Balance Sheet of Households and Nonprofit Organizations

Bond and exchange rates

Δ lvl	Interest Rate: Federal Funds (Effective)
Δ lvl	Interest Rate: U.S. Treasury Bills, Sec. Mkt., 1-Mo.
Δ lvl	Interest Rate: U.S. Treasury Bills, Sec. Mkt., 3-Mo.
Δ lvl	Interest Rate: U.S. Treasury Bills, Sec. Mkt., 6-Mo.
Δ lvl	Interest Rate: U.S. Treasury Const. Maturities, 1-Yr.
Δ lvl	Interest Rate: U.S. Treasury Const. Maturities, 3-Yr.
Δ lvl	Interest Rate: U.S. Treasury Const. Maturities, 5-Yr.
Δ lvl	Interest Rate: U.S. Treasury Const. Maturities, 7-Yr.
Δ lvl	Interest Rate: U.S. Treasury Const. Maturities, 10-Yr.
Δ lvl	Interest Rate: 30-Year Conventional Mortgage Rate
Δ lvl	Bond Yield: Moody's AAA Corporate
Δ lvl	Bond Yield: Moody's BAA Corporate
lvl	Spread: 3m – fed funds
lvl	Spread: 6m – fed funds
lvl	Spread: 1y – fed funds
lvl	Spread: 3y – fed funds
lvl	Spread: 5y – fed funds
lvl	Spread: 7y – fed funds
lvl	Spread: 10y – fed funds
lvl	Spread: AAA – fed funds
lvl	Spread: BAA – fed funds
lvl	Spread: BAA – AAA
Δ ln	Foreign Exchange Rate: Canadian Dollars to One U.S. Dollar
Δ ln	Foreign Exchange Rate: Japanese Yen to One U.S. Dollar
Δ ln	Foreign Exchange Rate: U.S. Dollars to One British Pound
Δ ln	Foreign Exchange Rate: Swiss Francs to One U.S. Dollar

Prices

$\Delta \ln$	Producer Price Index: Crude Materials for Further Processing (1982=100, SA)
$\Delta \ln$	Producer Price Index: Finished Consumer Foods (1982=100, SA)
$\Delta \ln$	Producer Price Index: Finished Goods (1982=100, SA)
$\Delta \ln$	Producer Price Index: Intermediate Materials: Supplies & Components (1982=100, SA)
$\Delta \ln$	Cpi-U: All Items (82-84=100, SA)
$\Delta \ln$	Cpi-U: Housing (82-84=100, SA)
$\Delta \ln$	Cpi-U: Transportation (82-84=100, SA)
$\Delta \ln$	Cpi-U: Commodities (82-84=100, SA)
$\Delta \ln$	Cpi-U: Durables (82-84=100, SA)
$\Delta \ln$	Cpi-U: Nondurables (82-84=100, SA)
$\Delta \ln$	Cpi-U: All Items Less Food (82-84=100, SA)
$\Delta \ln$	Cpi-U: All Items Less Shelter (82-84=100, SA)
$\Delta \ln$	Spot Oil Price: West Texas Intermediate
$\Delta \ln$	Personal Consumption Expenditures: Chain-type Price Index (2005=100, SA)
$\Delta \ln$	Gross Domestic Product: Implicit Price Deflator (2005=100, SA)

Stock market

$\Delta \ln$	S&P Composite Index Level
$\Delta \ln$	Dow Jones Industrial Average

References

- Aiolfi, M., Capistrán, C., & Timmermann, A. (2010). Forecast combinations. In: Clements, M. & Hendry, D., (Eds.), *Forecast Handbook*, Oxford.
- Bai, J., & Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70, 191-221.
- Bai, J., & Ng, S. (2006). Confidence intervals for diffusion index forecasts and inference for factor-augmented regressions. *Econometrica*, 74, 1133-1150.
- Bai, J., & Ng, S. (2007). Determining the number of primitive shocks in factor models. *Journal of Business & Economic Statistics*, 25, 52-60.
- Boivin, J., & Ng, S. (2005). Understanding and comparing factor-based forecasts. *International Journal of Central Banking*, 1, 117-151.
- Brunnermeier, M. K., & Julliard, C. (2008). Money illusion and housing frenzies. *Review of Financial Studies*, 21, 135-180.
- Campbell, S. D., Davis, M. A., Gallin, J., & Martin, R. F. (2009). What moves housing markets: A variance decomposition of the rent-price ratio. *Journal of Urban Economics*, 66, 90-102.
- Case, K. E., & Shiller, R. J. (1989). The efficiency of the market for single family homes. *American Economic Review*, 79, 125-137.
- Case, K. E., & Shiller, R. J. (1990). Forecasting prices and excess returns in the housing market. *American Real Estate and Urban Economics Association Journal*, 18, 253-273.
- Chamberlain, G., & Rothschild, M. (1983). Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica*, 51, 1281-1304.
- Cochrane, J. H. (2007). Financial markets and the real economy. In: Mehra, R., *The equity premium*, North Holland Handbook of Finance Series, North Holland, Amsterdam.
- Cochrane, J. H. (2011). Discount rates. *Journal of Finance*, 66, 1047-1108
- Connor, G., & Korajczyk, R. A. (1986). Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of Financial Economics*, 17, 255-289.
- Doz, C., Giannone, D., & Reichlin, L. (2011a). A quasi maximum likelihood approach for large approximate dynamic factor models. *Review of Economics and Statistics*, forthcoming.
- Doz, C., Giannone, D., & Reichlin, L. (2011b). A two-step estimator for large approximate dynamic factor models based on Kalman filtering. *Journal of Econometrics*, forthcoming.

- Fama, E. F., & Schwert, G. W. (1977). Asset returns and inflation. *Journal of Financial Economics*, 5, 115-146.
- Fama, E. F. (1981). Stock returns, real activity, inflation and money. *American Economic Review*, 71, 545-565.
- Favilukis, J., Van Nieuwerburgh, S., Ludvigson, S. (2012). The macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium. Working paper.
- Forni, M., Hallin, M., Lippi, M., & Reichlin, L. (2000). The generalized dynamic factor model: Identification and estimation. *Review of Economics and Statistics*, 82, 540-554.
- Forni, M., Hallin, M., Lippi, M., & Reichlin, L. (2004). The generalised dynamic factor: Consistency and rates. *Journal of Econometrics*, 119, 231-255.
- Forni, M., Hallin, M., Lippi, M., & Reichlin, L. (2005). The generalized dynamic factor model: One sided estimation and forecasting. *Journal of the American Statistical Association*, 100, 830-840.
- Gallin, J. (2008). The long-run relationship between house prices and rents. *Real Estate Economics*, 4, 635-658.
- Geweke, J. (1977). The dynamic factor analysis of economic time series. In: Aigner & Goldberger, Latent Variables in Socioeconomic Models, North Holland, Amsterdam.
- Gospodinov, N. & Ng, S. (2011). Commodity prices, convenience yields and inflation. *Review of Economics and Statistics*, forthcoming.
- Goyal, A., & Welch, I. (2003). Predicting the equity premium with dividend ratios. *Management Science*, 49, 639-654.
- Goyal, A., & Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21, 1455-1508.
- Hallin, M., & Liska, R. (2007). Determining the number of factors in the generalized factor model. *Journal of the American Statistical Association*, 102, 603-617.
- Jungbacker, B., & Koopman, S. J. (2008). Likelihood-based analysis for dynamic factor models. Working paper.
- Kim, C. J., & Nelson, C. R. (1999). *State-Space Models with Regime Switching*. The MIT Press.
- Ludvigson, S. C., & Ng, S. (2007). The empirical risk-return relation: A factor analysis approach. *Journal of Financial Economics*, 83, 171-222.
- Ludvigson, S. C., & Ng, S. (2009). Macro factors in bond risk premia. *Review of Financial Studies*, 22, 5027-5067.
- Ludvigson, S. C., & Ng, S. (2011). A factor analysis of bond risk premia. In: Ullah &

- Giles, Handbook of Empirical Economics and Finance.
- Lustig, H., Roussanov, N., & Verdelhan, A. (2010). Countercyclical currency risk premia. Working paper.
- Otrok, C., & Whiteman, C. H. (1998). Bayesian leading indicators: Measuring and predicting economic conditions in Iowa. *International Economic Review*, 39, 997-1014.
- Patton, A., Politis, D. & White, H. (2009). Correction to “Automatic Block-Length Selection for the Dependent Bootstrap” by D. Politis and H. White, *Econometric Reviews*, 28, 372-375.
- Rapach, D. E., & Strauss, J. K. (2009). Differences in housing price forecastability across US states. *International Journal of Forecasting*, 25, 351-372.
- Sargent, T. J., & Sims, C. A. (1977). Business cycle modeling without pretending to have too much a priori economic theory. In Sims, New Methods in Business Cycle Research, Federal Reserve Bank on Minneapolis, 45-110.
- Stock, J. H., & Watson, M. W. (1998). Diffusion indexes. NBER Working paper 6702.
- Stock, J. H., & Watson, M. W. (2001). A comparison of linear and nonlinear univariate models for forecasting macroeconomic time series. In: Engle, R., White, H. (Eds.), Cointegration, Causality and Forecasting. A Festschrift in honour of Clive W.J. Granger, Oxford University Press, Oxford. pp. 1-44.
- Stock, J. H., & Watson, M. W. (2002a). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97, 1167-1179.
- Stock, J. H., & Watson, M. W. (2002b). Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics*, 20, 147-162.
- Stock, J. H., & Watson, M. W. (2005). Implications of dynamic factor models for VAR analysis. Working paper 11467, National Bureau of Economic Statistics.
- Stock, J. H., & Watson, M. W. (2006). Forecasting with many predictors. In: Elliott, Graham, Granger, Clive, & Timmermann, Handbook of Forecasting, 20, 147-162.
- Stock, J. H., & Watson, M. W. (2008). Forecasting in dynamic factor models subject to structural instability. In: Castle & Shephard, The Methodology and Practice of Econometrics: A Festschrift in Honour of Professor David F. Hendry, Oxford University Press.
- Thom, R. (1985). The relationship between housing starts and mortgage availability. *Review of Economics and Statistics*, 67, 693-696.
- Timmermann, A. (2006). Forecast combinations. In: Elliott, G., Granger, C.W.J., Timmermann, A. (Eds.), Handbook of Economic Forecasting, vol. 1. Elsevier, Amsterdam.

Table 1. Summary statistics for estimated factors

i	$\text{AR1}(\hat{f}_{i,t})$	R_i^2	$\sum R_i^2$
1	0.74	28.0%	28.0%
2	0.79	12.5%	40.5%
3	0.47	9.3%	49.8%
4	0.11	5.9%	55.7%
5	0.69	5.4%	61.1%
6	0.46	4.9%	66.0%
7	0.41	3.8%	69.8%
8	0.40	2.1%	72.0%
9	-0.34	2.0%	74.0%
10	0.37	1.8%	75.8%
11	0.24	1.8%	77.6%

$\text{AR1}(\hat{f}_{i,t})$ is the first-order autocorrelation coefficient of the i th estimated factor, while R_i^2 is the proportion of the total variance explained by the i th estimated factor as determined by the i^{th} eigenvalue divided by the sum of eigenvalues.

Table 2. In-sample results: slope coefficients, explanatory power, and bootstrap confidence intervals.

	$\hat{f}_{1,t}$	$\hat{f}_{2,t}$	$\hat{f}_{3,t}$	y_t	$\bar{R}^2$
			$h = 1$		
OLS estimate	0.28	-0.30	-0.65	0.19	53.7%
Bootstrap C.I.	[0.01; 0.74]	[-0.79; -0.02]	[-1.66; -0.44]	[0.03; 0.46]	[42.1%; 60.5%]
			$h = 2$		
OLS estimate	0.78	-0.54	-1.08	0.20	49.3%
Bootstrap C.I.	[0.14; 1.76]	[-1.12; -0.10]	[-2.89; -0.78]	[-0.14; 0.84]	[37.2%; 59.7%]
			$h = 4$		
OLS estimate	1.29	-0.99	-1.23	0.90	53.2%
Bootstrap C.I.	[0.10; 3.27]	[-1.83; -0.24]	[-2.68; -0.94]	[0.38; 2.24]	[43.2%; 65.1%]
			$h = 8$		
OLS estimate	1.84	-2.31	-1.58	1.54	41.0%
Bootstrap C.I.	[-0.55; 6.28]	[-4.75; -0.66]	[-4.73; -0.85]	[0.84; 3.86]	[34.6%; 63.3%]

This table reports results of predictive regressions for the h -quarter ahead real house price growth: $y_{t+h} = \alpha + \beta_1 \hat{f}_{1,t} + \beta_2 \hat{f}_{2,t} + \beta_3 \hat{f}_{3,t} + \gamma y_t + \varepsilon_{t+h}$. For each regression, the table reports OLS estimates of the slope coefficients, 90% bootstrap confidence intervals, the adjusted R^2 -statistic, and 90% bootstrap confidence intervals for the adjusted R^2 -statistic. Coefficients in bold indicate statistical significance.

Table 3. Out-of-sample results: MSFE performance relative to the AR(SIC) benchmark.

Row	Model	Benchmark	Horizon			
			$h = 1$	$h = 2$	$h = 4$	$h = 8$
1	Three-factor model (SIC)	AR(SIC)	0.92	0.97	0.93	0.93
2	Fcst combi (SIC weights)	AR(SIC)	0.98	1.00	0.96	0.92
3	Fcst combi (MSFE weights)	AR(SIC)	0.98	1.00	0.95	0.92
4	Fcst combi (disc. MSFE weights)	AR(SIC)	0.98	1.00	0.94	0.91

We use factor forecast combination using three weighting schemes. The first scheme is based on the Schwartz information criterion (SIC). The second scheme is based on the past MSFE performance. The third scheme is based on past discounted MSFE performance. We also report results for the augmented three-factor model. We compare with an autoregressive benchmark where the number of lags is recursively chosen based on the SIC criterion. The table reports the ratio of the mean-squared-forecast-error (MSFE) between the given model and the autoregressive benchmark, AR(SIC).

Table 4. Out-of-sample results: MSFE performance relative to other benchmarks.

Row	Model	Benchmark	Horizon			
			$h = 1$	$h = 2$	$h = 4$	$h = 8$
1	Three-factor model (SIC)	AR(1)	0.75	0.77	0.78	0.98
2	Three-factor model (SIC)	Mean	0.65	0.61	0.51	0.78
3	Three-factor model (SIC)	Price-rent ratio	0.60	0.55	0.45	0.72

The table reports the ratio of the mean-squared-forecast-error (MSFE) between the three-factor model and various benchmarks.

Figure 1: Real house price growth rate.

Figure 2: Model fit.

Figure 3: R^2 between factor 1 and each individual series.
Documents/Papers/WP/house prices/WP/graphics/RsquarePanelTr1.pdf

Figure 4: R^2 between factor 2 and each individual series.
Documents/Papers/WP/house prices/WP/graphics/RsquarePanelTr2.pdf

Figure 5: R^2 between factor 3 and each individual series.
Documents/Papers/WP/house prices/WP/graphics/RsquarePanelTr3.pdf

- 2012-10: Peter Exterkate: Model Selection in Kernel Ridge Regression
- 2012-11: Torben G. Andersen, Nicola Fusari and Viktor Todorov: Parametric Inference and Dynamic State Recovery from Option Panels
- 2012-12: Mark Podolskij and Katrin Wasmuth: Goodness-of-fit testing for fractional diffusions
- 2012-13: Almut E. D. Veraart and Luitgard A. M. Veraart: Modelling electricity day-ahead prices by multivariate Lévy
- 2012-14: Niels Haldrup, Robinson Kruse, Timo Teräsvirta and Rasmus T. Varneskov: Unit roots, nonlinearities and structural breaks
- 2012-15: Matt P. Dziubinski and Stefano Grassi: Heterogeneous Computing in Economics: A Simplified Approach
- 2012-16: Anders Bredahl Kock and Laurent A.F. Callot: Oracle Inequalities for High Dimensional Vector Autoregressions
- 2012-17: Eric Hillebrand, Huiyu Huang, Tae-Hwy Lee and Canlin Li: Using the Yield Curve in Forecasting Output Growth and Inflation
- 2012-18: Eric Hillebrand and Tae-Hwy Lee: Stein-Rule Estimation and Generalized Shrinkage Methods for Forecasting Using Many Predictors
- 2012-19: Bent Jesper Christensen, Morten Ørregaard Nielsen and Jie Zhu: The impact of financial crises on the risk-return tradeoff and the leverage effect
- 2012-20: Hendrik Kaufmann, Robinson Kruse and Philipp Sibbertsen: On tests for linearity against STAR models with deterministic trends
- 2012-21: Andrey Launov, Olaf Posch and Klaus Wälde: On the estimation of the volatility-growth link
- 2012-22: Peter O. Christensen and Zhenjiang Qin: Information and Heterogeneous Beliefs: Cost of Capital, Trading Volume, and Investor Welfare
- 2012-23: Zhenjiang Qin: Heterogeneous Beliefs, Public Information, and Option Markets
- 2012-24: Zhenjiang Qin: Continuous Trading Dynamically Effectively Complete Market with Heterogeneous Beliefs
- 2012-25: Heejoon Han and Dennis Kristensen: Asymptotic Theory for the QMLE in GARCH-X Models with Stationary and Non-Stationary Covariates
- 2012-26: Lei Pan, Olaf Posch and Michel van der Wel: Measuring Convergence using Dynamic Equilibrium Models: Evidence from Chinese Provinces
- 2012-27: Lasse Bork and Stig V. Møller: Housing price forecastability: A factor analysis