

Forecasting Real Estate Prices*

Eric Ghysels

UNC[†]

Alberto Plazzi

University of Lugano and SFI[‡]

Walter Torous

UCLA[§]

Rossen Valkanov

UCSD[¶]

This Version: July 1, 2012

Abstract

This chapter reviews the evidence of predictability in US residential and commercial real estate markets. First, we highlight the main methodologies used in the construction of real estate indices, their underlying assumptions and their impact on the stochastic properties of the resultant series. We then survey the key empirical findings in the academic literature, including short-run persistence and long-run reversals in the log changes of real estate prices. Next, we summarize the ability of local as well as aggregate variables to forecast real estate returns. We illustrate a number of these results by relying on six aggregate indexes of the prices of unsecuritized (residential and commercial) real estate and REITs. The effect of leverage and monetary policy is also discussed.

Keywords: real estate, predictability, market efficiency, REIT.

JEL Classification: G14, G19, R30, R31, R33

*Comments Welcome. Prepared for the *Handbook of Economic Forecasting*: Vol II, G. Elliott and A. Timmermann (Eds.), Elsevier, forthcoming. We thank the participants at the 2011 Saint Louis Fed conference for suggestions and comments. We are especially grateful to the Editors, G. Elliott and A. Timmermann, and two anonymous referees for many suggestions that have greatly improved the article. All remaining errors are our own.

[†]Department of Economics and Kenan-Flagler Business School, Gardner Hall, Chapel Hill, NC 27599-3305, phone: (919) 962-9810, e-mail: eghysels@unc.edu.

[‡]Institute of Finance, Via Buffi 13, Lugano, 6904, Switzerland, phone: +41 (0)58 666-4677, email: alberto.plazzi@usi.ch.

[§]UCLA Anderson School, 110 Westwood Plaza, Los Angeles, CA 90095-1481, phone: (310) 825-4059, fax: (310) 206-5455, e-mail: walter.n.torous@anderson.ucla.edu.

[¶]Rady School of Management, UCSD, Pepper Canyon Hall, 9500 Gilman Drive, La Jolla, CA 92093, phone: (858) 534-0898, e-mail: rvalkanov@ucsd.edu.

1 Introduction

The importance of real estate as an asset class cannot be overstated. Its total value in the US at the end of 2011 was about \$25 trillion, of which more than \$16 trillion was in residential properties.¹ By comparison, at the end of the same year, the capitalization of the US stock market was in the neighborhood of \$18 trillion. Moreover, recent history suggests that fluctuations in real estate prices, whether in bubble or burst mode, have the potential to buoy up or wreak havoc on the financial sector and the rest of the economy. Some of that impact is due to leverage and the fact that real estate is the easiest asset to borrow against, especially from a household's perspective. Indeed, in 2011, about \$12 trillion of outstanding mortgage debt had been issued against the value of residential real estate.² But the connection between real estate and the macroeconomy is not a bubbles-only phenomenon. Case, Quigley, and Shiller (2005) show that variations in real estate prices have had a significant effect on aggregate consumption in the US, in fact more significant than the stock market, even before the recent volatility in the residential market. Reinhart and Rogoff (2009) document this to be the case more universally across a number of countries and over longer time periods. From an economic perspective, understanding what drives real estate values is no less important than is understanding the pricing dynamics of other asset classes, such as stocks, bonds, commodities, and currencies.

The real estate market is different from other financial markets in several important aspects. It is characterized by extreme heterogeneity due to the location and physical attributes of a property. Participants in that market face large transaction costs, carrying costs, illiquidity, and tax considerations. They also face large search costs stemming from real estate's heterogeneity. Investors have just limited possibility of exploiting forecast decreases in property values, because of the impossibility to short sale a specific asset and the absence of liquid real estate futures contracts.³ These large frictions suggest that the real estate market might generally not be efficient in the sense that other financial markets are (e.g., Fama (1970)). But before we can talk about market efficiency, which implies that some investors are able to take advantage of profit opportunities, we first must

¹The source is the Flow of Funds Accounts. This estimate obtains summing the value of Households real estate (Table B.100, line 4) with that of Nonfinancial Corporate Business (Table B.102, line 3).

²The source is the Flow of Funds Accounts, Total Mortgages (Table L.217), obtained as the sum of Home and Multifamily residential.

³There are some recently launched indexes that track the performance of residential and commercial mortgage backed securities. These indices allow investors to take short positions in assets that are undoubtedly correlated with the aggregate real estate portfolio. Shorting the construction sector of the economy is another indirect way of shorting real estate.

investigate whether price changes are in fact statistically predictable.

The presence of frictions does not imply that predicting real estate returns is an easy task. In practice, the opposite is true. An illustration of this fact can be gleaned from the transcripts of the Federal Open Market Committee's (FOMC) 2006 discussions, which were held at the peak of the recent housing bubble. This was a time when a growing consensus amongst economists that residential prices were inflated coincided with a growing uncertainty about their future direction. The transcripts reveal that most FOMC participants shared the opinion that we were in for a "a soft landing or a period of stabilization after several years of strong price appreciation."⁴ Now, with the benefit of a hindsight and the Great Recession behind us, we know that this prediction was considerably off the mark. Long-horizon forecasts can be equally challenging to make. One such forecast was formulated by Mankiw and Weil (1989) who argued that the rise of housing prices in the 1970s and 1980s was mostly due to the Baby Boom generation entering the residential market. Based on these findings and reasoning that future demand for housing will decrease over the next twenty years, the authors predicted that "real housing prices will fall substantially - indeed, real housing prices may well reach levels lower than those experienced at any time in the past forty years." Now, twenty or more years after Mankiw and Weil (1989) formulated this forecast, we have observed that the trends and volatility in the housing market were driven by factors other than demographic fundamentals.

In this chapter, we review the literature on return predictability in real estate markets. Many of the papers on this topic involve the use of indices at the city, regional, or national level rather than individual property prices. This is due to one obvious reason: real estate transactions are very infrequent. Hence, as a starting point, we discuss the construction and underlying assumptions behind some of the most widely used residential and commercial real estate indices. The distinction between residential and commercial properties is important as they tend to have different dynamics and return properties (Geltner and Miller (2006)). The difference is not surprising as a household's decision to purchase a home—presumably driven not only by investment considerations but also by the need to consume a housing unit—is quite different from that of an investor looking to purchase a retail property (Flavin and Yamashita (2002)). Perhaps the most transparent residential index is the median sales price, versions of which are provided by the Census Bureau and the National

⁴Excerpts from the transcript of the Meeting of the Federal Open Market Committee, May 10, 2006, available at <http://www.federalreserve.gov/monetarypolicyfiles/FOMC20060510meeting.pdf>.

Association of Realtors (NAR). While it is easy to construct and interpret, it does not adjust for the quality of properties that are on the market and thus confounds fluctuations in prices with fluctuations in real estate attributes. The fact that it is not a “constant quality” index makes it less desirable than some of the alternatives which specifically adjust for property attributes. Examples of constant-quality indices include, for residential properties, the Case-Shiller and the Federal Housing Finance Agency (FHFA) repeat-sales prices and, for commercial properties, the National Council of Real Estate Investment Fiduciaries’s Transaction-Based (TBI) hedonic prices. We discuss these and other indices in some detail in Section 2, because their statistical properties are determined as much by assumptions behind their construction as by market forces. Understandably, much creative energy and papers have been devoted to this topic. Without reliable indices, empirical research in real estate is virtually impossible.

To get a glimpse into the aggregate real estate data, in Figures 1 and 2 we plot three residential and three commercial indices that have been widely used in the literature and whose properties we will analyze in this Chapter. It is immediately clear that the three time-series in Figure 1 do not exhibit the same dynamics despite the fact that they are all intended to measure the same price appreciation of houses in the US. For instance, the growth rate (not the level) of the Case-Shiller index has a serial correlation of 0.939. Some of that serial correlation is due to the way the index is constructed and some of it is undoubtedly due to economic frictions. The same statistic for the growth rate of the Census median price is -0.517. Similarly dramatic differences are evident in Figure 2, where the price of a real estate investment trusts (REIT) portfolio exhibits volatility that dwarfs that of the other two indices. Before we can use these time-series, we have to understand how they are constructed, what they measure, and whether they are suitable for forecasting. In Section 2.1, we discuss the various types of real estate indices, provide details behind their constructions and, in Section 2.2, we present their summary statistics.

Predictive regressions in the real estate literature in many respects mirror those in other asset classes. The forecasted quantity, often future price changes or returns, is regressed on a set of predetermined variables, which are chosen to test a set of economic hypotheses. The presence of return forecastability is interpreted either as evidence of market inefficiencies (Fama (1970)) or of time-varying risk premia in otherwise efficient markets. However, predictive regressions are inherently reduced-form expressions and cannot identify the economic reasons underlying the

forecastability without further modeling restrictions (Fama and French (1988a)).

It is useful to divide predictive regressions into three categories based on the predictors and maintained hypotheses. First, if the predictor is lagged returns as in Gau (1984), Gau (1985), Linneman (1986), Guntermann and Smith (1987), Rayburn, Devaney, and Evans (1987), Case and Shiller (1989), McIntosh and Henderson (1989), Gyourko and Voith (1992), Kuo (1996), Hill, Sirmans, and Knight (1997), Hill, Sirmans, and Knight (1999), Gu (2002), and Schindler (2011), this is a test of weak-form market efficiency. However, simple serial dependence tests in real estate are complicated by the fact that the price changes of some indices are serially correlated by construction, making it difficult to disentangle spurious correlations from actual market inefficiencies. The evidence of weak-form market efficiency is presented in Section 3.1.

Second, valuation ratios – such as the rent-price ratio or price-income ratio – are often used as predictors (e.g., Hamilton and Schwab (1985), Meese and Wallace (1994), Geltner and Mei (1995), Capozza and Seguin (1996), Lamont and Stein (1999), Malpezzi, Himmelberg, Mayer, and Sinai (2005), Campbell, Davis, Gallin, and Martin (2009), Gallin (2008)). We review predictive regression with valuation ratios in Section 3.2. Such regressions are motivated either by the valuation ratios' ability to detect deviations and slow adjustments toward an equilibrium or because they proxy for time variation in expected returns. Overall, the evidence supports the view that valuation ratios are not able to capture all the variation in the conditioning set.

Third, a richer set of hypotheses can be tested by including property- and/or region-specific economic variables, whose aim is to proxy for demand and supply shocks in the real estate market. Such predictors, used, amongst others, by Rosen (1984), Linneman (1986), Skantz and Strickland (1987), Case and Shiller (1990), Abraham and Hendershott (1996), Pace, Barry, Gilley, and Sirmans (2000), MacKinnon and Al Zaman (2009), and Plazzi, Torous, and Valkanov (2010), include demographic variables, income variables, construction costs, and zoning restrictions. One can argue that these regressions account more fully for the heterogeneity in real estate investments. The evidence from these regressions is reviewed in Section 3.3.

Real estate data present some unique challenges in forecasting settings. First, the predictive results have to take into account high transaction costs, which in real estate can be 6 percent, or even higher, of the property value. In addition to the statistical significance, the coefficient estimate

must be large enough to cover those costs. Second, the available real estate data is relatively short in its duration and is observed at a monthly or quarterly frequency.⁵ Sparse datasets are available from the 1970s, but most empirical work is done with series starting from the 1980s or later. The lack of longer and higher-frequency data renders estimation and hypothesis testing difficult. Third, the in-sample fit of predictive regressions is often traced to dichotomous variables for geographical location, coastal proximity, or whether a commercial property is of a certain type (apartments, retail space, offices, or industrial buildings). While these fixed-effects are important in accounting for the heterogeneity of the asset, they are not predictors in the usual sense of the word. They do not change over time and cannot be the source of time-series predictability. Fourth, the predictability evidence is mainly based on in-sample statistics. It is rarely evaluated with mean squared prediction errors (MSPE) or other out-of-sample analysis (West (2006)), mainly because of the severe data limitations. Finally, and related to the previous point, parameter stability and the robustness of the forecasting model are rarely investigated (e.g., Rossi (2012)).

A real-estate-related market that does not suffer from the high transaction costs and infrequent observations issues is that of publicly-traded real estate investment trusts. REITs are exchange-traded funds that derive most of their income from real estate investments and whose returns provide a remarkably clean venue for testing whether or not real estate returns are forecastable. The REIT market has been given particular attention in the real estate literature and we are also devoting special attention to it in Section 4. The focus on REITs is attributable to the fact that their returns have less measurement error and are observable at higher frequency than other real estate investments. Hence, econometric issues arising in the estimation and forecasting of returns can largely be addressed.⁶ However, empirical work with REITs does have its limitations. For instance, investing in a REIT is not the same as investing in the underlying commercial property market. The risk-return characteristics of the investments might be different. Ross and Zisler (1991) document that REITs have the risk-return profile of small-cap stocks and co-move more with the stock market rather than the underlying real estate market.

We supplement the summary of existing findings with our own set of predictive regressions estimated with three residential and three commercial real estate indices at the national (US) level

⁵A notable exception is Eichholtz (1997) whose bi-annual residential index of Amsterdam properties spans the period 1628 to 1973.

⁶These issues stem from the fact that the predictor is often a near-integrated process whose innovations are correlated with the innovations of returns (Cavanagh, Elliott, and Stock (1995), Stambaugh (1999)).

with monthly and quarterly data from 1991 to the end of 2010. We choose to work mainly with aggregate indices as they are available over a long time span and allow us to keep a common set of predictors. Some cross-sectional results are provided using metropolitan level data for residential properties. Following the literature, we run various specifications of the predictive regressions. For all non-REIT indices, the data restricts us to stay with in-sample comparisons. However, with REITs, we are able to estimate a more complete (and interesting) predictive system, to adjust the estimates for known small-sample biases, to impose relevant economic restrictions, to look for predictability at various horizons, and out-of-sample. These results are discussed in Sections 3 and 4.

The data and empirical methods that we survey are remarkably diverse. As a way of providing a bird's eye view of this body of work, in Table 1 we summarize most of the covered papers along with the main findings pertaining to returns forecastability. The papers differ not only in their econometric approach, but also in the type of properties they investigate (residential, commercial, REITs), the geographical coverage of the data, the time span, and the conditioning variables. Despite all these differences, a common set of findings emerge from this literature, most of which we are able to observe in our aggregate predictive regressions. These stylized facts can be summarized as follows:

- Price changes of repeat-sales and hedonic indices are very positively serially correlated at monthly and quarterly frequency, whereas median price indices exhibit negative serial correlation. The serial correlation of REIT returns is similar to that of small-cap stocks.
- Transaction costs and other frictions are too large for the serial correlation to translate into economic gains for the non-REIT and REIT data.
- Valuation ratios, such as the rent-price ratio or the income-price ratio, have some predictive power, in-sample. It is mostly attributable to time varying expected returns rather than to exploitable market inefficiencies.
- Variables, such as construction costs, demographic changes, and regulatory restrictions, have a sizeable impact on future real estate returns. However, the out-of-sample properties of these forecasts are largely unexplored.

- For REITs, there is weak in-sample evidence of return predictability and stronger evidence for rent-growth predictability. The evidence for out-of-sample predictability is not strong.
- Leverage is positively related to future returns.

We should point out that several of these results obtain in pooled regressions with more cross-sectional than time-series observations. Hence, their “predictive” nature should be taken with some caution. This is a general theme in real estate research, as the lack of good time-series data prevents us from applying the standard forecasting toolbox. For instance, it precludes the use of what some might consider to be the ultimate forecast evaluation tool: out-of-sample predictive measures, such as the MSPE comparison. The MSPE analysis is mostly asymptotic in nature and, as West (2006) points out, is not useful for studies with only a handful of observations for out of sample evaluation. However, the lack of a long time-series data also raises a host of new interesting questions, such as how to use the richness of the cross-section in formulating and evaluating forecasts.

There is also a growing literature investigating the portfolio implications of investing in housing, as in Ross and Zisler (1991), Goetzmann (1993), Flavin and Yamashita (2002), and Cauley, Pavlov, and Schwartz (2007). Recent papers by Lustig and Van Nieuwerburgh (2005), Campbell and Cocco (2007), and Piazzesi, Schneider, and Tuzel (2007) investigate the connection between housing, consumption, and asset pricing. Becker and Shabani (2010) examine the implications of large household debt (a mortgage) on investment decisions. As a next step, it would be interesting to incorporate predictable real estate returns in the optimal investment decision of households. The portfolio setting would provide a natural measure of the economic importance of real estate return predictability in absolute terms and relative to the predictability of stock and bond returns. A first step in that direction is the work of MacKinnon and Al Zaman (2009) and Plazzi, Torous, and Valkanov (2010).

Following the recent real estate crisis, two research areas have generated considerable interest. First, there has been renewed attention on the role of leverage on house price dynamics. Theoretical work by Stein (1995), McDonald (1999), Spiegel (2001), and Ortalo-Magné and Rady (2006) and the empirical findings of Linneman and Wachter (1989), Genesove and Mayer (1997), and Lamont and Stein (1999), Brown (2000), Aoki, Proudman, and Vlieghe (2004), and Favilukis, Ludvigson, and Van Nieuwerburgh (2011) suggest that leveraged properties are more sensitive to

economic shocks. The main amplifying channel in those papers is due to the fact that a household's ability to borrow is directly tied to asset values. While not all of the papers we discuss in that section contain direct evidence of real estate predictability, they all suggest that leverage is an important determinant of house price dynamics. On a related topic, Mian and Sufi (2009), (2010) and Favilukis, Kohn, Ludvigson, and Van Nieuwerburgh (2012) analyze the effect of the recent credit expansion on real estate prices. We review this literature in Section 5.1. Second, the recent real estate crisis has focused attention on the effect of monetary policy on real estate prices which we survey in Section 5.2. Section 6 concludes and offers directions for current and future research.

2 The Real Estate Data

In this section, we summarize the availability of US real estate data.⁷ The focus is on price indices which can naturally be categorized, based on the construction methodology, into four groups: median-price indices, repeat-sales indices, hedonic indices, and stock-market-based indices. In each category, the indices can track either residential or commercial properties.⁸ In subsection 2.1, when introducing the various indices, we pay particular attention to the way they are constructed and the effect of the construction on subsequent forecasting regressions. Subsection 2.2 presents summary statistics for several well-known indices that will be used in the rest of this chapter.

2.1 Real Estate Index Definitions

2.1.1 Median Price Indices

Median price indices track the price at which the median priced home within a particular area trades in a given period. At present, residential median-price indices are provided for U.S. single-family homes by the Federal Housing Finance Agency, the Census Bureau (new homes)⁹, and by the National Association of Realtors (existing homes)¹⁰. They are all available at monthly

⁷While we try to cover as much ground as possible, the increasing interest in real estate and the lower cost of information acquisition have resulted in an increase of available data sources.

⁸A third useful categorization is to decompose a property value into land and improvements. While we do not analyze land data in our article, aggregate, state-level, and MSA-level price indices for land are provided by the Lincoln Institute following the work of Davis and Heathcote (2007). See <http://www.lincolninst.edu/>.

⁹See <http://www.census.gov/const/www/newressalesindex.html>.

¹⁰See <http://www.realtor.org/>.

frequency, but the sample spans vary.

The appeal of median price indices stems from the ease with which they can be computed and their unambiguous interpretation. However, they ignore potentially important changes in the characteristics of the dwellings being sold. In particular, if high-quality and low-quality homes are put on the market at different times, the corresponding median prices may exhibit spurious time series fluctuations due mainly to differences in the quality of the properties. Moreover, it is reasonable to expect the mix of homes sold to be correlated with local economic conditions as more expensive homes will tend to be put on the market in expansionary times. These considerations suggest that a median price index provides not only a noisy but also a systematically biased estimate of the behavior of home prices in a particular market. Because of this, attempts have been made to keep the quality of the median house constant through time when constructing these indices. Corrections also include stratification methods which adjust for compositional changes in the transactions, as in Prasad and Richards (2008).

By far, the two most popular methods to explicitly control for quality and infrequent trading of real estate properties are repeat-sales regressions and hedonic models, which we review next.

2.1.2 Repeat-Sales Indices

Repeat-sales indices use information about homes that transact at least twice during the sample period to infer market-wide price movements. Let $P_{i,t}$ be the price of home i at the end of period t , and $p_{i,t}$ its logarithmic transformation. The standard repeat-sales approach models $p_{i,t}$ as the sum of two components:

$$p_{i,t} = p_{m,t} + e_{i,t} \quad (1)$$

where $p_{m,t}$ denotes the aggregate real estate index – an equally-weighted portfolio of properties – and $e_{i,t}$ is a property-specific mean-zero stochastic drift. Shocks to $e_{i,t}$, denoted by $\varepsilon_{i,t} = e_{i,t} - e_{i,t-1}$, are assumed to be i.i.d. both cross-sectionally as well as over time, with finite variance σ_ε^2 .

We are interested in obtaining estimates of $p_{m,t}$ using a sample of I individual property transaction prices, but we don't have observations for the same property in each period. Instead, we have information for home i , $i = \{1, \dots, I\}$, on the date of its initial purchase t_i , the date of its first sale $t_i + T_i$, with $T_i \geq 1$, and the corresponding prices. The subscript i captures the fact that the

transaction dates t_i and $t_i + T_i$ are dwelling-specific. Then, following (1), the log return during the $[t_i; t_i + T_i]$ period can be expressed as:

$$p_{i,t_i+T_i} - p_{i,t_i} = p_{m,t_i+T_i} - p_{m,t_i} + \sum_{\tau=t_i+1}^{t_i+T_i} \varepsilon_{i,\tau}. \quad (2)$$

Motivated by this expression, the standard repeat-sales regression (RSR) approach of Bailey, Muth, and Nourse (1963) consists of estimating via ordinary least squares the following cross-sectional regression

$$y_i = \beta X_i + u_i \quad (3)$$

where the dependent variable y_i is the holding period return for home i , or $y_i = p_{i,t_i+T_i} - p_{i,t_i}$. The regressor X_i is a dummy variable that contains values for each time period, except for the first. It equals 1 on the first sale date $t_i + T_i$, -1 on the purchase date t_i , and 0 otherwise. If the purchase period t_i coincides with the first period of the sample, then the purchase date dummy is omitted. If there are a total of $T + 1$ periods, then the $T \times 1$ vector $\hat{\beta}$ provides an estimate of the log price of the aggregate index, or $\hat{p}_{m,t} = \hat{\beta}_t$. The value of the log price in the initial period is normalized to zero.

The OLS estimator of (3) is, however, not efficient. In particular, the variance of the error term $u_i = \sum_{\tau=t_i+1}^{t_i+T_i} \varepsilon_{i,\tau}$ increases linearly with the interval of time between the two transaction dates, or $\sigma_{u_i}^2 = T_i \sigma_\varepsilon^2$. As a result, the OLS estimator overweighs the information on transactions that occur after longer time intervals, ignoring the larger noise embedded in their price changes. In this context, the best linear unbiased estimator (BLUE) is a GLS estimator of (3) where each observation is weighted by the inverse of the square root of its holding period. The resultant error terms are now i.i.d. and the system satisfies the Gauss-Markov conditions. This GLS estimator coincides with the maximum likelihood estimator when we assume normality of the underlying $\varepsilon_{i,s}$ (Goetzmann (1992) and references therein).

Case and Shiller (1987) extend model (1) to allow for the presence of noise in individual home prices. Formally, the log price of property i is expressed as

$$p_{i,t} = p_{m,t} + e_{i,t} + n_{i,t} \quad (4)$$

where $n_{i,t}$ is a normal i.i.d. noise factor with finite variance σ_n^2 which captures imperfections in the housing market. The variance σ_n^2 is constant across properties because it is determined by market-wide conditions. The three components are assumed to be uncorrelated amongst each other at all leads and lags. Iterating equation (4) over the transaction period $[t_i; t_i + T_i]$, we obtain that $\sigma_{u_i}^2$ now equals the sum of a fixed component, $2\sigma_n^2$, plus a component which is linearly increasing in the length of the holding period, $T_i\sigma_\epsilon^2$. The weighted repeat sales (WRS) method of Case and Shiller (1987) adapts the GLS estimator to account for the presence of this constant term in the variance of the error. Its construction is based on a three-step procedure. In the first step, regression (3) is estimated by OLS and the corresponding residuals \hat{u}_i are stored. The second step consists of a weighted least square regression of these residuals squared on a constant and on the time interval between transactions. The constant term of this regression represents an estimate of $2\sigma_n^2$ whereas the slope is an estimate of σ_ϵ^2 . In the third step, a GLS regression of (3) is run where each observation is weighted by the inverse of the square root of the corresponding fitted value from the second step.

Using this methodology, Case and Shiller (1987) construct real estate price indices for Atlanta, Chicago, Dallas, and San Francisco/Oakland relying on nearly 40,000 pairs of transactions over the 1970-1986 period. Compared to median-based indices, the resultant series do not exhibit marked seasonal patterns and display considerable cross-sectional and time-series fluctuations. Further, the weighting implied by the WRS has a substantial effect on the quarter-to-quarter changes in the index compared to the RSR approach. The improvement is largely attributable to the common component in the error variance, σ_n^2 , being quite substantial, on the order of 6% to 7%. By looking at the ratio between the standard deviation of the estimated index and the average standard error of the estimates, they also show that their WRS index captures quite precisely the level of aggregate prices and its annual differences. Quarterly differences, on the other hand, are quite noisy and poorly estimated.

A variant of the Case and Shiller (1987) methodology has been proposed by Goetzmann and Spiegel (1995). They document that including an intercept term in the matrix X of dummy variables helps reduce biases in the estimation. This fixed “non-temporal” component in housing returns most likely relates to property-specific improvements occurring at the time of a sale, which can be as high as 2% to 3% of the investment. Alternative repeat-sales methodologies include

shrinkage-type estimators and Bayesian approaches (e.g., Kuo (1996) and Goetzmann (1992)). Goetzmann (1992) compares the performance of various RSR estimators using simulation on a cross-section of common stocks during a given year. He finds that there seems to be little, if any, advantage to using anything more sophisticated than the GLS estimator when focusing on monthly data and the number of repeat sales observations is large enough relative to the number of intervals estimated.

An appealing feature of the repeat-sales estimator is the fairly limited amount of variables that are required to construct the index, consisting at the very least of prices changes and dates of individual property transactions. Measures of homes characteristics and quality are not directly used in the estimation but may be needed to identify and exclude properties which have undergone major quality changes—such as renovations, expansions, or re-zoning—between transactions. The estimation procedure is computationally tractable, and standard econometric procedures can be used to construct the relevant statistics.

On the other hand, repeat-sales estimators make use of just a limited number of transactions as the information on homes that transacted once is neglected. Also, homes that are sold repeatedly may not be representative of the population as a whole, thus giving rise to a selection bias problem (Clapp, Giaccotto, and Tirtiroglu (1991), Gatzlaff and Haurin (1998), Quigley (1995), Korteweg and Sorensen (2011)). From a statistical perspective, the estimation may be inaccurate because of the singularity or near-singularity of the matrix X . This will occur when no or very few transactions are available in a given period. The practical solution in such a case is to omit the redundant columns, and to calculate the index over longer time intervals. Single-period returns are then assigned the average return during those periods. The accuracy of the index increases at lower frequencies, but autocorrelation is induced in the higher frequency returns. This issue has no clear solution and tends to be more relevant as we attempt to construct indices in thin markets where the number of properties is large relative to the turnover.

An additional concern is represented by spurious autocorrelation in returns arising from overlapping information. Due to the presence of the house-specific noise component, the estimates of p_m may exhibit serial correlation in first differences even if house prices truly follow a random walk. The sign of this serial correlation is not clear and depends on the timing of the sales of the

homes relied upon, but it tends to be negative over short time intervals.¹¹ Longer (one-year) returns are instead more precisely estimated and generally display positive autocorrelation. The autocorrelation properties of the index returns are clearly of importance when analyzing predictability, an issue we will return to in Section 3.1.¹² Lastly, the β estimates and thus the whole time series of the index may change as new information becomes available and the coefficients in (3) are re-estimated. These revisions may be substantial, on the order of one to two percentage points on an annual basis (Abraham and Schauman (1991)). Moreover, they tend to be insensitive to sample size, with systematic and persistent dynamics (Clapp and Giaccotto (1999), Clapham, Englund, Quigley, and Redfearn (2006)). The revision of the index also renders out-of-sample tests hard to evaluate, as the original series is no longer available.

For the residential real estate market, the most well-known repeat-sales indices are the S&P/Case-Shiller Home Price Indices and the HPI Index constructed by the Federal Housing Finance Agency.¹³ The S&P/Case-Shiller Home Price Indices are based on the repeat-sales methodology as modified by Case and Shiller (1987) (see Standard and Poor's (2008)). The index tracks monthly changes in the value of single-family homes both nationally as well as in 20 individual metropolitan areas. The indices are calculated monthly using a three-month moving average and published with a two month lag. The national index is a quarterly indicator for the nine U.S. Census divisions, and captures approximately 75% of the U.S. residential housing stock by value.¹⁴ For the national index, for most of the MSAs indices, and for the Composite 10 index the data begins in 1986, while all remaining metropolitan indices and the Composite 20 begin in 2000. To account for sample selection, sales that occur within six months of one another are excluded owing to the likelihood that the homes have been renovated.

The methodology behind the repeat-sales indices provided by the Federal Housing Finance Agency (FHFA) is a variant of Case and Shiller (1987). The difference is that the second step also involves a quadratic term in the regression of squared residuals on the time interval between trans-

¹¹Webb (1981a),(1981c), and (1981b) show that under some conditions the autocorrelation in return errors approaches -0.5 as the number of observations goes to infinity.

¹²Another concern is that, as noted by Goetzmann (1992), RSR methods estimate the average cross-sectional log return (geometric average), which is lower than the log of the arithmetic average return by Jensen's inequality. This issue is not alleviated by augmenting the number of observations.

¹³Another popular series is the Conventional Mortgage Home Price Index (CMPHI) jointly created by Freddie Mac and Fannie Mae based on mortgages purchased or securitized. See <http://www.alliemae.org/cmhpi.html>.

¹⁴Time-series data and further information on the index construction can be found at the website <http://www.standardandpoors.com> and <http://www.macromarkets.com/index.shtml>.

actions. The indexes are based on data of conventional conforming mortgage transactions obtained from Freddie Mac and Fannie Mae. The HPI provides a broader geographic coverage with respect to the Case-Shiller index owing to the national operations of the two government sponsored housing enterprises. However, this comes at a cost as mortgage transactions on attached and multi-unit properties, properties financed by government insured loans, and properties financed by mortgages exceeding the conforming loan limits determining eligibility for purchase by Freddie Mac or Fannie Mae (such as sub-prime mortgages) are excluded. Further details about its construction are provided by Calhoun (1996). Monthly indices for the U.S. and Census divisions based on sales price data are available since January 1991. Quarterly indices estimated using both sales prices and appraisal data for the U.S., Census divisions, and metropolitan areas start in the first quarter of 1975.

Competing repeat-sales indexes for the aggregate and local US residential market starting in 1975 are also available from CoreLogic. They are constructed based on a broad universe (about 50 million) of mortgages, and cover about 98 percent of all U.S. Zip codes. Unfortunately, the data are not publicly available. For a recent application using the CoreLogic series, see Favilukis, Kohn, Ludvigson, and Van Nieuwerburgh (2012). Recently, repeat-sale indices have also been introduced into commercial real estate markets. Prominent among these are the Moodys/REAL commercial property price index (CPPI) and the CoStar commercial repeat sales index (CCRSI).

It is important to emphasize that these indices only track price appreciation. That is, they only account for changes in prices and ignore any intermediate cash flow over the period the home is being held. These cash flows include explicit or implicit rent, tax effects, and maintenance costs. Clearly, true measures of (excess) returns to real estate must reflect all inflows and outflows arising from the trading and management of a dwelling. An average implied rent and a rent-to-price ratio series for the Case-Shiller and FHFA indices has been constructed by the Lincoln Institute based on the methodology of Davis, Lehnert, and Martin (2008).¹⁵

¹⁵The data are downloadable at <http://www.lincolnst.edu/resources/>.

2.1.3 Hedonic Indices

Repeat sales models measure movements in property prices in a given location over time. They do not shed light on what specific factors determine these prices at a given point in time. To answer this question, we turn our attention to hedonic pricing models. In hedonic models, the price of a property is expressed as a function of a set of characteristics which determine its quality (such as square footage, number of bedrooms, etc.) and other factors (such as proximity to a school).¹⁶ This relation may arise as the equilibrium outcome of a competitive market with heterogeneous goods whose characteristics enter the agent's utility function (Rosen (1974)).

An important classification of hedonic models pertains to the functional form relating the property price and its characteristics. The standard semi-log formulation assumes a linear specification of the type

$$p_{i,t} = \beta Z_i + \delta D + \epsilon_{i,t} \quad (5)$$

where $p_{i,t}$ is the log transaction price of property i in period t , Z_i is a $C \times 1$ vector of property attributes (also known as hedonic variables) including a constant term, and D is a $T - 1 \times 1$ vector of time dummies, one for each period except the first. The OLS estimates of (β, δ) in (5) are obtained by pooling the information of all transactions and have an immediate interpretation. The estimates of β measure the marginal utility an investor derives from having one additional unit of a characteristic, also known as a shadow price. The parameter estimates of δ capture, by contrast, the period-specific change in log price once the effect of property characteristics has been accounted for. Similar to repeat-sales models, the vector $\hat{\delta}$ is then regarded as an estimate of the log price of the quality-adjusted aggregate index, or $\hat{p}_{m,t} = \hat{\delta}_t$. The value of the log price in the initial period is again normalized to zero.

Semi-log hedonic models are often preferred for their ease of estimation. Standard errors and statistical tests are easily computed. Linear models are, however, clearly prone to model misspecification. Several researchers have explored alternative specifications which allow for greater flexibility through nonparametric functional forms or second-order expansions (see Halvorsen and Pollakowski (1981), Wallace (1996), and Clapp (2004)). These models have been found to provide superior out-of-sample predictive performance compared to linear ones (Pace (1993)), but their

¹⁶For an extensive discussion of hedonic models, see Hill (2011). An early application of hedonic models to commercial real estate is Hoag (1980) who investigates the risk and return characteristics of industrial real estate.

estimation requires the availability of large datasets and shadow prices are not easily obtainable.

Another key element in the implementation of hedonic models is the choice of the appropriate set of characteristics. The most commonly used characteristics are lot size, square footage, number of bedrooms, number of bathrooms, and age (Wallace (1996), Sirmans, MacDonald, Macpherson, and Zietz (2006)). Others include garage space and the presence of air conditioning, a swimming pool, and a fireplace. In general, this list is dictated by data availability. Data on properties characteristics, along with owners' own assessment of their values, can be found in the Panel Study of Income Dynamics (PSID) which starts in 1968. Moreover, several variables that may affect pricing such as the amount of traffic noise and sunlight exposure are not directly measurable or observable. This renders hedonic models prone to both omitted variable and selection bias, as missing observations for some characteristics may lead to data censoring toward, for example, high quality buildings. As in any regression-type approach, the maintained assumption for consistency of the estimates is that the included variables are not correlated with omitted determinants.¹⁷ A perhaps comforting result is the evidence that shadow prices for the same characteristics resulting from the estimation of the semi-log model on different databases appear to be rather stable (Sirmans, MacDonald, Macpherson, and Zietz (2006)). In addition, location identifiers (such as zip codes, location dummies for proximity to the ocean or nearby lakes) are usually included in the regression in order to account for unobserved heterogeneity, as in Campbell, Giglio, and Pathak (2011).

Repeat sales models can be viewed a special case of hedonic pricing models. To see this, consider a property i which sold twice, say at times t_1 and t_2 , and apply the hedonic pricing model, expression (5), at each of these time points assuming the shadow prices of the property's attributes as well as the sizes of the attributes themselves do not change between sales:

$$\begin{aligned} p_{i,t_1} &= \beta Z_i + \delta_{t_1} D_{t_1} + \epsilon_{i,t_1} \\ p_{i,t_2} &= \beta Z_i + \delta_{t_2} D_{t_2} + \epsilon_{i,t_2}. \end{aligned}$$

¹⁷Shiller (2008) argues that hedonic models are subject to the risk that researchers "cherry pick" the functional form and characteristics to obtain the desired results. This argument, however, applies broadly to all empirical studies.

Subtracting these expressions gives

$$\begin{aligned} p_{i,t_2} - p_{i,t_1} &= \delta_{t_2} D_{t_2} - \delta_{t_1} D_{t_1} + \epsilon_{i,t_2} - \epsilon_{i,t_1} \\ &= \delta_{t_2} D_{t_2} - \delta_{t_1} D_{t_1} + \nu_i. \end{aligned}$$

If we assume a sale occurred at t_2 and assign $D_{t_2} = -1$ and a purchase at t_1 and set $D_{t_1} = +1$, the latter expression above corresponds to the repeat sales model given in expression (3).

Several studies compare the relative performance of hedonic and repeat-sales models. Meese and Wallace (1997) exploit the fact that repeat-sales estimators can be viewed as constrained versions of a dynamic hedonic model in which it is assumed that (i) homes that sold twice are representative of the whole market and (ii) the shadow prices of the attributes are constant over time and therefore cancel out in the construction of the index. They reject both of these assumptions using data on transactions prices and characteristics for 50,000 homes located in the cities of Oakland and Fremont, California. In addition, they find that repeat-sales indices tend to be very volatile. They attribute this behavior to sample selection bias, non-constancy of the characteristics' shadow prices, and sensitivity to small-sample problems which all make these approaches less suitable to study efficiency in local markets. Surprisingly, they find that the time series properties of the readily-available median sales price index were very close to those of a hedonic Fisher Ideal index. Similar conclusions are reached by Clapp and Giaccotto (2002) looking at the empirical distribution of prediction errors using data for Dade County, Florida.

As in the case for repeat-sales models, hedonic models provide estimates of the average log return, and not of the log return of the average property. They rely on the assumption that the set of houses that transact is representative of the market as a whole. If instead the sample of house sold varies with economic conditions, the resultant indices may be systematically biased. The magnitude of this bias can be analyzed by comparing hedonic models with price indices based on censored regression procedures, as in Gatzlaff and Haurin ((1997), (1998)).

An alternative way of constructing hedonic-based indices which does not require data on properties characteristics is to take advantage of appraisal valuations. Appraisals are estimates of the current value of a property provided either by the owner (so called "internal appraisals") or by a professional agent ("external appraisals"). The key insight here is that while an appraisal value may

represent a noisy estimate of a property's true market value, it serves as a valuable hedonic variable, summarizing a building's characteristics which are either observable, such as its size, or are unobservable, such as its quality. For the U.S. commercial real estate market, a popular index which is based primarily on appraisal values is the National Property Index (NPI) constructed by the National Council of Real Estate Investment Fiduciaries, NCREIF.¹⁸ NCREIF assets are institutional-grade commercial properties managed by investment fiduciaries on behalf of tax-exempt investors, mostly pension funds. The commercial properties are acquired in the private market for investment purposes only.¹⁹ Based on the information provided by its members, NCREIF constructs quarterly indices for the aggregate commercial real estate as well as indices disaggregated by property type and region. The indices are value-weighted by each property market value, and include cash flows from net operating income and capital expenditures. The series for the U.S., industrial properties, retail properties, and offices start in the first quarter of 1978, while the index for apartments is available from 1984.

A well-known drawback of using appraisal valuations is that the resultant returns respond with a lag to changes in actual market values and are much smoother (Fisher (2005)). The Transaction-Based Index (TBI) constructed by the MIT Center for Real Estate uses the information on transaction prices of properties sold from the NCREIF database to provide a more timely measure of market movements.²⁰ The index is based on the two-stage methodology of Fisher, Geltner, and Pollakowski (2007) which combines the information of infrequent transaction prices with that of frequent appraisal valuations. In the first stage, quarterly transaction data are used to estimate a hedonic price model in which corresponding transaction prices are regressed against properties' lagged appraisal values as well as several dummy variables controlling for time, property type, and location. The estimated coefficients from this regression are then used in a second stage to construct predicted prices based on the appraisal values and other characteristics of those properties that did not transact in a given quarter.²¹ In order to construct the aggregate TBI Index, the first-stage estimates are then applied to a representative property mirroring the average char-

¹⁸See <http://www.ncreif.org/data.aspx>.

¹⁹When a property is sold or is subject to a change of use, it exits the database. Due to changes in its composition and the type of assets included, the NPI index may therefore not be representative of the commercial real estate market as a whole.

²⁰See <http://web.mit.edu/cre/research/credl/tbi.html>.

²¹The methodology also accounts for transaction sample selection bias in the first stage using a Heckman (1979) two-step approach and applies Bayesian noise filtering technique to reduce the effect of noise in the quarterly series due to the limited number of transactions. See Fisher, Geltner, and Pollakowski (2007) for further details on this estimation procedure.

acteristics of the data. The methodology can also be used to construct pseudo–market prices for individual properties, as in Plazzi, Torous, and Valkanov (2011). The nation-wide index is available quarterly from 1984:Q1, while property-specific indices for apartments, industrial properties, retail properties, and offices start in 1994:Q1.²²

2.1.4 Hybrids

A combination of two or more types of indices might attenuate the deficiencies in the individual approaches. Along those lines, Case, Pollakowski, and Wachter (1991), Case and Quigley (1991), Quigley (1995), and Meese and Wallace (1997) combine repeat-sales and nonparametric hedonic methodologies in the construction of hybrid indices. The specification of Case and Quigley’s (1991) hybrid model is appealingly simple. It involves estimating the repeat-sales model (3) and the hedonic model (5) jointly in a GMM system of equations.

Hill, Sirmans, and Knight (1997) improve upon the estimation of Case and Quigley (1991) hybrid model. More specifically, they use hedonic regressions to estimate the effect of depreciation (the shadow price of a building’s age) and impose a first-order autoregressive process for the error term to capture sluggish adjustments to economic shocks. They then jointly estimate a repeat-sales regression consistent with this error structure via maximum-likelihood and document substantial efficiency gains in terms of lower standard errors and narrower interval estimates for the resultant index. Hybrid indices seem to offer improvements over either the repeat-sales or the purely hedonic models (Case, Pollakowski, and Wachter (1991) and Meese and Wallace (1997)), which illustrates well the fact that the adoption of one particular model to the exclusion of all others is likely to result in the suboptimal use of information. The intuition for this results is analogous to that in the forecast combinations literature (Timmermann (2006)).

2.1.5 Stock Market-Based Indices

Institutions and individuals can take positions in the commercial real estate market by investing in publicly-traded REIT companies. Market-based indices can be obtained from the trading of individual REIT stocks. These indices are usually constructed as value-weighted averages of firm-

²²The starting date for the property-specific indices is motivated by the need of a sufficient number of transactions to estimate the model parameters separately within each property.

specific REIT returns. The two standard data sources here are the CRSP/Ziman Real Estate Data Series and the FTSE NAREIT US Real Estate Index Series. Both indices track the performance of the US market and provide disaggregated information across REIT types (equity, mortgage, and hybrids). The CRSP Index is available from 1979, while the NAREIT data starts in 1972. The CRSP Index also provides separate indices for Apartments, Industrial and Offices, and Retail.

A few caveats are, however, in order when using REIT data. First, the overall value of the 163 REITs traded at the end of 2010 was about \$366 billion, and thus represent quite a small fraction of the approximately \$10 trillion estimated value of non-residential real estate market. Hence, REITs may not constitute a representative sample of the U.S. commercial real estate market as a whole. Second, the number of traded REITs varies considerably over time, from about 100 trusts in the early 1980s to slightly less than 200 during the mid 2000s. Third, the market is characterized by a few large companies and many smaller REITs. This description is consistent with the fact that in 2009 the average market cap of REITs was \$1.37 billion, while the median market cap was only \$0.618 billion. An investment in REITs exposes investors to the risks inherent in small-cap stocks. Finally, many REIT companies have a significant amount of debt and thus the return series reflect the profit for the equity stake of investing in real estate. Since equity is nothing but a call option on the value of assets (Merton (1974)), higher debt levels amplify the effect of shocks to the value of the asset (property). Hence, we expect REITs returns to exhibit higher mean and volatility compared to those of the commercial real estate market.

2.1.6 Other Methods

The non-observability of the true underlying price process has also prompted some researchers to apply filtering techniques to extract the information of true prices embedded in noisy transaction prices. Engle, Lilien, and Watson (1985) use an EM algorithm (based upon Kalman filtering and smoothing) to estimate unobservable rent-price ratios, by relying on hedonic prices and a present value model between prices and future rents. For forecasting real house prices in the U.K., a Kalman filter model with time-varying coefficients has been used by Brown, Song, and McGillivray (1997). Giaccotto and Clapp (1992) use Monte Carlo simulation to show that Bayesian-type techniques based on a Kalman filter should be preferred by appraisers to estimate current true prices.

Finally, the increasing availability of large databases has lead researchers to explore the use of spatial econometric techniques that control for geographical and temporal dependence in real estate prices.²³ Spatial dependence refers to the fact that properties which are in geographical proximity to each other will tend to be subject to similar shocks. Moreover, the location of a dwelling may play an important role in its pricing due to the presence of factors such as proximity to schools, parks, and malls. We would then expect the error components in hedonic regressions to be more correlated the closer the two properties are to each other. In contrast, temporal dependence refers to the fact that the parameters of hedonic models (the attribute prices) may change over time. This is similar to modeling time-varying coefficients in standard regression analysis. Error terms which refer to transactions occurring in distant periods are then likely to be less correlated. Explicitly accounting for these two sources of correlation helps reduce the bias and improves the efficiency of the estimators. Application of spatial models to housing can be found in Can (1992), Pace, Barry, Clapp, and Rodriguez (1998), Pace, Barry, Gilley, and Sirmans (2000), Caplin, Chopra, Leahy, LeCun, and Thampy (2008), Nappi-Choulet Pr. and Maury (2009).

2.2 Summary Statistics

Residential: For residential properties, we have three data sources—one median price index and two repeat-sales indices. First, the Census Bureau provides median and average prices of US residential properties at monthly frequency. Second, Standard and Poor’s/Case-Shiller (Case-Shiller) construct repeat-sales prices that are available at various frequencies and aggregation levels. A national index is constructed quarterly and a monthly version for the 10 largest metropolitan areas (C10) is also available. In addition, we compute a monthly equally-weighted average across all available areas at a point in time and label it as “EW”. Third, the Federal Housing Finance Agency’s (FHFA) repeat-sales House Price Index is available monthly (Purchase Only) and quarterly (All Transactions) at the national level. All three series track price appreciation in nominal, as opposed to real, house prices. We compute log price changes of all indices, which cannot be interpreted as returns in the usual sense of the term since the levels do not account for rents. Rent-to-price data for the Case-Shiller and FHFA data are from the Lincoln Institute at quarterly frequency. This data is also used to obtain growth in rents. Unfortunately, we do not have access to residential hedonic

²³See Anselin (1988) for a book treatment of spatial regressions.

price data. To facilitate the comparison across data sources, we sample all series starting in 1991 (or later, depending on the series availability). The exception is REITs, for which we use a sample starting from January 1980.

In Panel A of Table 2, we report summary statistics – annualized means, annualized standard deviations, skewness, and AR(1) coefficient– for log price changes of all indices as well as the rent growth rate and the rent-to-price ratios. The average growth rate of the Census median sales index and of the repeat-sales series are all in the neighborhood of 3%. However, the standard deviation of the median sales index, at 8.8%, is significantly higher than the 2.0% to 5.1% observed for the Case-Shiller and FHFA indices.²⁴ All indices are negatively skewed, but the skewness in repeat-sales indices is larger in absolute value.

From a time-series perspective, the most significant difference between the median and the repeat-sales indices is the level of time-dependence. Changes in log levels of the Census series are significantly negatively correlated, with an AR(1) coefficient of -0.522. The Case-Shiller C10 and EW monthly log changes exhibit an AR(1) coefficient of 0.938 and 0.929, respectively. To a significant extent, this dependence reflects the fact that, as explained above, the Case-Shiller indices are constructed as a three-month moving average of an underlying series. Indeed, if we take the quarterly index, the AR(1) coefficient is significantly lower (0.613). The FHFA monthly and quarterly log changes also have large positive AR(1) coefficients of 0.756 and 0.708, respectively.

The average rent growth rate for the repeat-sale indices is 3%. The series exhibit very little volatility and high serial correlation (0.900). The rent-to-price ratio is 4.5% for both indices. Its standard deviation is very small compared to that of the log price changes, but it is extremely persistent, with AR(1) coefficients of 0.987 (Case-Shiller) and 0.994 (FHFA). Given that these estimates are downward-biased (Andrews (1993)), there is little doubt that these series are close to non-stationary. The high level of persistence in the rent-to-price ratio is similar to that observed in valuation ratios (dividend-price, earnings-price) of the US stock market (Welch and Goyal (2008)). The average return from residential real estate can be computed by adding the average price appreciation and the average rent-to-price ratio. Over our sample, it is 7.6% for both repeat-sales indices.

Commercial: Commercial properties naturally fall into one of four categories: apartments (Apt),

²⁴It is worth noticing that over the 1963-2010 period for which Census data are available, the return standard deviation is even higher at 13%.

industrial properties (Ind), offices (Off), and retail properties (Rtl). Indices are available for each of these categories as well as for the overall commercial real estate market. The first source of commercial real estate values is the NPI from the National Council of Real Estate Investment Fiduciaries (NCREIF). An alternative hedonic index, based on the work by Fisher, Geltner, and Pollakowski (2007), is the transactions-based index (TBI).

Repeat-sales commercial real estate indices are relatively new. One such index, the Moodys/REAL commercial property price index (CPPI) provides monthly data at the aggregate level and quarterly series by property type from 2001 to the present.²⁵ While this time span is too short for forecasting exercises, we include the CPPI for completeness and provide summary statistics. We have no doubt that repeat-sales indices will play a growing role in commercial real estate.

REIT is a value-weighted index of all publicly-traded REITs in the CRSP-Ziman database. Using monthly returns with and without dividends, we construct its dividend-price ratio and dividend growth rate (see Appendix A.2 for details). Since REIT is a stock-market based index, it presents a unique opportunity to investigate the performance and predictability of commercial real estate returns without the complications inherent in hedonic and repeat-sales indices. It is therefore not surprising that many academic papers have been written on REIT return predictability and so we will also devote special attention to this market.

Panel B of Table 2 contains summary statistics for the NPI, TBI, CPPI, and REIT indices. Whenever available, we also report summary statistics for rent growth rates (or dividend growth, in the case of REITs) and rent-to-price ratios (dividend-price ratios, for REITs). With the exception of CPPI, all return series include cash flows distributions (net rents or dividends).

The TBI index shows a higher average return than the TBI (9.0% versus 6.8%) and a higher standard deviation (9.6% versus 5.2%). Since the two indexes are based on the same data, the increased volatility of the TBI is attributable to the reduced impact of the smoothness coming from appraisal valuations. Indeed, the AR(1) coefficient of the NPI is 0.80, about one order of magnitude larger than that for the TBI (0.087). The low serial correlation in the return series makes the TBI particularly appealing from an economic perspective. The rent growth rate of the

²⁵The CPPI tracks same-property realized round-trip price changes based on transactions data provided by Real Estate Analytics, LLC (REAL). The RCA database aims at collecting price information for every commercial property transaction in the U.S. over \$2,500,000 in value. Thus, it reflects a more extensive set of properties than those in the NCREIF portfolio. See Geltner and Pollakowski (2007) and <http://mit.edu/cre/research/credl/rca.html>.

TBI exhibits moderate serial dependence (AR(1) coefficient of 0.365) whereas its rent-to-price ratio is extremely persistent (AR(1) coefficient of 0.971). The changes in the log CPPI have a very low average mean of 1.1%. This is mostly due to the sample over which these statistics were calculated.²⁶ The CPPI series are also quite volatile and exhibit moderate dependence (AR(1) of 0.452).

The average REIT returns, at the bottom of Panel B, Table 2, have a mean of about 10.4% and a standard deviation of 17.9%. These numbers are higher than for the other commercial real estate indices. By comparison, during the common 1991-2010 period which is used in our forecasting regressions, the return mean and volatility for the All-property index are respectively 0.105 and 0.202. Since REIT is a stock index, it is useful to compare its return and volatility with that of the market-wide portfolio return. The CRSP value-weighted portfolio return has an average of 10.7% and a standard deviation of 16.3% over a similar period, which implies a higher Sharpe ratio than that of REITs. REIT returns have a relatively higher serial correlation of 0.146, which is in line with that of small-cap stocks in the US stock market.

Conditioning Variables: The conditioning variables we select proxy for time variation in the state of the economy and thus in the prevailing investment opportunity set. These variables have also been shown to successfully capture time variation in expected returns of the aggregate U.S. stock market and bond returns. These include the lagged aggregate stock market (Campbell and Vuolteenaho (2004)), its dividend-price ratio (Fama and French (1988b), Lettau and Van Nieuwerburgh (2008)), the relative 3-month Treasury bill calculated as the current rate minus its twelve-month moving average (Hodrick (1992)), the inflation rate (Fama and Schwert (1977)), the term spread as difference between the 5-year and 3-month log yields (Fama and French (1989)), the Cochrane and Piazzesi (2005) tent-shaped combination of forward rates, and industrial production growth (Fama (1990)). Summary statistics, reported in the bottom panel of Table 2, show a wide range of persistence with AR(1) coefficients ranging from 0.29 for Industrial Production growth to as high as 0.92 for the Term Spread. Details on the data source and construction are provided in Appendix A.2.

²⁶Unfortunately, a longer time span is not available.

3 Forecasting Real Estate Returns

The extensive predictability literature in finance and real estate considers variations of the following linear predictive regression:

$$r_{t+1} = \alpha + \beta' X_t + \epsilon_{t+1} \quad (6)$$

where r_{t+1} is a return (or price change) and X_t is a vector of variables, observable at time t . Predictability in r_{t+1} may arise because of two distinct economic reasons. First, it might be due to market inefficiency if some available information is not incorporated in prices in a timely manner by market participants (e.g., Fama (1970)). Second, predictability might be due to time-variation in expected returns (e.g., Campbell and Shiller (1988)). Unfortunately, the existence of predictability in a reduced-form regression (6) does not allow us to trace its economic provenance. Also, the existence of predictability does not necessarily imply that the market is inefficient in the usual sense of the term (Fama (1970)). For a market to be inefficient, investors should be able to exploit some of the serial dependence. This point is discussed in detail by Case and Shiller (1989) in the context of residential real estate.

Linear models are deceptively simple. An extensive literature has investigated their statistical properties (estimation and inference) and out-of-sample predictive performance (Rapach and Zhou (2012) in this Handbook). Statistical complications arise because the predictor X_t is often persistent and its innovations are correlated with ϵ_{t+1} , which induces bias in the estimation of β (Stambaugh (1999)). Moreover, excessive noise in the returns series renders hypothesis testing unreliable. We will revisit some of these issues below.

In this section, we review the literature on real estate predictability. We also report estimates from our own predictive regressions using the indices introduced above. The discussion is organized around the kind of predictive information that is included in X_t and the implied hypotheses.

3.1 Serial Dependence in Real Estate Returns and Weak-Form Market Efficiency

We start off with the simplest information set X_t , that of past returns, r_t . In this case, regression (6) tests for serial correlation in returns and weak-form market efficiency. Several studies

in the real estate literature find that returns (or price changes) exhibit positive serial correlation, including Gau (1984), Gau (1985), Linneman (1986), Guntermann and Smith (1987), Rayburn, Devaney, and Evans (1987), Case and Shiller (1989), McIntosh and Henderson (1989), Gyourko and Voith (1992), Kuo (1996), Hill, Sirmans, and Knight (1997), Hill, Sirmans, and Knight (1999), Gu (2002), and Schindler (2011). However, the evidence on whether this predictability can be exploited for financial gains is less clear.

In one of the earliest papers of weak-form efficiency for the US real estate market, Gau (1985) investigates the persistence in monthly returns to commercial real estate in Vancouver during the 1971-1980 period. Rather than using simple returns, he works with abnormal returns, defined as returns adjusted for various sources of systematic risk. The cross-sectional risk-adjustments alter the unconditional mean of the returns series, but have little effect on their dynamics. Gau (1985) finds that the forecasting errors from predicting abnormal returns using past price information are too small to be exploitable by a trading strategy. Linneman (1986) uses hedonic prices, also adjusted for risk, to test for market efficiency in the Philadelphia residential market. He finds evidence of serial dependence in the data, but concludes that the predictability is insufficient “to cover the high transaction costs associated with transacting real estate.” Guntermann and Smith (1987) apply a portfolio approach to uncover the autocorrelation of aggregate unanticipated total returns to residential real estate in 57 MSAs using the Federal Housing Administration data over the 1968-1982 sample. Their study is one of the first to explicitly take into account rental income in computing returns. They document positive predictability over horizons of one to three years, and negative autocorrelation at the four- to ten-year horizon. This pattern is consistent with short-run momentum and long-run reversal. Consistent with Linneman (1986), the persistence is not large enough for various trading rules to appear profitable once transaction costs are considered. Rayburn, Devaney, and Evans (1987) and McIntosh and Henderson (1989) reach similar conclusions using different datasets and methodologies.

In an influential study, Case and Shiller (1989) test for weak-form efficiency in four U.S. single-family markets: Atlanta, Chicago, Dallas, and San Francisco/Oakland. They do so using their weighted repeat sales (WRS) index (Case and Shiller (1987)). To reduce errors-in-variables issues, they randomly partition their sample of transactions into two groups and obtain two corresponding WRS indices for each of the four cities. They then regress quarterly observations on the annual

return in one index on the one-year lagged annual return of the other index. This approach, which can be seen as instrumental variables (IV), produces consistent, albeit biased, estimates of the autoregressive coefficient.²⁷ Case and Shiller (1989) document substantial predictability in real and excess returns to housing, with predictive R^2 ranging from 0.11 to as high as 0.48 corresponding to average trading profits between 1% and 3%. They also find it much harder to forecast individual properties returns using the city-wide index, due to the large amount of noise-to-signal ratio in such data. Out-of-sample performance deteriorates considerably due to measurement error in estimating the aggregate index using only a subset of the sample. Moreover, the random partitioning approach implies that the estimates and forecasts will change if we alter the partitioning of the data.

Hill, Sirmans, and Knight (1999) and Schindler (2011) provide some additional evidence on the Case and Shiller (1989) data. Hill, Sirmans, and Knight (1999) reject the hypothesis of a random walk in these price series using the methodology in Hill, Sirmans, and Knight (1997). The test is based on the idea that a random walk process for prices would induce heteroskedasticity in repeat-sales indices. They also show that the GLS procedure of Case and Shiller (1989) can be improved upon by accounting for a stationary, autoregressive component in house prices. Schindler (2011) provides recent evidence of predictability in the Case-Shiller real and nominal log nominal price changes, computed for the national and 20 metropolitan areas indices. Not surprisingly, he finds strong evidence for dependence in the price changes, with some indices exhibiting strong positive autocorrelation even at 24 monthly lags. Perhaps more surprising is his finding that, after comparing different buy-and-hold and dynamic trading strategies, the author concludes that in some markets the persistence in the data is large enough to be exploitable. It is worth mentioning that the markets with the largest gains from the trading strategies – Los Angeles, Las Vegas San Diego, and San Francisco – are also the markets that have exhibited the largest bubbles. Hence, it is not clear whether these strategies would fare equally well out-of-sample.

Gu (2002) studies the autocorrelation properties of quarterly returns to the Conventional Mortgage Home Price Index (CMHPI) for all U.S. states, the District of Columbia, nine Census Divisions, and an aggregate index for the US during the 1975–1999 period. He finds that the degree of persistence and the sign of correlation varies geographically as well as over time. The findings

²⁷To reduce the bias arising from noisy instruments, Kuo (1996) proposes an alternative Bayesian approach to estimate an AR(2) model based on repeat-sales. His setup explicitly models the unknown true indices as random variables and hence does not necessitate partitioning the repeat-sales sample. The corresponding posterior means of the AR(2) coefficients suggest that repeat-sales indices are more dependent than what found in previous studies.

in Gu (2002) also point to the difficulty of comparing the results of weak-form efficiency studies in real estate using information from different markets, datasets (some aggregated, others not), sample periods, and methodologies.

A growing literature uses regime switching models to capture real estate price dynamics.²⁸ In a regime switching model, the time-series properties of a series depend on the realization of an underlying state variable.²⁹ Moreover, it is reasonable to expect substantial variation in regimes across different areas and property types owing to the reliance of real estate to local economic and geographic conditions. Shifts in house price dynamics may arise, for example, because of a changing relationship between housing, income, and interest rates (see e.g. Muellbauer and Murphy (1997), Boz and Mendoza (2010), and Favilukis, Ludvigson, and Van Nieuwerburgh (2011)) or from the interaction between credit-constrained households, lenders, and developers (Spiegel (2001)).³⁰

Theoretically, the entire conditional density of a process may depend on the current state realization. In practice, for tractability, the empirical applications have mainly adopted first-order autoregressive regime switching models:

$$r_t = c_{s_t} + \phi_{s_t} r_{t-1} + \epsilon_t \quad (7)$$

where $\epsilon_t \sim N(0, \sigma_{s_t}^2)$. In this context, returns follow an AR(1) process whose intercept, slope, and error variance depend on the realization of the regime variable s_t . Thus, the data-generating process is subject to jumps in the unconditional mean, persistence, and level of volatility. The choice of the number of regimes n that better describes the data is usually chosen by likelihood ratio tests or information criteria. However, due to the availability of relatively short time-series, most of the existing studies rely on low-order models with two or three regimes (Crawford and Fratantoni (2003)). The fact that prices of direct investments in real estate are available at low

²⁸Following the influential work of Hamilton (1989), regime switching models have been extensively used in the macroeconomic and finance literature to capture nonlinearities in exchange rates (Engel and Hamilton (1990)), interest rates (Gray (1996), Garcia and Perron (1996), Bansal and Zhou (2002)), stock returns (Perez-Quiros and Timmermann (2001), Granger and Hyung (2004), Guidolin and Timmermann (2008)), GDP growth (Diebold and Rudebusch (1999)), and GNP growth (Hamilton (1989)).

²⁹Using Monte Carlo simulations, Van Norden and Vigfusson (1998) provide evidence that regime-switching tests for bubbles suffer from a downward bias distortion even with relatively large samples – i.e., they reject the null of no bubble too often – but display considerable power in detecting nonstationarities. This makes them suitable to capture the persistent real estate cycles.

³⁰Similarly, Guirguis, Giannikos, and Anderson (2005) find evidence of parameters instability in the relationship between real estate prices and the fraction of the population aged between 25 and 35, real disposable income, stock of owner occupied dwellings, the expected nominal capital gains, the nominal post-tax mortgage interest rate. They interpret this fact as evidence of structural changes and suggest the use of time-varying coefficient methods coupled with Kalman filter.

frequency – monthly, at the very best – also complicates the detection of short-lived regimes. The estimates are obtained using maximum likelihood (Hamilton (1994)).

Among the empirical studies, Crawford and Fratantoni (2003) compare the in-sample and out-of-sample forecasting performance of regime switching models to that of ARIMA processes and GARCH models on the annualized growth rates in the state-level OFHEO quarterly return series from 1979 until 2001. Their approach parallels that of Perez-Quiros and Timmermann (2001) for the stock market. The authors document considerable heterogeneity in the time-series properties of residential returns across states. For example, past returns explain only 6% of the variance of returns to Ohio but about 75% in the case of California. Their data exhibits a strong degree of persistence and nonlinearity in the volatility process which is modeled as an EGARCH. They also observe that the time-series patterns seem to be better captured by a two-states regime model, which deliver much lower in-sample RMSEs and R^2 . Out-of-sample, however, ARMA models display better forecasting properties (lower MSFE), probably due to the tendency of regime switching models to overfit in-sample.³¹ Interestingly, this last finding is consistent with Gu (2002)'s analysis on CMHPI data.

Another stream of papers makes a connection between a high degree of positive serial correlation in returns and bubbles in real estate markets. For example, Gyourko and Voith (1992) analyze autocorrelation in median house prices for 56 MSAs during 1971-1989. They find significant differences in autocorrelations across areas, which they interpret as evidence of market inefficiency. At the same time, they argue that a global component drives residential prices and that “the national economy strongly influences local housing markets.” Glaeser, Gyourko, and Saiz (2008) use the FHFA/OFHEO data to show that large price increases in residential properties were almost entirely experienced in cities where housing supply is more inelastic. Hence, they argue that boom-bust housing cycles are largely driven by housing supply rather than demand shocks.

The stylized facts that emerge from this literature can therefore be summarized as follows: (i) most residential indices exhibit serial correlation in log price changes; (ii) the serial correlation is positive at horizons up to a few years; (iii) at longer horizons, we observe a reversal, or a negative serial correlation in returns; (iv) the economic significance of this serial correlation and whether it

³¹Perez-Quiros and Timmermann (2001) document that even when the data are generated by a regime switching process, simple autoregressive models may provide better short-term forecasts.

can be exploited by market participants is still an open question.

To illustrate these findings, we run simple autoregressive tests, which are mostly descriptive in nature. The goal is to replicate some of the key dynamics outlined above for returns in the aggregate residential and commercial real estate markets. As we don't have access to the disaggregated data used in the construction of the indices, we cannot replicate the more involved estimation specifications. We consider the following long-horizon regression:

$$r_{t+1:t+T} = \alpha(T) + \beta(T) r_{t-T+1:t} + \varepsilon_{t+1:t+T} \quad (8)$$

where $r_{t+1:t+T}$ is the log change of a price index over T periods. For $T = 1$, equation (8) collapses to an AR(1) model. Long-horizon return regressions (8) are used frequently in empirical finance to investigate the behavior of equity returns at various horizons (e.g., Fama and French (1988a)). Versions have also been used in the residential real estate literature by Guntermann and Smith (1987) and in the commercial real estate literature by Plazzi, Torous, and Valkanov (2010). Under the null hypothesis that the price process has no predictable component, $\beta(T)$ should be zero at all horizons. Deviations from the null hypothesis imply that there are predictable dynamics at different horizons.

In Figure 3, we provide the results from this regression estimated for four indices, two residential and two commercial. We choose indices that exhibit various degrees of serial correlation. For residential, we use the median and FHFA/OFHEO indices³² whereas for commercial, the comprehensive TBI and REIT indices. For all regressions, we plot the least squares estimates of $\beta(T)$, $T = 1, \dots, 48$ months with a solid line along with 2 times the Newey and West (1987) standard errors (dotted line), based on T lags. The Census median index exhibits negative serial correlation at short horizons. As the horizon increases, the correlation turns positive between 12 and 24 months and then reverts to zero. With the exception of the short-horizon negative correlation, we cannot reject the null that the Census median forecast changes are uncorrelated. The picture is quite different for the repeat-sales FHFA/OFHEO data. We observe a strong degree of serial correlation at short and medium horizons. The $\beta(T)$ estimates are positive and significant for horizons up to 30 months. Then, the serial correlation turns negative but insignificant.

³²The Case-Shiller is even more serially correlated than the FHFA/OFHEO index. This is due to its construction, which takes a three-month average of an underlying repeat-sales index.

The TBI index is not serially correlated at the very short horizon. The estimates of $\beta(T)$ turn significantly positive for horizons between 6 and 18 months. Then, similarly to the FHFA/OFHEO index, we observe a reversal and a negative, albeit not significant, correlation in long-horizon returns. Finally, the REIT index exhibits a positive but insignificant correlation at the one-month horizon, which is typical for small-cap stocks (Campbell, Lo, and MacKinlay (1997)). As the horizon increases, we notice a negative drift in the estimates which turns significant after about 30 months. This is consistent with the model in Fama and French (1988a) who show that in the presence of a small mean reverting component in returns, the estimate of $\beta(T)$ will be negative at long horizons. This evidence supports the presence of a small predictable component in REIT prices.

The results in Figure 3 show the main salient findings in the real estate literature, namely positive serial correlation in the price changes at short horizons of up to 2-3 years and reversals at horizons beyond 3 years. In the case of the FHFA/OFHEO and TBI indices, the positive serial correlation is due partly to the construction of the index, as discussed in section 2.1, and partly to market inefficiencies in real estate markets. The construction of an index that captures a quality-adjusted price without introducing artificial dynamics remains an important topic of research. The current indices, especially those that are filtered, are appropriate for capturing the state of the real estate market. However, their excessive serial correlation does not make them suitable for forecasting exercises. Moreover, the majority of the studies cited above reach the conclusion that high transaction costs render the observed predictability hard to exploit by market participants.

3.2 Predictability Based on Valuation Ratios

Valuation ratios, such as the dividend-price, the book-to-market, and the earnings-price, have a long-standing tradition as predictors of equity returns (see Rapach and Zhou (2012) in this Handbook and references therein). Analogous ratios have been used in the real estate literature, some of which are the rent-to-price (Hamilton and Schwab (1985), Meese and Wallace (1994), Geltner and Mei (1995), Capozza and Seguin (1996), Campbell, Davis, Gallin, and Martin (2009), Himmelberg, Mayer, and Sinai (2005), Gallin (2008), Plazzi, Torous, and Valkanov (2010)), the loan-to-value (Lamont and Stein (1999)), and the price-to-income ratio (see, e.g., Malpezzi (1990),

(1999)). The economic reason for the use of ratios as predictors of future returns is straightforward and hinges on the plausible assumption that the variables used to form the ratios are co-integrated in logs (Engle and Granger (1987)). To see that, consider the log rent-price ratio. If log rents and log prices are co-integrated, then the log rent-price ratio must be a mean-reverting process. If at time t , the ratio is higher than its unconditional mean, the mean reversion implies that either expected returns would be high or that expected growth in rents would be low, or a combination of the two (Campbell, Davis, Gallin, and Martin (2009), Plazzi, Torous, and Valkanov (2010)). Similar logic applies to most valuation ratios.

To understand the appeal of the rent-price ratio to forecast either future rent growth or future returns, it is useful to introduce some notation. Let H_t denote rents net of all operating expenses of a property and P_t denotes its current price. Then the rent-price ratio is H_t/P_t and we denote its log transformation by $hp_t \equiv \ln(H_t) - \ln(P_t)$. It is also known as the capitalization, or cap, rate of the property. We will use the terms rent-price ratio and cap rate interchangeably from now on. Following Campbell and Shiller (1988), Campbell, Davis, Gallin, and Martin (2009) and Plazzi, Torous, and Valkanov (2010) show that hp_t can be expressed as

$$hp_t = k + E_t \left[\sum_{j=0}^{\infty} \rho^j r_{t+1+j} \right] - E_t \left[\sum_{j=0}^{\infty} \rho^j \Delta h_{t+1+j} \right] \quad (9)$$

where r_{t+1+j} is the future return of the property, Δh_{t+1+j} is the future growth in its rents, and k and ρ are linearization constants. In other words, fluctuations in the log cap rate must be able to predict either future returns, or future growth in rents, or both.

Expression (9) is the basis for the following long-horizon regressions:

$$r_{t+1:t+T} = \beta_r(T)hp_t + \tau_{t+1:t+T}^r \quad (10)$$

$$\Delta h_{t+1:t+T} = \beta_d(T)hp_t + \tau_{t+1:t+T}^d \quad (11)$$

where $r_{t+1:t+T} \equiv \sum_{j=1}^T r_{t+j}$ and $\Delta h_{t+1:t+T} \equiv \sum_{j=1}^T \Delta h_{t+j}$ approximate, respectively, the first and second term in (9) for a large T .³³ These approximations are appealing because $r_{t+1:t+T}$ and $\Delta h_{t+1:t+T}$ represent the log T -period return and rent growth, respectively. The forecasting regres-

³³The constant ρ is usually close to unity.

sion are estimated at various horizons, ranging from one period ($T = 1$) to several years ahead. The framework above has been used to investigate predictability of residential, commercial, and REIT returns. Given the extensive literature about REITs, we devote more attention to that literature in a separate section.

One might wonder why should the cap rate alone predict future real estate returns. Shouldn't other variables, such as construction costs, local economic conditions, zoning laws, and demographic trends be part of the set of explanatory variables? The assumption that the cap rate is the only conditioning variable is equivalent to assuming that all other economic factors are successfully summarized by that one quantity. In other words, this ratio captures all relevant economic fluctuations and it is the sole state variable. To the extent that some of the current economic information is not embedded in that ratio, the model will be misspecified.

Most of the literature focuses on return (or price appreciation) predictability in expression (10). A notable exception is the work of Hamilton and Schwab (1985). They find a significant negative relation between the rent-price ratio (not in logs) and future growth in rents. Using aggregate REIT data, our regressions support their findings (section (4)).

Meese and Wallace (1994), Capozza and Seguin (1996), Gallin (2008), and Campbell, Davis, Gallin, and Martin (2009) are some of the studies testing for return predictability within the above framework while Himmelberg, Mayer, and Sinai (2005) offer a slightly broader approach. Gallin (2008) estimates equations (10) and (11) using quarterly repeat-sales index data from 1970:Q1 to 2005:Q4. The two equations are estimated separately at four-years-ahead horizons ($T = 16$). He finds that the rent-price ratio has a positive relation with future returns and a negative relation with future rent growth rates, as predicted by expression (9). In his work, the coefficient $\beta_d(T)$ in the rent growth regression is statistically significant, whereas $\beta_r(T)$ in the return regression is not. Hence, from a statistical perspective, the evidence of rent-price predicting future returns is tenuous. Meese and Wallace (1994) formulate a different test which they carry out with transaction level data for Alameda and San Francisco counties in Northern California. They find evidence of short-run violations but long-run consistency with the present value relation and argue that high transaction costs might be the reason behind the differences across horizons. Capozza and Seguin (1996) point out that the predictive power of the cap rate is best observed once they account for other cross-sectional differences in rental versus owner-occupied housing. They use a pooled

sample of 64 metropolitan areas across the US from 1960 to 1990 and most of the data is from the decennial Census of Housing and Population.

In a recent work, Campbell, Davis, Gallin, and Martin (2009) also use expression (9) as a starting point of their analysis. Rather than assuming that future returns are an adequate proxy for expected returns, the authors follow Campbell (1991) and use vector autoregression (VAR) to forecast the future quantities in (9). One of the forecasting variables is the log rent-price ratio. Based on the VAR estimates, they document predictable variations in expected returns and expected rent growth for 23 metropolitan markets, four regional markets, and the national housing market over the 1975-2007 period with quarterly data. Consistent with the results of Gallin (2008), the rent-price ratio explains a larger fraction of the variability of expected returns than of expected rent growth. This is true for the entire sample and even during the boom subsample of 1997-2007. However, the statistical significance of the return predictability is not compelling. To illustrate the degree of real estate predictability by the log rent-price ratio, we run short-horizon equivalents of regression (10). The data is sampled at quarterly frequency and $T = 1$, which implies that we forecast returns one quarter ahead. We include the lagged return in addition to the lagged log rent-price ratio in the regressions because of the high degree of serial correlation documented in the previous section. The regressions are estimated for residential properties with the Census, Case-Shiller, and FHFA/OFHEO databases, and for commercial properties with the NCREIF, TBI, and REIT indices.

The residential results are presented in Table 4, Panel B.³⁴ When included as a stand-alone variable, the log rent-price ratio is a rather weak predictor of future returns for all three series. The coefficient is positive for the Census and FHFA/OFHEO series but is far from being statistically significant, and is essentially zero for the Case-Shiller index. In the subsequent specification, we add lagged returns as a control variable. The results from the Census data are now the most discouraging, from a predictability perspective. The point estimate for lagged log rent-price ratio is 0.008 with a Newey-West t -statistic of 0.255. This lack of predictability might be due to the fact that the Census median returns series do not adjust for the quality of properties. Some evidence pointing in that direction is presented in the next subsection. The Case-Shiller and FHFA/OFHEO estimates are positive, as suggested by Malpezzi (1999), but still statistically insignificant. The

³⁴We cannot run the residential regressions at monthly horizons, because the rent-price ratio data is only available quarterly.

large R^2 s of 0.559 and 0.409 are mostly due to the serial correlation in returns, which is captured by the lagged return term, r .

In Table 5, we display the predictive regressions results for commercial real estate returns based on quarterly observations over the same 1991:Q2-2010:Q4 period. The specifications are directly comparable with those for residential properties in Table 4, and the results are very much in agreement. More precisely, we observe a positive relation between the log rent-price ratio and future commercial real estate returns. While the point estimates are slightly larger for all three indices, and especially for the TBI, the Newey-West t -statistics are in the range of -0.121 to 1.494 when no other conditioning variables are included and in the range of 0.964 to 1.532 when controlling for lagged returns. The inability to reject the null of no predictability might be due to a lack of power of our test, especially given the presence of noise in the returns series. We will be able to explore the lack of power direction a bit further in the case of REITs, as we have market-based monthly observations over a longer time span. In sum, these results suggest that while there is a positive relation between the log rent-price ratio and future returns, as suggested by expression (6), it is not statistically significant in our samples. Admittedly, short-horizon predictability, even if it were present, is hard to detect in real estate indices that are so serially correlated.

The price-income ratio is suggested by Malpezzi (1999) as another predictor of real estate price changes. The underlying assumption behind this approach is the presence of an equilibrium relationship between house prices and household income. A temporary deviation of the price-income ratio above its long-run mean implies that either future prices will have to come down, the income level has to increase, or both. Malpezzi tests this prediction on repeat-sales and hedonic indexes on residential housing data at the MSA level. He documents that a one-unit increase in the distance between the price-income ratio and its equilibrium level is associated with a 2.7% drop in housing prices in the subsequent year. He also finds that other variables, such as mortgage rates, population growth and housing regulations, have an effect on future prices beyond the lags of the price-income. For instance, the stringency of the regulatory environment raises prices consistently, suggesting that supply shocks are important in determining house price dynamics (Glaeser, Gyourko, and Saiz (2008)). An inelastic housing supply is also suggested by the evidence that higher growth rates of population and income are associated with future price changes. The effect of supply shocks, however, hinges crucially on the assumption about the ability of households to move. If

this can happen freely, shocks to supply do not have an impact on house prices (Van Nieuwerburgh and Weill (2010)). Finally, higher mortgage rates predict lower price changes.

The vast majority of studies on real estate predictability estimate regressions (10) and/or (11) by OLS, equation by equation. However, if future returns and future rent growth are correlated (and there is no reason to believe that they are not), then equation-by-equation OLS regressions suffer from an omitted-variables bias. This argument has been made by Lettau and Ludvigson (2001) and Kojien and Van Binsbergen (2010) in the context of stock returns. The reason is that, in the presence of a correlation, the rent-to-price ratio is not sufficient to capture the time variation in the predicted variables and may understate their degree of time-variation. The same point has been made by Plazzi, Torous, and Valkanov (2010) in the case of commercial real estate predictability. To circumvent this issue, Plazzi, Torous, and Valkanov (2010) assume a first-order autoregressive process for the unobserved expected returns and expected rent growth and link it to the reduced form OLS regression coefficient in the rent-to-price ratio regression. They estimate the parameters of the underlying processes on a panel of quarterly returns to apartments, industrial properties, offices, and retail properties for 53 MSAs over the 1994-2003 period. Because of the added structure in the predictive regressions, they are able to identify and estimate substantial variations in both expected returns and expected rent growth in the time-series and across property types (in the 2%-20% range).

While the use of valuation ratios in predictive regressions has its appeal, an obvious drawback is that the predictive ratio might not be capturing all time variation in the conditioning information set. This is yet another reason to suspect that the results for such regressions may suffer from omitted-variable bias. But even if a ratio were a sufficient proxy for the time variation in economic conditions, the reduced form version of regressions (10) and (11) do not allow us to understand the economic forces that are behind the predictive relation. Is the forecastability due to demographic changes or other demand shocks? Or is it driven by supply shocks, such as tighter real estate regulations and zoning laws? What role, if any, do slowly-adjusting construction costs play? Those questions can only be answered if additional conditioning variables are introduced in the predictive regressions.

3.3 Predictability Based on Economic Variables

There is considerable evidence that economic variables, other than past returns or valuation ratios, are associated with future appreciations in property values, as shown by Rosen (1984), Linneman (1986), Skantz and Strickland (1987), Case and Shiller (1990), Abraham and Hendershott (1996), Pace, Barry, Gilley, and Sirmans (2000), MacKinnon and Al Zaman (2009), and Plazzi, Torous, and Valkanov (2010), among others. The empirical framework in most of these papers is the predictive regression (6) with conditioning information X_t that includes demographic variables (population growth, percentage of population within a certain age), income and employment variables, construction costs, housing starts, tax rates, zoning restrictions, and other regulatory variables. The selection of the conditioning information is dictated by the data, the level of aggregation, and the methodology.

Linneman (1986) uses data from the Annual Housing Survey for the Philadelphia residential market in 1975 and 1978 to test whether a wide set of property characteristics are associated with future changes in property values. He regresses 1975 residential values on a broad set of structural and neighborhood characteristics and finds that the 1975 residuals are significantly correlated with 1978 property prices.

Linneman's (1986) work is one of the first to document persistence in changes in residential values. However, it is not a predictive model in the strict sense of the term. In regression (6), we are looking to relate systematic changes in X_t with future returns. Most of the property characteristics in his hedonic model are fixed-effects and do not change with time. In an interesting study, Skantz and Strickland (1987) investigate house price dynamics following an unexpected disaster, Houston's widespread flood on July 26, 1979, which affected several subdivisions of the city. They find that prices in the flooded subdivision did not decrease immediately after the flood, which suggests that residential values already reflected the higher flood risk. They also document that house prices started to adjust downwards a year after the event, mainly because of an increase in flood insurance premia. The fact that, net of insurance premia costs, home prices are not affected by a natural disaster is a compelling evidence in favor of market efficiency.

Case and Shiller (1990) test for predictability in the excess total returns (returns minus the Tbill rate) of four metropolitan areas (Atlanta, Chicago, Dallas, San Francisco) with a number of

conditioning variables including lagged returns. The return data is computed from repeat-sales indices, rental data from the 1970 Census, and residential rent from the Bureau of Labor Statistics. The larger information set allows the authors to test a stronger version of the market efficiency hypothesis than in Case and Shiller (1989). The dataset is at quarterly frequency, from 1970 to 1986. The conditioning variables, the majority of which are available at metropolitan level, are: the rent-price ratio, the mortgage payment-income ratio, construction costs-price ratio, employment growth rate, real per capita income growth, growth rate in construction costs, percentage change in adult population (between ages 25 and 44), the percentage change in marginal tax rate, and housing starts divided by the population. The forecasting variable of interest is the excess return over the next four-quarters. Case and Shiller (1990) show that the economic predictors are able to capture a significant fraction of the fluctuations in future real estate returns; their fully-specified predictive regressions have an R^2 of 0.336 to as much as 0.615 (their Tables 7, 8, and 9). From the included variables, real per capita income growth and the increase in the adult population are strongly positively related with future annual excess returns. The sign is consistent with economic intuition, namely, improved economic conditions and demographic booms both put demand pressure on house prices. Two measures of fundamental value, the price-rent ratio and construction costs divided by price, also forecast future returns with a positive sign. To put things in perspective, the rent-price ratio alone has an in-sample R^2 of 0.109 (Tables 8 and 9). Variables such as the growth rate of employment, marginal individual income tax rate, and houses starts are also found to be important predictors. Overall, the addition of control variables increases the in-sample R^2 from about 10% to as much as 60%, thus rejecting the hypothesis of a semi-strongly efficient market. Their analysis is, however, in-sample and the pooled regression (across cities) use overlapping observations and do not explicitly account for cross-sectional correlation in the residuals.

Abraham and Hendershott (1996) explore the systematic predictability patterns of residential properties and link them to supply shocks. They express the growth in real estate prices as a linear function of lagged state variables such as the growth in real construction costs, the growth in real income per working age adult, the growth in employment, and the change in real after-tax rate. The error term of this regression is interpreted as a deviation from a market equilibrium and, therefore, a proxy for a bubble as it reflects adjustment dynamics and data noise. Working on the repeat-sales house price indices from Freddie Mac-Fannie Mae for 30 MSAs, they find that coastal

and inland cities are more pronounced in those markets. Their results imply that predictability in the residential market is to some degree related to measures of local supply elasticity such as the availability of desirable land.

The papers thus far have focused on predictability in the residential market. Similar studies have been conducted with commercial real estate data. For instance, MacKinnon and Al Zaman (2009) analyze the predictability of returns to direct real estate investment, proxied by the TBI index, and REITs in the context of a long-horizon asset allocation problem.³⁵ Their in-sample analysis reveals that TBI returns are predictable by variables such as lagged REIT returns, bond returns, and the T-bill interest rate, while cap rates and employment growth show only limited predictive power. This result is noteworthy especially because, as we saw in Table (2), price changes in the TBI index are close to uncorrelated. The mean-reversion in direct real estate investment makes it a less risky asset for long-horizon investors. As a result, for reasonable target values of the total portfolio expected return, the fraction of wealth invested in commercial properties is found to be no lower than 17%. Interestingly, in their application, the correlation between REITs and TBI returns is high enough to make REITs a redundant asset class when investors have access to direct investment in commercial properties.

The interaction between prices, vacancy rates, transactions volume, housing stock, and rents have also been investigated. For example, Wheaton and Torto (1988) consider deviations from equilibrium and adjustment dynamics between vacancy and rents in the office market. They find a statistically significant and qualitatively strong relationship between current excess vacancy and future real rent changes, even when accounting for the possibility in trends in the structural vacancy rates. Zhong-Guo (1997) uses a vector autoregression (VAR) error correction model to analyze the joint relationship between sales volume and median sales prices from the National Association of Realtors. Granger-causality tests indicate that sales affect price significantly, but home prices affect sales only weakly. DiPasquale and Wheaton (1994) show that the Tax Reform Act of 1986 had economically significant supply and demand effects on the rental properties markets. Based on their estimates, they predicted that, over the next ten years, real rents should have increased by 8 percent as a direct result of the legislation. These numbers are economically meaningful, especially given the low variability of rents over the decades preceding the 1986 legislation.

³⁵For a recent investigation on the European market, see Fugazza, Guidolin, and Nicodano (2007).

Evidence for predictability in commercial real estate returns and rent growth is also provided by Plazzi, Torous, and Valkanov (2010) using three principal components extracted from the level and change in population, employment, per-capita income, and construction costs. In particular, either the first or first and second principal components – which load on the level and growth in population and income as well as on construction costs – are found to be significant at the 1% level or better across all four property types. When the cap rate and a coastal dummy are also included, the adjusted R^2 in the returns predictive regressions range from 17% (for offices) to 37% for retail properties. In the rental growth predictive regressions, the adjusted R^2 are from 8% for offices to 14% for apartments. Interestingly, the cap rate remains significant even after the inclusion of these principal components, which suggests that the valuation ratio it is truly capturing time-varying dynamics rather than mere cross-sectional differences. Consistent with the findings in Abraham and Hendershott (1996), Plazzi, Torous, and Valkanov (2010) document cross-sectional differences in predictability depending on density and land-use restrictions. Evidence of return predictability is drawn primarily from locations characterized by lower population density and less stringent land-use restrictions. By contrast, rent growth predictability is more likely observed in locations characterized by higher population density and more severe land-use restrictions.

We revisit the predictability of aggregate real estate indices by conditioning variables other than lagged returns or log rent-price ratio. Panels A and B of Table 4 present various specifications of regression (6), estimated with monthly and quarterly residential data. In Panel A, we observe that the inclusion of the stock market's dividend-price ratio (dp_m), which is a proxy for the state of the equity market, has a negative effect on future residential returns. The effect is statistically significant for the Case-Shiller and FHFA/OHFEO indices. The relative T-Bill rate and the Cochrane and Piazzesi (2005) bond factor are significant for the repeat-sales indices. In the fourth, most comprehensive specification, the inclusion of lagged returns and all other economic variables reveals that, at monthly horizons, the data is simply too serially correlated for any of the additional predictors to be statistically significant.

In the quarterly predictability regressions, Panel B of Table 4, we observe that several economic variables are statistically significant in explaining future fluctuations in real estate price changes. More specifically, for the Case-Shiller index, the stock-market's dividend-price ratio, the relative T-Bill rate, inflation, and the Cochrane and Piazzesi (2005) bond factor are significant at the 10%

confidence level or better. The results for the FHFA/OFHEO are similar, albeit less significant. In all specifications, the log rent-price ratio is insignificant. Also, the Census median price index remains the least forecastable of the three indices.

In Table 5, we present the equivalent results for commercial properties at quarterly frequency. As in the previous table, the stock market's log dividend price ratio is negative and significantly related to future returns of the NCREIF and TBI indices. In the third specification of the regressions, higher industrial production growth leads to higher future changes in the same two indices. In the case of NCREIF, the term spread is also statistically significant but its point estimate is negative. The NCREIF is the most forecastable index, with as much as 70.9 percent of its changes being predicted, in-sample, by the economic variables. For the TBI, the explanatory power drops to 25.5 percent. REIT returns are the least predictable, as the joint explanatory power of all predictor yields an adjusted R^2 of 15.4 percent.

A recurring theme in the extant real estate literature is that the predictability of returns varies across geographic regions (e.g., Case and Shiller (1989), Gyourko and Voith (1992), Abraham and Hendershott (1996), Gu (2002), Crawford and Fratantoni (2003), Fratantoni and Schuh (2003), Capozza, Hendershott, and Mack (2004), and Hill (2004)). In a particularly exhaustive study of 62 metropolitan areas from 1979 to 1995, Capozza, Hendershott, and Mack (2004) note that “the dynamic properties of housing markets are specific to the given time and location being considered.” The economic sources of heterogeneity in predictive regressions are the same ones that determine house price dynamics, namely, demographic changes, regulations and zoning restrictions, local economic conditions, as well as heterogeneous responses to global macro-economic shocks. While differences in datasets, variable definitions, and methodologies make it hard to compare results across studies, quantifying the predictive ability of economic variables across metropolitan areas is of clear interest.

To illustrate the cross-sectional differences, we run predictive regressions similar to the ones discussed in Tables 4 and 5, but with MSA-level (rather than national) indices. To do so, we compute quarterly log price changes of 25 MSA regions from FHFA/OHFEIO over the 1991-2010 period. For each region, we regress the one-period-ahead price changes on the same set of conditioning variables that were used in Table 4. In other words, we run 25 MSA-level regressions, whose coefficient estimates, t -statistics, and R^2 s are comparable to those for the aggregate FHFA/OFHEO

index, reported in the very last column of Table 4. Rather than tabulating a large number of statistics, in Figure 4, we summarize the average predictive coefficients on each variable across regions (top-left panel), the average t -statistics (top right), the number of significant coefficients across regions (bottom left), and the average R^2 (bottom right).

The average coefficient on lagged returns is about 0.2, with an average t -statistic of nearly 2, and is statistically significant in 13 out of the 25 metropolitan areas. Interestingly, the most strongly significant predictors in the top-right panel – the lagged return, the market dividend-price ratio, and the RTB – are the same ones that were significant in the aggregate regressions (Table 4). Here, an additional predictor, the CP factor, displays a comparable importance. The same four variables appear as the most frequently significant across metropolitan areas, with the addition of inflation, suggesting that their average t -statistics are not driven by a few outliers. The fact that inflation and industrial production do not reach (on average) statistical significance might indicate the need for MSA-specific measures of economic activity. Our regressions, which are designed to capture common movements across MSAs price changes (because the predictors are the same across MSAs), show that at least some of the time-variation in residential returns is attributable to systematic, market-wide fluctuations. Abraham and Hendershott (1996), Capozza, Hendershott, and Mack (2004), and Del Negro and Otrok (2007) report similar findings. The average R^2 is 49%, in line with the 48% value documented for the national series. There is, however, considerable dispersion in the individual regressions R^2 (bottom-right plot), ranging from 25% for St. Louis, MO-IL to as high as 79% for Los Angeles-Long Beach-Glendale, CA.

As a more direct test of the presence of common factors in the cross-section of 25 MSAs, we extract the first 10 principal components of their covariance matrix. In the top panel of Figure 5, we plot the fraction of the total variance explained by each of these components. Strikingly, the first principal component explains slightly less than 70% of the covariance, while the other components are much less important. This evidence supports the findings in Figure 4 and the assertion that macro-economic fluctuations are behind some of the time-variation, at least over our sample period. Of course, a significant part of the unexplained variance is due to local factors. The bottom panel in the figure displays the fraction explained by the first component over the 2001 to 2010 period, estimated with a 40-quarter rolling window basis. Interestingly, the common component increased significantly during the 2008-2010 period, undoubtedly as a result of the bust

of the residential real estate bubble and the subsequent financial crisis.

4 REITs

A real estate corporation is considered to be a REIT for tax-purposes if it distributes at least 90 percent of its taxable income as dividends.³⁶ This unambiguous link between cash flows from commercial real estate and dividends makes REITs particularly suitable for predictability tests. In addition, REITs are traded on the US stock exchange, are relatively liquid, and have small transaction costs relative to other real estate investments.

We devote particular attention to REITs for two additional reasons. First, the data is available at higher frequency and over a longer time span than other datasets. This allows us to correct the parameter estimates for well-known, small-sample biases and to cast the predictive system in a GMM framework. Cross-equations restrictions, imposed by the model, may provide further efficiency gains in the estimation. These improvements will ultimately allow us to verify whether more precise parameter estimates and better in-sample fit obtain by using longer and less noisy series. Second, the longer sample allows us to investigate the (pseudo) out-of-sample performance of the forecasts. Thus far, most of the discussion had focused on estimation and in-sample performance. However, as is well-known in the stock market forecasting literature, in-sample fit does not necessarily translate into successful out-of-sample performance (Welch and Goyal (2008)). The disconnect between in-sample and out-of-sample results is likely due to parameter estimation error, structural breaks (Rossi (2012) in this Handbook) or more general model misspecifications. We investigate the out-of-sample predictability of REITs by also incorporating some new insights.

The literature on REIT predictability is voluminous. A partial list of the works on the topic include Liu and Mei (1992), Mei and Liu (1994), Li and Wang (1995), Mei and Gao (1995), Nelling and Gyourko (1998), Liu and Mei (1998), Brooks and Tsolacos (2001), Brooks and Tsolacos (2003), Serrano and Hoesli (2007), and Serrano and Hoesli (2010). The empirical approach in

³⁶In order to qualify as a REIT, a company must comply with certain provisions within the U.S. Internal Revenue Code. As required by the tax code, a REIT must: Be an entity that is taxable as a corporation; Be managed by a board of directors or trustees; Have shares that are fully transferable; Have no more than 50 percent of its shares held by 5 or fewer individuals during the last half of the taxable year; Invest at least 75 percent of its total assets in qualifying real estate assets, which include interests in real property, interests in mortgages on real property, or shares in other REITs; Derive at least 75 percent of its gross income from real estate related services, such as rents from real property or interest on mortgages financing real property; Have no more than 25 percent of its assets consist of stock in taxable REIT subsidiaries; Pay annually at least 90 percent of its taxable income in the form of shareholder dividends.

most of these studies can be framed into the following predictive system:

$$r_{t+1} = \mu_r + \beta_r x_t + \tau_{t+1}^r \quad (12)$$

$$\Delta d_{t+1} = \mu_d + \beta_d x_t + \tau_{t+1}^d \quad (13)$$

$$x_{t+1} = \mu_x + \phi x_t + \tau_{t+1}^{dp} \quad (14)$$

The one-period regressions are sometimes augmented by investigating the predictability of long-horizon returns $r_{t+1:t+T} \equiv T^{-1} \sum_{j=1}^T r_{t+j}$ and dividend growth $\Delta d_{t+1:t+T} \equiv T^{-1} \sum_{j=1}^T \Delta d_{t+j}$, as in equations (10) and (11) above. The conditioning variable x_t is usually the log dividend-price ratio. Multiple predictors are also considered, in which case β_r , β_d , and x_t are vectors.

In many instances, the predictor x_t is highly persistent and its innovations are correlated with the return innovations. In other words, ϕ is close to unity and the correlation between τ_{t+1}^{dp} and τ_{t+1}^r is non-zero. If that were the case, Cavanagh, Elliott, and Stock (1995) and Stambaugh (1999) have shown that the OLS estimator of β_r is biased. The larger (in absolute value) is the correlation between the innovations, the larger is the bias. This result is of relevance in real estate predictive regressions, as the predictor is often a valuation ratio, the T-bill rate, or an yield spread, which are all persistent. Several approaches of dealing with the small-sample bias have been suggested in the literature. We discuss the Stambaugh (1999) bias-correction below.

Liu and Mei (1992) is one of the earlier studies to investigate the predictability of US REIT returns. Using monthly data from 1972 to 1989, they document that the log cap rate on REITs and the T-Bill, and to a lesser extent the dividend-price ratio of the aggregate stock market, forecast one-month ahead REIT returns. The corresponding adjusted R^2 for REITs at 0.175 is found to be much higher than the 0.087 value for the overall market, and comparable to the 0.165 R^2 of a portfolio of small-cap stocks. This makes economic sense as REITs fall, on average, in the small-caps category of stocks. In a follow-up study, Mei and Liu (1994) find that market-timing trading strategies based on the predictors in Liu and Mei (1992) outperform a buy-and-hold portfolio strategy for REIT firms in an out-of-sample test. In particular, they report an R^2 of the conditioning model relative to the unconditional mean of 0.136 for a portfolio of real estate building companies stocks, and lower R^2 s at 0.109 and 0.083 in the case of the mREIT and eREIT portfolios, respectively. Similar predictability is however not present for common stocks. In a recent work, Serrano and

Hoesli (2010) document that REITs returns in more mature and well-established REIT markets – most notably the US, the Netherlands, and Australia – are more predictable than are local stock market returns. Trading strategies based on the forecasting regressions outperform a buy-and-hold benchmark portfolio in all ten countries, and the gains exceed the transaction costs in about half of them. In sum, these studies suggest that predictability in real estate returns might be exploitable.

Several papers have gone beyond the simple univariate, linear forecasting models of equations (12-14). For example, the results in Serrano and Hoesli (2007) indicate that neural networks models outperform linear regressions (time varying coefficient regressions and VAR systems) in forecasting equity REIT returns, when the models jointly include stock, bond, real estate, size, and book-to-market factors. Similar findings based on European data can be found in Brooks and Tsolacos (2003). Although there is no consensus in the literature as to what forecasting technique works best, increasing the complexity of the model seems in general to provide some in-sample improvements in forecasting accuracy.

Performance continuation and reversals are also related to the time series properties of asset returns and have been the subject of many studies in the financial economics literature. For securitized real estate, Mei and Gao (1995) examine serial persistence of weekly returns and argue that a contrarian-based strategy is profitable only if transaction costs are ignored. Analogous results based on monthly observations can be found in Nellling and Gyourko (1998), who document that the extant predictability is strongest after 1992 following the major reforms in the REITs market. Using a filter-based rule, Cooper, Downs, and Patterson (1999) show that a contrarian strategy is in many cases more profitable than its associated execution costs. Graff and Young (1997) use different frequencies and find positive momentum effects with yearly data, evidence of performance reversals with monthly data, and no evidence of momentum or reversals with quarterly data. Finally, Stevenson (2002) provides international evidence of momentum effects over short and medium term horizons, as well as little support for price reversals.

To highlight and update the results we just discussed, we test for predictability using the monthly equity REIT series from the CRSP/Ziman database during the full 1980-2010 period. We construct monthly dividends using total and without-dividends returns as in Fama and French (1988b). The dividend series is then calculated as sum of the current and past 11 months dividends. The log dividend-price ratio is defined as the log dividend minus the log of the level in the current

month. The left-hand side variables are the real log return and dividend growth, deflated by the CPI Index.³⁷ We also look at the results for the REIT portfolios of companies that hold mainly apartment buildings (Apt), industrial buildings and offices (Ind & Off), and retail properties (Rtl), as they may exhibit different predictability properties.

Table 6 presents OLS estimates of equations (12) and (13) for all REITs, Apt, Ind & Off, and Rtl. The predictive regressions are estimated over horizons of $T = \{1, 3, 6, 12\}$ months. Some studies compute long-horizon returns by overlapping one-period returns whereas others consider non-overlapping windows. We provide both sets of results. Panel A displays the estimates for the return regressions (12) and Panel B displays those for the dividend-growth regression (13). It is well-known that the persistence of the dividend-price ratio (ϕ in (14) close to one) and a negative correlation of the predictor innovations with those of returns (correlation of τ_{t+1}^r and τ_{t+1}^{dp}) induce significant small-sample bias in the OLS estimates in regression (12) (Stambaugh (1999)). Therefore, the displayed coefficient estimates are adjusted for the small-sample bias using Stambaugh's (1999) correction. In the non-overlapping results, the t -statistics are calculating with Newey and West (1987) standard errors with four lags. In the overlapping data, the number of lags equals the number of overlapping observations T .

In Panel A of Table 6, the positive OLS estimates of β_r at various horizons show that the log dividend-yield is positively related to future REIT returns. This is true for the entire REIT portfolio and especially for retail properties. Interestingly, industrial and office properties exhibit little predictability at any horizon. Plazzi, Torous, and Valkanov (2010) document a similar finding, namely, that industrial and office property returns are the least predictable, using a different dataset. As the horizon increases, we observe an increase in the t -statistics. For overlapping returns, they reach customary significance levels at the yearly horizon. This result supports the claim in previous studies that returns are predictable at long horizons. However, we observe that the non-overlapping long-horizon results are not significant which raises the possibility that the high t -statistics might be the product of distorted inference, due to the severe serial correlation, induced by the overlap (Valkanov (2003)).

The OLS estimates of β_d in Panel B are negative, as implied by the log-linearized expression (9). The point estimates are larger in magnitude than those in Panel A and, more importantly, they

³⁷Working with nominal or excess returns yield very similar results.

are statistically significant, with t -statistics for overlapping REITs in the -3.5 range. For industrial and office properties, the predictability is even more significant as the t -statistics are about -5 . This finding mirrors the Plazzi, Torous, and Valkanov's (2010) results that those properties exhibit the most predictable rent growth rates. It is comforting to observe that estimates and t -statistics using overlapping and non-overlapping dividend growth data yield very similar results. The R^2 in the overlapping and non-overlapping regressions is 0.2 for all REITs and as high as 0.4 to 0.5 for industrial and office properties. The strong ability of the dividend-price ratio to predict REITs dividend growth rates is in sharp contrast with the findings for the aggregate stock market (Lettau and Van Nieuwerburgh (2008)). This difference may be partly attributable to the strict payout policy that makes REITs dividends less prone to artificial smoothing, or to the nature of REIT cash flows.

The mixed predictability results of returns might be due to the lack of statistical power of our tests. Recent work by Lettau and Van Nieuwerburgh (2008) suggests that estimation and inference in long-horizon regressions can be improved by noting that the predictive system (12)–(14) implies the following time-series restriction relating one-period and T -period slope coefficients:

$$\beta_r(T) = \beta_r \left(\frac{1 - \phi^T}{1 - \phi} \right) \quad \text{and} \quad \beta_d(T) = \beta_d \left(\frac{1 - \phi^T}{1 - \phi} \right). \quad (15)$$

In addition, the Campbell and Shiller (1988) present-value identity imposes a restriction between the predictive regression coefficients:

$$\beta_r - \beta_d = 1 - \rho\phi \quad (16)$$

where ρ is the log-linearization coefficient. This expression is nothing but a restatement of the fact that, owing to the variability of dp , we must observe either return predictability, dividend growth predictability, or both. These observations suggest that more efficient estimates of β_r and β_d can be obtained through a GMM estimator which imposes the above restrictions across different horizons T , as in Lettau and Van Nieuwerburgh (2008) and Plazzi, Torous, and Valkanov (2010).

In Table 6 we report the one-month coefficients β_r and β_d , estimated with GMM using horizons of $\{1, 3, 6, 12\}$ months. The estimates of β_r in Panel A remain positive and statistically insignificant for all REITs and for the various property types. The estimates of β_d (Panel B) are negative

and smaller in magnitude than those in the unrestricted OLS regressions. However, their t -statistics have increased both with overlapping and non-overlapping data. The OLS and GMM results point to the same three conclusions. First, there is a positive but largely insignificant in-sample relation between the log dividend-price ratio and future returns. Second, the log dividend-price ratio forecasts dividend growth rates with a negative sign and the estimates are statistically significant. Finally, industrial and office properties seem to exhibit the least returns predictability and the most dividend growth predictability. This last finding merits further attention as it might provide further insights about the underlying economic sources behind the trade-off in predictability between returns and dividend growth rates.

The empirical discussions were thus far framed around in-sample fit in predictive regressions. However, the ultimate test of such regressions resides in whether they can produce accurate out-of-sample forecasts. Recently, the stock market predictability literature has found that significant predictors do not necessarily produce accurate out-of-sample forecasts. For instance, Welch and Goyal (2008) have documented that stock return predictors do not outperform the simple unconditional mean in a (pseudo) out-of-sample comparison. As a measure of relative forecasting accuracy, Welch and Goyal (2008) use the out-of-sample (as opposed to in-sample) R^2 defined as:

$$R_{OOS}^2 = 1 - \frac{\sum_{t=1}^T (r_t - \hat{r}_{t|t-1})}{\sum_{t=1}^T (r_t - \bar{r}_{t|t-1})} \quad (17)$$

where $\hat{r}_{t|t-1}$ and $\bar{r}_{t|t-1}$ are the predicted value and the historical average return, respectively, estimated using information up to and including time $t - 1$.

We take this opportunity to look at the out-of-sample predictive power of conditioning variables in REIT regressions. In addition to the log dividend-price ratio, we consider the same commonly-used predictors that were discussed in Tables 3 and 4.³⁸ We conduct the out-of-sample comparison by splitting the sample in two periods. A first 180-month period, corresponding to the 1980-1995 sample, is used to estimate first the forecasting model. Then, the estimates from that sample are taken to form the first forecasts of returns and dividend growth for January, 1996. Subsequently, we include the January 1996 observations to re-estimate the models and formulate forecasts for February, 1996 and so on until December, 2010. We do this at horizons of $T = \{1, 3, 6\}$ months.

³⁸For the Cochrane and Piazzesi (2005) factor, we construct out-of-sample estimates by re-estimating the coefficients using only information available up to time $t - 1$.

Given the limited time-series, longer horizons are not possible.

In Panel A of Table 7, we present the monthly in-sample (column “IS”) and out-of-sample (column “OOS unc”) R^2 s. The “unc” stands for “unconstrained” to differentiate those results from some constrained ones, discussed below. From the definition of R^2_{OOS} in equation (17), positive values imply that the model delivers a lower mean square forecasting error (MSFE) than the unconditional mean. Looking down the IS column for real returns, we notice that several predictors have in-sample predictive power. However, the OOS unc column reveals that, with the exception of the lagged stock market return, the OOS R^2 are all negative. In other words, the MSFEs of all but one forecasting variables are higher than that of the unconditional mean. A combined forecast, obtained by equally-weighting all 8 predictions, does well in-sample but its out-of sample performance barely beats the unconditional mean (last row in Panel A).

Campbell and Thompson (2008) show that imposing economic constraints on the OOS forecasts of stock market returns results in significant improvements in their OOS R^2 . Following their work, we consider several non-linear constraints. The first one (column “ $\text{sign}(\beta_r)$ ”) restricts the estimated coefficient in a given period to have the expected sign.³⁹ The second constraint (column “OOS $\hat{r} > 0$ ”) imposes non-negativity of the forecasted real return. In column “OOS both”, the forecast must satisfy both constraints. In periods when a constraint is violated, the historical average is used instead as a forecast.⁴⁰ The OOS results from the three constraints are displayed in Table 7. Imposing the $\text{sign}(\beta_r)$ constraint brings several of the negative OOS R^2 s into positive territory and increases the OOS R^2 s of the combined forecast. The $\hat{r} > 0$ and joint constraints lead to modest improvements in the predictions. In Panels B and C, we present the results at quarterly and semi-annual horizons and observe several differences with respect to the monthly results. The log dividend-price ratio, which was insignificant at monthly frequency, is more significant in-sample. In Panel C, its in-sample R^2 s is as high as 3.862 percent. Similar increases in IS R^2 s occur for dp_m , and TSP . However, this in-sample fit does not translate into OOS performance. Indeed, the MSFE of dp is 15.211 percent higher than that of the unconditional mean. Some predictors, such as the CP factor exhibit improvement in OOS forecasting power. Imposing the constraints does not help the log dividend-yield but does lead to a better performance of IP .

³⁹For the dividend-price ratio, the Campbell and Shiller (1988) decomposition implies a positive coefficient. For the other variables, we impose the same positivity constraint, although the guidance from economic theory is less clear. The sign constraints are broadly consistent with the prior of counter-cyclical risk premia.

⁴⁰Other types of constraints could be imposed to enhance forecastability. See, for example, Pettenuzzo, Timmermann, and Valkanov (2011).

The results in Tables 6 and 7 all point toward weak evidence of in-sample predictability at long horizons. Econometric refinements, designed to reduce bias and increase efficiency of the estimates, do not alter significantly this conclusion. The OOS forecasting exercise suggests that most of the predictors yield MSFEs similar to or higher than the unconditional mean. This is true whether or not we impose economic constraints. The in-sample significance of our predictive results appears somewhat smaller than documented in previous studies. This difference may be partly due to the fact that our sample includes the particularly volatile 2007 to 2010 period and to the downward bias adjustment. Perhaps the most novel result in the tables is that the dividend growth of REITs is forecastable in-sample, although the OOS results are, once again, much weaker.

5 Real Estate, Leverage, and Monetary Policy

5.1 The Effect of Leverage

The relation between real estate prices and leverage is a natural one, especially in the residential market. While in principle any asset could be used as collateral, housing is by far the easiest asset to borrow against. At the peak of the recent US residential bubble, about \$12 trillion of outstanding mortgage debt had been issued against the value of properties, which at the time were worth in the neighborhood of \$25 trillion (Case (2007)). In the UK, Aoki, Proudman, and Vlieghe (2004) report that about 80 percent of all household borrowing is secured on housing. High levels of home borrowing are bound to have an effect on house price dynamics and some of the studies that analyze this relationship are Linneman and Wachter (1989), Stein (1995), Lamont and Stein (1999), Spiegel (2001), Aoki, Proudman, and Vlieghe (2004), Ortalo-Magné and Rady (2006), Lustig and Van Nieuwerburgh (2005), Favilukis, Ludvigson, and Van Nieuwerburgh (2011).

Stein (1995) proposes a static fully-rational model that explains the joint relationship between house prices, trading volume, and leverage. In his framework, households must repay any existing loan and make a down-payment prior of moving to a new home. If households experience a negative exogenous shock to their income, the home-equity portion for families with a high level of leverage may not be enough to meet the down-payment of a larger house. These households may

therefore decide not to move, thus creating a decrease in demand which further depresses prices.⁴¹ On the contrary, following a positive income shock, these financially constrained families prefer to move to their desired location promptly, thus increasing prices and volume. Stein's (1995) model implies that the mix of leverage and liquidity constraints amplifies the effect of changes in asset values on the demand for housing. In cities where a large fraction of families are highly leveraged, the impact of fundamental shocks on house prices will be significantly higher, thus giving rise to self-reinforcing effects. As a result, home prices would display pronounced boom-to-bust cycles which may appear incompatible with an efficient capital market dynamics.⁴²

McDonald (1999), Spiegel (2001), and Ortalo-Magné and Rady (2006) are three additional papers that investigate the theoretical links between leverage and real estate prices. Spiegel (2001) shows that the presence of credit constraints can lead to construction cycles. As an implication of his model, he shows that leverage and developer construction decisions forecast time variation in expected housing returns. Ortalo-Magné and Rady (2006) propose a life-cycle model in which households differ in their income and thus ability to afford down-payment on a home. Young agents are constrained in their ability to borrow to purchase their first "starter" home. Moreover, changes in the price of starter homes shifts the agents' demand for trade-up homes, thus establishing a link between the price of a starter home and the price of trade-up homes. While the modeling assumptions are different from Stein's (1995), the two papers rely on the same amplifying effects of leverage and emphasize the role of down-payments and liquidity constraints.

The empirical predictions of Stein (1995) are investigated by Lamont and Stein (1999) whose main focus is on the effect of leverage on future house price fluctuations. Working on a sample of 44 metropolitan areas available at annual frequency between 1984 and 1994, they look at the fraction of all owner-occupants with an outstanding mortgage balance to house value ratio greater than 80 percent. This measure is meant to capture the relative presence of "constrained mover" families which play a destabilizing effect in Stein's (1995) model. Lamont and Stein (1999) find that leverage plays an important role in predicting future real estate prices in two distinct ways. First, lenders and borrowers may be willing to take on high leveraged positions just if they fore-

⁴¹ Alternatively, they may also try to list their home for a relatively high price as this represents a low-cost alternative. This would explain why even in falling markets there is some inertia in the decrease of prices.

⁴² The model's main predictions hold as long as the market for renting does not represent a costless alternative than buying a new home. Tax and moral hazard reasons suggest that renting is not a perfect substitute for direct ownership, so that the most efficient way to consume new housing is owner-occupied. The magnitude of these effects can be big as a large fraction of all home sales are to repeat buyers.

see house prices to rise. Consistent with this expectation hypothesis, high leverage today is found to positively correlate with future price appreciation. In their sample, real estate prices are predictable by lagged price changes, with a positive sign, and negatively by the lagged price-income ratio. These effects suggest that housing prices are driven by local economic conditions, short-run momentum, and long-run reversal to fundamentals.

A second important finding of Lamont and Stein (1999) is the economically and statistically large effect of leverage when it is interacted with changes in income. While shocks to income seem to be fully absorbed by real estate prices within four years, there are considerable cross-sectional differences in the response of house prices depending on the initial distribution of debt levels. In high-leverage cities, housing prices react quickly to an income shock and overshoot in the short-run. This effect peaks in the fourth year, when prices start to mean-revert to their new long-run level. By contrast, low-leverage cities display a more gradual response to the same economic shock and a smooth transition to the new equilibrium level. These effects are consistent with the main prediction of Stein (1995). Even if the level of leverage in an area is ultimately an endogenous variable, their analysis suggests a causal relationship which runs from changes in leverage to house price fluctuations. An empirical limitation of their findings, however, is that they use homeowners' estimates of their home value rather than true market prices.

Several other methods have been used to quantify the effect of leverage on the demand for residential and commercial properties. Linneman and Wachter (1989) document that wealth and income constraints lead to a lower probability of home ownership. Using data from the Federal Reserve Board's 1977 Survey of Consumer Credit and the 1983 Survey of Consumer Finances, they show that mortgage market innovations tend to relax financial constraints. However, the authors do not investigate price effects as a function of the constraints. Genesove and Mayer (1997) use data from the Boston condominium market in the early 1990's to show that an owner's level of leverage determines his behavior as a seller and might have an impact on transaction prices. They find that leveraged owners tend to set a higher asking price, their properties stay longer on the market, and conditional on selling, the transaction price is higher than that of less leveraged owners. These findings are broadly consistent with the predictions of Stein (1995) and Ortalo-Magné and Rady (2006). Brown (2000) compares the performance of highly-leveraged mortgage REITs (typically 80 percent leverage or more) to those of less leveraged properties held by property management

companies. In the period of the late 1980s and early 1990s, which was characterized by large declines in commercial real estate values, he finds that the leveraged mortgage REITs were net sellers of highly leveraged assets, whereas equity REITs were net buyers. Moreover, the returns of mortgage REITs were significantly more negative than those of equity REITs. The overall evidence of the 1989 and 1990 period suggests that high levels of debt forced mortgage REITs to sell assets at fire sale prices, thus resulting in large losses.

Recently, Favilukis, Ludvigson, and Van Nieuwerburgh (2011) formulate an elegant general equilibrium overlapping-generation model in which heterogeneous agents derive utility from a durable consumption good and housing and face stochastic life-cycle earnings. Two frictions are at play in their model. First, the households face borrowing constraints in the form of a down-payment as well as transaction costs. A “financial market liberalization”, FML, corresponds to a relaxation of these constraints. Second, financial markets are incomplete, which implies that both idiosyncratic and aggregate risks cannot be perfectly insured. This rich framework represents an ideal laboratory to investigate the ability of FML and foreign capital inflows of producing the swings in housing markets, aggregate quantities, and consumption that are observed in U.S. data. The model is indeed able to generate endogenous swings in rent-to-price ratios largely as a response to a financial market liberalization. A positive shock to FML – i.e. an easing of collateral requirements and lower transaction costs – increases risk-sharing as allows households to borrow more aggressively against unexpected shocks to their wealth. This, in turn, decreases risk premia and increase prices. The opposite pattern is observed during a reversal of FML. The authors also show that low interest rates, which are the result of an inflow of foreign capital into the domestic bond market, play a comparatively small role in driving house prices fluctuations. The reason is that the decrease in interest rates is almost offset by an increase in risk premia due to the fact that domestic savers are pulled out of the bond market by foreign purchases, and are therefore more exposed to systematic equity and housing risk. The demand pressure on Treasuries from capital inflows keeps interest rates low thus balancing the decrease in precautionary savings from FML, but does not constitute the ultimate cause of the housing boom and bust.

The model in Favilukis, Ludvigson, and Van Nieuwerburgh (2011) is also capable of generating plausible amounts of real estate predictability for both excess returns and rent growth. High rent-price ratios forecast higher future (excess) returns, as a result of the improved risk-sharing above

mentioned. As a point in case, the R^2 from predicting one-year excess returns implied by the model is about 3%, which increases to 10% at the five-year horizon. High rent-price ratios also forecast higher future rental growth rates, contrary to the main prediction of standard present-value models. The authors argue that this empirical regularity can only be understood in a general equilibrium context, in which a negative economic shock can simultaneously drive discount rates up and residential investment down. As a result of a shrinking supply of housing, future rental growth rates increase. This positive correlation between risk premia and expected rental growth is at the root of the positive coefficient in the predictive regression.

Favilukis, Kohn, Ludvigson, and Van Nieuwerburgh (2012) provide further evidence in line with the predictions of the model by taking an empirical look at the determinants of residential real estate dynamics in the U.S. and several other countries. They document that a measure of credit supply of the banking sector is by far the most powerful explanatory variable of contemporaneous real house price growth during the 2000-2010 decade internationally, and for the U.S. also in the longer 1991-2010 period. This variable is meant to capture an easing of credit constraints, such as an increase in loan-to-value ratios and in the availability of new mortgage contracts, a reduction of fees, and laxer documentation requirements. They show that credit supply is also a powerful *forecaster* of both future house prices and rent-to-price ratios, with R^2 s exceeding 0.40 in predictive regressions of one- to four-quarter ahead. These results still hold when credit supply is purged from endogenous shocks by controlling for expected economic conditions, which suggests the existence of a causality link running from exogenous shocks in the tightness of credit supply to subsequent movements in house prices. In contrast, capital flows, interest rates, and GDP growth are found not to be statistically significant determinants of current and future real house price dynamics, both in the U.S. as well as internationally.

A similar argument is found in Mian and Sufi (2009), who investigate the links between credit supply, leverage, and housing prices in the 2007 subprime crisis. Using zip-code level data, they find that mortgage origination during the 2002-2005 period was stronger in zip codes with a larger fraction of sub-prime applicants. The same areas, however, exhibit a sharp decrease in income and employment growth in both absolute and relative terms with respect to other zip codes in the same county. Thus, more mortgages originated in zip codes with worsening relative prospects. Mian and Sufi (2009) also test whether the increase in credit has been driven by expectations of future house

price appreciation. Their identification strategy exploits the index of housing supply elasticity based on land-topology metrics developed by Saiz (2010). House price growth in elastic MSAs is expected not to exceed the inflation rate for house construction costs, as supply can quickly accommodate any increased demand (Glaeser, Gyourko, and Saiz (2008)). Historically, the bulk of house price appreciations have been concentrated in MSAs with inelastic supply. Based on this evidence, one would expect not to observe any increase in lending in elastic MSAs areas, as these do not have good prospects for house price appreciations. Contrary to this prediction, even zip codes in elastic MSAs with a greater fraction of sub-prime lenders experienced an increase in the relative fraction of originated mortgages sold for securitization and a positive relative mortgage origination growth. These results are supportive of a supply-based hypothesis, one in which the increased in credit supply through more relaxed lending practices, a decrease in the price of debt (the subprime-prime spread), and an increase in securitization in high subprime zip codes all lead to a reduction in denial rates and to an increase in house prices. Moreover, 2002-2005 was the only period in the last eighteen years when house prices were negatively correlated to income growth, and thus to fundamentals. Also, it was a period of great expansion of the sub-prime lending industry.

In a companion paper, Mian and Sufi (2010) document that, from 2002 to 2006, the average homeowner extracted about 25 to 30 cents of every dollar increase in his home equity. This home-equity-based borrowing channel is much stronger for more financially constrained households, such as younger, low quality borrowers, and high credit card users. They estimate that, on aggregate, about 53% of the increase in homeowners' debt from 2002 to 2006, or about \$1.25 trillion, was due to home-equity borrowing. Defaults of these existing homeowners represented 39% of total new defaults in the U.S. during the same period. Their results point at the importance of a feedback effect, resulting from an increase in collateral values and impacting the availability of new credit, and are consistent with the aggregate evidence of a wealth effect in Case, Quigley, and Shiller (2011).

Overall, these studies corroborate the claim that exogenous economic shocks lead to larger price fluctuations in leveraged properties. The work of Lamont and Stein (1999) and Brown (2000) suggest that the amplification mechanism derives from the fact that, in real estate, the ability to borrow is directly linked to asset values. An exogenous decrease in those values can lead to a

reduction in asset demand, as the borrowing capacity has decreased. However, more research is needed to document whether positive and negative economic shocks lead to an asymmetric response in prices. The analysis in Mian and Sufi (2009) and Favilukis, Kohn, Ludvigson, and Van Nieuwerburgh (2012) leads to the conclusion that the supply of (subprime) credit in the mid 2000s fueled to a large extent the subsequent house price increases.

5.2 Monetary Policy and Real Estate

Under the premise that monetary policy can stimulate economic activity by reducing borrowing costs, and given that the real estate market is heavily reliant on credit, it is reasonable to ask what is the effect of monetary policy shocks on real estate prices and sales volume. An emerging literature is looking at this question, with recent contributions by Fratantoni and Schuh (2003), Aoki, Proudman, and Vlieghe (2004), Iacoviello (2005), Ahearne, Ammer, Doyle, Kole, and Martin (2005), Del Negro and Otrok (2007), Hamilton (2008), and Iacoviello and Neri (2010).

At the aggregate level, Iacoviello (2005) shows that house prices and GDP respond negatively to tight monetary policy in the US with quarterly data from 1974 to 2003. To explore the role of borrowing constraints in the observed link between house prices and monetary policy, Iacoviello uses a new-Keynesian model with collateral constraints tied to real estate values for firms, as in Kiyotaki and Moore (1997). The importance of collateral to lower borrowing costs is also considered by Aoki, Proudman, and Vlieghe (2004) who note that house prices may have a direct effect on consumption via the credit market. Their theoretical model is supported by the empirical evidence in Case, Quigley, and Shiller (2005), (2011) that price changes in real estate have a large impact on aggregate consumption. In the Aoki, Proudman, and Vlieghe (2004) model, a house represents collateral for homeowners, and borrowing on a secured basis against ample housing collateral is generally cheaper than borrowing against little collateral or on an unsecured basis, such as a personal loan or credit card. Therefore, an increase in housing prices makes more collateral available to homeowners, which in turn encourages them to borrow more to finance their desired level of consumption and housing investment. Looking at structural changes in the UK's retail financial market, they also show that cheaper access to home equity means that, for a given house price increase, more borrowing will be devoted to consumption relative to housing investment. The

response of consumption to an unanticipated change in interest rates will therefore be larger, and the response of house prices and housing investment will be smaller. Ahearne, Ammer, Doyle, Kole, and Martin (2005) find that the links between monetary policy shocks and real estate prices are observed more generally across eighteen developed economies.

Due to the geographical heterogeneity of real estate, it is unlikely that a monetary policy shock will have the same effect on all regional markets. In fact, it has been documented that the impact of monetary actions varies across regions (Carlino and DeFina (1998), (1999)) and is a function of local economic conditions. It is therefore possible that the response of real estate prices to monetary policy shocks may differ across regions in magnitude and duration. Based on this premise, Fratantoni and Schuh (2003) quantify the importance of regional heterogeneity in housing markets with respect to monetary policy shocks. They start off with a model in which the monetary authority sets the global monetary policy and the mortgage rate serves as the central channel for monetary transmission. At the regional level, income, housing investment, and housing prices are determined by households and firms. Fratantoni and Schuh (2003) contrast the standard method of estimating the effect of monetary policy with one in which regional heterogeneity matters. Following the standard approach first, the authors estimate a structural VAR model of log non-housing deflator, log housing deflator, log per-capita non-housing GDP, log real per-capita housing investment, FED funds rate, and a nominal rate on a 30-year mortgage during the 1966-1998 period. Their VAR estimates suggest that a monetary shock has a significantly larger and more rapid effect on the housing than non-housing sector and that monetary policy accounts for a large fraction of fluctuations in the housing sector. However, since housing supply and demand are determined locally, this motivates their regional heterogeneous-agent VAR (HAVAR) analysis.

A HAVAR model involves the estimation of regional VARs which are then aggregated at a national level. The authors show that aggregating these VARs induces non-linearities in the form of time variation (aggregation across heterogeneous regions) and state dependence, due to prevailing economic conditions at the time of the monetary policy intervention. They estimate their model on MSA-level data of housing starts (from the Bureau of Labor Statistics), housing prices (from the FHFA repeat sales transactions), and state-level income (from the Census Bureau). Using a balanced panel of 27 MSAs over 1986:Q3 to 1996:Q2, they estimate unrestricted VARs via OLS and look at the effect of monetary policy tightening, defined as a transitory 100-basis points shock

to the FED funds rate, on regional income, housing investment, and house appreciation. The resulting impulse-response functions show that the magnitude and duration of the regional responses vary widely across areas. For instance, monetary tightening is moderately less effective when the economy is experiencing a coastal housing boom. Their key finding is that there are economically and statistically significant differences between the dynamic responses of the HAVAR model they develop and the conventional VAR approach. Peak responses to monetary shocks can vary by more than 1%, and mean lags, by more than one year, depending on local conditions.

In a complementary work, Del Negro and Otrok (2007) use a dynamic factor model to decompose house price movements into regional shocks and a common national component. The motivation for the decomposition is the considerable empirical heterogeneity in the growth rate of house prices across states. They use the FHFA/OFHEO data during the 1986-2005 period. The dynamic factor model is estimated via Bayesian methods and the common component is identified from the idiosyncratic movements. Del Negro and Otrok (2007) find that, historically, fluctuations in housing prices have been mainly driven by local factors. Interestingly, growth rates of housing are far less correlated across states than are growth rates in real per capita income. This heterogeneity is due to the fact that states have different exposures to the common business cycle. However, in the more recent period, many states display an increase in house prices due to the national factor. The authors investigate whether this increased correlation may be due to monetary policy intervention. Interestingly, the impact of monetary policy shocks on house prices is found to be fairly small. Thus, the authors conclude that the Fed expansionary policy was not behind the recent boom in house prices.

Hamilton (2008) uses a new, high-frequency measure of monetary policy shocks—daily innovations in the 1- and 2-month-ahead futures contracts—and traces out its effect on the housing market. More specifically, he estimates the response of new homes sold (reported by the Census Bureau) following monetary policy shocks. The main finding in the empirical analysis, which is carried out at the aggregate level, is that sales respond with a considerable lag to monetary shocks. He attributes the delay of several months to heterogeneity in search time across households. It would be interesting to extend Hamilton's (2008) approach and investigate the impact of Fed policy shocks on real estate prices. However, it is fair to say that this, and the other papers in this literature, suggest that monetary policy actions have an important role to play in understanding real estate

price fluctuations.

6 Concluding Remarks

How difficult is it to predict changes in real estate prices? In this chapter, we revisit this question from the perspective of the academic real estate literature. However, before we can even tackle the question of predictability, we must step back and address an even more basic issue: constructing reliable real estate price indices in light of the fact that the underlying asset is extremely heterogeneous, faces high transaction costs, and is inherently illiquid. Various indices have been proposed and all have limitations. Recently much interest has focused on repeat-sales indices, both for commercial and residential properties. However, while these indices may be suitable for capturing the current state of the market, they might not necessarily be ideal for forecasting future *market* prices.

Researchers have explored a variety of predictors of real estate price changes beyond simple autoregressive models. A partial list of the most successful forecasting variables include: valuation ratios such as the rent-to-price and income-to-price ratios; local economic variables, such as the employment rate, income growth, construction costs; demographic trends such as population growth; local space market variables such as housing starts, vacancy rates, and transactions volume; proxies for zoning restrictions; measures of leverage and monetary policy action. The interpretation and success of the predictive models varies considerably as do the datasets on which they have been tested. From a statistical perspective, evaluating the accuracy of real estate forecasts is a challenge: the lack of a sufficiently long time-series of price data has prevented researchers from conducting meaningful out-of-sample MSFE comparisons. Most of the reported evidence is in-sample and exploits cross-sectional differences. However, from an economic perspective, whether the profits from exploiting predictable patterns in real estate prices are enough to cover transaction and search costs is still an unresolved issue.

A notable exception is the REITs market, for which exchange-traded returns are available for a relatively long timespan. In-sample results highlight some predictability of REIT returns, especially at yearly horizons. Significantly more predictability is observed in the growth rates of their cashflows. Out-of-sample predictions of REIT returns mimic those of the aggregate stock market, where unconstrained forecasts perform relatively poorly, although economic constraints may pro-

vide some improvements. Linking REITs returns to those of the underlying properties may be an interesting application to trace the origin of the extant predictability.

Appendix

A Data sources

A.1 Real Estate Data

Residential: The Census Median and Average Sales Prices of New Homes Sold are obtained from the Census.⁴³ The S&P/Case-Shiller indexes are the NSA series from Macromarkets.⁴⁴ FHFA/OFHEO NSA Purchase only monthly and All-Transactions quarterly indexes are from the Federal Housing Finance Agency website.⁴⁵

The 25 metropolitan areas in the FHFA/OFHEO Purchase Only Indexes tape are: Atlanta-Sandy Springs-Marietta, GA; Baltimore-Towson, MD; Chicago-Joliet-Naperville, IL (MSAD); Cleveland-Elyria-Mentor, OH; Dallas-Plano-Irving, TX (MSAD); Denver-Aurora-Broomfield, CO; Edison-New Brunswick, NJ (MSAD); Houston-Sugar Land-Baytown, TX; Los Angeles-Long Beach-Glendale, CA (MSAD); Miami-Miami Beach-Kendall, FL (MSAD); Minneapolis-St.Paul-Bloomington, MN-WI; Nassau-Suffolk, NY (MSAD); New York-White Plains-Wayne, NY-NJ (MSAD); Oakland-Fremont-Hayward, CA (MSAD); Philadelphia, PA (MSAD); Phoenix-Mesa-Glendale, AZ; Pittsburgh, PA; Riverside-San Bernardino-Ontario, CA; St. Louis, MO-IL; San Diego-Carlsbad-San Marcos, CA; Santa Ana-Anaheim-Irvine, CA (MSAD); Seattle-Bellevue-Everett, WA (MSAD); Tampa-St. Petersburg-Clearwater, FL; Warren-Troy-Farmington Hills, MI (MSAD); Washington-Arlington-Alexandria, DC-VA-MD-WV (MSAD).

Commercial: NCREIF indexes are from the NCREIF website.⁴⁶ The Transaction-Based Index (TBI) and the Moody's/REAL Commercial Property Price Index (CPPI) are from the MIT center for real estate.⁴⁷ REIT indexes are from the CRSP/Ziman All-Equity tape.

⁴³<http://www.census.gov/const/uspricemon.pdf>

⁴⁴http://www.macromarkets.com/csi_housing/sp_caseshiller.asp

⁴⁵<http://www.fhfa.gov/Default.aspx?Page=87>

⁴⁶<http://www.ncreif.org/tbi-returns.aspx>

⁴⁷<http://web.mit.edu/cre/>

A.2 Financial and Macro Variables

The aggregate stock market is the value-weighted NYSE/AMEX/NASDAQ index. The dividend-price ratio is constructed following Fama and French (1988b) as sum of current and previous 11-month (or 3 quarters, for quarterly data) dividends over the current price index. The 3-month rate is the Fama Risk-free rate, average of bid and ask. The inflation rate is the return to the CPI index. The Cochrane and Piazzesi (2005) factor is constructed using Fama-Bliss Discount Bond Yields. The source for all these series is the Center for Research in Security Prices (CRSP). Industrial Production is obtained from the Federal Reserve Bank of St. Louis.

References

- Abraham, Jesse M., and Patric H. Hendershott, 1996, Bubbles in Metropolitan Housing Markets, *Journal of Housing Research* 7, 191–207.
- Abraham, Jesse M., and William S. Schauman, 1991, New Evidence on Home Prices from Freddie Mac Repeat Sales, *Real Estate Economics* 19, 333–352.
- Ahearne, Allan, John Ammer, Brian Doyle, Linda Kole, and Robert Martin, 2005, House prices and monetary policy: A cross-country study, Board of Governors of the Federal Reserve System: International Finance Discussion Paper 841.
- Andrews, D.W.K, 1993, Exactly median-unbiased estimation of first order autoregressive unit root models, *Econometrica* 61, 139–165.
- Anselin, Luc, 1988, *Spatial Econometrics: Methods and Models* (Kluwer Academic Publishers: Dordrecht, The Netherlands).
- Aoki, Kosuke, James Proudman, and Gertjan Vlieghe, 2004, House prices, consumption, and monetary policy: a financial accelerator approach, *Journal of Financial Intermediation* 13, 414–435.
- Bailey, Martin J., Richard F. Muth, and Hugh O. Nourse, 1963, A Regression Method for Real Estate Price Index Construction, *Journal of the American Statistical Association* 58, pp. 933–942.
- Bansal, Ravi, and Hao Zhou, 2002, Term Structure of Interest Rates with Regime Shifts, *The Journal of Finance* 57, 1997–2043.
- Becker, T., and R. Shabani, 2010, Outstanding debt and the household portfolio, *Review of Financial Studies* 23, 2900–2934.
- Boz, Emine, and Enrique G. Mendoza, 2010, Financial innovation, the discovery of risk, and the u.s. credit crisis, NBER Working Paper 16020.
- Brooks, Chris, and Sotiris Tsolacos, 2001, Forecasting real estate returns using financial spreads, *Journal of Property Research* 18, 235–248.

- , 2003, International evidence on the predictability of returns to securitized real estate assets: econometric models versus neural networks, *Journal of Property Research* 20, 133–155.
- Brown, David T., 2000, Liquidity and Liquidation: Evidence from Real Estate Investment Trusts, *The Journal of Finance* 55, 469–485.
- Brown, Jane P., Haiyan Song, and Alan McGillivray, 1997, Forecasting UK house prices: A time varying coefficient approach, *Economic Modelling* 14, 529–548.
- Calhoun, Charles A., 1996, OFHEO House Price Indexes: HPI Technical Description, Discussion paper Office of Federal Housing Enterprise Oversight, Washington, D.C.
- Campbell, Davis, Morris Davis, Joshua Gallin, and Robert Martin, 2009, What Moves Housing Markets: A Variance Decomposition of the Price-Rent Ratio, *Journal of Urban Economics* 66, 90–102.
- Campbell, J, and J Cocco, 2007, How do house prices affect consumption? evidence from micro data, *Journal of Monetary Economics* 54, 591–621.
- Campbell, J. Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.
- Campbell, John Y., Stefano Giglio, and Parag Pathak, 2011, Forced Sales and House Prices, *American Economic Review* 101, 2108–2131.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press: Princeton).
- Campbell, John Y., and Robert J. Shiller, 1988, The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors, *Review of Financial Studies* 1, 195–228.
- Campbell, John Y., and Samuel B. Thompson, 2008, Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?, *Review of Financial Studies* 21, 1509–1531.
- Campbell, John Y., and Tuomo Vuolteenaho, 2004, Bad beta, good beta, *American Economic Review* 94, 1249–1275.

- Can, A., 1992, Specification and Estimation of Hedonic Housing Prices Models, *Regional Science and Urban Economics* 22, 453–474.
- Caplin, Andrew, Sumit Chopra, John Leahy, Yahn LeCun, and Trivikraman Thampy, 2008, Machine learning and the spatial structure of house prices and housing returns, Working paper.
- Capozza, Dennis R., Patric H. Hendershott, and Charlotte Mack, 2004, An Anatomy of Price Dynamics in Illiquid Markets: Analysis and Evidence from Local Housing Markets, *Real Estate Economics* 32, 1–32.
- Capozza, D. R., and P. Seguin, 1996, Expectations, efficiency, and euphoria in the housing market, *Regional Science and Urban Economics* 26, 369–386.
- Carlino, Gerard A., and Robert H. DeFina, 1998, The Differential Regional Effects of Monetary Policy, *Review of Economics and Statistics* 80, 572–587.
- , 1999, Do States Respond Differently to Changes in Monetary Policy?, *Business Review* July/August, 17–27.
- Case, Bradford, Henry Pollakowski, and Susan Wachter, 1991, On Choosing Among House Price Index methodologies, *AREUEA Journal* 19, 286–307.
- Case, Bradford, and John M. Quigley, 1991, The Dynamics of Real Estate Prices, *Review of Economics and Statistics* 73, 50–58.
- Case, Karl E., 2007, The Value of Land in the United States: 1975-2005, in Gregory Ingram, and Yu-Hung Hong, ed.: *Land Policies and Their Outcomes* (Lincoln Institute of Land Policy).
- , J. M. Quigley, and R. Shiller, 2005, Comparing wealth effects: The stock market versus the housing market, *Advances in Macroeconomics* 5, 1.
- , 2011, Wealth Effects Revisited 1978-2009, NBER Working Paper, No. 16848.
- Case, Karl E., and Robert J. Shiller, 1987, Prices of Single-Family Homes Since 1970: New Indexes for Four Cities, *New England Economic Review*, September/October, 46–56.
- , 1989, The Efficiency of the Market for Single-Family Homes, *The American Economic Review* 1, 125–137.

- , 1990, Forecasting Prices and Excess Returns in the Housing Market, *Real Estate Economics* 18, 253–273.
- Cauley, S., A. Pavlov, and E. Schwartz, 2007, Homeownership as a constraint on asset allocation, *Journal of Real Estate Finance and Economics* 34, 283–311.
- Cavanagh, C., G. Elliott, and J. Stock, 1995, Inference in models with nearly integrated regressors, *Econometric Theory* 11, 1131–1147.
- Clapham, Eric, Peter Englund, John M. Quigley, and Christian L. Redfearn, 2006, Revisiting the Past and Settling the Score: Index Revision for House Price Derivatives, *Real Estate Economics* 34, 275–302.
- Clapp, John M., 2004, A Semiparametric Method for Estimating Local House Price Indices, *Real Estate Economics* 32, 127–160.
- , and Carmelo Giaccotto, 1999, Revisions in Repeat-Sales Price Indexes: Here Today, Gone Tomorrow?, *Real Estate Economics* 27, 79–104.
- , 2002, Evaluating House Price Forecasts, *Journal of Real Estate Research* 24, 1–26.
- , and Dogan Tirtiroglu, 1991, Housing Price Indices Based on All Transactions Compared to Repeat Subsamples, *Real Estate Economics* 19, 270–285.
- Cochrane, John H., and Monika Piazzesi, 2005, Bond Risk Premia, *American Economic Review* 95, 138–160.
- Cooper, Michael, David H. Downs, and Gary A. Patterson, 1999, Real Estate Securities and a Filter-Based, Short-Term Trading Strategy., *Journal of Real Estate Research* 18, 313–334.
- Crawford, Gordon W., and Michael C. Fratantoni, 2003, Assessing the Forecasting Performance of Regime-Switching, ARIMA and GARCH Models of House Prices, *Real Estate Economics* 31, 223–243.
- Davis, Morris, Andreas Lehnert, and Robert F. Martin, 2008, The Rent-Price Ratio for the Aggregate Stock of Owner-Occupied Housing, *Review of Income and Wealth* 54, 279–284.

- Davis, Morris A., and Jonathan Heathcote, 2007, The price and quantity of residential land in the united states, *Journal of Monetary Economics* 54, 2595 – 2620.
- Del Negro, Marco, and Christopher Otrok, 2007, 99 Luftballons: Monetary policy and the house price boom across U.S. states, *Journal of Monetary Economics* 54, 1962–1985.
- Diebold, Francis X., and Glenn D. Rudebusch, 1999, *Business cycles: durations, dynamics, and forecasting* . chap. Measuring Business Cycles: A Modern Perspective, pp. 117–143 (Princeton University Press: Princeton, New Jersey).
- DiPasquale, Denise, and William C. Wheaton, 1994, Housing Market Dynamics and the Future of Housing Prices, *Journal of Urban Economics* 35, 1–27.
- Eichholtz, Piet, 1997, A long run house price index: The heregracht index, 1628-1973, *Real Estate Economics* 25, 175–192.
- Engel, Charles, and James D. Hamilton, 1990, Long Swings in the Dollar: Are They in the Data and Do Markets Know It?, *The American Economic Review* 80, 689–713.
- Engle, Robert, and Cliff Granger, 1987, Co-integration and error-correction: Representation, estimation, and testing, *Econometrica* 55, 251–276.
- Engle, R., D. Lilien, and M. Watson, 1985, A dymimic model of housing price determination, *Journal of Econometrics* 28, 307–326.
- Fama, Eugene, 1970, Efficient Capital Markets: A Review of Theory and Empirical Work, *Journal of Finance* 25, 383–417.
- , and Kenneth French, 1988a, Permanent and Temporary Components of Stock Prices, *The Journal of Political Economy* 96, 246–273.
- Fama, E., and G.W. Schwert, 1977, Asset returns and inflation, *Journal of Financial Economics* 5, 115–146.
- Fama, Eugene F., 1990, Stock returns, expected returns, and real activity, *The Journal of Finance* 45, pp. 1089–1108.

- , and Kenneth R. French, 1988b, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3 – 25.
- , 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23 – 49.
- Favilukis, Jank, David Kohn, Sidney Ludvigson, and Stijn Van Nieuwerburgh, 2012, *International Capital Flows and House Prices: Theory and Evidence* . chap. 8, p. forthcoming (Housing and the Financial Crisis, NBER Book Series).
- Favilukis, Jank, Sidney Ludvigson, and Stijn Van Nieuwerburgh, 2011, Macroeconomic implications of housing wealth, housing finance, and limited risk-sharing in general equilibrium, NBER working paper.
- Fisher, Jeffrey, 2005, US Commercial real estate indices: the NCREIF property index, Discussion paper BIS Papers No 21.
- Fisher, Jeff, David Geltner, and Henry Pollakowski, 2007, A Quarterly Transactions-based Index of Institutional Real Estate Investment Performance and Movements in Supply and Demand, *The Journal of Real Estate Finance and Economics* 34, 5–33.
- Flavin, Marjorie, and Takashi Yamashita, 2002, Owner-occupied housing and the composition of the household portfolio, *American Economic Review* 92, 345–362.
- Fratantoni, Michael, and Scott Schuh, 2003, Monetary Policy, Housing, and Heterogeneous Regional Markets, *Journal of Money, Credit and Banking* 35, 557–589.
- Fugazza, Carolina, Massimo Guidolin, and Giovanna Nicodano, 2007, Investing for the Long-run in European Real Estate, *The Journal of Real Estate Finance and Economics* 34, 35–80
10.1007/s11146-007-9002-5.
- Gallin, J., 2008, The long-run relationship between house prices and rents, *Real Estate Economics* 36, 635–68.
- Garcia, Ren, and Pierre Perron, 1996, An Analysis of the Real Interest Rate Under Regime Shifts, *Review of Economics and Statistics* 78, pp. 111–125.

- Gatzlaff, Dean H., and Donald R. Haurin, 1997, Sample Selection Bias and Repeat-Sales Index Estimates, *The Journal of Real Estate Finance and Economics* 14, 33–50.
- , 1998, Sample Selection and Biases in Local House Value Indices, *Journal of Urban Economics* 43, 199–222.
- Gau, George W., 1984, Weak Form Tests Of The Efficiency Of Real Estate Investment Markets, *Financial Review* 19, 301–320.
- , 1985, Public Information and Abnormal Returns in Real Estate Investment, *Real Estate Economics* 13, 15–31.
- Geltner, D., and J. Mei, 1995, The present value model with time-varying discount rates: Implications for commercial property valuation and investment decisions, *Journal of Real Estate Finance and Economics* 11, 119–135.
- Geltner, David, and Henry Pollakowski, 2007, A Set of Indexes for Trading Commercial Real Estate Based on the Real Capital Analytics Transaction Prices Database, MIT/CREDL.
- Geltner, David M., and Norman G. Miller, 2006, *Commercial Real Estate Analysis and Investments* (Cincinnati: South-Western Educational Publishing).
- Genesove, D., and Mayer, 1997, Equity and time to sale in the real estate market, *The American Economic Review* 87, 255–269.
- Giaccotto, Carmelo, and John Clapp, 1992, Appraisal-Based Real Estate Returns under Alternative Market Regimes, *Real Estate Economics* 20, 1–24.
- Glaeser, Edward L., Joseph Gyourko, and Albert Saiz, 2008, Housing Supply and Housing Bubbles, *Journal of Urban Economics* 64, 198 – 217.
- Goetzmann, William, 1992, The Accuracy of Real Estate Indices: Repeat Sale Estimators, *Journal of Real Estate Finance and Economics* 5, 5–53.
- Goetzmann, W., 1993, The single family home in the investment portfolio, *Journal of Real Estate Finance and Economics* 6, 201–222.

- Goetzmann, William N., and Matthew Spiegel, 1995, Non-Temporal Components of Residential Real Estate Appreciation, *Review of Economics and Statistics* 77, 199–206.
- Graff, Richard A., and Michael S. Young, 1997, Serial Persistence in REIT Returns, *Journal of Real Estate Research* 14, 183–214.
- Granger, Clive W.J., and Namwon Hyung, 2004, Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns, *Journal of Empirical Finance* 11, 399 – 421.
- Gray, Stephen F., 1996, Modeling the conditional distribution of interest rates as a regime-switching process, *Journal of Financial Economics* 42, 27–62.
- Gu, Anthony Y., 2002, The Predictability of House Prices, *Journal of Real Estate Research* 24, 213–233.
- Guidolin, Massimo, and Allan Timmermann, 2008, International asset allocation under regime switching, skew, and kurtosis preferences, *Review of Financial Studies* 21, 889–935.
- Guirguis, Hany S., Christos I. Giannikos, and Randy I. Anderson, 2005, The US Housing Market: Asset Pricing Forecasts Using Time Varying Coefficients, *The Journal of Real Estate Finance and Economics* 30, 33–53.
- Guntermann, Karl L., and Richard L. Smith, 1987, Efficiency of the Market for Residential Real Estate, *Land Economics* 63, pp. 34–45.
- Gyourko, Joseph, and Richard Voith, 1992, Local market and national components in house price appreciation, *Journal of Urban Economics* 32, 52 – 69.
- Halvorsen, Robert, and Henry O. Pollakowski, 1981, Choice of functional form for hedonic price equations, *Journal of Urban Economics* 10, 37–49.
- Hamilton, Bruce W., and Robert M. Schwab, 1985, Expected appreciation in urban housing markets, *Journal of Urban Economics* 18, 103–118.
- Hamilton, James D., 1989, A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle, *Econometrica* 57, pp. 357–384.

- , 1994, *Time Series Analysis* (Princeton University Press: Princeton, New Jersey).
- , 2008, Daily monetary policy shocks and new home sales, *Journal of Monetary Economics* 55, 1171–1190.
- Heckman, James, 1979, Sample Selection Bias as a Specification Error, *Econometrica* 47, 153–161.
- Hill, Robert, 2011, Hedonic Price Indexes for Housing, OECD Statistics Working Papers, 2011/01, OECD Publishing, available at <http://dx.doi.org/10.1787/5kghzxp6g6f-en>.
- Hill, R.Carter, C.F. Sirmans, and John R. Knight, 1997, Estimating Capital Asset Price Indexes, *Review of Economics and Statistics* 79, 226–233.
- , 1999, A random walk down main street?, *Regional Science and Urban Economics* 29, 89–103.
- Hill, Robert J., 2004, Constructing Price Indexes across Space and Time: The Case of the European Union, *The American Economic Review* 94, pp. 1379–1410.
- Himmelberg, C., C. Mayer, and T. Sinai, 2005, Assessing high house prices: bubbles, fundamentals, and misperceptions, *Journal of Economic Perspectives* 19, 67–92.
- Hoag, J. W., 1980, Towards indices of real estate value and return, *Journal of Finance* 35, 569–580.
- Hodrick, RJ, 1992, Dividend yields and expected stock returns: alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–386.
- Iacoviello, Matteo, 2005, House prices, borrowing constraints, and monetary policy in the business cycle, *American Economic Review* 95, 739–764.
- , and Stefano Neri, 2010, Housing market spillovers: Evidence from an estimated dsge model, *American Economic Journal: Macroeconomics* 2, 125–164.
- Kiyotaki, Nobuhiro, and John Moore, 1997, Credit cycles, *Journal of Political Economy* 105, 211–248.

- Koijen, Ralph, and Jules H. Van Binsbergen, 2010, Predictive regressions: A present-value approach, *Journal of Finance* 65, 1439–1471.
- Korteweg, Arthur, and Morten Sorensen, 2011, Estimating loan-to-value and foreclosure behavior, NBER Working Paper No. 17882.
- Kuo, Chiong-Long, 1996, Serial correlation and seasonality in the real estate market, *The Journal of Real Estate Finance and Economics* 12, 139–162.
- Lamont, Owen, and Jeremy Stein, 1999, Leverage and House-Price Dynamics in U.S. Cities, *RAND Journal of Economics* 30, 498–514.
- Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *The Journal of Finance* 56, 815–849.
- Lettau, Martin, and Stijn Van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies* 21, 1607–1652.
- Li, Yuming, and Ko Wang, 1995, The Predictability of REIT Returns and Market Segmentation, *Journal of Real Estate Research* 10, 471–482.
- Linneman, Peter, 1986, An empirical test of the efficiency of the housing market, *Journal of Urban Economics* 20, 140–154.
- , and Susan Wachter, 1989, The Impacts of Borrowing Constraints on Homeownership, *Real Estate Economics* 17, 389–402.
- Liu, Crocker H., and Jianping Mei, 1992, The predictability of returns on equity REITs and their co-movement with other assets, *The Journal of Real Estate Finance and Economics* 5, 401–418.
- Liu, C. H., and J. Mei, 1998, The predictability of international real estate markets, exchange rate risks, and diversification consequences, *Real Estate Economics* 26, 3–39.
- Lustig, Hanno N., and Stijn G. Van Nieuwerburgh, 2005, Housing collateral, consumption insurance, and risk premia: An empirical perspective, *The Journal of Finance* 60, 1167–1219.
- MacKinnon, Gregory H., and Ashraf Al Zaman, 2009, Real Estate for the Long Term: The Effect of Return Predictability on Long-Horizon Allocations, *Real Estate Economics* 37, 117–153.

- Malpezzi, S., 1990, Urban Housing and Financial Markets: Some International Comparisons, *Urban Studies* 27, 971–1022.
- Malpezzi, Stephen, 1999, A Simple Error Correction Model of House Prices, *Journal of Housing Economics* 8, 27–62.
- Mankiw, Gregory, and David Weil, 1989, The Baby Boom, the Baby Bust, and the Housing Market, *Regional Science and Urban Economics* 19, 235–258.
- McDonald, John F., 1999, Optimal leverage in real estate investment, *Journal of Real Estate Finance and Economics* 18, 239–252.
- McIntosh, Willard, and Glenn V. Henderson, 1989, Efficiency of the office properties market, *The Journal of Real Estate Finance and Economics* 2, 61–70.
- Meese, Richard, and Nancy Wallace, 1994, Testing the present value relation for housing prices: Should i leave my house in san francisco?, *Journal of Urban Economics* 35, 245 – 266.
- Meese, Richard A., and Nancy E. Wallace, 1997, The Construction of Residential Housing Price Indices: A Comparison of Repeat-Sales, Hedonic-Regression, and Hybrid Approaches, *The Journal of Real Estate Finance and Economics* 14, 51–73.
- Mei, Jianping, and Bin Gao, 1995, Price reversal, transaction costs, and arbitrage profits in the real estate securities market, *The Journal of Real Estate Finance and Economics* 11, 153–165.
- Mei, Jianping, and Crocker H. Liu, 1994, The predictability of real estate returns and market timing, *The Journal of Real Estate Finance and Economics* 8, 115–135.
- Merton, Robert, 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449–470.
- Mian, Atif, and Amir Sufi, 2009, The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis, *Quarterly Journal of Economics* 124, 1449–1496.
- , 2010, House Prices, Home-Equity-Based Borrowing, and the U.S. Household Leverage Crisis, *American Economic Review* forthcoming.

- Muellbauer, John, and Anthony Murphy, 1997, Booms and Busts in the UK Housing Market, *The Economic Journal* 107, pp. 1701–1727.
- Nappi-Choulet Pr., Ingrid, and Tristan-Pierre Maury, 2009, A Spatiotemporal Autoregressive Price Index for the Paris Office Property Market, *Real Estate Economics* 37, 305–340.
- Nelling, Edward, and Joseph Gyourko, 1998, The Predictability of Equity REIT Returns., *Journal of Real Estate Research* 16, 251–268.
- Newey, Whitney K., and Kenneth D. West, 1987, A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703–708.
- Ortalo-Magné, François, and Sven Rady, 2006, Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints, *Review of Economic Studies* 73, 459–485.
- Pace, R.Kelley, Ronald Barry, J. Clapp, and M. Rodriguez, 1998, Spatiotemporal Autoregressive Models of Neighborhood Effects, *Journal of Real Estate Finance and Economics* 17, 14–33.
- Pace, R.K., Ronald Barry, O. Gilley, and C.F. Sirmans, 2000, A method for Spatial-Temporal Forecasting With an Application to Real Estate Prices, *International Journal of Forecasting* 16, 229–246.
- Pace, R. Kelley, 1993, Nonparametric methods with applications to hedonic models, *The Journal of Real Estate Finance and Economics* 7, 185–204.
- Perez-Quiros, Gabriel, and Allan Timmermann, 2001, Business cycle asymmetries in stock returns: Evidence from higher order moments and conditional densities, *Journal of Econometrics* 103, 259 – 306.
- Pettenuzzo, D., A. Timmermann, and R. Valkanov, 2011, Return predictability under equilibrium constraints on the equity premium, UCSD working paper.
- Piazzesi, M., M. Schneider, and S. Tuzel, 2007, Housing, consumption and asset pricing, *Journal of Financial Economics* 83, 531–569.
- Plazzi, Alberto, Walter Torous, and Rossen Valkanov, 2010, Expected Returns and Expected Growth in Rents of Commercial Real Estate, *Review of Financial Studies* 23, 3469–3519.

- , 2011, Exploiting Property Characteristics in Commercial Real Estate Portfolio Allocation, *Journal of Portfolio Management* 35, 39–50.
- Prasad, Nalini, and Anthony Richards, 2008, Improving Median Housing Price Indexes through Stratification, *Journal of Real Estate Research* 30, 45–72.
- Quigley, J. M., 1995, A Simple Hybrid Model for Estimating Real Estate Price Indexes, *Journal of Housing Economics* 4, 1–12.
- Rapach, D., and G. Zhou, 2012, Forecasting Stock Returns, in A. Timmermann, and G. Elliott, ed.: *Handbook of Economic Forecasting*, vol. II. forthcoming.
- Rayburn, William, Michael Devaney, and Richard Evans, 1987, A Test of Weak-Form Efficiency in Residential Real Estate Returns, *Real Estate Economics* 15, 220–233.
- Reinhart, Carmen, and Kenneth Rogoff, 2009, *This Time is Different: Eight Centuries of Financial Folly* (Princeton University Press: Princeton, New Jersey).
- Rosen, Kenneth T., 1984, Toward a Model of the Office Building Sector, *Real Estate Economics* 12, 261–269.
- Rosen, Sherwin, 1974, Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition, *Journal of Political Economy* 82, 34–55.
- Ross, S., and R. Zisler, 1991, Risk and return in real estate, *Journal of Real Estate Finance and Economics* 4, 175–190.
- Rossi, B., 2012, *Handbook of Economic Forecasting*. chap. Advances in Forecasting under Instabilities (Elsevier-North Holland).
- Saiz, Albert, 2010, The Geographic Determinants of Housing Supply, *The Quarterly Journal of Economics* 125, 1253–1296.
- Schindler, Felix, 2011, Predictability and Persistence of the Price Movements of the S&P/Case-Shiller House Price Indices, *The Journal of Real Estate Finance and Economics* forthcoming, 1–47.

- Serrano, Camilo, and Martin Hoesli, 2007, Forecasting EREIT Returns., *Journal of Real Estate Portfolio Management* 13, 293–309.
- , 2010, Are Securitized Real Estate Returns more Predictable than Stock Returns?, *Journal of Real Estate Finance and Economics* 41, 170–192.
- Shiller, Robert, 2008, Derivative Markets for Home Prices, NBER w.p. 13962.
- Sirmans, G., Lynn MacDonald, David Macpherson, and Emily Zietz, 2006, The Value of Housing Characteristics: A Meta Analysis, *The Journal of Real Estate Finance and Economics* 33, 215–240.
- Skantz, Terrance R., and Thomas H. Strickland, 1987, House Prices and a Flood Event: An Empirical Investigation of Market Efficiency, *Journal of Real Estate Research* 2, 75–83.
- Spiegel, Matthew, 2001, Housing Return and Construction Cycles, *Real Estate Economics* 29, 521–551.
- Stambaugh, Robert, 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.
- Standard, and Poor's, 2008, S&P/Case-Shiller Home Price Indices Index Methodology, Discussion paper New York: Standard and Poor's, available at <http://www.standardandpoors.com/home/en/us>.
- Stein, Jeremy C., 1995, Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects, *The Quarterly Journal of Economics* 110, 379–406.
- Stevenson, Simon, 2002, Momentum Effects and Mean Reversion in Real Estate Securities., *Journal of Real Estate Research* 23, 47–64.
- Timmermann, A., 2006, *Handbook of Economic Forecasting* . chap. Forecast Combinations, pp. 135–178 (North-Holland).
- Valkanov, Rossen, 2003, Long-horizon regressions: Theoretical results and applications, *Journal of Financial Economics* 68, 201–232.
- Van Nieuwerburgh, Stijn, and Pierre-Olivier Weill, 2010, Why Has House Price Dispersion Gone up?, *Review of Economic Studies* 77, 1567–1606.

- Van Norden, Simon, and Robert Vigfusson, 1998, Avoiding the Pitfalls: Can regime-Switching Tests Reliability Detect Bubbles?, *Studies in Nonlinear Dynamics & Econometrics* 3, 1–22.
- Wallace, Nancy E., 1996, Hedonic-based price indexes for housing: theory, estimation, and index construction, *Economic Review* 3, 34–48.
- Webb, Cary, 1981a, A Discrete Random Walk Model for Price Levels in Real Estate, Working Paper, Department of Mathematics, Chicago State University.
- , 1981b, The Expected Accuracy of a Real Estate Price Index, Working Paper, Department of Mathematics, Chicago State University.
- , 1981c, Trading Activity and the Variance of an Index Number, Working Paper, Department of Mathematics, Chicago State University.
- Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.
- West, K.D., 2006, *Handbook of Economic Forecasting* . chap. Forecast Evaluation, pp. 101–133 (North-Holland).
- Wheaton, William C., and Raymond Torto, 1988, Vacany Rates and the Future of Office Rents, *AREUEA Journal* 19, 430–436.
- Zhong-Guo, Zhou, 1997, Forecasting Sales and Price for Existing Single-Family Homes: A VAR Model with Error Correction., *Journal of Real Estate Research* 14, 155.

Table 1: Summary of Literature

The table summarizes many of the papers that document predictability in real estate returns (residential, commercial, and REITs). In the first column, in parenthesis after the year of the paper, we report the market for which the analysis is conducted followed by the sample period. The conditioning variables used in each group of papers are listed in the middle column. The last column reports whether the results support strong (Y), weak (N), or mixed (Y/N) evidence of return predictability.

Papers	Conditioning Variables	Predictability evidence
Real Estate (Residential and Commercial)		
Gau (1984, 1985; Vancouver, 1971-1980); Linneman (1986; Philadelphia, 1975-1978); Guntermann and Smith (1987; 57 MSAs, 1968-1982); Rayburn, Devaney, and Evans (1987; Memphis, 1970-1984)	lagged returns/abnormal returns	Y/N (insufficient to cover transaction costs)
Case and Shiller (1989; 4 cities, 1970-1986); Hill, Sirmans, and Knight (1999; CS (1989) data); Kuo (1996, CS (1989) data); Schindler (2011; CS national and 20 MSAs; 1987-2009)	lagged returns/abnormal returns	Y
Capozza, Hendershott and Mack (2004; US metropolitan areas, 1979-1995)	geographic, demographic, and economic variables	Y/N (income, population growth, and construction costs are important factors)
Gyourko and Voith (1992; aggregate indices, 1971-1989); Glaeser, Gyourko, and Saiz (2008; aggregate indices, 1982-2007)	long-term mean reversion in prices	Y/N (less mean reversion in markets with more elastic housing supply)
Gu (2002; US indices, 1975-1999); Crawford and Fratantoni (2003; US indices, 1979-2001)	lagged returns/regime switching	Y/N (instability across locations and time periods)
McIntosh and Henderson (1989; Dallas-Forth Worth; 1979-1985)	lagged returns	N
Capozza and Seguin (1996; US indices, 1960-1990), Meese and Wallace (1994; county data, 1970-1988)	rent-price	Y
Campbell, Davis, Gallin, and Martin (2009; US indices, 1975-2007), Gallin (2008; US index, 1970-2005)		
Geltner and Mei (1995; US indices, 1975-1992)		
Plazzi, Torous, and Valkanov (2007, 21 MSAs, 1985-2000); Ghysels, Plazzi, and Valkanov (2010; 53 MSAs, 1994-2001)	rent-price, geographic, demographic, and economic variables	
Favilukis, Kohn, Ludvigson, and Van Nieuwerburgh (2012; US, 1991-2010); Mian and Sufi (2009; US, 2002-2009)	rent-price, credit supply	Y
Lamont and Stein (1999; US metropolitan areas, 1984-1994)	loan-to-value	Y
Malpezzi (1999; US metropolitan areas, 1979-1996)	price-to-income	Y
Abraham and Hendershott (1996, US indices, 1977-1992)	geographic, demographic, and economic variables	Y/N (mostly pronounced in areas with more elastic housing supply)
Zhong and Guo (1997; US indices, 1970-1994)	sales volume	Y
MacKinnon and Al Zaman (2009; US index, 1984-2007)	REIT returns, bond returns, rent-price	Y/N
REITs		
Liu and Mei (1992, 1994; US, 1972-1989);	dividend-price ratio and other macro variables	Y
Mei and Gao (1995; US, 1962-1990); Nelling and Gyourko (1998; US, 1975-1995);	returns	Y/N (insufficient to cover transaction costs)
Cooper, Downs, and Patterson (1999; US, 1973-1995); Serrano and Hoesli (2010; 10 countries, 1990-2007);	returns	Y
Stevenson (2002; 11 countries, 1977-2000); Graff and Young (1997; US, 1987-1996)		
Serrano and Hoesli (2007; US, 1978-2006)	returns to other asset classes	Y

Table 2: Summary Statistics

Annualized mean, annualized standard deviation, skewness, and first-order autoregressive coefficient for returns/price changes, rent/dividend growth, and rent/dividend-price ratio of aggregate real estate indices and conditioning variables. For residential real estate (Panel A), the indices are: the Census Median and Average; the Case-Shiller aggregate US, Composite 10 (C10), and equally-weighted average of available MSA indices (EW); the FHFA/OFHEO quarterly All-Transactions and monthly Purchase Only indices. Rent growth and the rent-price ratio for the Case-Shiller and OFHEO indices are from the Lincoln Institute. For commercial real estate (Panel B), the indices are from NCREIF, TBI, CPPI, and CRSP/Ziman REIT for the aggregate market (All) and separately for apartments (Apt), industrial properties (Ind), offices (Off), and retail properties (Rtl). In Panel C, the Financial and Macro Variables are the CRSP Value-Weighted NYSE/AMEX/NASDAQ index, the 3-month Treasury bill minus its twelve-month moving average (RTB), the return on the CPI index (CPI), the term spread as difference between the 5-year and 3-month yields (TSP), the Cochrane-Piazzesi (2005) interest rate factor (CP), and industrial production growth (IPG). Begin date reports the first return observation. The monthly (M) or quarterly (Q) frequency of each index is also denoted.

	Begin date	Returns/Price changes				Dividend/Rent Growth				Dividend/Rent-to-Price Ratio			
		Mean	Std	Skew	AR(1)	Mean	Std	Skew	AR(1)	Mean	Std	Skew	AR(1)
Panel A: Residential Real Estate													
Census													
Median (M)	Jan 1991	0.034	0.088	-0.425	-0.522	-	-	-	-	-	-	-	-
Average (M)	Jan 1991	0.032	0.080	0.136	-0.192	-	-	-	-	-	-	-	-
Case-Shiller													
US (Q)	1991:Q1	0.029	0.051	-1.362	0.613	0.030	0.007	-0.941	0.900	0.045	0.013	-0.989	0.987
C10 (M)	Mar 1991	0.034	0.033	-0.773	0.938	-	-	-	-	-	-	-	-
EW (M)	Mar 1991	0.029	0.029	-1.594	0.929	-	-	-	-	-	-	-	-
OFHEO													
US (Q)	1991:Q1	0.032	0.026	-0.911	0.756	0.030	0.007	-0.941	0.900	0.045	0.009	-0.765	0.994
US (M)	Feb 1991	0.031	0.020	-1.165	0.708	-	-	-	-	-	-	-	-
Panel B: Commercial Real Estate													
NCREIF													
All (Q)	1991:Q1	0.068	0.052	-1.888	0.801	-	-	-	-	-	-	-	-
Apt (Q)	1991:Q1	0.081	0.051	-2.402	0.821	-	-	-	-	-	-	-	-
Ind (Q)	1991:Q1	0.071	0.053	-1.747	0.809	-	-	-	-	-	-	-	-
Off (Q)	1991:Q1	0.060	0.065	-1.563	0.742	-	-	-	-	-	-	-	-
Rtl (Q)	1991:Q1	0.074	0.045	-0.938	0.750	-	-	-	-	-	-	-	-
TBI													
All (Q)	1991:Q1	0.090	0.096	-0.850	0.087	0.021	0.053	-0.388	0.365	0.056	0.024	-0.825	0.971
Apt (Q)	1994:Q2	0.098	0.090	-0.430	0.256	0.020	0.104	3.097	0.206	0.052	0.032	-0.218	0.980
Ind (Q)	1994:Q2	0.097	0.115	-0.463	0.053	0.018	0.106	-0.044	0.163	0.058	0.025	-0.310	0.948
Off (Q)	1994:Q2	0.095	0.090	-0.670	0.415	0.021	0.098	-0.446	0.569	0.047	0.027	-0.242	0.971
Rtl (Q)	1994:Q2	0.089	0.089	1.120	0.184	0.027	0.062	0.912	0.421	0.065	0.027	-0.340	0.967
CPPI													
All (Q)	Jan 2001	0.011	0.075	-1.354	0.452	-	-	-	-	-	-	-	-
Apt (Q)	2001:Q1	0.034	0.114	-1.324	0.193	-	-	-	-	-	-	-	-
Ind (Q)	2001:Q1	0.022	0.113	-1.793	0.256	-	-	-	-	-	-	-	-
Off (Q)	2001:Q1	0.022	0.109	-1.716	0.006	-	-	-	-	-	-	-	-
Rtl (Q)	2001:Q1	0.031	0.093	-1.249	0.376	-	-	-	-	-	-	-	-
REITs													
All (Q)	Jan 1980	0.104	0.179	-1.730	0.146	0.018	0.077	-1.700	-0.050	0.067	0.051	0.384	0.957
Apt (Q)	Jan 1980	0.114	0.190	-1.081	0.090	0.053	0.185	0.884	-0.090	0.072	0.083	0.452	0.970
Ind & Off (Q)	Jan 1980	0.066	0.225	-1.816	0.111	-0.039	0.191	-1.357	-0.099	0.071	0.076	1.173	0.940
Rtl (Q)	Jan 1980	0.121	0.200	-1.818	0.144	0.025	0.124	-4.494	-0.078	0.065	0.055	0.453	0.947
Panel C: Stock, Tbill, Inflation, Industrial Production													
CRSP VW (M)	Jan 1980	0.107	0.163	-1.096	0.114	0.050	0.047	2.159	0.082	0.027	0.038	0.580	0.992
RTB (M)	Jan 1980	-0.021	0.038	-0.046	0.885	-	-	-	-	-	-	-	-
CPI (M)	Jan 1980	0.034	0.012	-0.486	0.555	-	-	-	-	-	-	-	-
TSP (M)	Jan 1980	0.015	0.010	-0.410	0.924	-	-	-	-	-	-	-	-
CP (M)	Jan 1980	0.013	0.020	0.340	0.783	-	-	-	-	-	-	-	-
IPG (M)	Jan 1980	0.019	0.025	-1.096	0.294	-	-	-	-	-	-	-	-

Table 3: Correlation Matrix

Correlation matrix of quarterly returns to the Census Median, Case-Shiller Composite 10, FHFA/OFHEO US Monthly, NCREIF All, TBI All, and REIT All indices during the common 1991:Q2-2010:Q4 sample period.

	Census Median	CS10	FHFA/OFHEO	NCREIF	TBI	REIT
Census Median	1					
CS10	0.148	1				
FHFA/OFHEO	0.029	0.803	1			
NCREIF	0.080	0.425	0.335	1		
TBI	0.029	0.340	0.329	0.604	1	
REIT	0.334	0.374	0.181	0.280	0.220	1

Table 4: Forecasting Residential Real Estate

OLS slope estimates of the regression of returns on the Census Median, Case-Shiller Composite 10, and FHFA/OFHEO US Monthly indices on a constant (not reported) and the following lagged conditioning variables: the return and rent-price ratio of the index (r and hp respectively), the return and dividend-price ratio of the aggregate stock market (r_m and dp_m respectively), and the financial and macro variables as defined in Table 2. The rent-price ratio used for the Census Median in Panel B is the average between the Case-Shiller and FHFA/OFHEO rent-price ratios. In Panel A, the horizon is monthly from February 1991 until December 2010. In panel B, the horizon is quarterly from 1991:Q2 until 2010:Q4. In parenthesis below the estimates, Newey and West (1987) HAC t -statistics based on four lags are reported.

Panel A: Monthly												
	Census				Case-Shiller				FHFA/OFHEO			
r	-0.569			-0.573	0.938			0.863	0.725			0.593
	(-11.894)			(-12.211)	(26.222)			(21.141)	(10.592)			(7.097)
r_m	0.058			0.061	0.017			0.007	0.017			0.013
	(1.067)			(1.203)	(0.885)			(1.063)	(1.372)			(1.755)
dp_m	-0.005			-0.010	-0.016			-0.003	-0.007			-0.004
	(-1.115)			(-1.393)	(-5.542)			(-1.592)	(-5.182)			(-3.395)
RTB				-0.055	-0.073			0.491	0.041			0.234
				(-0.227)	(-0.300)			(2.876)	(0.990)			(3.210)
CPI				0.013	0.003			0.259	0.200			0.237
				(0.019)	(0.004)			(1.431)	(3.554)			(1.911)
TSP				0.224	0.359			-0.118	0.062			-0.163
				(0.788)	(1.032)			(-0.819)	(1.399)			(-2.080)
CP				-0.153	-0.137			0.251	0.029			0.186
				(-0.691)	(-0.550)			(2.801)	(0.880)			(3.788)
IPG				0.370	0.464			0.076	-0.029			-0.081
				(0.768)	(1.279)			(0.717)	(-0.535)			(-1.058)
R^2	0.317	0.005	0.004	0.332	0.881	0.250	0.279	0.893	0.501	0.136	0.223	0.543

Panel B: Quarterly												
	Census				Case-Shiller				FHFA/OFHEO			
hp	0.011	0.008		-0.005	0.000	0.017		0.015	0.006	0.009		0.000
	(0.379)	(0.255)		(-0.153)	(-0.007)	(1.342)		(0.941)	(0.239)	(0.883)		(-0.015)
r				-0.523				0.480				0.535
				(-7.262)				(9.834)				(3.267)
dp_m				-0.008				-0.034				-0.012
				(-0.769)				(-6.521)				(-1.805)
r_m				0.005				0.008				-0.012
				(0.070)				(1.211)				(-0.710)
RTB				-0.050				1.460				0.747
				(-0.093)				(1.882)				(2.417)
CPI				0.125				0.579				-0.121
				(0.121)				(1.996)				(-0.645)
TSP				0.001				-0.417				-0.624
				(0.002)				(-0.651)				(-1.876)
CP				0.011				0.709				0.567
				(0.029)				(1.525)				(2.640)
IPG				0.065				-0.035				-0.067
				(0.306)				(-1.043)				(-0.423)
R^2	0.001	0.273	0.002	0.001	0.334	0.000	0.559	0.255	0.268	0.670	0.002	0.409

Table 5: Forecasting Commercial Real Estate

OLS slope estimates of the regression of returns on the NCREIF All, TBI All, and REIT All indices on a constant (not reported) and lagged conditioning variables. Variables definition follows from Table 4. The horizon is quarterly from 1991:Q2 until 2010:Q4. In parenthesis below the estimates, Newey and West (1987) HAC t -statistics based on four lags are reported.

	NCREIF					TBI					REITs				
hp/dp	-0.003	0.011			-0.007	0.031	0.031			0.003	0.070	0.109			0.136
	(-0.121)	(0.985)			(-1.139)	(0.917)	(0.964)			(0.123)	(1.494)	(1.532)			(1.620)
r		0.819			0.363		0.057			-0.254		0.257			0.016
		(13.330)			(2.066)		(0.460)			(-1.871)		(1.727)			(0.113)
r_m			-0.053		-0.026		-0.035			-0.031			0.016		0.025
			(-3.553)		(-2.057)		(-1.545)			(-1.236)			(0.405)		(0.506)
dp_m			0.059		0.005		0.093			0.000			0.365		0.260
			(1.328)		(0.397)		(1.137)			(0.002)			(1.428)		(1.229)
RTB			0.698		0.570				1.098	1.481				1.630	0.729
			(2.076)		(2.405)				(1.512)	(1.793)				(1.231)	(0.518)
CPI			0.630		0.357				0.509	0.551				3.521	3.885
			(2.407)		(2.033)				(1.009)	(1.032)				(1.486)	(1.544)
TSP			-1.265		-0.309				-0.640	-0.658				0.073	0.918
			(-2.635)		(-0.870)				(-0.857)	(-0.997)				(0.033)	(0.444)
CP			0.099		0.073				0.193	0.334				0.748	-0.850
			(0.282)		(0.391)				(0.348)	(0.534)				(0.640)	(-0.573)
IPG			0.759		0.499				1.116	1.182				1.262	1.748
			(3.101)		(1.826)				(2.643)	(2.848)				(0.950)	(1.228)
R^2	0.001	0.650	0.371	0.588	0.767	0.024	0.027	0.074	0.253	0.317	0.022	0.082	0.088	0.179	0.270

Table 6: Forecasting REITs

Slope coefficients of the regressions of T -month real returns (Panel A) and real dividend growth (Panel B) on the aggregate CRSP/Ziman All-Equity REIT index (All) and separately for apartments (Apt), industrial properties and offices (Ind & Off), and retail properties (Rtl) on a constant (not reported) and lagged log dividend-price ratio. OLS results are shown for separate monthly horizons. The table also reports the one-month two-stage GMM estimates which impose the present value constraint (equation 16) across equations and the short-long horizon relationship (equation 15) across horizons $T = \{1, 3, 6, 12\}$. The coefficients are bias-adjusted following Stambaugh (1999). The t -statistics are Newey and West (1987) based on four lags for non-overlapping and T lags for overlapping returns. The sample is monthly from December 1980 until December 2010.

Panel A: Real Returns												
Estimation, Horizon	All			Apt			Ind & Off			Rtl		
	β_r	$t(\beta_r)$	R^2									
Overlapping												
OLS, $T=1$	0.001	0.035	0.002	0.000	0.013	0.001	-0.012	-0.791	0.001	0.003	0.189	0.002
OLS, $T=3$	0.031	0.820	0.009	0.022	0.967	0.010	-0.020	-0.547	0.001	0.039	0.998	0.012
OLS, $T=6$	0.117	1.490	0.039	0.080	1.652	0.040	0.000	-0.005	0.000	0.129	1.660	0.045
OLS, $T=12$	0.280	1.684	0.092	0.200	1.968	0.110	0.031	0.182	0.002	0.321	1.872	0.119
GMM, $T = \{1, 3, 6, 12\}$	0.013	1.126	-	0.010	1.421	-	-0.004	-0.328	-	0.017	1.435	-
Non-Overlapping												
OLS, $T=1$	0.001	0.041	0.002	0.000	0.014	0.001	-0.012	-0.841	0.001	0.003	0.235	0.002
OLS, $T=3$	0.014	0.345	0.009	0.009	0.338	0.009	-0.031	-0.865	0.001	0.026	0.614	0.014
OLS, $T=6$	0.024	0.309	0.013	0.023	0.508	0.017	-0.080	-1.139	0.004	0.041	0.520	0.019
OLS, $T=12$	0.062	0.504	0.027	0.057	0.671	0.037	-0.112	-0.644	0.003	0.126	0.837	0.049
GMM, $T = \{1, 3, 6, 12\}$	0.009	1.397	-	0.003	1.607	-	-0.001	-0.128	-	0.008	1.410	-
Panel B: Real Dividend Growth												
Estimation, Horizon	All			Apt			Ind & Off			Rtl		
	β_d	$t(\beta_d)$	R^2									
Overlapping												
OLS, $T=1$	-0.015	-3.232	0.029	-0.026	-2.266	0.042	-0.046	-4.125	0.071	-0.020	-2.298	0.026
OLS, $T=3$	-0.051	-3.465	0.094	-0.081	-2.912	0.118	-0.125	-5.628	0.212	-0.063	-2.765	0.078
OLS, $T=6$	-0.119	-3.361	0.162	-0.182	-3.101	0.201	-0.255	-5.231	0.325	-0.137	-2.671	0.139
OLS, $T=12$	-0.261	-3.551	0.230	-0.384	-3.483	0.320	-0.541	-4.994	0.447	-0.263	-2.458	0.193
GMM, $T = \{1, 3, 6, 12\}$	-0.018	-9.716	-	-0.031	-7.156	-	-0.045	-9.366	-	-0.025	-6.828	-
Non-Overlapping												
OLS, $T=1$	-0.015	-2.753	0.029	-0.026	-2.097	0.042	-0.046	-4.062	0.071	-0.020	-2.210	0.026
OLS, $T=3$	-0.049	-3.063	0.095	-0.082	-2.002	0.145	-0.127	-4.198	0.210	-0.059	-2.149	0.079
OLS, $T=6$	-0.120	-3.586	0.170	-0.205	-2.514	0.280	-0.271	-4.760	0.388	-0.154	-2.228	0.173
OLS, $T=12$	-0.230	-4.629	0.213	-0.396	-3.354	0.417	-0.488	-4.547	0.525	-0.295	-2.367	0.231
GMM, $T = \{1, 3, 6, 12\}$	-0.019	-6.696	-	-0.035	-5.677	-	-0.055	-9.179	-	-0.023	-3.227	-

Table 7: Forecasting REITs In-Sample and Out-Of-Sample

In-sample (IS) and Out-Of-Sample (OOS) R^2 , as defined in equation (17), for the predictive regression of real returns and real dividend growth to the aggregate CRSP/Ziman All-Equity REIT index at the monthly (Panel A), quarterly (Panel B), and semi-annual (Panel C) frequency. The forecasters are defined as in Table 5. “Combined” refers to the average forecast across all forecasters. Specification ‘unc’ is the unconstrained regression; specification “ $\text{sign}(\beta_r)$ ” and “ $r > 0$ ” replace the forecast with the unconditional mean when the slope coefficient or the forecast, respectively, are negative; specification “both” imposes both constraints. The OOS results are based on a 180-month burn-in period. The full sample is monthly observations from December 1980 until December 2010.

Forecaster	Real returns					Real Dividend-Growth	
	IS	OOS unc	OOS $\text{sign}(\beta_r)$	OOS $\hat{r} > 0$	OOS both	IS	OOS unc
Panel A: Monthly							
dp	0.181	-0.686	-0.571	-0.321	-0.205	2.866	-1.873
r_m	4.760	4.022	4.022	1.885	1.885	0.572	0.985
dp_m	0.005	-0.637	-0.572	-0.324	-0.260	0.008	-1.160
RTB	0.071	-2.005	0.000	-1.565	0.000	0.058	-2.015
CPI	0.641	-3.105	0.000	-3.075	0.000	0.047	-1.846
TSP	0.313	-0.488	-0.488	-0.406	-0.406	0.559	-10.616
CP	0.250	-0.385	-0.318	0.211	0.278	1.739	-4.108
IP	2.805	-0.227	0.327	-0.985	-0.411	2.320	-1.920
Combined	2.093	0.187	0.620	-0.226	0.241	6.862	0.291
Panel B: Quarterly							
dp	0.919	-6.189	-6.160	-4.337	-4.308	9.366	-15.520
r_m	0.470	1.463	2.218	0.770	1.526	0.230	-0.319
dp_m	0.121	-2.406	-2.406	-1.284	-1.284	0.002	-1.442
RTB	0.031	-3.792	0.000	-3.704	0.000	0.258	-0.290
CPI	0.690	-6.142	-0.445	-6.373	-0.445	0.268	0.297
TSP	2.156	-1.007	-1.007	0.749	0.749	0.602	-15.706
CP	0.798	0.638	0.649	0.638	0.649	2.536	-3.415
IP	2.240	1.059	2.156	-0.162	1.150	3.116	-3.049
Combined	2.093	-1.319	-0.166	-1.270	-0.045	6.862	-0.266
Panel C: Semi-Annual							
dp	3.862	-15.211	-15.185	-9.784	-9.757	16.217	-40.390
r_m	0.480	0.377	0.611	0.377	0.611	0.016	-1.645
dp_m	0.450	-3.825	-3.825	-2.314	-2.314	0.002	-2.137
RTB	0.001	-4.650	0.000	-4.650	0.000	0.526	-0.418
CPI	0.001	-3.694	-0.017	-3.694	-0.017	0.658	0.912
TSP	2.955	-0.984	-0.984	-0.522	-0.522	1.076	-19.220
CP	0.193	1.216	1.095	1.216	1.095	5.412	-5.349
IP	1.240	-1.267	2.143	-1.888	1.469	5.580	-5.200
Combined	2.093	-2.625	-1.340	-2.164	-0.890	6.862	-2.492

Figure 1: Residential Real Estate Indices

Time series plot of the Census Median, Case-Shiller Composite 10, and OFHEO. All series are sampled quarterly and normalized at one in 1991:Q1.




Figure 2: Commercial Real Estate Indices

Time series plot of the NCREIF All, TBI All, and CRSP/Ziman REIT All indices. All series are sampled quarterly and normalized at one in 1984:Q4.




Figure 3: Long-Horizon Return Autocorrelations

The thick solid line reports the OLS slope coefficients $\beta(T)$ in the regression of T -month overlapping returns on a constant and their T -period lagged value, $r_{t+1:t+T} = \alpha + \beta(T) r_{t-T+1:t} + \epsilon_{t+1:t+T}$. The dashed lines denote the point estimate plus or minus two Newey and West (1987) HAC standard errors based on T lags. The plots are (in clockwise order, starting from top left) for the Census Median index quarterly sampled, FHFA/OFHEO Quarterly series, REIT All quarterly sampled, and TBI All. The X-axis reports the horizon T in months. The sample period is as in Table 2.




Figure 4: Predicting Residential Real Estate MSA Returns

Summary of the predictive regression of each of the 25 MSAs OFHEO Purchase Only quarterly return series on a constant, its own (lagged) return, the aggregate OFHEO rent-to-price ratio, and macro and financial variables as defined in Table 4. The regressors are all jointly included as in the largest specification of Tables 4 and 5. For each regressor (except the constant) across MSA, the Figure shows in the top left panel the average slope coefficient, in the top right panel the average t -statistic, and in the bottom left panel the number of t -statistics greater than two in absolute value. The bottom right panel displays the R^2 for each regression. The list of MSAs is reported in Appendix A.1. The sample is quarterly observations from 1991:Q2 until 2010:Q4.




Figure 5: Principal Component Analysis by MSA

The top plot reports the percentage explained by the first 10 principal components extracted from the covariance matrix of returns to the 25 MSAs OFHEO Purchase Only quarterly series. The bottom plot shows the percentage explained by the first component based on 40-quarter rolling windows. The full sample is quarterly observations from 1991:Q2 until 2010:Q4.

