Forecasting US Commercial Property Price Indexes using Dynamic Factor Models.[☆,☆☆]

Alex van de Minne[a,*], Marc Francke[b,c], David Geltner[a]

[a] *MIT Center for Real Estate, 105 Massachusetts Ave, Cambridge, MA 02139.*
[b] *Faculty of Economics and Business, University of Amsterdam, Plantage Muidergracht 12, 1001 NL Amsterdam.*
[c] *Ortec Finance, Naritaweg 51, 1043 BP Amsterdam.*

Abstract

The general purpose of a dynamic factor model (DFM) is to summarize a large number of time series into a few common factors. Here we explore a number of DFM specifications applied to 80 granular, non-overlapping indexes of commercial property prices in the US, quarterly from 2001 to 2017. We examine the nature and the structure of the factors and the index forecasts that can be produced using the DFMs. We consider specifications of 1, 2, 3 and 4 common factor trends. As a major motivation for the use of DFMs is their ability to improve out-of-sample forecasting of systems of numerous related series, we apply the DFM estimated factor returns in an Autoregressive Distributed Lag (ARDL) model to forecast the individual real estate price series. We compare the forecasted residuals to a conventional Autoregressive (AR) forecast model as a "benchmark" for two markets: Boston apartments and Dallas commercial. The results show that the ARDL model predicts the crisis and subsequent recovery really well, whereas the "benchmark" model typically follows the previous price trend. We find that the DFM forecasts are most precise with only one or two factors. The two prominent factors may reflect general economic conditions and the rental housing market, respectively.

☆MIT/CRE Price Dynamics Research Platform Working Paper

☆☆This paper was funded by the Real Estate Research Institute. We would also like to thank Real Capital Analytics for providing to us the data.

*Corresponding author

Email addresses: m.k.francke@uva.nl (Marc Francke), dgeltner@mit.edu (David Geltner), avdminne@mit.edu (Alex van de Minne)

Keywords: structural time series, EM algorithm.
JEL-codes: R32, C01.

1. Introduction

In recent years, an explosion in the amount of economic and financial data has prompted econometricians to develop or adapt new methods to efficiently summarize the information in large databases with many time series. Among these methods, **dynamic factor models** (DFMs), have seen rapid growth and become very popular among macro-economists. These models summarize a large number of related time series into a small number of factors common to the original series. In a DFM we describe temporal variation in a set of N observed variables that are related or reflect a common system (like GDP, interest rates, employment and such macro-economic variables, or in our case asset price indexes of numerous separate commercial real estate space markets) using linear combinations of $M << N$ "hidden" common factors which are identified and estimated from the data. As a result we can summarize big quantities of N time series, into a small number of M common factors. Because we may then be better able to forecast a small number of common factors than the large number of original series separately, DFMs may be able to improve forecasting of the series.

Dynamic factor analysis is already widely used in practice. A famous example is the Chicago FRB National Activity Indicator. It is based on 85 monthly series describing the US economy, covering production, income, employment, personal consumption, housing, sales, inventories and orders. Applications of DFMs abound in the empirical economic literature as well. A few examples include asset pricing models (Ross, 1976), consumer theory (Gorman, 1981; Lewbel, 1991), the assessment of performance and risk measurement in finance (Campbell et al., 1997), and policy related questions (Bernanke et al., 2005; Stock and Watson, 2005; Favero et al., 2005; Del Negro and Otrok, 2007). Factor models have also been applied to commercial real estate in an investment context (see Naranjo and Ling, 1997), though the new breed of *dynamic* factor models has not yet to our knowledge been applied to commercial property price forecasting.

The advent of high quality granular indexes of commercial property asset transaction prices makes the time ripe to explore the application of the new DFM methodologies to commercial property, with a consideration of issues of interest to the real estate investment industry. Granular indexes track individual local property asset markets defined by the space markets. For example, a granular index typically tracks prices in a single metro area or part of a metro such as CBD or suburbs and for a single property usage type sector such as apartments or office. (In our data we combine three core non-residential commercial sectors together: office, retail, and industrial property.) Such "micro" or granular level tracking of asset prices is important because the space markets are segmented, not integrated, which can result in different pricing and different risk and return behavior at the granular level.

But this means that there are dozens of commercial property price series in a complete database. While distinct and therefore important to track separately, these dozens of series are also interrelated. They are all part of the overall commercial property capital market and none of the underlying space markets are isolated from the US macro-economy. This is exactly the type of database in which the DFM methodology has been found to be useful in other fields such as macro-economics. There, it is also known that DFMs can improve existing forecasting models. The main problem with the "standard" forecast methodologies for multiple, interrelated time series, such as vector autoregression (VAR) models, is that they become intractable with such a large number of time series. (Bernanke et al., 2005).

We believe this is the first study to explore DFM based forecasting to commercial property price index returns. In the first step, we reduce a large set of (N) price index returns series into a small set of (M) factors. This is of interest in its own right, as it provides insights about how commercial real estate markets co-move. It tells you which property price indexes tend to move with which factors, which can be of interest to investors for portfolio management and diversification. In the second step we forecast the M factor returns by using univariate autoregressive models on all the the M factors. In the third step we individually forecast the N price index returns using Autoregressive Distributed Lag (ARDL) models, which include lagged price index returns and the dynamic factors.

We focus on two distinct markets: Boston apartments and Dallas commercial. Not only are the property types different, the first market is a typical "supply constraint" market, whereas the second is not. We are specifically interested in how well the benchmark and ARDL with factors predicted the crisis and subsequent recovery. Our results show that the benchmark model almost completely "misses" both events. Indeed, the benchmark model predicts a price increase (decrease) during the crisis (recovery). In contrast, the ARDL model with factors, predicts both events with impressively high accuracy.

Another benefit of using a DFM is that we do not require any other (third party) data. Instead of explaining real estate returns using additional macro economic variables, we use the factor trends returns from the DFM. These factor return are estimated directly from the panel of indexes. Assuming that the indexes co-move and using that information is not controversial within real estate (Francke et al., 2017).

Having reliable asset price forecasts is obviously important from a practical perspective. Price forecasts are important for tactical level portfolio management (where to buy and where to sell in the intermediate term). If (or when) property price derivatives become important, price forecasting will be essential for pricing real estate derivates (Geltner and Fisher, 2007).

Real Capital Analytics provided to us 80 non-overlapping commercial real estate property price indexes. The indexes are quarterly between 2001Q1 and 2017Q2. More specific, we observe 40 regions and have a commercial (combined office, retail and industrial) and apartment index per region. The regions include 26 metro markets (such as Jacksonville, Boston, Chicago...) with many of those broken into central versus non-central areas.

Our findings are of interest not only for the overall forecast performance results, but also for the structural insights they provide about commercial property asset price dynamics. Generally we find that more factors increase the DFM model fit. However, the forecasts of property price index returns are better when only including 1 or 2 factors. In fact, this is a known phenomena in DFM literature (Eickmeier and Ziegler, 2008).

This paper proceeds as follows. Section 2 provides the DFM used in this research, and the forecast models. Section 3 provides a description of the data. The estimated factor trends and forecast results are shown and discussed in Section 4. Finally Section 5 concludes.

2. Model

2.1. The Dynamic Factor Model

Dynamic factor models fall in the realm of structural time series (Harvey, 1989). A structural time series model is a model in which the trend, error terms, plus other relevant components (like seasonality), are modeled explicitly. For example, Francke et al. (2017) estimate a repeat sales model for commercial real estate, in which the variance parameters are allowed to be time-varying (following a random-walk). Other examples of the use of structural time series in real estate applications include Goetzmann (1992), Schwann (1998), Francke and De Vos (2000), and Francke (2010).

In this paper, we only consider a structural time series of the form:

$$\text{data} = \text{trends} + \text{noise}. \tag{1}$$

In math, suppose we have univariate series of log commercial real estate price index returns, denoted by Δp_t, observed in period t, with $t = 1, \ldots, T$. We model log index returns, instead of the log indexes themselves, to account for issues regarding non-stationarity; log price indexes are typically integrated of order 1.

The most simple univariate structural time series model, without any explanatory variables, is given by

$$\Delta p_t = \mu_t + \epsilon_t, \tag{2}$$

$$\mu_t = \mu_{t-1} + \eta_{t-1}. \tag{3}$$

This model is also known as a random walk plus noise model. The term μ_t represents the unknown trend of the index returns at period t. The components ϵ_t and η_t are the error components. It is assumed that $\epsilon_t \sim N(0, r)$ and $\eta_t \sim N(0, q)$. The signal-to-noise parameter r/q determines the smoothness of the trend component, where smaller values give a more smooth trend. A large signal-to-noise ratio ($r/q \to \infty$) coincides with time fixed effects.

Next, suppose we have N price indexes. We could individually estimate a trend component per price index return series. However, this results in N different trends that all have to be interpreted independently, and interactions between the estimated trends are ignored. This is especially problematic for large N. The dynamic factor model aims to overcome these disadvantages by reducing the N univariate trends to M common factors, where $1 \leq M << N$. To illustrate, suppose we assume $M = 2$ factors (μ_m, with $m = 1, \ldots, M$) and $N = 4$ returns (Δp_i, with $i = 1, \ldots, N$). The full model is now provided by (while ignoring the distributions of the error terms and initial conditions for now)

$$\begin{bmatrix} \Delta p_1 \\ \Delta p_2 \\ \Delta p_3 \\ \Delta p_4 \end{bmatrix}_t = \begin{bmatrix} \gamma_{11} & \gamma_{12} \\ \gamma_{21} & \gamma_{22} \\ \gamma_{31} & \gamma_{32} \\ \gamma_{41} & \gamma_{42} \end{bmatrix} \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}_t + \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \end{bmatrix} + \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \\ \epsilon_3 \\ \epsilon_4 \end{bmatrix}_t , \tag{4}$$

$$\tag{5}$$

$$\begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}_t = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}_{t-1} + \begin{bmatrix} \eta_1 \\ \eta_2 \end{bmatrix}_{t-1} . \tag{6}$$

The observation errors are given by

$$\begin{bmatrix} \epsilon_1 \\ \epsilon_2 \\ \epsilon_3 \\ \epsilon_4 \end{bmatrix}_t \sim N \left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} , \begin{bmatrix} r_{11} & r_{12} & r_{13} & r_{14} \\ r_{21} & r_{22} & r_{23} & r_{24} \\ r_{31} & r_{32} & r_{33} & r_{34} \\ r_{41} & r_{42} & r_{43} & r_{44} \end{bmatrix} \right) , \tag{7}$$

and the innovation errors by

$$\begin{bmatrix} \eta_1 \\ \eta_2 \end{bmatrix}_t \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix} , \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix} \right) , \tag{8}$$

where α is a constant per observed time series (Δp_{it}) and the $(M \times N)$ matrix Γ contain the so-called factor loadings γ_{ij}. The factor loadings can be interpreted as the correlation of the price index returns with each factor return μ_m. For example, the predicted values of our first time series are $\Delta \hat{p}_{1t} = \hat{\mu}_{1t} * \hat{\gamma}_{11} + \hat{\mu}_{2t} * \hat{\gamma}_{12} + \hat{\alpha}_1$. Thus, the general idea is that the observations (Δp_{it}) are modeled as a linear combination of factor returns (μ_m) and factor loadings (γ_{ij}) plus some offsets (α_i).

A more general (vectorized) formulation of the dynamic factor model is given by

$$\mathbf{\Delta p}_t = \mathbf{\Gamma} \boldsymbol{\mu_t} + \boldsymbol{\alpha} + \boldsymbol{\epsilon_t}, \tag{9}$$

$$\boldsymbol{\mu_t} = \boldsymbol{\mu_{t-1}} + \boldsymbol{\eta_{t-1}}, \tag{10}$$

where $\mathbf{\Delta p}_t$ is a N-vector of log price index returns, and $\boldsymbol{\mu_t}$ a M-vector of factors. We assume that $\boldsymbol{\epsilon}_t \sim N(0, \mathbf{R})$, $\boldsymbol{\eta}_t \sim N(0, \mathbf{Q})$ with initial condition for the factors $\boldsymbol{\mu}_0 \sim N(\mathbf{m}_0, \mathbf{V}_0)$. The unknown parameters in the model are elements of $\mathbf{\Gamma}, \mathbf{R}, \mathbf{Q}, \boldsymbol{\mu}, \boldsymbol{\mu}_0$ and $\mathbf{V}_0$ and are usually referred to as hyperparameters.

The hyperparameters can be estimated using the Kalman filter by maximum likelihood based methods or full Bayesian Markov Chain Monte Carlo algorithms. In our application we use the Kalman filter, where we estimate the hyperparameters using the Expectation-Maximization (EM) algorithm. The advantage of this methodology is that it does not require long time series (which we do not have, see Section 3). It also converges relatively fast for large numbers of time series, compared to some of the other methodologies. Also note that the model is fully Gaussian, which is easy to process in a Kalman filter framework.

2.2. Identification of the Dynamic Factor Model

The issue is that if one does not constrain $\mathbf{\Gamma}$, $\boldsymbol{\alpha}$ and $\mathbf{Q}$ in Eqs. (9) – (10), the parameters in the DFM are not identified (see Harvey, 1989, Chapter 4.4). There is substantial literature on how to make the system identifiable, see for example Harvey (1989); Geweke and Zhou (1996); Aguilar and West (2000); Forni et al. (2000); Stock and Watson (2002); Bai and Wang (2015); Aßmann et al. (2016), among many others. We use the constraints introduced by Zuur et al. (2003), who largely follow Harvey (1989) (Chapter 5.8.1), but with one crucial difference to make the estimation more robust in the EM framework (see further below).

First, set γ_{ij} in Γ to zero if $j > i$. Thus, the top right corner of matrix Γ fully consists of zeros. In our previous example, that implies we set γ_{12} to zero. More general,

$$\Gamma = \begin{bmatrix} \gamma_{11} & 0 & \cdots & \cdots & 0 \\ \gamma_{21} & \gamma_{22} & \ddots & & \vdots \\ \vdots & \vdots & \ddots & 0 & \vdots \\ \gamma_{M-1,1} & \gamma_{M-1,2} & \cdots & \gamma_{M-1,M-1} & 0 \\ \gamma_{M,1} & \gamma_{M,2} & \cdots & \gamma_{M,M-1} & \gamma_{M,M} \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ \gamma_{N,1} & \gamma_{N,2} & \cdots & \gamma_{N,M-1} & \gamma_{N,M} \end{bmatrix}. \tag{11}$$

Secondly, we set $\mathbf{Q}$ equal to a diagonal matrix, implying uncorrelated innovation errors. In our previous example, with $M = 2$, the variance-covariance matrix of the innovation errors is given by

$$\mathbf{Q} = \begin{bmatrix} q_{11} & 0 \\ 0 & q_{22} \end{bmatrix}. \tag{12}$$

We also have to constraint $\boldsymbol{\alpha}$ in a meaningful way. Harvey (1989) suggests to set the first M values to zero. In our previous example, this would mean, $\boldsymbol{\alpha} = [0, 0, \alpha_3, \alpha_4]'$. However, Zuur et al. (2003) found that the EM estimates are not robust to this constraint, and takes long to converge. We therefore simply demean our data, and omit $\boldsymbol{\alpha}$ altogether. We further standardize the time series by dividing by its own standard deviation (Bernanke et al., 2005). This step is not necessary for identification, but does increase the efficiency (and speed of convergence) of the EM algorithm. Thus, the index return p_{it} for series i at time t in the measurement equation (9) is replaced by

$$\Delta \mathrm{p}^*_{it} = \big(\Delta \mathrm{p}_{it} - \Delta \bar{\mathrm{p}}_{i.}\big)/\sigma_i, \tag{13}$$

where $\Delta \mathrm{p}^*_{it}$ is the transformed index return for series i at time t, and $\Delta \bar{\mathrm{p}}_{i.}$ and σ_i are the average and standard deviation of index return series i.

Finally, following Zuur et al. (2003) the initial condition of the state vector is specified by a diffuse prior, $\boldsymbol{\mu}_0 \sim N(0, \kappa\mathbf{I})$, where κ is large, and $\mathbf{I}$ denotes the identity matrix. Implementing these restrictions and initial conditions within the EM algorithm is not trivial (Wu et al., 1996). Please consult Zuur et al. (2003) for the exact algorithm used in this paper, which we will omit here for space.[1]

Note that by implementing these restrictions a unique solution for the factor loadings exists, with one huge caveat. The caveat is that the factors depend on the ordering of the series; the first factor is determined by the price index returns of the first series, the second factor by the index returns of the first two series, etc. Other solutions therefore also exists (Harvey, 1989, p. 450) by simply ordering the series differently.

However, once the parameters have been estimated, a factor rotation can be applied to the estimated factor loadings and factors. Many techniques to do the rotation exist in literature (see Harman, 1976; Basilevsky, 1994; Browne, 2001; Bernaards and Jennrich, 2005, for overviews), like oblimin,

[1] Zuur et al. (2003) also show how to handle missing values, and how to enter covariates in the DFM. However, both are not needed in our research.

othomax, quterartimax, etc. We use the varimax rotation. Varimax is so called because it maximizes the sum of the variances of the squared factor loadings. More specifically, we introduce rotation $(M \times M)$ matrix $\mathbf{H}$, which is interacted with the factors $(\boldsymbol{\mu_t})$ and the process errors $(\boldsymbol{\eta_t})$;

$$\boldsymbol{\Delta}\mathbf{p}_t^* = \boldsymbol{\Gamma}\mathbf{H}^{-1}\boldsymbol{\mu_t} + \boldsymbol{\epsilon_t} \tag{14}$$

$$\mathbf{H}\boldsymbol{\mu_t} = \mathbf{H}\boldsymbol{\mu_{t-1}} + \mathbf{H}\boldsymbol{\eta_{t-1}} \tag{15}$$

The varimax rotation subsequently seeks a a rotation matrix $\mathbf{H}$, that creates the largest difference between loadings. As varimax is (arguably) the most widely used technique for rotation. Therefore, most statistical software packages will have built-in functions for this.

Note that we have no constraints on the variance-covariance $\mathbf{R}$. In our application we used a simple diagonal and equal variance-covariance structure. We did also test diagonal and unequal and equal variance / equal covariance structures. The results remain similar and are omitted here for the sake of brevity. Obviously, these results are available upon request. (Technically we could also specify $\mathbf{R}$ to be unconstrained - i.e. unequal variance and covariance - however we found that computing time increased considerably by doing so.)

2.3. Forecasting

Forecasters can predict time series by fitting small-scale time series models, such as (vector) autoregressive models. These models have been shown to perform fairly well in the past. However, in this day and age there is a lot of information available to us which potentially could be useful for predicting.

However, it is not feasible to include every potential relevant variable simultaneously in a forecasting equation. This becomes especially problematic with large N and/or small T. This is where factor indexes come into place. Again, the general idea of factor models is that the bulk of variation of many variables can be explained by a small number of factors or shocks. Factor models exploit the variables' co-movement and efficiently reduce the dimension of the dataset to just a few underlying factors. More specifically, a relatively small number of factors are entered into a forecasting equation to predict our returns, and only a few parameters need to be estimated.

Unsurprisingly, forecasting time series (like GDP) using factor models, have gained in popularity in the last decade or two. For example Eickmeier and Ziegler (2008) summarize 52 papers where factor models are used for forecasting inflation and output for a variety of countries. Most of these (and other) studies use the factors in an Autoregressive Distributed Lags (ARDL) framework (Eickmeier and Ziegler, 2008; Barhoumi et al., 2013).

The ARDL with (estimated) factors $\hat{\mu}_{mt}$ and price returns Δp_{it}, is given by

$$\Delta p_{it} = \beta_0 + \sum_{m=1}^{M} \sum_{k=1}^{r_m} \beta_{mk} \hat{\mu}_{m,t-k} + \sum_{k=1}^{p} \phi_k \Delta p_{i,t-k} + \zeta_{it}, \tag{16}$$

where we assume that the error term is normally distributed, $\zeta_{it} \sim N(0, \sigma_\zeta^2)$.

Conditional on the number of lags (p, and r_m, for $i = 1, \ldots, M$), the model can be estimated by Ordinary Least Squares (OLS). Note that we estimate Eq. (16) for every price index return series in the data separately, so N-times. The number of lags is usually determined by the Akaike Information Criteria (AIC).

In order to use Eq. (16) as a forecasting model, we need forecasts of the factors $\boldsymbol{\mu_t}$ as well. We forecast the factor returns using an autoregressive model for each factor individually,

$$\mu_{mt} = \phi_0 + \sum_{k=1}^{p} \phi_k \mu_{m,t-k} + \zeta_{mt}. \tag{17}$$

We use Eqs. (16) and (17) to forecast our price index returns Δp_{it} "within sample". For example, say we have data until 2017. First, we leave out all 2017 data, and subsequently forecast - using data up to 2016 - the returns in 2017.[2] We re-do this for multiple periods, like 2016, 2015, and so on. Next, we compare the forecasts to the actual returns and compute the "forecast residuals". We compare the results to a *benchmark model*. This benchmark model is an autoregressive model, so equivalent to Eq. (17). After computing the forecast residuals in similar fashion for this benchmark model, we can evaluate whether or not forecasting improves by including the factors.

As noted earlier, this type of three step forecasting procedure (step 1; estimate DFM, step 2; forecast factors, and step 3; use factors in ARDL

[2] In reality, the indexes and factor indexes would revise over time. Thus, we would have to re-estimate the indexes and factor indexes with data up to only 2016 as well. However, the indexes are given to us. Thus we are not able to do this. Also, estimating all factor indexes takes approximately a day on our desktop computer. Re-estimating all the factor indexes would therefore become intractable.

framework) has gained in popularity, and its easy to see why. However, there are also some critiques to this approach. For example, Boivin and Ng (2006) show that the forecast performance of factor models may worsen if one (or more) of the factors that are included are irrelevant for the variable of interest - also known as the oversampling problem. This is partly resolved by dropping irrelevant factors if the AIC deteriorates for example. Another issue - and related to the first - is that we have to forecast the factor indexes themselves as well. This creates extra uncertainty in the forecasts. It should also be noted that there is evidence that the DFM forecast models improve with more N and more T (Stock and Watson, 2002; Bai and Ng, 2002).

We use the ARDL forecast model, because; (1) it is the most widely used forecast model for this purpose, (2) its simplicity and (3) the theory behind it (Stock and Watson, 1998).[3] However, note that scholars have recently built DFM models which can forecast themselves, by including autoregressive terms in the state equation of the DFM. See for example the "Factor Augmented Vector Autoregressive Model" by Bernanke et al. (2005). These types of models are out of the scope of this research.

3. Data

We obtain our price indexes for a selection of cities from Real Capital Analytics (RCA). RCA is recognized as one of the most respected commercial real estate data firms in the world. RCA focuses on commercial and residential real estate which were bought and sold by institutional investors. With their transaction data, RCA estimates repeat sales indexes world wide, using the repeat sales methodology developed in Francke et al. (2017). For this study we use 80 non-overlapping Commercial Property Price Indexes (CPPIs), which covers commercial (= combination of office, retail and industrial properties) and residential real estate in the United States on very granular (metro) level. RCA also has "all types" (apartments + commercial), state, regional and country level indexes (plus some other specialty indexes), but these are left outside of the scope of this paper, as they are overlapping. For most metros, we observe 2 CPPIs (one for residential and one commercial).

[3]We also used VAR models to forecast. The general results still hold true. Using a VAR might be preferable if one is concerned for biases caused by endogeneity. Even though these results are not given in this paper, they are availably upon request.

However, in some metros we are able to identify sub-markets, like Manhattan, the Burroughs and the suburbs, in the New York Metro metro area. The indexes are quarterly and run from 2001 to mid 2017. Our total panel data is therefore $N = 80$ and $T = 66$. For a full overview of all the indexes we used, please consult the website of RCA.

A graphical representation of all indexes is given in Figure A.1a, where 2006Q4 = 100. Figure A.1b gives the (log) return of all our indexes, and Figure A.1c gives the "standardized" returns, using Eq. (13). Note that these standardized returns are our left-hand side variable in the DFM. Finally, Figure A.1d gives a histogram of the average quarterly price growth per market.

[Place Figure A.1 about here]

Overall, the indexes do seem to co-move. Indeed, the indexes go up prior to the crisis, then crash and subsequently recover starting in approximately 2009. Even though it is not presented here, there is a lot of first order autocorrelation in the returns, which indicates predictability. This is not an uncommon find in real estate literature (see Case and Shiller, 1989; Quan and Quigley, 1991; Barkham and Geltner, 1995; Geltner et al., 2003, among others). This autoregressive representation (or 'inertia') is inherent to the price formation process in real estate and does not imply arbitrage opportunities. More specifically: (1) Participants in real estate markets have incomplete information about the attributes of the purchase, (2) some period of costly search must be incurred by both buyers and sellers, due to the heterogeneity of real estate and (3) trades are decentralized, i.e. market prices are the outcome of pairwise negotiations (Case and Shiller, 1989; Quan and Quigley, 1991). Most of the autocorrelation is gone after a year.

Still, there is a big variety in both exact timing and in overall growth rate. For example, we find that most markets in the major metro areas have had the highest price growth; New York, Los Angeles and San Fransisco and had an overall price growth of almost 2% or more per quarter. Also cities like Miami, Portland and Seattle performed relatively well. On the other hand, cities like Las Vegas, Chicago, Dallas, Atlanta and Philadelphia (among others) did poorly. Usually, when apartments had a relative high price growth in a region, the commercial properties did the same.

You can also see big differences in volatility. High growth markets also tend to be the most volatile, see for example New York and Los Angeles.

More interesting is that some low overall growth markets like Las Vegas and Phoenix also had relative high volatility. Some other markets with low volatility include Raleigh/Durham and Minneapolis.

4. Results

The factor index returns and corresponding factor loadings are discussed in Section 4.1. Subsequently we discuss the forecast results in Section 4.2.

4.1. Factors and Factor Loadings

In our research we limit ourselves to a maximum of four factors. One consideration is computing time; Estimating 5 factors takes almost a day. Also, it is quite well established that more factors isn't per se better when it comes to forecasting. Our results in Section 4.2 confirm this as well. The reason is that we need to forecast those factors as well, which adds uncertainty (Section 2.3). We also become subject to the aforementioned 'over-sampling'. Admittedly, the model fit (AIC-wise) does seem to improve with higher M. Still, it should be stressed that there is no consensus in literature to what the best model fit statistic is in the first place, and other statistics might give different results. Other than the AIC, we also have the criteria developed in; Forni et al. (2000); Bai and Ng (2002); Breitung and Pigorsch (2013), for example. Finally, estimating many factors defeats the purposes of DFM analysis, which is; *Summarizing large quantities of data into a few common factors*. Table B.1 gives some diagnostics of all M, with $M \leq 4$. The results are ordered by model fit (using the penalized AIC), from best to worst.

[Place Table B.1 about here]

The specification with 4 hidden trends, gives the best overall model fit, even after penalizing the AIC for the increasing amount of parameters. Interestingly, is that for forecasting purposes we find that less factor trends typically works better, see Section 4.2. This - again - buttresses earlier findings in other literature, that more factor trends are not per se better for forecasting (Eickmeier and Ziegler, 2008).

Next we will turn our attention to the estimated factor returns $\hat{\mu}_{tm}$. Using Eq. (13) we modeled the standardized returns, and not the returns themselves. Since it is difficult to interpret the standardized factors, we first "unstandardize" the factor returns. Next we will explain how this "unstandardizing" works.

First, the fitted value of every index can be computed as follows. Let's take Boston (commercial) as an example, using a DFM with 1 factor return (M = 1). The fitted price return for said market is than given by

$$\Delta\hat{p}_{\mathbf{Boston},t} = \big(\hat{\mu}_t \times \hat{\gamma}_{\mathbf{Boston}} + \Delta\bar{p}_{\mathbf{Boston}}\big)\sigma_{\mathbf{Boston}}$$

. Note that we get unique solution because for every market i, the standard deviation of the returns (σ_i), the factor loading ($\hat{\gamma}_i$) and the average return ($\Delta\bar{p}_i$) is different. With two factor returns, we get two unique factor loadings, etc. We simply take the average of each "unstandardized" factor returns to get a single measure for each factor. This is given by;

$$r_{mt} = \frac{1}{N}\sum_{i=1}^{N}\big(\hat{\mu}_{mt}\hat{\gamma}_{mi}\big)\sigma_i, \qquad (18)$$

where r_{mt} is the unstandardized factor return m at time t, σ^2 are the variances of the index returns of the different N markets and $\hat{\mu}_{mt}$ is the estimated factor m at time t from the DFM. Next we sum the returns (and exponentiate them) to get the unstandardized factor indexes.

[Place Figure A.2 about here]

Note, that we do not add back the average price growth trends ($\Delta\bar{p}_i$) to the factor returns, as we cannot attribute individual market growth to a specific factor. Thus all indexes start and stop at 100. (I.e. the indexes remain demeaned.)

Next we focus on the *factor loadings* ($\hat{\gamma}_{mi}$). "Unstandardizing" the factor loadings can be done with basic algebra as well, which is equivalent to simply running a multivariate OLS model per market with the "unstandardized" data (market index left, factor indexes right) plus a constant. Figure A.3 give the kernal distributions of these factor loadings per M. These results are also only for the DFM with equal variance and covariance.

[Place Figure A.3 about here]

With M = 1, note that the factor loadings are positive for all index return series. In general most factor loadings are positive, even with higher M. We noted earlier in Section 3, that finding co-movement in real estate literature is not controversial (Geltner and Mei, 1995; Francke and van de

Minne, 2017; Francke et al., 2017). The spatial distribution of these factor loadings is given in Appendix C. With M = 1, a higher "unstandarized" factor loading simply means the index had a higher estimated (return) variance. The factor loadings with M = 1 are relatively high in Florida. Typically, the factor loadings are higher for commercial real estate. Factor loadings for apartments in Boston, Washington DC, California (except for Los Angeles) and Denver are especially on the low side.

The first factor of the DFM with M = 2 had a high pre-crisis run up, see Figure A.2. In contrast, the second factor remained relatively stable, starting at 100 in 2001, and ending at 102 in 2007. The crash is similar %-wise between the two factors. However, the first factor definitely moved first. The second factor did experience a more steep recovery, and shows less volatility overall.

The size of the factor loadings are relatively similar between the two factors, see Figure A.3. Figures A.8a – A.8b gives the spatial distribution of these factor loadings. It is evidents that the first and second factor are typically a mirror image of each other. Indeed, if the loading is high on the first factor, it is typically low (sometimes even negative) on the second factor and vice verse. Never do we observe a very high loading on both. If we look at markets that are mostly impacted by the first factor (red), they are almost exclusively apartments. Examples are apartments in California (not including San Fransisco), Chicago, Boston, New York, Philadelphia, Washington DC and Miami. Most of these are also big cities. For commercial, only the Fort Myers metro area (including Naples and Sarasota) loads almost exclusively on the first factor. The second factor does load onto *both* sectors evenly. For apartments we observe; Seattle, Denver, Dallas, Raleigh/Durham, Nashville, Washington DC and Atlanta. For commercial we observe; Seattle, Denver, Dallas, Boston and Dallas. Thus, if the second factor loads onto commercial in a specific market, so does apartments in some cases.

The factors with M = 3 are given in Figure A.2c. The third (first) factor is reminiscent of the first (second) factor of the DFM with M = 2. There are some differences. The first factor goes down more pre-crisis compared to the second factor in the DFM with M = 2. The second factor is new. Interestingly, the shape is a very typical boom-bust cycle, one that we also observed in the DFM with M = 1. The third factor does go down a full year before the other two. The timing of the crash between the second and third factor are similar though.

As with the previous examples, the factor loadings are similar in size (Figure A.3). Also similar to previous examples, is that the markets are typically not highly loaded on all factors, but rather one or sometimes two. First, apartments almost never load highly on the second factor (except perhaps Charlotte). In contrast, it is the most dominant factor component for commercial real estate in most markets. Apartments load more aggressively on the low volatile first factor trend. Almost none of the markets load mainly on the third factor, apart from apartments in North-Carolina and Portland (OR).

The DFM with M = 4 gives some familiar results, see Figure A.2. We see again the typical boom-bust factor (factor 2) and the factor trend that goes down at the beginning of the sample, but also recovers more swiftly (factor 3). What is "new" is that the third factor from the DFM with M = 3, is "broken up" into two separate factors (factor 1 and 4). They are very similar. (Which can be interpreted as a hint to stop increasing M.) The main exception being that the fourth factor shows a bit of the boom-bust cycle, whereas the first does not.

As previously, the second factor (typical boom-bust cycle) is mainly loaded on commercial properties. The third factor is again the most dominant factor in apartments. Although quite surprisingly, some commercial property markets also load on this factor highly; Boston, Denver, Seattle and Dallas. The "new" factors are more equally loaded on the markets. Although, the fourth factor seems to impact apartments slightly more.

Because the factors are estimated simultaneously, it cannot be said which factor trend is "dominant". Still, a feel for this can be developed by comparing the different factor trends in Figure A.2, and see which pattern prevails. Arguably, there are 2 factors that are "dominant". The typical boom-bust cycle, which we observe in the DFM with $M = 1$, and the second factor in the DFM with $M = 3$ and M = 4, is one. We found that commercial properties typically load on this factor. The second factor we found multiple times was factor 2 in the DFM with $M = 2$, the first factor in the DFM with $M = 3$, and the third factor in the DFM with $M = 4$. This factor slows down at the beginning of the sample, then drops during the GFC, followed by a steep recovery. Apartments typically load on this one.

We noted earlier that we cannot attribute the individual price growths to the factor trends. One exception being the model with 1 factor trend (M = 1). With M = 1, you can add the average return of all the market indexes to Eq. (18), to get *the* common trend (= average growth + average cyclicality from DFM). We have done so in Figure A.4. This can thus be interpreted as *the* US common trend running through all properties. It is therefore useful - as a robustness check - to compare the index with the readily available "all types US index" from Real Capital Analytics As the methodology to estimate both indexes, data (panel versus micro transaction data) and the exact interpretation are completely different, some differences are expected.[4] Still, both indexes have many similarities. Indeed, you can observe the boom-bust cycle in both cases (although slightly attenuated in the DFM case), and both indexes end at around 200 index points. Both indexes also show a brief slowdown, a year before the crisis started. The correlations between the two indexes are therefore unsurprisingly high, with 0.99 between the index levels and 0.97 between the index returns. We also add to Figure A.4 simply the mean of all market indexes in our data per period.

[Place Figure A.4 about here]

4.2. Forecasting

In this Section we add the estimated factors in an ARDL model to forecast individual index returns, as is outlined in Section 2.3, and compare the performance to a benchmark autoregressive model. We test our forecasting models on two distinct markets; Boston apartments and Dallas commercial. Not only are the property types different, the first location is typically "supply constraint", whereas the second is not (Saiz, 2010). We will be forecasting 8 quarters out of sample, or 2 years. We focus on two specific points in time; the eve of the Great Financial Crisis and the subsequent recovery. Both periods are characterized by a sudden and structural break with the past, making predictions difficult. We can subsequently compare the forecasted returns with the actual returns.

We have many combinations of possible optimal lag lengths. We can have different lags for the (univariate) benchmark model, separate for each index. The same can be said for fitting the individual univariate factor returns that need to be forecasted. Finally, when estimating the ARDL, both the AR coefficient and all the individual factor returns can have its own lag length. In our application we use every lag between 1 and 8. The only restriction

[4] Also, hotels are also included in the all types index by RCA. These are omitted from our panel data. Hotels represent only a small proportion of total transaction volume, so the impact is not likely to be substantial.

we impose is that the amount of lags on the factor trends and the ARDL are similar, reducing the number of combinations considerably. Note that we still have four different DFM specifications (with M = 1 through 4). Thus the total number of possible models is; 4 models $\times$ 8 lags = 32. The results of the "best" benchmark model and the "best" ARDL with factor model is given in Figures A.5 – A.6, for Boston and Dallas respectively. With "best" we mean the model that resulted in the lowest forecast residuals.

[Place Figure A.5 about here]

[Place Figure A.6 about here]

The benchmark AR model gives a familiar picture. Past price growth is taken as the projection for the future. The sudden breaks are therefore "ignored". Arguably, only the recovery in Boston is picked up slightly using the benchmark model. (But too slow and too late.) Somewhat surprising, is that the best performing benchmark model during the crisis was an AR(1) specification for both markets. In other words, when allowing for more lags, the benchmark model predicted even more price growth. During the recovery, the best benchmark model had 5 and 2 lags, for Boston and Dallas respectively.

In contrast to the benchmark model, the "best" ARDL with factor trends predicts really well. The crash is clearly visible in both markets, albeit slightly delayed in Dallas. The recovery is equally impressively forecasted. Typically, the "best" ARDL models has only 1 or 2 factors and 7 to 8 lags. The only "outlier" being the model that best predicted the crash in Dallas: Here the ARDL with 3 factor trends and 1 lag resulted in the best results.

In general, by looking at the results of other markets as well, we find that the best forecasts are rendered by low M high lag length ARDL models. These results are omitted to conserve space but are available upon request. Indeed, in none of our 80 markets do we find that the ARDL with 4 factor trends (with any lag length) results in the "best" forecasts. As mentioned before, this is probably caused by by the fact that we have to forecast multiple factor trends alongside our ARDL, which can result in more uncertainty.

It should be noted that we only present the "best" model based on comparing the forecasts with the actual realized returns. But not all ARDL specifications predicted the crisis so nicely. In fact, "only" approximately 1 in 5 ARDL models predicted a downturn. The best benchmark model also outperformed some of the ARDL models.[5] Still, an overall rule of thumb is that if one picks an ARDL with only 1 or 2 factor trends, it will typically outperform the "standard" benchmark model to a large extent. Interestingly enough, is that based on the AIC, the ARDL with more factors almost exclusively resulted in better model fit, even after penalizing for the amount of parameters. The AIC is not the only tool to measure model fit, and future research could benefit from implementing these different metrics. See Forni et al. (2000); Bai and Ng (2002); Breitung and Pigorsch (2013) among others.

Obviously, the results are based on the history of the RCA CPPIs, and there is no guarantee that we will be able to forecast the next big crash. Still, the results are impressive, especially considering we only used the factors and not any other data source, like local GDP and interest rates. (which is typical when forecasting real estate returns.)

5. Conclusion

In the last few years, the growth in the available amount of economic and financial data has prompted econometricians to develop or adapt new methods enabling them to summarize efficiently the information contained in large databases containing many time series. Of these methods, dynamic factor models (DFMs), has seen rapid growth and is very popular among macro-economists. These models can be used to summarize the information contained in a large number of economic (time) series into a small number of factors common to the original set of (time) series.

In this study we applied a dynamic factor model to commercial property prices. This is the first time that we can do such a comprehensive study in commercial real estate, due to the availability of very granular Commercial Property Price Indexes (CPPIs). We use 80 different CPPIs provided by Real Capital Analytics. These indexes are on a quarterly basis and run from 2001Q2 to 2017Q2. We estimate the DFM with 1, 2, 3 and 4 factor trends. Typically, we find 2 factor trends that keep re-emerging. The first is a typical boom-bust cycle. The second trend was already going down pre-crisis. Its recovery was also more swift. Commercial real estate tends to load highly

[5] During the crash there is also a clear case of "overfitting" with the ARDL model with high M and high lag length. This is evident from "extreme" high or low forecast, like predicting a 400% price increase in 2 years after 2007.

onto the first trend, whereas apartments are more affected by the second trend described here.

In a second step we use the estimated factor returns in an Autoregressive Distributed Lag (ARDL) model to forecast the individual price returns 2 years into the future. For the Boston apartment and Dallas commercial markets, we compare the forecast residuals to a "benchmark" Autoregressive (AR) model. Our findings suggests that the ARDL with factors greatly outperforms the "benchmark" model. The "benchmark" model almost completely misses the crisis and subsequent recovery, whereas the ARDL with factors is very accurate. Our results indicate that the forecasts generally improve when the factors are used. Interestingly, using less factors generally results in lower forecast error.

Improving forecasts is important in practice, as it is used in asset-allocation models, and for the pricing of real estate derivatives. Apart from forecasting per se, the DFM analysis also provides interesting insight about the structure of real estate returns, in particular regarding co-movements across markets that are important for portfolio diversification.

The goal of our forecasting exercise was not per se to find the "best" forecasting model. For example, we did not use any other explanatory variables. This can be seen as one of the pros of the DFM framework. (We do not *need* any other variables.) However, one can still use other regressors in the DFM and in the ARDL, and it *could* improve the forecasts. Also, we found that using the AIC - even after penalizing it for the amount of parameters - would typically, not pick the "best" model. Other Information Criteria exist, but are out of the scope of this research. Finally, other specifications / more modern versions of the DFM exists, which we did not pursue in this current study.

Aguilar, O., and M. West, 2000, Bayesian dynamic factor models and portfolio allocation, *Journal of Business & Economic Statistics* 18, 338–357.

Aßmann, C., J. Boysen-Hogrefe, and M. Pape, 2016, Bayesian analysis of static and dynamic factor models: An ex-post approach towards the rotation problem, *Journal of Econometrics* 192, 190–206.

Bai, J., and S. Ng, 2002, Determining the number of factors in approximate factor models, *Econometrica* 70, 191–221.

Bai, J., and P. Wang, 2015, Identification and bayesian estimation of dynamic factor models, *Journal of Business & Economic Statistics* 33, 221–240.

Barhoumi, K., O. Darné, and L. Ferrara, 2013, Dynamic factor models: A review of the literature, *OECD Journal of Business Cycle Measurement and Analysis* 2013.

Barkham, R., and D. M. Geltner, 1995, Price discovery in American and British property markets, *Real Estate Economics* 23, 21–44.

Basilevsky, A., 1994, *Statistical factor analysis and related methods: Theory and applications* (Wiley-Blackwell).

Bernaards, C. A., and R. I. Jennrich, 2005, Gradient projection algorithms and software for arbitrary rotation criteria in factor analysis, *Educational and Psychological Measurement* 65, 676–696.

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.

Boivin, J., and S. Ng, 2006, Are more data always better for factor analysis?, *Journal of Econometrics* 132, 169–194.

Breitung, J., and U. Pigorsch, 2013, A canonical correlation approach for selecting the number of dynamic factors, *Oxford Bulletin of Economics and Statistics* 75, 23–36.

Browne, M. W., 2001, An overview of analytic rotation in exploratory factor analysis, *Multivariate Behavioral Research* 36, 111–150.

Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The econometrics of financial markets*, volume 2 (princeton University press Princeton, NJ).

Case, K. E, and R. J. Shiller, 1989, The efficiency of the market of single-family homes, *The American Economic Review* 79, 125–137.

Del Negro, M., and C. Otrok, 2007, 99 luftballons: Monetary policy and the house price boom across US states, *Journal of Monetary Economics* 54, 1962–1985.

Eickmeier, S., and C. Ziegler, 2008, How successful are dynamic factor models at forecasting output and inflation? A meta-analytic approach, *Journal of Forecasting* 27, 237–265.

Favero, C. A., M. Marcellino, and F. Neglia, 2005, Principal components at work: the empirical analysis of monetary policy with large data sets, *Journal of Applied Econometrics* 20, 603–620.

Forni, M., M. Hallin, M. Lippi, and L. Reichlin, 2000, The generalized dynamic-factor model: Identification and estimation, *Review of Economics and statistics* 82, 540–554.

Francke, M. K., 2010, Repeat sales index for thin markets: a structural time series approach, *Journal of Real Estate Finance and Economics* 41, 24–52.

Francke, M. K., and A. F. De Vos, 2000, Efficient computation of hierarchical trends, *Journal of Business and Economic Statistics* 18, 51–57.

Francke, M. K., D. M. Geltner, A. M. Van de Minne, and R. White, 2017, Real estate index revisions in thin markets, *MIT Working Paper* .

Francke, M. K., and A. M. van de Minne, 2017, The hierarchical repeat sales model for granular commercial real estate and residential price indices, *The Journal of Real Estate Finance and Economics* 55(4), 511–532.

Geltner, D., and J. Mei, 1995, The present value model with time-varying discount rates: Implications for commercial property valuation and investment decisions, *The Journal of Real Estate Finance and Economics* 11, 119–135.

Geltner, D. M., and J. Fisher, 2007, Pricing and index considerations in commercial real estate derivatives, *Journal of Portfolio Management* 99.

Geltner, D. M., B. D. MacGregor, and G. M. Schwann, 2003, Appraisal smoothing and price discovery in real estate markets, *Urban Studies* 40, 1047–1064.

Geweke, J., and G. Zhou, 1996, Measuring the pricing error of the arbitrage pricing theory, *Review of Financial Studies* 9, 557–587.

Goetzmann, W. N., 1992, The accuracy of real estate indices: Repeats sale estimators, *Journal of Real Estate Finance and Economics* 5, 5–53.

Gorman, W. M., 1981, Some engel curves, *Essays in the theory and measurement of consumer behaviour in honor of sir Richard Stone* .

Harman, H. H., 1976, *Modern factor analysis* (University of Chicago Press).

Harvey, A., 1989, *Forecasting Structural Time Series Models and the Kalman Filter* (Cambridge University Press, Cambridge).

Lewbel, A., 1991, The rank of demand systems: theory and nonparametric estimation, *Econometrica: Journal of the Econometric Society* 711–730.

Naranjo, A., and D. C. Ling, 1997, Economic risk factors and commercial real estate returns, *The Journal of Real Estate Finance and Economics* 14, 283–307.

Quan, D. C., and J. M. Quigley, 1991, Price formation and the appraisal function in real estate markets, *The Journal of Real Estate Finance and Economics* 4, 127–146.

Ross, S. A., 1976, The arbitrage theory of capital asset pricing, *Journal of economic theory* 13, 341–360.

Saiz, A., 2010, The geographic determinants of housing supply, *The Quarterly Journal of Economics* 125, 1253–1296.

Schwann, G. M., 1998, A real estate price index for thin markets, *Journal of Real Estate Finance and Economics* 16, 269–287.

Stock, H., J, and M. W. Watson, 2002, Macroeconomic forecasting using diffusion indexes, *Journal of Business & Economic Statistics* 20, 147–162.

Stock, J. H., and M. W. Watson, 1998, Diffusion indexes, Technical report, National bureau of economic research.

Stock, J. H., and M. W. Watson, 2005, Implications of dynamic factor models for VAR analysis, Technical report, National Bureau of Economic Research.

Wu, L. S., J. S. Pai, and J. R. M. Hosking, 1996, An algorithm for estimating parameters of state-space models, *Statistics & probability letters* 28, 99–106.

Zuur, A. F., R. J. Fryer, I. T. Jolliffe, R. Dekker, and J. J. Beukema, 2003, Estimating common trends in multivariate time series using dynamic factor analysis, *Environmetrics* 14, 665–685.

Appendix A. Figures

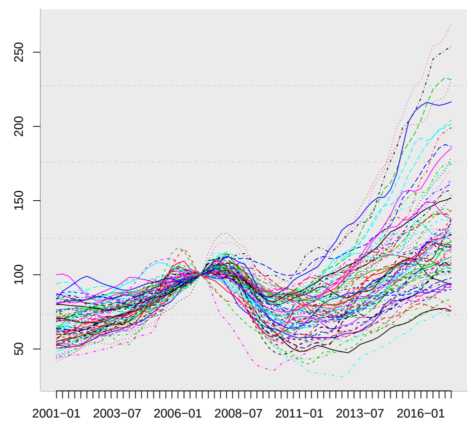

(a) Price levels.

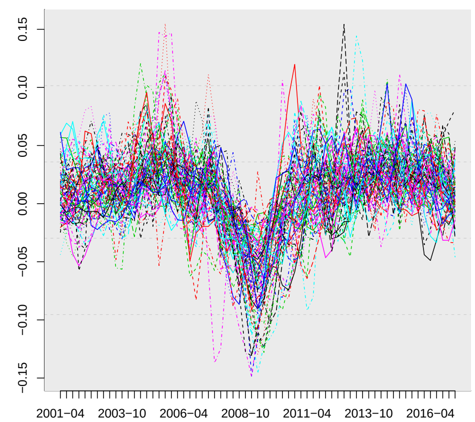

(b) Price returns.

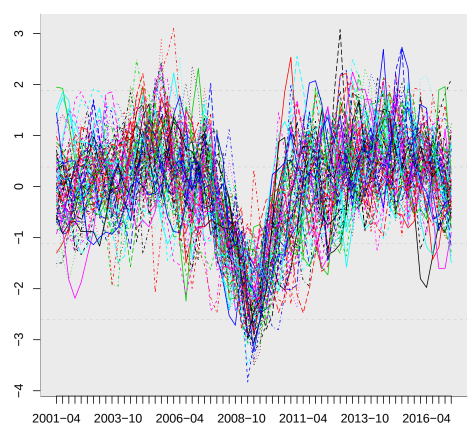

(c) Standardized returns.

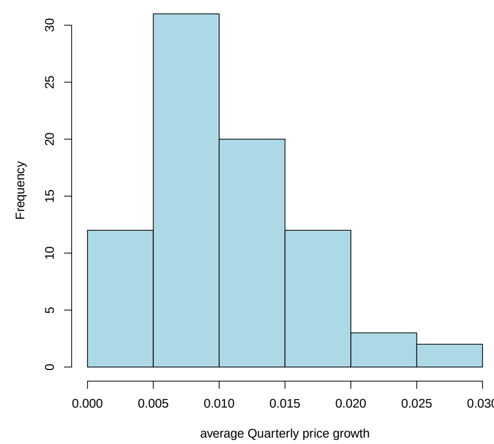

(d) Price Growth.

Figure A.1: Price indexes and returns.

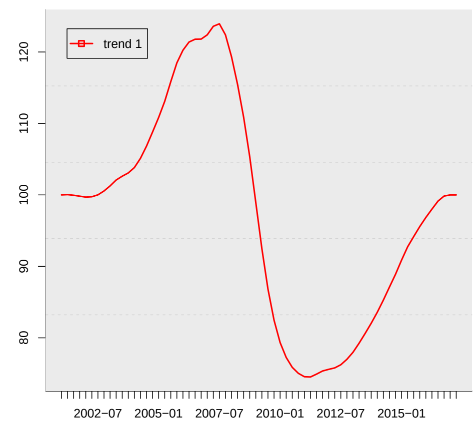

(a) Factor index, $M = 1$.

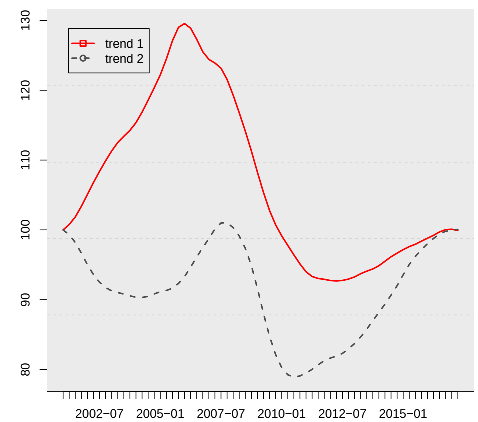

(b) Factor index, $M = 2$.

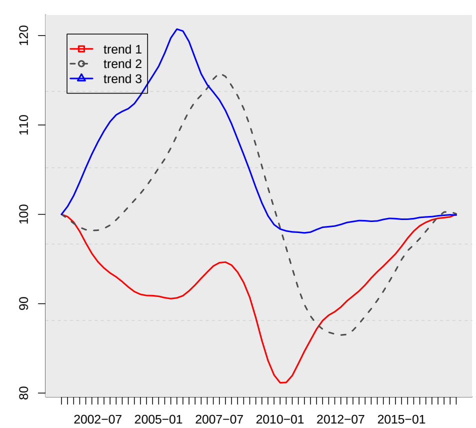

(c) Factor index, $M = 3$.

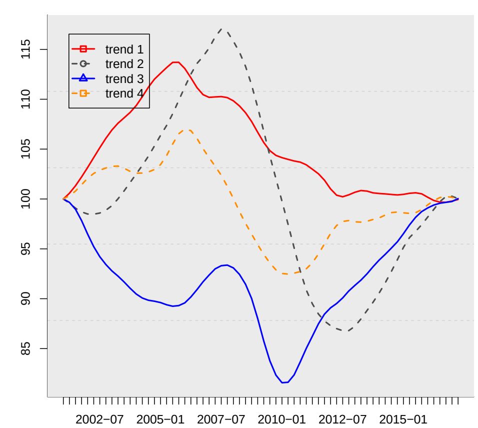

(d) Factor index, $M = 4$.

Figure A.2: Unstandardized factor indexes with diagonal and unequal variance-covariance structure.

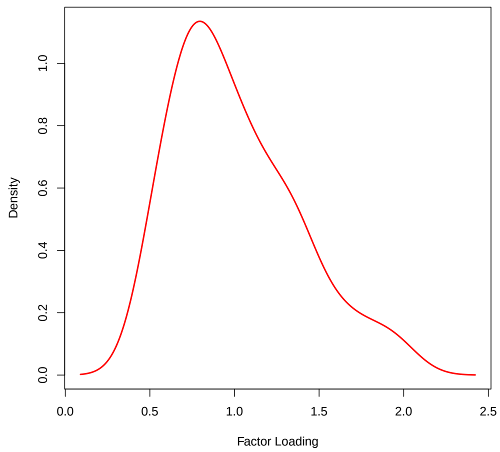

(a) Factor loadings, $M = 1$.

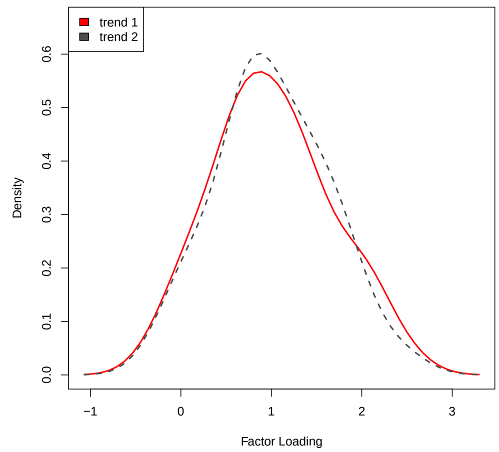

(b) Factor loadings, $M = 2$.

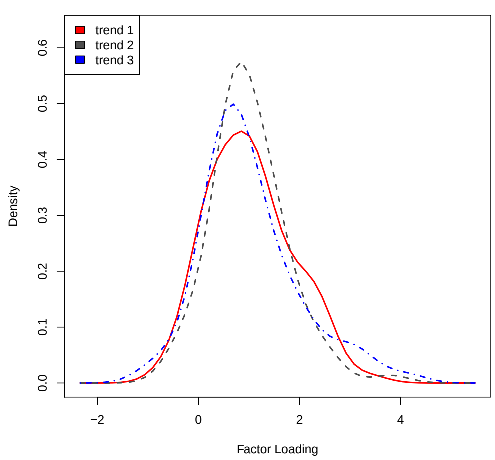

(c) Factor loadings, $M = 3$.

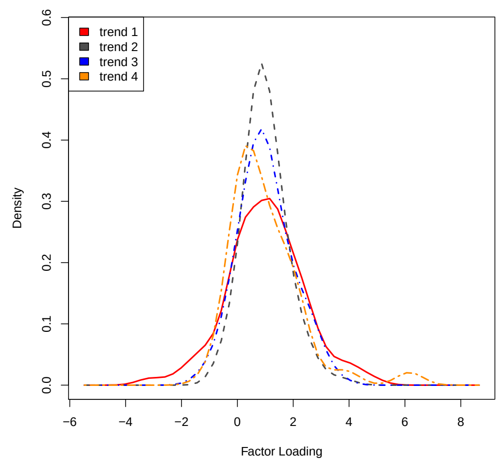

(d) Factor loadings, $M = 4$.

Figure A.3: Kernal distributions of the (unstandardized) factor loadings with diagonal and unequal variance-covariance structure.

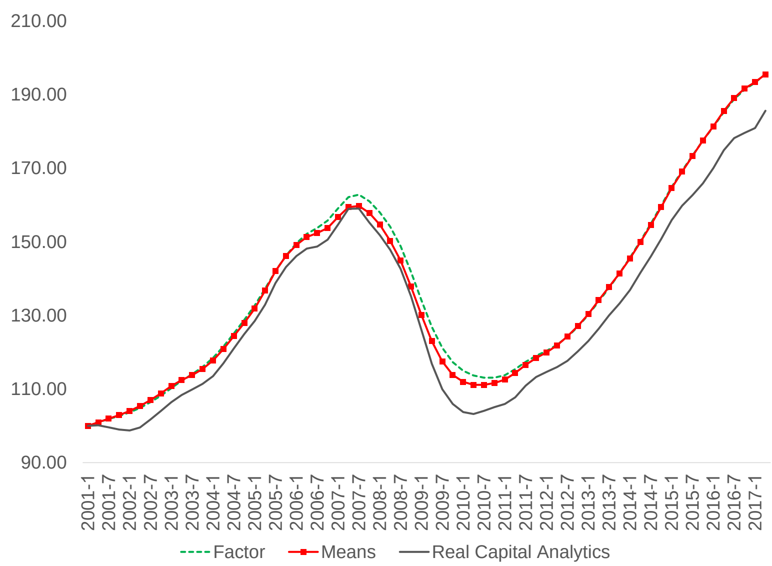

Figure A.4: Factor "index" with M=1 (**Factor**), the mean index values of our data (**Means**), and the 'all types US index' published by **Real Capital Analytics**.

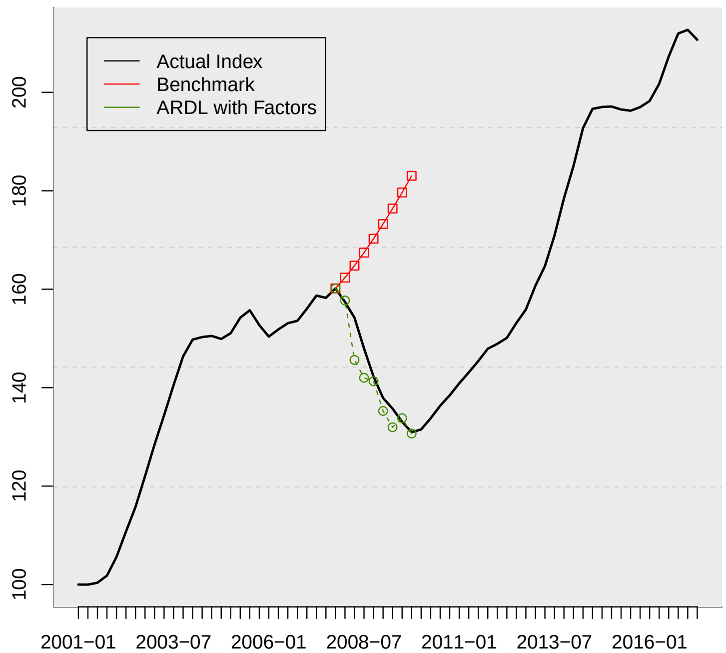

(a) Crisis.

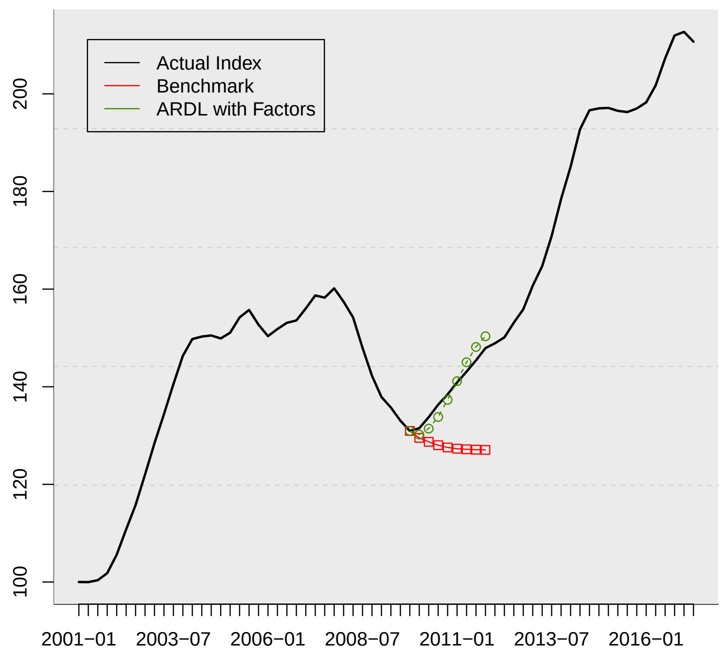

(b) Recovery.

Figure A.5: Boston apartment 2Yr forecasts.

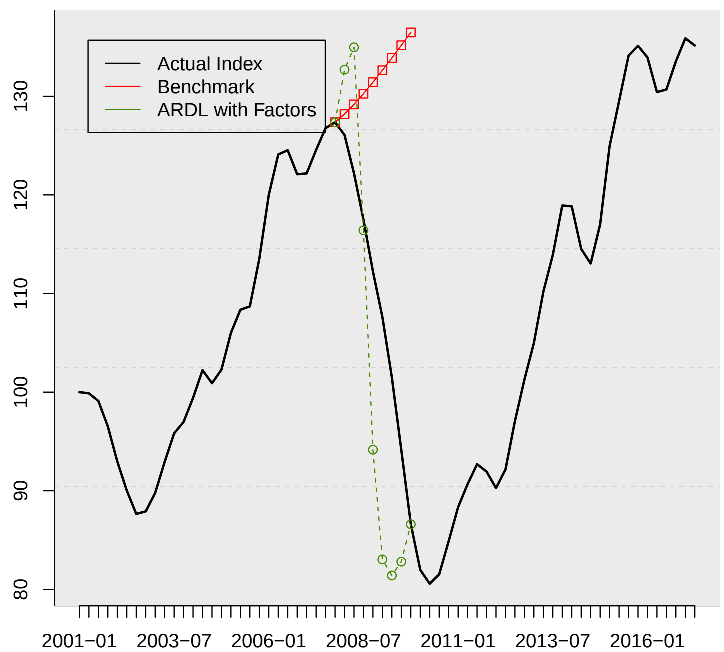

(a) Crisis.

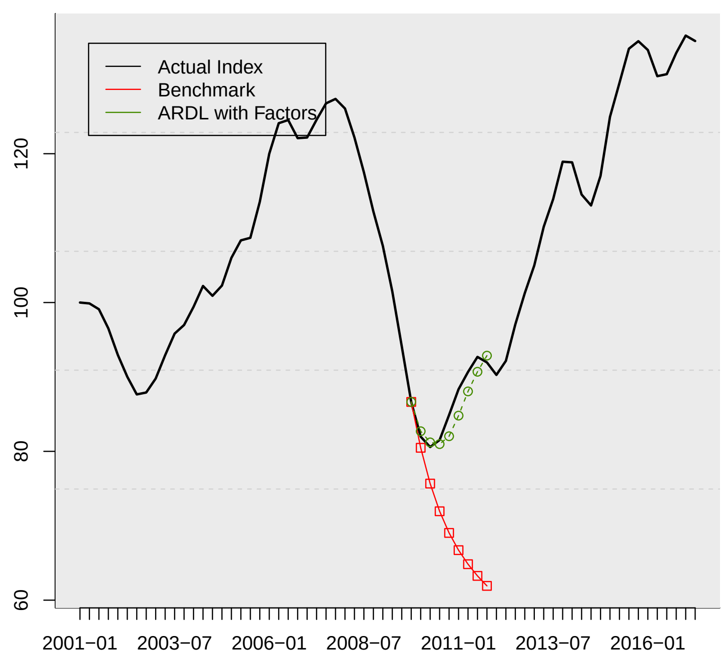

(b) Recovery.

Figure A.6: Dallas commercial 2Yr forecasts.

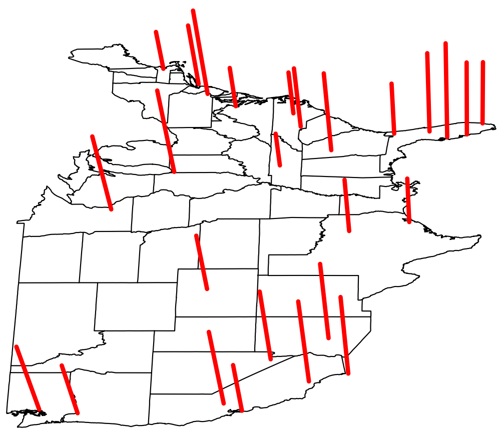

(a) Commercial.

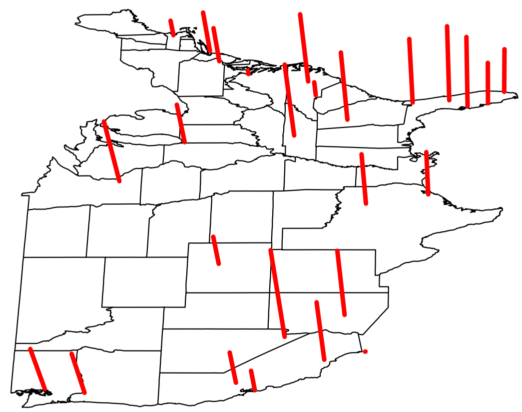

(b) Apartments.

Figure A.7: Factor Loadings with M=1.

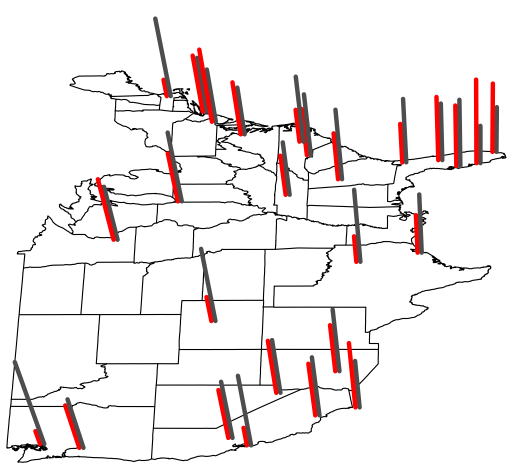

(a) Commercial.

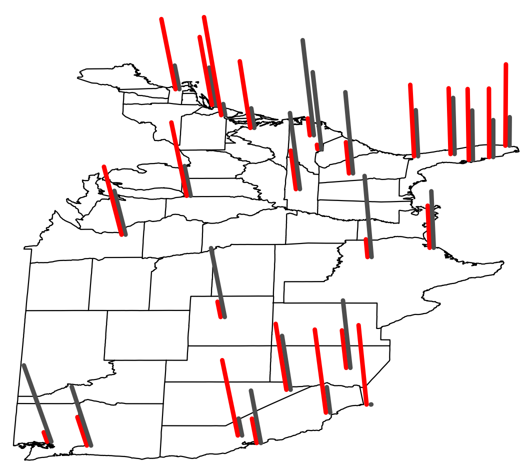

(b) Apartments.

Figure A.8: Factor Loadings with M=2.

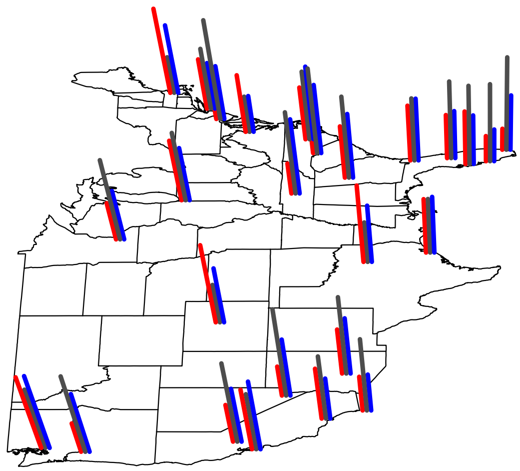

(a) Commercial.

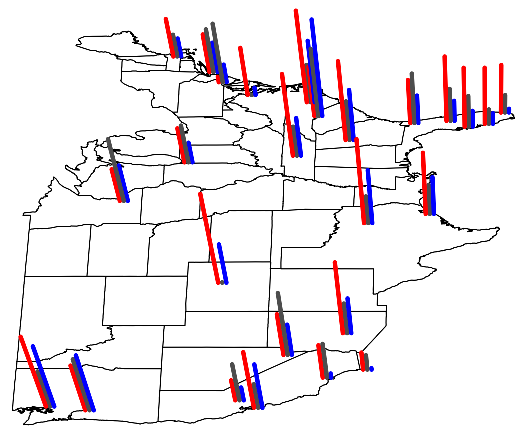

(b) Apartments.

Figure A.9: Factor Loadings with M=3.

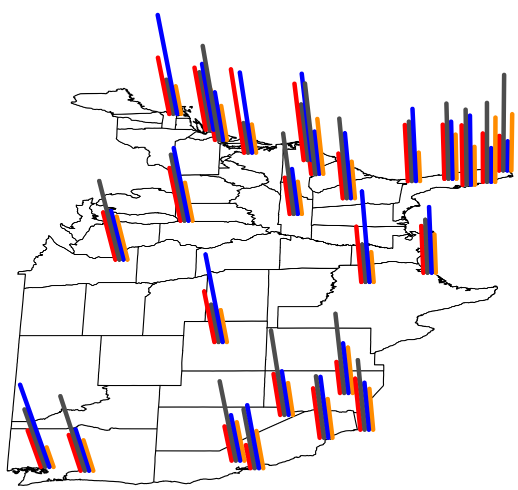

(a) Commercial.

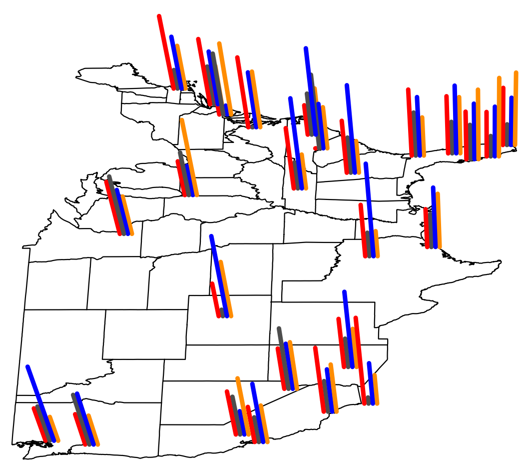

(b) Apartments.

Figure A.10: Factor Loadings with M=4.

Appendix B. Tables

Table B.1: Model fit of different specifications of the DFM, ordered by model fit.

R	**M**	**logLik**	**K**	**AICc**
diagonal and equal	4	-4,992.7	315	10,656.2
diagonal and equal	3	-5,130.8	238	10,760.5
diagonal and equal	2	-5,347.5	160	11,025.3
diagonal and equal	1	-5,588.4	81	11,341.5

R gives the design of the variance-covariance matrix of the observation errors, M is the number of factors, logLik denotes the Log likelihood, K is the number of parameters, AICc is the weighted or penalized AIC.

Appendix C. Spatial Distribution of Factor Loadings

The lowest factor loading per factor index is given a 0, and the highest factor loading is given a 1. Whenever a factor loading is between the min and max, the relative position between the two is given. Some markets are omitted from these figures, as they are a "rest" category, like "South East Rest". Also, if we have multiple markets within 1 metro we take the average. This is the case for Chicago, New York, San Fransisco, Miami and Los Angeles.

[Place Figure A.7 about here]

[Place Figure A.8 about here]

[Place Figure A.9 about here]

[Place Figure A.10 about here]