

10-28-2014

Using Cash Flow Dynamics to Price Thinly Traded Assets: The Case of Commercial Real Estate

Walter I. Boudry

Cornell University School of Hotel Administration, wb242@cornell.edu

Crocker H. Liu

Cornell University School of Hotel Administration, chl62@cornell.edu

Tobias Mühlhofer

University of Texas

Walter N. Torous

Massachusetts Institute of Technology

Follow this and additional works at: <http://scholarship.sha.cornell.edu/workingpapers>

 Part of the [Real Estate Commons](#)

Recommended Citation

Boudry, W. I., Liu, C. H., Mühlhofer, T., & Torous, W. N. (2014). *Using cash flow dynamics to price thinly traded assets: The case of commercial real estate*[Electronic version]. Retrieved [insert date], from Cornell University School of Hotel Administration site: <http://scholarship.sha.cornell.edu/workingpapers/15>

This Working Paper is brought to you for free and open access by the School of Hotel Administration Collection at The Scholarly Commons. It has been accepted for inclusion in Working Papers by an authorized administrator of The Scholarly Commons. For more information, please contact hlmdigital@cornell.edu.

Using Cash Flow Dynamics to Price Thinly Traded Assets: The Case of Commercial Real Estate

Abstract

Previous studies of share repurchase have primarily focused on examining announcement effects and long-term operating performance in order to distinguish among the diverse possible hypotheses for repurchase. One of the most important rationales they have studied is the over-investment hypothesis: firms repurchase in order to avoid investing in negative net present value projects. While the recent empirical analyses have presented some indirect evidence in support of the over-investment hypothesis, this study examines this rationale for repurchase from a unique perspective, empirically showing that project returns have an important influence on the decision to repurchase shares. Our sample of firms consists of 125 real estate investment trusts (REITs) in order to utilize a time series of real estate capitalization rates (property ROAs) from market transactions on different property types. These cap rates proxy for a REIT's project opportunity set. Using a both Logit and Tobit models that corrects for other possible buyback rationales, we show that during periods with relatively low cap rates, REITs are more likely to both repurchase shares and repurchase larger amounts of shares than when cap rates are high.

Keywords

thinly traded assets, asset pricing, panel vector autoregression, commercial real estate

Disciplines

Real Estate

Comments

Required Publisher Statement

Copyright held by the authors.

Using Cash Flow Dynamics to Price Thinly Traded Assets: The Case of Commercial Real Estate

Walter I. Boudry
School of Hotel Administration
Cornell University

Crocker H. Liu
School of Hotel Administration
Cornell University

Tobias Mühlhofer*
McCombs School of Business
University of Texas

Walter N. Torous
MIT

October 28, 2014

Abstract

We propose a technique to infer cash flow yields for investment assets whose trades are infrequent, but for which cash flow data is available. We construct a *Self-Propagating* Rolling-Window Panel VAR framework, adapted from a Dynamic Gordon Growth Model setup. We use this framework to estimate yields and volatility in yields for untraded commercial properties as out-of-sample predictions from our VAR based on these properties' cash flow data. We find that our predicted cash flow yields closely resemble ex-post realized transaction yields, and that these predicted yields even outperform appraisals in this respect. We find that this paradigm provides a good representation of commercial real estate yields, and propose that investors can readily apply this algorithm to infer values of untraded investment assets.

Keywords: Thinly traded assets, asset pricing, panel vector autoregression, commercial real estate.
JEL Codes: G12, R33

*The emails of the authors are: wb242@cornell.edu, chl62@cornell.edu, tobias.muhlhofer@austin.utexas.edu, wtorous@mit.edu respectively. Contact author: Tobias Mühlhofer, McCombs School of Business, University of Texas at Austin, Department of Finance, 1 University Station B6600, Austin, TX 78712. Telephone: 512-850-5710, E-mail: tobias.muhlhofer@austin.utexas.edu. We wish to thank the National Council of Real Estate Investment Fiduciaries and Bob White of Real Capital Analytics for the use of their data.

1 Introduction

This study focuses on the measurement of returns and risk in a market where assets frequently suffer from thin trading. This issue is important since thin trading can distort return and risk measurement leading to incorrect conclusions about the pricing and thus the performance of these assets. We define a thin market as one in which a limited number of buyers and sellers exist. Since there are fewer transactions, assets are less liquid and prices exhibit higher volatility, with larger spreads typically occurring between bids and asks.

Even though a thin asset market exists, the underlying price discovery process is not necessarily ad-hoc in nature if the asset in question generates regular cash flows. In this case, one can explicitly model the price formation process using the asset's fundamental valuation equation. Using this logic, price volatility and, as a result, return volatility can be assessed vis-à-vis variations in the observable cash flows. This result is based on the premise that cash flow fundamentals are the most important source of pricing an asset and its risk.

We contribute to the literature by developing a technique to infer cash flow yields for assets that are not continuously traded, but for which continuous cash flow data exists. This cash flow yield can then be applied to cash flow levels for inferring prices directly. The empirical setting in which we develop this technique is Commercial Real Estate, where thin trades are common, but continuous cash flow data through rents exists. The starting point for the methodology we develop is the Dynamic Gordon Growth Model, in the style of Shiller (1992) and Campbell and Shiller (1988b). We apply a version of the Campbell-Shiller Vector Autoregression (VAR) to a large panel dataset of commercial-property cash-flow data. We develop a rolling-window panel VAR technique to estimate cash flow yields for untraded commercial properties, as out-of-sample predictions from this VAR. We then use these predicted yields in the coefficient estimation for the next iteration of the rolling VAR, joining them with new cash-flow data (and the small number of transactions available) to generate the next set of estimated yields, which we then re-use in the subsequent iteration. By this method, we estimate close to 200,000 property yields across a panel dataset. Since each subsequent VAR iteration uses the predictions generated in the previous run, we term this procedure a *Self-Propagating* Rolling-Window Panel VAR. We find that our predicted

yields explain between 75% and 93% of variation in ex-post realized transaction yields. In contrast, appraisal-based yields explain only half to two-thirds as high a fraction of this variation as our predictions.

Prior finance studies have focused on adjusting systematic risk to account for thin trading. More specifically, when shares are thinly traded, various authors have proposed downward corrections in Beta estimates for this bias (see, for example, Cohen, Hawawini, Maier, Schwartz and Whitcomb (1983), Dimson (1979), Scholes and Williams (1977)). In the current study we use a different approach, focusing on cash flow dynamics to account for firm (property) risk. This is not a new concept, although, as far as we are aware, so far this approach has not been used to explicitly infer cash flow yields of untraded investment assets. In the literature so far, for example, Da and Warachka (2009) show that changes in expected cash flows partially drive the cross-sectional variation in stock returns using an analyst's earnings beta which they show accounts in part for the value premium, size premium, and long-term return reversals. More recently, Driessen, Lin and Phalippou (2011) use cross-sectional cash flows to estimate abnormal performance and risk exposure of non-traded assets. A related strand of literature concentrates on earning betas. Beaver, Kettler and Scholes (1970), Beaver and Manegold (1975), Gonedes (1975), Ismail and Kim (1989) show that accounting information including accounting earnings beta and cash flow beta are related to a firm's beta. Mandelker and Rhee (1984) further find that the degrees of operating and financial leverage explain a large portion of the variation in beta. Thus firm risk (beta) varies the higher the fixed operating costs and fixed financial costs in addition to the volatility in earnings and cash flows.

The genesis for using the variation in cash flow fundamentals to assess pricing risk stems from the log-linear dividend ratio model of Campbell and Shiller (1988b), Campbell (1991), and Shiller (1992). The model characterizes the relation between asset prices in the next period and changes in rational expectations of future dividend growth and future asset returns. The model thus allows both expected future cash flows and expected returns (discount rates) to influence asset prices. Kallberg, Liu and Srinivasan (2003) as well as Mühlhofer and Ukhov (2011) find that this model is consistent with REIT pricing. If the cash flows for REITs are similar to those for underlying

properties, there is a good chance that the model would also hold for direct real estate. Showing this is one research objective of our study. Plazzi, Torous and Valkanov (2010) use a complementary approach using the dynamics of commercial property cap rates rather than dividend growth to study risk.

We apply our methodology to a large panel data set of commercial-property rental cash flows and, where available, transaction prices.¹ We proceed by first establishing the feasibility of our technique in a single time-series rather than a panel, using index data for the Los Angeles Office market as a test case, to establish that properties are, in fact, priced according to a Dynamic Gordon Growth Model paradigm. We find that on a single time series, our rolling VAR generates predicted yields that closely resemble ex-post realized yields. Having established this, we then examine the feasibility of the procedure in a panel setting, using appraisal-based yields to provide continuously-populated data and thus a more controlled environment. In this setting we closely examine the VAR coefficients and find that they match what would be expected if the formation of yields follows a Dynamic Gordon Growth Model. In this setting, too, we find that the predicted yields from a rolling panel VAR resemble appraisals. Finally, having established the feasibility of our approach in these two controlled settings, we proceed to developing and applying our *Self-Propagating* Rolling-Window Panel VAR procedure and test its performance vis-a-vis appraisals and realized transaction cap rates, finding, as stated above, that the predicted cap rates from this model explain a substantial amount of the variation in transaction cap rates, and outperform appraisals in this respect. We further explore whether a distinction in the price formation process exists between the most liquid submarkets and the rest of the market, and find against this hypothesis. The price formation process that we model through our procedure seems universal throughout this universe.

While we apply this methodology to Commercial Real Estate data, this procedure should be applicable to any investment asset which is thinly traded, but which has consistent cash flow information. Other examples of such asset classes might include natural-resource extraction sites (such as mines or oil and gas wells), as well as thinly-traded fixed-income securities with variable cash flows, such as municipal revenue-based bonds. The reason why we choose Commercial Real

¹This dataset comes from the National Council of Real Estate Investment Fiduciaries (NCREIF) and tracks a property portfolio worth \$344 billion and comprising over 7,000 properties, as of the end of 2013.

Estate to develop and demonstrate this procedure is in large part an econometric one. In this setting we have reliable cash flow data, as well as enough transactions to monitor the predictive ability of our technique. At the same time, transactions are still few enough that this can be considered a very thinly traded market, of the type for which this procedure is designed. Within this market, we also have some heterogeneity in the amount of cash flow volatility, due to different lease lengths in different sectors (Apartments tend to be year to year, while especially Industrial leases on purpose-built property can be as long as ten years or more). This heterogeneity allows us to make a statement on the relative importance of informative cash flow variation as opposed to cash flow stickiness for the success of this procedure.

We further test the importance of informative cash flow variation, by applying our procedure to the stock market, in which, as is well documented, we should expect cash-flow (i.e. dividend) variation to be less informative in the price-formation process. We apply our Rolling-Window Panel VAR to the panel of stocks in the CRSP universe, and, as expected, find dramatically reduced explanatory power of our procedure, in predicting dividend yields. This result strongly supports the hypothesis that informativeness of cash-flow variation plays an important role in the effectiveness of a Dynamic Gordon Growth Model in capturing asset-price movements.

The rest of this study proceeds as follows. Section 2 builds our methodology; Section 3 describes our data in detail; Section 4 presents our results; Section 5 concludes.

2 Methodology

2.1 Underlying Theory

While commercial property is thinly traded and therefore pricing observations for this asset class are difficult to find, the income cash flows generated by commercial real estate are readily observable. The methodology we propose for this study is based upon the use of this cash flow data and an analysis of a local sub-market's particular price formation process, in order to infer prices for properties in periods when they are not traded.

Specifically, like any investment asset, commercial property is priced by an investor as the

present discounted value of all future cash flows the asset entitles him to. Given the long life of commercial property, basic finance theory tells us that such an investment can be priced as an infinite series of cash flows, which leads to reduced form expressions such as the Gordon Growth Model:

$$P_t = \frac{CF_{t+1}}{r - g} \quad (1)$$

Here, P_t is the price of the property at time t , CF_{t+1} is the income cash flow produced by the property over the next period, r is the risk-adjusted discount rate, and g is the expected cash-flow growth rate. This well known expression shows the role that the real estate capital market plays in the price formation process. In particular, the capital market performs a risk assessment and growth projection for a given cash flow series in order to assign a price to a property. The observed cap rate (the denominator of Equation 1) represents the ex-post result of this price formation process. Consequently, if the price formation process is known ex-ante, or alternatively if one can infer what cap rate is used for a particular submarket at a given point in time, then we can infer property prices, even when a property is not traded.

In order to infer cap rates, it is necessary to understand how both r and g are determined. The finance literature (for example Shiller (1992), or Campbell and Shiller (1988a)) argues that, given the persistence in the cash flows of many investment assets, cash flow growth expectations can be modeled by analyzing the time series of past cash flow growth rates. The market's inferred growth expectations can then be written as a function of this time series. The discount factor, r is made up of the risk-free rate plus a risk premium. The risk premium, in turn, is a function of expected cash flow risk. Relying on the same arguments as above, this risk should be persistent, and so the market's expectations of cash flow risk (and therefore the risk premium) can be modeled as a function of past cash flow volatility levels, i.e. simply the squares of past changes in cash flow. This should especially be the case for property types with long -term contractual leases such as office buildings.

Campbell and Shiller (1988a) and Shiller (1992) conduct a log-linearization of a dynamic version

of Equation 1, and show that this can be estimated through a Vector-Autoregression (VAR). Vector autoregressions are a useful and flexible way of analyzing economic relations in time series data. More specifically, a VAR allows for the mutual impact of the variables and is thus well suited for inter-dependent economic time series. In other words, the technique is useful in examining complex relationships among variables when the variables are serially correlated. Typically, VARs have little serial correlation in the residuals. This is helpful for separating out the effects of economically unrelated influences in the VAR. All variables in a VAR are treated equally by including for each variable, an equation explaining its evolution based on its own lags and the lags of all other variables in the model. Thus, the VAR recognizes that variables can have an impact on other variables. VAR generalizes easily to more variables and more lags of variables. In the current setting, we use the VAR model to reveal the evolution of the yield, long term interest rates, net-operating-income growth as well as the dynamic interactions between these variables. For the purposes of our study, the parameter matrix of this VAR would show exactly how a particular submarket uses past cash flow information in order to construct cap rates and therefore prices. In principle, we would use properties for which we have both cash flows and realized prices to *train* these VARs for each market. Having done this, we can infer cap rates and prices for properties that are not traded by inserting their cash flows into the VAR and simply calculating fitted values.²

Given an observable series of underlying cash flows, the task that must be accomplished in order to infer the price of an untraded asset is to model the cash-flow yield; in commercial property this is the capitalization rate, while in equity securities this is the dividend yield. We therefore use as our starting point the log-linearization of the Gordon Growth Model of Campbell and Shiller (1988a) and Shiller (1992), who formulate an estimable dynamic version of the GGM as follows:

$$\delta_t = \sum_{j=1}^{\infty} \rho^j E_t [r_{t+j} - \Delta d_{t+j}] + C \quad (2)$$

In the above notation, δ is the log cash flow yield, ρ is the log of the time-varying risk premium, r is the log of the risk-free interest rate, and Δd is the growth rate of log cash flows. The above

²It should be clear that in this procedure, the properties used to *train* the VAR are the same properties on which the predictions are made, as the training happens in the earlier time periods of the same sample.

equation can be estimated through a VAR system that contains the three state variables δ , r , and Δd . The one-lag version of this VAR system can be written as follows:

$$\begin{bmatrix} \delta_t \\ r_t \\ \Delta d_t \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} \delta_{t-1} \\ r_{t-1} \\ \Delta d_{t-1} \end{bmatrix} + \begin{bmatrix} u_{1,t} \\ u_{2,t} \\ u_{3,t} \end{bmatrix} \quad (3)$$

In this setup, the time-varying risk premium (which is difficult to observe directly), is omitted. The matrix of coefficients shows the nature of the price-formation process at work in the market being modeled. Economically, if a_{12} (the coefficient on the risk-free rate in the yield equation) has a positive sign, and a_{13} (the coefficient on cash flow growth in the same equation) has a negative sign, this constitutes empirical evidence that prices are formed by investors through a Gordon-Growth-Model process. Such an observation justifies the use of this methodology to model inferred cash flow yields.

The system in Equation (3) can be written more compactly in matrix form as

$$z_t = Az_{t-1} + v_t \quad (4)$$

where z_τ is the observed vector of state variables at time τ , A is the matrix of coefficients, and v_t is the vector of error terms. The inferred k -period forward predicted cash-flow yields that we are attempting to generate in this project can then be obtained by multiplying the time- t realization of the state variables in the system, by the VAR-coefficient matrix A , k times, or

$$E[z_{t+k}] = A^k z_t \quad (5)$$

2.2 In-Sample and Basic Rolling Out-of-Sample Estimation

Initially, to establish the feasibility of our approach, we conduct this procedure in a single time series at a market level. We choose the market for Los Angeles Office properties (see specifics on data below). For δ we use a market-wide capitalization rate based on index data for prices and cash

flows. For cash flows, we use index-level Net Operating Incomes (NOIs). From this latter source, we also construct Δd , as a first difference of log NOI per square foot. The variable for r becomes the log of the long-term interest rate. To account for possible seasonality effects in our quarterly data, we estimate VARs with up to four lags. In this stage, we estimate the VAR over the entire time series of data, and the inferred yields become simply the fitted values from this VAR.

Since the main focus of our study is to attempt to infer prices of untraded assets, the primary portion of our empirical analysis focuses on a property-by-property level. This means we are dealing with panel data consisting of a cross section of properties over time, and so we modify the VAR approach for this setting³. In the panel setting, we stack time series observations for individual properties and then we modify Equation (3) to look as follows:

$$\begin{bmatrix} \delta_{i,t} \\ r_t \\ \Delta d_{i,t} \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} \delta_{i,t-1} \\ r_{t-1} \\ \Delta d_{i,t-1} \end{bmatrix} + \begin{bmatrix} u_{1,i,t} \\ u_{2,i,t} \\ u_{3,i,t} \end{bmatrix} \quad (6)$$

The above representation shows that we model yield for property i at time t , through lagged values of the yield for the same property, the common long-term interest rate (which is identical for all properties at a certain time), and the lagged cash flow growth for that property. Economically, the coefficient matrix estimated in this way characterizes the common price formation process applied to all properties in the sample used, with the cross section giving additional statistical power to the estimation of the coefficients.

Initially we estimate a single VAR system or a single set of rolling VAR systems⁴ across the entire panel. We then run separate VAR estimations for separate market segments, by property type, in order to allow the particulars of the price formation process to vary in the cross section. All data is de-meanned.

It is well known that in panel data studies, unobserved systematic cross-sectional heterogeneity can lead to biased OLS standard errors and thus to incorrect statistical inferences. For this reason,

³VARs are seldom used for panel data in the literature, perhaps because few settings warrant this from an economic perspective. For one example of such a procedure in use, however, see Love and Zicchino (2006).

⁴See discussion of out-of-sample estimation below.

such studies often use fixed effects and clustering of standard errors at various levels of panel observation.⁵ We elect not to incorporate such techniques in the procedures we use. This is because, while we do present VAR coefficients and hypothesis tests associated with them in our study, these purely serve to convey to the reader the basic interrelationships between the state variables involved in our study, and are not important for our primary result. Our primary interest lies in the out-of-sample predictions from the panel VARs we estimate, which should not be biased by this omission. On the contrary, it has been shown (for example, Wooldridge (2010)), that in the presence of unobserved systematic cross-sectional heterogeneity in panel data, OLS coefficients may become inefficient. If this is the case, this should negatively impact, rather than overstate, the predictive power of our panel VARs. While Petersen (2009) argues in this context that the use of random-effects estimation would help this situation, we elect to omit this for two reasons. First, as far as we know, the statistical properties of a GLS-based random-effects estimator in the context of a panel VAR are not sufficiently well understood. Second, a meaningful run of such a procedure would require a large amount of data, both in the cross section and the time series. Therefore, this would preclude us from being as parsimonious in our approach as we currently are, which would exclude a large part of the early portion of the data sample.

In keeping with the asset pricing literature, we begin by estimating a single VAR over all our data and computing our inferred yields as fitted values from this VAR. We do this to assess the quality of the inferred yields we produce, in a better-understood controlled setting. However, since the focus of our study is to infer prices for untraded assets, we then expand this procedure to generate out-of-sample estimates of yields. In particular, we run the panel VAR described in Equation (6) above, over a 40-quarter (10-year) rolling window and then conduct out-of-sample predictions. Specifically, in each iteration for a quarter t , we estimate our VAR using data from time $t - 39$ to t . By using Equation (5)⁶, we then compute predicted yields for each property i for quarter $t + 1$ as our inferred yields; to assess the quality of the yields produced, we then compare these inferred yields to the ex-post realized yields for each respective property i at time $t + 1$.

⁵Love and Zicchino (2006) also do this in their panel VAR study, although they have a different motivation for doing so, as they have economic interest in these coefficients.

⁶We adjust this analogously to Equation (6), such that z_τ contains the entire cross-section at time τ .

Given that we have an unbalanced panel, we choose to include in each VAR estimation properties for which we have continuous data for at least $2L + 1$ quarters, ending at time $t + 1$ (where L is the lag order of the VAR). This avoids estimating low-power time-series coefficients by using a panel that is too shallow to offer true time-series insights. On the other hand this is traded off against the survivorship bias that would be induced if we were to require that a property exist for, say, the entire 40 quarters that we use for each run. Given that we also have a cross section of properties to give us statistical power, it is not necessary for each single property to exist that long. Once again, we consider lags up to four quarters.⁷

To assess the quality of our predicted yields, we look at two measures. The first is a ratio of the standard deviation of our series of predicted yields, divided by that of the actual realized yields. This shows the fraction of total volatility that we capture; if our VAR system resembles the market's actual price formation process, this should be high (in theory close to 1, if we had actual prices in our data). For the in-sample estimation, this suffices to assess the quality of the estimated yields. For the out-of-sample procedure, there is one additional measure we must consider in addition to this. Supposing that we produced VARs that were extremely noisy, but whose estimated yields in no way resembled ex-post realized yields; in that case, the previously discussed measure would be high (and could even be above one), but this would still not indicate a production of high-quality yield forecasts. Therefore, in this case we examine the above measure jointly with the correlation coefficient between the predicted and the ex-post realized yields.

At this stage, for δ (which we call *yield* in the tables) we use the natural log of the ratio of property-level Net Operating Income divided by the appraised value of the property (or transaction price in quarters in which the property trades). Since appraisals tend to be smoothed, our goal from this line of research is to generate out-of-sample yields which are more volatile than the ex-post realized yields (i.e. the ratio discussed above should be greater than one). We plan to accomplish this by including additional measures of cash flow and modeling time-varying risk premia in future versions of the paper. For r (which we call *lt.rate* in the tables) we use the long-term interest rate,

⁷Economically, estimation windows between five and ten years which we use in this study, also seem warranted, as they match with the majority of holding periods for commercial properties. Hochberg and Mühlhofer (2011), for example, report that in the NCREIF data set the mean holding period is just below five years, with the third quartile at approximately 6 and a quarter.

and for Δd (which we call *noi.growth* in the tables) we use the difference in log of Net Operating Income per square foot for the respective property.

2.3 Self-Propagating Rolling-Window VAR Procedure

Having established step by step the feasibility of our approach, we next make the final refinement to our procedure, by implementing what we call a *Self-Propagating* Rolling-Window Panel VAR. In a rolling VAR procedure, in each VAR estimation run, it is necessary to have a full set of populated data fields, to *train* the VAR (i.e. to generate a coefficient matrix A). In the procedure described above, we use appraisals to generate the yields (cap rates) that are input into each VAR estimation, to generate the coefficients, and with which the coefficients matrix is multiplied to generate predictions.

We now move away from such a heavy reliance on appraisals, to a more transaction-based approach. The commercial property market is too thinly traded to estimate a VAR such as Equation 6 purely using cap rates based on transaction prices. That is the primary motivation of this study. Instead, we use our rolling-window VAR to populate the data with VAR-predicted cap rates where a genuine cap rate does not exist. The primary difficulty in this context, of course, becomes estimating the VAR system which is then used to predict cap rates, vis-a-vis this lack of data. Specifically, we proceed as follows.

Consider an unbalanced panel dataset of all the state variables, for a set of properties, $i \in 1, 2, 3, \dots, I$, over time, $t \in 1, 2, 3, \dots, T$. While transactions and therefore prices may be scarce in this context, it should be noted that observations for all other state variables in the VAR system are continuously populated. We generate a *mixed cap rate* for each property, i at time t , which we initially populate with a local transaction cap rate. That means, for each property, i , in whose market, m , one or more transactions occur at time t , we set *mixed.cap.rate* $_{i,t}$ equal to the average of transaction cap rates in market m at that time. We define a property's local market m , by the interaction of its Core-Based Statistical Area (CBSA, a classification of urban areas, roughly equivalent to the traditional MSA designations) and property type (i.e. Apartment, Industrial, Office, Retail). Given that our data consists of exclusively institutional-grade real estate, we believe

that a local market cap rate defined in this way constitutes a reasonable proxy for a property's own cap rate. Applying this procedure fills in cap rates for some panel observations but still leaves the vast majority of cap rate observations blank. In practice, out of 54,120 quarter-CBSA-type combinations, 5,192, or 10.1% of markets have any transactions and are therefore populated this way. This shows how thinly traded this market is, and the scope for improvement through our procedure. It will be the task of our Self-Propagating Rolling VAR to fill in the rest.

We specify a time window, w , over which to run each iteration of the rolling-window VAR. For the initial VAR run, which uses observations for all properties present in the dataset from $t = 1$ to w , we have no choice but to fill in the missing cap rates using appraisals as property valuations, since we need continuous data to estimate the VAR system. After estimating the coefficient matrix for the initial VAR run, however, we then fill in all unpopulated cap rates (i.e. cap rates for properties in whose markets no local cap rate was available) for time $t = w + 1$ as predictions from the panel VAR we just estimated. We then move the estimation window up by one period and re-estimate the VAR system from $t = 2$ to $t = w + 1$. It should be noted that the previously unpopulated cap rate observations at $w + 1$ are now populated with predictions from the previous VAR, which are then matched up with new data for the other state variables (cash flows and interest rates), all of which is used in the estimation. The one-period forward predictions from this new VAR run are then used to populate the previously-unpopulated cap rate observations for $t = w + 2$, which are then combined with new data for the other state variables in that period to be used for the VAR run from $t = 3$ to $t = w + 2$, to generate predicted cap rates for $w + 3$, and so forth. It is thus apparent how this VAR generates the data necessary for its next run, and therefore why we refer to it as *self-propagating*.

It should be apparent that, due to the structure of the panel VAR, through this procedure we generate a specific new cap rate for each individual property simultaneously, rather than an overall market cap rate. To see how this works, consider a panel of four properties ($i = 1, \dots, 4$), for which we have estimated the coefficient matrix, as illustrated in Equation 6. The property-specific predictions of the state variables are then calculated by multiplying the time- t matrix of property-

by-property values of state-variable realizations with the transcript of the coefficient matrix⁸, as follows:

$$\begin{bmatrix} \delta_{1,t} & r_t & \Delta d_{1,t} \\ \delta_{2,t} & r_t & \Delta d_{2,t} \\ \delta_{3,t} & r_t & \Delta d_{3,t} \\ \delta_{4,t} & r_t & \Delta d_{4,t} \end{bmatrix} \begin{bmatrix} a_{1,\delta} & a_{1,r} & a_{1,\Delta d} \\ a_{2,\delta} & a_{2,r} & a_{2,\Delta d} \\ a_{3,\delta} & a_{3,r} & a_{3,\Delta d} \end{bmatrix} = \begin{bmatrix} \delta_{1,t+1} & r_{t+1} + \epsilon_1 & \Delta d_{1,t+1} \\ \delta_{2,t+1} & r_{t+1} + \epsilon_2 & \Delta d_{2,t+1} \\ \delta_{3,t+1} & r_{t+1} + \epsilon_3 & \Delta d_{3,t+1} \\ \delta_{4,t+1} & r_{t+1} + \epsilon_4 & \Delta d_{4,t+1} \end{bmatrix} \quad (7)$$

Of primary interest, in this case, are the predictions for δ in each property. As described above, we then take these cap rate values and fill them in for properties at time $t + 1$ where this field is unpopulated (i.e. no local cap rate exists). Before doing this, of course, they must be re-measured and exponentiated, as the VAR is run with de-measured logs of variable realizations.

There may be a concern that using previously made predictions in a new estimation run may cause the new predictions (which would cumulate errors on top of errors) to become extremely noisy. Our results show, however, that this is not the case. There may be two reasons for this. First, in the panel, some new cap rate data is used, since local transactions occur, which supplements the information (or non-information) in the predictions. Second, all other state variables contain new information which is processed in the new estimation run.

As previously stated, we are dealing with an unbalanced panel of properties in this study. As before, we choose to include properties that, for a given number of lags L included in the VAR, have at least $2L + 1$ time-series observations, including their quarter of sale. Further, for the vast majority of properties that first appear in the dataset at a time beyond the initial window from 1 to w (where we automatically filled unpopulated *mixed.cap.rate* observations with appraisal cap rates), we populate any empty *mixed.cap.rate* observations in the property's first $2L + 1$ quarters of existence with appraisal-based cap rates. There are two points to be considered with regards to this rule. First, if we did not do this, we could only use properties for our estimation, which in their first $2L + 1$ quarters had local transactions in their market. This would be a non-randomly drawn sample

⁸In this coefficient matrix, the second subscript refers to the equation from which the coefficient came, so, for example $a_{2,\delta}$ refers to the second coefficient (i.e. the one on interest rate) from the δ equation.

and would therefore introduce selection bias. Second, the vast majority of NCREIF properties was actually purchased within a year or less of its cash flows being reported in the dataset. Within that time period, an appraisal would rely very heavily on the just-concluded transaction price, and thus essentially this way we introduce near-transaction-based cap rates into the estimation. Overall, the worry that we fill in much of our data this way, leaving no room for our self-propagating VAR would be unfounded: in practice, our self-propagating VAR is left with the task of estimating 195,734 cap rates.

In order to be as parsimonious as possible, and to use as few appraisals as possible, we try to use a short window length w . For estimations conducted on the entire dataset, we use a window size of 20 quarters (5 years). With much shorter window sizes, estimates become exceedingly noisy. Further, it may be economically justifiable to estimate the VAR over a sample that covers a large part of the previous property cycle.⁹ It should be noted here that our rolling-window procedure allows the VAR coefficients to vary over time. Again this should be economically justified, as investors may make different use of the underlying data in different economic times.

Besides estimation over the entire dataset, we also conduct this procedure on each property sector alone. Economically, it is conceivable that the price formation process (and therefore the VAR coefficients) for the different property-type submarkets might differ amongst each other. Due to a smaller cross-section of properties in this setting, we use a window size w of 30 quarters. We use four lags throughout this part.

We expand the set of state variables used in this procedure to include, in addition to the variables described above, the log of the square of quarterly NOI growth. We do this in an attempt to capture conditional cash flow volatility, and therefore get at a partial indicator of discount-factor information. In contrast to the Campbell-Shiller setting where an attempt is made to distinguish between cash-flow and discount-factor information, in this case this inclusion should be beneficial. Since we are trying to predict as much as possible of cap rates through cash flow dynamics, it seems warranted to try to capture at least some discount factor information. For illustrative purposes, we also run the estimation described in the previous section, with this additional variable.

⁹Further, as stated above, this fits well within the holding period range for the investors in this dataset.

It should be noted that the algorithmic nature of this procedure is set up to easily allow replicability. Therefore this estimation technique should be suitable for investors wishing to estimate current cap rates for properties they are examining, without resorting to appraisals.

We construct the following statistics to examine the predictive power of our self-propagating VAR. First, we compare our predictions to appraisals. For this, we use the ratio of standard deviation of the *mixed cap rates* we predict over standard deviation of appraisal-based cap rates, as well as their correlation. The same arguments as above apply for the interpretation of these statistics.

Next, we compare our self-propagating VAR's ability to predict eventual transaction cap rates (when the property is sold) to that of appraisals. To do this, we first use appraisal cap rates as the *predicted* series and report the ratio of standard deviations of predicted over transaction cap rates, as well as their correlation. Finally, we also report an out-of-sample R-squared (see Welch and Goyal (2008)) defined as

$$R_{OOS}^2 = 1 - \frac{\sum_{t=w}^T (\delta'_{t+1} - \delta_{t+1})^2}{\sum_{t=w}^T (\bar{\delta}_t - \delta_{t+1})^2} \quad (8)$$

In this expression δ_{t+1} is the ex-post realized transaction cap rate at time $t+1$, δ'_{t+1} is the predicted cap rate and $\bar{\delta}_t$ is the historical average cap rate over the rolling window ending at time t . This figure compares the sum-squared prediction error to the prediction error that would be obtained by using the historical mean as the best predicted cap rate. The more of an improvement the prediction offers over the historical mean, the closer to 1 this statistic gets. If the prediction does no better than the historical mean, the statistic is zero, and if it does worse, the statistic is negative.

Because a property is not appraised in a transaction quarter, we use the appraisal from the previous quarter as a comparison. To make a valid comparison, we then do the same thing with the predicted *mixed.cap.rate* from our self-propagating rolling VAR from the quarter before a transaction. We compute the ratio of standard deviations, the correlation, and the out-of-sample R^2 , this time using the predictions from our self-propagating VAR, compared to transaction cap rates. We compare these statistics to those calculated for appraisals. The most important comparison to be

made here will be in the out-of-sample R-squareds, as the best indicator of predictive power.

3 Data

Our study depends crucially on reliable, high-quality time series data on property-level cash flows. We therefore use the NCREIF property database for this information. In its property database, the National Council of Real Estate Investment Fiduciaries (NCREIF) collects individual-property-level data for institutional-grade commercial properties held by private entities (primarily pension funds). This collection effort is undertaken with the primary objective of compiling the National Property Index (or NPI) series, which constitute the de-facto industry standard commercial property indices in the United States. As of 2012, the value of the property portfolio in NCREIF's universe stood at \$320 billion. While this is difficult to estimate, many institutional investors consider this portfolio to be the vast majority of non-owner-occupied institutional-grade commercial real estate held by private entities. Each landlord of a property in NCREIF's universe reports a variety of information to NCREIF on a quarterly basis, including essential property characteristics, and (very importantly) operating cash flows, appraisals, and transaction values. While membership in NCREIF (and thus reporting of data to NCREIF) is voluntary, inclusion in NCREIF's database is considered desirable and prestigious on the part of private managers. NCREIF's stated policy is to only report data on high-grade institutional-quality commercial real estate, as this is the type of real estate its NPI tracks. As a result, inclusion of one's property transactions in NCREIF's database and indices is viewed as confirming a level of quality on the included investor. Therefore, most eligible managers choose to become members of NCREIF, and thus subject themselves to quarterly reporting of transactions. NCREIF membership constitutes a long-term contract and commitment, and once included, it is not possible for an investor to report performance only in certain quarters and not in others; the investor is contractually obligated to report all cash flows and transactions going forward. Data reported by NCREIF members to NCREIF is protected by a strict non-disclosure agreement.¹⁰ Thus, manipulating performance numbers is viewed as ineffective, as it cannot help the investor signal quality beyond membership itself. As a result, NCREIF members are both

¹⁰As academic researchers, we are given access to NCREIF's raw data under the same non-disclosure agreement.

willing and able to fully and confidentially report this data to NCREIF. This arrangement gives us the opportunity to examine cash flows and valuations (including transactions where available) in a very large set of commercial properties in some detail and with a high level of confidence in the accuracy of the data. The NCREIF sample runs from Q1 1978 through Q2 2012.

For the single-market test run of our VAR on Los Angeles Office property, we use price index data from RCA to construct cap rates.

Table 1 shows summary statistics for the NCREIF property database. We show time-series distributional statistics for the number of properties in the dataset, as well as time-series distributions partitioned into region or type. Lastly, we show statistics on cross-sectional distributions of Net Operating Income (NOI) per square foot, subdivided by property type, at three points in the dataset: first, the end of 2011 (the last full year of data we have), and then ten and twenty years earlier. The number of properties in our sample changes over time with only 26 properties remaining in our database over the entire 1991 to 2011 time period. In terms of the number of properties by property type, industrial properties constitute 41 percent of our sample followed by office properties (26%), apartments (16.5%) and retail properties (15%). Hotels represent only 1.5% of our sample. From a regional standpoint, 35% of properties are located in the West with an additional 28%, 21%, and 16% located in the South, East and Midwest respectively. When property size (square feet) is used in lieu of the total number of properties a slightly different perspective emerges. All property types have similar aggregate square feet except for hotels. A similar situation arises when different regions are examined with respect to property size. All four regions have similar square feet in aggregate. When NOI per square foot is examined over the 1991, 2001, and 2011 period, industrial properties tend to have the lowest mean return and standard deviation of the five property types followed by apartments. All other property types have higher mean returns and standard deviations.

4 Results

4.1 Primary Results: Commercial Property Markets

Figure 1 shows the actual log capitalization rate for Los Angeles office properties relative to the fitted values for cap rates from our two lag VAR model. The predicted cap rate tends to co-move with the actual cap rate albeit it appears to lag the actual cap rate by a quarter. This is especially evident in the latter half of the sample. The fitted cap rate tends to be a little smoother than the actual cap rate which is not surprising since this is an in-sample fitted value from a VAR. With out-of-sample rolling predictions we could possibly get at something more volatile, if the VAR does a better job than the appraiser.

Table 2 reports the Campbell-Shiller VAR coefficients for our full-sample panel. The state variables used are *yield*, the log of the capitalization rate, *lt.rate*, the log of the long-term interest rate, and *noi.growth*, the change in quarterly log NOI per square foot. The first panel shows the coefficients for the VAR using one lag. The second panel reports the VAR using two lags. Although we run these VARs for up to four lags, we omit tabulating the coefficients for our higher-lag VARs to save space. The VARs are estimated through equation-by-equation OLS. Table 2 shows that a two-lag VAR model does not offer any additional information over a one lag VAR model as evidenced by the incremental R^2 . Economically, this is an interesting finding, in that the commercial property market is thought to be fairly informationally inefficient. However, we do find in this case that new cash-flow and interest rate information is integrated by the market very quickly. It may thus be the case, that the market's apparent inefficiency is more a product of space markets (which determine rental cash flows) than capital markets, which price existing cash flow information.

Observe that a large portion of the variation in the yield (30.8 percent) is accounted for by the combination of a one period lag in the yield, long term interest rate, and NOI growth. All three lag variables are statistically significant at the 1 percent level with the coefficient corresponding to the interest rate being positive while the coefficient associated with NOI growth is negative. Since the interest rate and NOI growth represent the denominator ($r - g$) in the Gordon growth model, this suggests that our VAR does a good job in simulating this growth model. Recall that both r and g

are linear since we are using logs, which is the idea underlying the Campbell-Shiller method. If the signs were not consistent with the growth model, then something in addition to the Gordon growth model would also be useful in determining prices. It appears that prior information embedded in each of our variables is useful for prediction purposes. While we previously noted that using a two lag VAR model provides little, if any, additional explanatory power, we should note that the Gordon growth model continues to hold when a two lag VAR model is used. Table 2 also reveals that the combination of a one period lag in the yield, long term interest rate, and NOI growth also accounts for a sizeable portion of the variation in the long term interest rate (83 percent) and also the growth rate in NOI (18.5 percent). Evidence of persistence in the yield which represents the ratio of cash flow to the price of commercial real estate and the persistence of NOI growth are necessary conditions in our attempt to show that an investor can focus on cash flow dynamics to account for firm (property) risk.

Figure 2 shows impulse response functions from the full-panel VAR with four lags. The first panel shows the response on all state variables obtained by shocking *yield*, the second the response obtained by shocking *lt.rate*, and the third the response obtained by shocking *noi.growth*. The figure is consistent with Table 2. By shocking yield (refer to the first graph on the left), we see that yield itself (i.e. cap rate) is positively persistent, with a slow decay pattern. Past yields contain information on the future direction of the yield. In contrast, the future interest rate is not responsive to yield shocks while NOI growth responds negatively to a yield shock. In other words, an increase in the yield (cap rate) forecasts a decline in the growth rate of net operating income. The middle graph shows that the expected yield should increase given an increase in the long term interest rate. The increase in yield given either a decline in NOI growth or an increase in the long term interest rate is consistent with the Gordon growth model relationship. In contrast, the expected growth in net operating income is invariant to an increase in the long term interest rate. The graph on the extreme right shows that the yield on commercial real estate should decline if NOI growth is expected to increase. Overall, seasonality might partly account for the reverting patterns exhibited in NOI growth. It should also be noted that the long-term interest rate is only affected by changes in itself, and not by changes in the other two state variables. This is in line

with the relative coefficient sizes in the *lt.rate* equation of Table 2, and is consistent with economic intuition. The risk-free interest rate should be fairly exogenous to commercial property yields and cash flow growth. The important relationships are again the response of yield to the long term interest rate and NOI growth.

Table 3 presents statistics comparing the predicted yields (log cap rates) from our VARs, with actual ex-post realized yields. Specifically, the table presents, for each VAR specification, the ratio of the standard deviations between the predicted and the realized series of yields. The first section presents this for the predictions (i.e. fitted values) from a set of full-sample panel VARs. The second section presents out-of-sample predictions for a set of rolling-window panel VARs. The predicted yields are for quarter $t + 1$ and are generated by a VAR using 40 quarters' worth of data, ending at t . The variable $pred.yield_{i,t}$, in the first section is the fitted value from the VAR for this variable. The realized yield at quarter t , for the same property i is $yield_{i,t}$. In the second section, for the rolling VAR $pred.yield_{i,t,t+1}$ is the prediction constructed in quarter t , for the yield for property i in quarter $t + 1$, while $yield_{i,t+1}$ is the ex-post realized yield for the same property at that time. For the out-of-sample predictions we present both the ratio of standard deviations of the two series, as well as their correlation coefficient. In parentheses there is the value of a t-statistic testing the hypothesis that the actual correlation between the two series is 0. The set of state variables for the VAR consists of $yield_{i,t}$, $lt.rate_t$, the log of the long-term interest rate, and $noi.growth_{i,t}$, the quarterly difference of log NOIs. All variables are de-meanned. The ratio of the standard deviations between the predicted and the realized series of yields is similar for both the full-sample panel VAR system and the rolling panel VAR system with the latter, only slightly lower since this represents out-of-sample forecasts (as opposed to an in-sample forecast in the former case). The 53 percent ratio of standard deviations in addition to the 61 percent correlation between the predicted yield versus the actual yield indicates that our VAR model which represents a dynamic linearized Gordon Growth Model does a good job in predicting the yield and in turn the volatility in the yield. As a benchmark with which to compare our results to, Campbell and Shiller report values around 15% (ratio of standard deviations) and 25% (correlation) respectively. Besides this, the results appear to be invariant to the number of lags used in the VAR system which suggests that a one lag VAR

model suffices. Using additional lags does not provide any new information.

Overall, we have now established in a more controlled environment that this type of modeling of cash flow dynamics should constitute a good methodology to infer cash flow yields (i.e. capitalization rates) when these are not observable.

We now turn to the results illustrating the performance of our full new technique, the *Self-Propagating* Rolling-Window Panel VAR estimation of yields for untraded properties. Figure 3 shows a time series of cross-sectional averages of realized transaction-based cap rates (the black solid line) and our predicted *mixed cap rates* (the red dashed line). Notice how closely the predictions in cap rate track the actual realized values, especially after about 1995. This should show how well our estimation procedure infers yields of untraded properties, as, when these properties do trade, the realized yields closely resemble our estimates. This is especially remarkable, given that our VAR runs more or less in the blind for much of the time over the majority of properties, as far as yield data is concerned, using its own estimates in each subsequent run, with only cash flow data being updated.

The early 1990s show a large amount of volatility, both in realized, as well as in predicted cap rates. This is likely for two reasons. First, in this time period, there truly was a large amount of cap rate volatility, as the market was recovering from the Savings-and-Loan crisis. Second, during this time the coverage of NCREIF's data was still much smaller than in the late 1990s and beyond, so that individual outliers in terms of cap rate would have made more of a difference on overall averages shown. At the same time, our VAR procedure has less data to work with during that time, leading to less power and therefore more noisy estimates. In the run-up of yields around the 2007-2008 financial crisis, a small lag of our predicted cap rates is apparent. It should be remembered that the predicted cap rate shown is that from the previous quarter as our headline results make a comparison with appraisals (see Section 2). In the quickly-increasing cap rates of that time, this lag likely would have made the most difference.

The subsequent figure (Figure 4) shows how our *Self-Propagating* VAR performs in the cross-section. The bars show the difference between the time-series average transaction cap rate for each CBSA, and the time-series average predicted cap rate for the same CBSA. With the exception

of one negative outlier, the vast majority of predicted average cap rates is within less than one percentage point of the realized cap rate, with all observations within two percentage points.¹¹

Table 4 presents our main results illustrating the performance of the *Self-Propagating* Rolling VAR. Panel A shows the results for the full time series. This means that, given that our data starts in 1980 (accounting for the lags needed in the VAR), and the first prediction for *All Types* (i.e. all property types combined) is made for the first quarter of 1985 (since we use a 20-quarter window size. For the individual property types, with a 30-quarter window size, the first prediction is made for the third quarter of 1987. The predictions analyzed are then quarterly, for each property that exists in the sample at a given time, until 2012.

The first column (“Predicted vs. Appraisal”) compares our predicted *mixed cap rates* to appraisals at the same time. We report the ratio of the standard deviation of our predicted series over that of the appraisal series, to show what fraction of the variation we capture. Since we are making out-of-sample predictions, in order to distinguish between a set of noisy predictions and one that actually captures the variation in appraisals, we also report the correlation between the two series and a t-test, testing the null hypothesis of zero correlation. For All Types, in Panel A, we find that the standard deviation of our predicted cap rates is .823 times that of appraisals, which means we capture a large amount of the variation in this series, and the statistically highly significant positive correlation of .3514 indicates that these predictions do resemble appraisal-based cap rates. It should be remembered that appraisals themselves have well-known problems, such as smoothing and temporal-lag bias. Therefore, matching appraisals as closely as possible should not be a primary goal for our predicted cap rates. The main point to be taken from these statistics is that our estimates remain well-controlled, in terms of their statistical properties. A fear that re-using previously generated predictions in subsequent iterations of the VAR would lead to extreme noise in our predictions (or perhaps a gradual dying off in their variance, as in an impulse-response function) is therefore unfounded, and we can be confident of the reasonable performance of our procedure. Performing this comparison with appraisals for this purpose has the added advantage

¹¹The one large negative outlier is Worcester, MA. Given that this is a small market, this result is likely driven by one highly idiosyncratic transaction. Given that our prediction over-states the actual cap rate, this would most likely have been some type of glamour sale, or a misrecorded transaction.

that we have observations for nearly the entire panel, while, as stated before, transactions tend to be scarce.

The second and third column (“Appraisal vs. Transaction”) and (“Predicted vs. Transaction”) considered jointly, on the other hand, show directly how well our procedure predicts actual market-transaction cap rates, and, by implication, how accurate its inferred cap rates for untraded properties would be. The “Appraisal vs. Transaction” column shows a benchmark of what is possible through appraisals, to which we can then compare our procedure in the third column. It should be noted that Column 2 purely reports pre-existing data from the NCREIF dataset. None of our estimation technique enters in a any way into any of the statistics reported there. In these comparisons, too, we report the ratios of standard deviations, correlation coefficients, and t-tests of these to assess the overall statistical properties of our predictions in this comparison. The salient result, however, is the Out-of-Sample R^2 which, it is generally argued (see for example Welch and Goyal (2008)), is one of the best ways to assess the predictive ability of an out-of-sample estimation, as it distinguishes noisy predictions, from predictions that genuinely capture the variation in the true underlying series.

Comparing the values of the ratios of σ , as well as the correlations between columns two and three, we find that our estimated cap rates once again have good statistical properties in their relation with transaction cap rates, when compared with appraisals. Most importantly, however, we find that the R_{OOS}^2 values show that appraisals capture .438 of the total variation in transaction cap rates, while our predictions capture .7511 of total variation in transaction cap rates, for all properties combined and over the entire time series. This means, our technique not only captures a very large portion of the variation of cap rates in itself, but it also substantially outperforms appraisals in this respect. This is in spite the fact that appraisers have substantial property-specific information available to them (which a quarter before the transaction should really not change substantially until the transaction is complete) while we limit our estimation to only the time series of cash flow information. Economically, besides showing the attractiveness of our procedure, this also shows how strongly yields (and therefore prices) in this thinly traded market are shaped by the time-series dynamics of cash flow information.

There may seem to be a discrepancy throughout Table 4 in that the correlations between appraisals and transaction cap rates are generally higher than those between our predicted cap rates and transactions, while the opposite is the case for the Out-of-Sample R^2 (R_{OOS}^2). The difference is that the correlation coefficient shows the degree of covariation between two series up to an arbitrary linear transformation on each. In contrast, the Out-of-Sample R^2 makes a direct comparison between the unadjusted raw series. Given that the arbitrary linear transformation in the correlation coefficient is difficult to get at in practice, it is generally argued that the R_{OOS}^2 offers a cleaner comparison and therefore shows the more relevant result in this setting. We report the ratio of standard deviations and the correlations primarily because this is how the statistical properties of VAR predictions are customarily compared to the realized series in this line of literature.

The lower sections of each panel of Table 4 show our technique applied to each subsample of the data by property type.¹² As stated before, this is an important potential distinction to make, as the price formation process (and therefore our VAR coefficients) might differ among property sectors. Overall, we see that the predictive ability of our procedure does better in Apartments, than in Industrial and Office. In Apartments, we find an R_{OOS}^2 of .844 (about nine percentage points above the statistic for the combined data) , while for Industrial this is .6857 (or about seven percentage points lower than for all properties), and for Office this is .7361, very close to the level for all properties. This is consistent with economic intuition, in that apartment leases are generally very short, while industrial and office leases are much longer. This means that economically there is much less stickiness in Net Operating Incomes (NOIs) for Apartments than for the other two property types, and that therefore these cash flows are more informative of current market lease levels, which may be used by investors in their forward projections used in Gordon-Growth-Model valuations. This may also be part of the reason why appraisals better predict Apartment cap rates than Industrial or Office. Statistically, it also means that Apartment cash flow data is more variable, and therefore gives the VAR more information to work with in predicting cap rates. Very importantly, though, we note that our primary result holds for all property-type submarkets, in that our predicted yields always capture a substantially higher part of variation in realized yields than

¹²While NCREIF does track a few hotel properties, the number of properties is too small for our procedure to work in this sector.

appraisals do, with differences in R_{OOS}^2 between 20 percentage points (Office) and 37 percentage points (Industrial).

For Retail, we also find a large outperformance in predictive ability by our estimates over appraisals. Our estimates here show an Out-of-Sample R^2 of .8495, the highest among all property types, lying at least by a very small margin over Apartments, while appraisals show the second lowest (.4832), only above Industrial. Economically, while retail leases are fairly long, a substantial portion of cash flows in Retail properties is paid through percentage rents, in which the landlord participates in tenants' revenues. This makes the cash flows from this property also highly variable, and less sticky, leading to better estimates of cash flows going forward. It seems to be the case that appraisers are less able to account for this in their cap rate formation. Overall, we find that, both in the entire sample, as well as in all property-type subsectors, our predicted cap rates capture a substantial portion of the variation of actual transaction cap rates, and substantially outperform appraisals in this respect. This lends strong support to the contribution made by the *Self-Propagating* VAR technique we develop in this study.

The subsequent panels of Table 4 show time-period subsamples of our results. It should be noted that the results shown in these panels come from the application of the VAR procedure to the full sample, and only report those estimates made for the date range specified. This makes intuitive sense, as the estimation procedure consists of a rolling-window VAR. Therefore, letting the estimation only start in time to produce the desired estimates would actually be detrimental, since, as explained in Section 2, we have to fill the first estimation window with appraisal-based cap rates. Letting the estimation start at the beginning, on the other hand, gives us estimates that are purely based on cap rates filled with the VAR's own estimates, as specified there. Thus, these are actually more *pure* estimates from the *Self-Propagating* VAR.

In Panel B, we show results from 1995 going forward. This cuts out the early part of the sample, where, as shown in Figure 3, the data quality improves to generate less noise in cap rates. Overall, this panel shows almost identical patterns to Panel A. We capture between approximately 67% and 85% percent of variation of transaction cap rates, with ratios of σ and correlations in reasonable ranges, as above. Similarly to the previous panel, we capture a substantially higher amount of

variation in transaction cap rates than appraisals do (between about 24 and 32 percentage points more). Here, Apartment and Retail estimations still do the best, with Industrial faring the worst and Office being close to the estimate for All Types. Qualitatively, this situation is maintained in Panel C, which shows statistics on estimates from 2005 forward. This panel does show a slight decrease in the overall fraction of variation predicted both by our estimation and by appraisals. The good predictive ability of our model with respect to Apartments is also reduced in this panel; this may be due to Apartments' being the most closely affected by the run-up and collapse of the housing market of this time period. However, we still capture a substantial fraction of variation in these cap rates ($R_{OOS}^2 = .6749$) and still outperform the predictive ability of appraisals by 21 percentage points. In this time period, Retail shows the best predictive ability ($R_{OOS}^2 = .8997$), which is consistent with the previous intuition about percentage rents and cash-flow informativeness.

Panel D reports results from the beginning of the financial crisis (2007Q3) going forward. Together with Panel C (which already largely contains this period), these results show that our technique works remarkably well, even in times of economic turmoil. Even during this time period, we still capture a substantial fraction of variability of transaction cap rates (.7314 for all types and for individual types between .6263 for Industrial and as high as .929 for Retail). In fact, we do slightly better in this time period than in the previous panel. For Apartments, the Out-of-Sample R^2 is higher again, at .7361, which puts it close to the predictive ability for All Types. Since apartments were the most affected by the fluctuations in single-family housing, it is consistent with economic intuition that these cap rates might be more difficult to model in this time period. In that sense, the ability of our technique to capture this large a fraction of cap rate variation for individual properties is remarkably good.

The Office submarket also shows an interesting development between Panels C and D. In Panel C, there seems to be a large amount of volatility in the market, with appraisals so noisy that the ratio of σ s is actually at about 1.02. In Panel D, this goes back below one, and the Out-of-Sample R^2 s rise substantially. This may have to do with the privatization of Equity Office Properties (EOP) by Blackstone, with the subsequent selling of EOP's portfolio. This caused a large shock to the office market. By the middle of 2007, however, most of this was complete and so cap rate

volatility would have declined again. However, despite this, our model still captures just under 70% of variation including this shock, and over 80% thereafter, even in a market with such long leases, in a time of general financial turmoil. This panel shows overall, that even in an economic crisis, in which market liquidity almost completely dried up, our estimation technique captures a substantial amount of the variation in actual transaction cap rates. Further, in a time of such low transaction volume, a technique such as this one is especially valuable, as it is otherwise especially difficult to accurately infer cash flow yields and therefore prices.

4.2 Most-Liquid versus Less-Liquid Markets

We now proceed to test the hypothesis that the price-formation process may differ between submarkets which are highly liquid and populated in large part by institutional investors, versus markets which are not. In order to do this, for each quarter and property type, we rank CBSAs by a liquidity measure. The measures we consider are number of transactions, square-footage transacted, and dollar-value transacted, each of which we derive from CoStar. Then, we define a dummy variable, which is equal to one if a CBSA is a top-10 liquid market for that quarter and type, and zero if it is not. We consider an analogous dummy for top-20 markets. This liquidity dummy, generally, also includes the largest cities, and the so-called “gateway” cities, so economic delineations of markets by these criteria would be largely equivalent to this.

We use the liquidity dummy ($liq_{m,t}$, which indicates whether market m is a top-10 or top-20 market in quarter t) in our VARs, to form a fully-interacted model, as follows:

$$\begin{bmatrix} \delta_{i,t} \\ r_t \\ \Delta d_{i,t} \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} & a_{17} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} & a_{27} \\ a_{31} & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} & a_{37} \end{bmatrix} \begin{bmatrix} \delta_{i,t-1} \\ r_{t-1} \\ \Delta d_{i,t-1} \\ liq_{m,t} \\ liq_{m,t}(\delta_{i,t-1}) \\ liq_{m,t}(r_{t-1}) \\ liq_{m,t}(\Delta d_{i,t-1}) \end{bmatrix} + \begin{bmatrix} u_{1,i,t} \\ u_{2,i,t} \\ u_{3,i,t} \end{bmatrix} \quad (9)$$

As before, the above shows a one-lag representation of the VAR system, which is easily extendable to the four periods we use in our estimation. Predicted cap rates across the panel can be computed in an analogous way as before, simply with a wider coefficient matrix and matrix of explanatory variables. We interact the dummy with the lagged state variables in the system, in order to allow the entire price-formation process to vary across the two types of market, rather than just the mean change in cap rate. Note that the liquidity dummy is *not* lagged, but simultaneous to the time period for which the estimation is conducted. This should be economically warranted, as the state of a market at the time a transaction occurs should be the relevant determinant of its price-formation process. We begin this analysis in 1995, as the CoStar data does not have sufficient coverage before this date.

The results for this procedure (not tabulated to save space) support the hypothesis that the price formation process is the same between the two types of markets and that therefore the more parsimonious procedure without the market split is more appropriate in this context. The ratios of standard deviations of our predicted cap rates versus actual transaction cap rates become slightly higher (by two to five percentage points), but the out-of-sample R-squareds become slightly *lower* (by one to two percentage points). This indicates that the additional variables in the model only add noise (in that a higher number of coefficients with the same amount of data will sacrifice statistical power), but do not add genuine explanatory power. The exact choice of liquidity variable among the alternatives presented above, as well as the delineation of top-10 versus top-20 does not qualitatively alter these results. These results therefore, once again, support the hypothesis that the price-formation process modeled by our procedure is universal throughout these markets, and does not vary by liquidity or institutional participation.

Overall, these results thus show the accuracy with which our *Self-Propagating* Rolling-Window Panel VAR predicts actual realized transaction yields. We consistently capture a large fraction of the variation in these yields and consistently substantially outperform the predictive ability of appraisals in this respect. It should be remembered that the yields we predict pertain to individual properties and therefore allow investors to accurately assess the specific value of their untraded investment assets.

4.3 The Informativeness of Cash Flows: Stock Markets

In our main results, we argue, in part, that the informativeness of cash-flow variation plays an important role in the success of our Rolling-Panel VAR procedure. We now further explore this hypothesis, by applying our procedure to a different asset class, the stock market. While in this market trades are abundant, and therefore it is not necessary to infer cash-flow yields of untraded stocks, this market nevertheless offers a suitable natural laboratory to explore the role that cash-flow informativeness plays in this context.

It is well documented that dynamic Gordon-Growth models do not do well in explaining stock price movements. This is known as the *Excess-Volatility Puzzle* (first documented by Shiller (1981) and LeRoy and Porter (1981)). A large part of this literature (see, for example Boudoukh, Michaely, Richardson and Roberts (2007), or Chen (2009)) focuses on the idea that dividend-smoothing makes the cash flows associated with stocks uninformative in the price-formation process. Therefore, if cash-flow informativeness plays an important role in the success of our methodology, we should see significantly diminished explanatory power for stock prices.

In order to test this, we use stock data from the Center for Research in Securities Prices (CRSP), and compute imputed dividend yields and dividend growth rates for all stocks in the sample, as ratios of returns including distributions, and returns without distributions.¹³ We then use this data in a Rolling-Panel VAR, equivalent to the out-of-sample estimation presented in Section 2.2. Once again, since stocks are continuously traded, the *Self-Propagating* procedure becomes unnecessary here. We maintain a quarterly frequency, a window size of twenty quarters, the same time-length inclusion criteria for each individual stock, as well as four lags in the VAR.

Table 5 shows the results from this estimation. To make the results comparable to the application for Commercial Property, we show the same sample splits, with the full sample covering the same time period as the complete NCREIF sample, and the time-period subsamples defined analogously as before. The table compares our estimated dividend yields to the ex-post dividend yields, and should therefore be compared to the right column of the different panels of Table 4.

When making this comparison, we see that overall the ratios of standard deviations between

¹³For this methodology, see, for example Mühlhofer and Ukhov (2011).

predicted and realized yields are somewhat lower (.57 to .64, versus .70 to .92) and correlations are somewhat higher, at .50 to .57, but comparable to the results from Table 3, (the non-*Self-Propagating* VAR), which compares estimated yields to appraisal yields. The most dramatic difference, however, can be found in our main statistic of interest, the Out-Of-Sample R-Squared. The values for stocks are dramatically lower in this case, lying only between .25 and .32, while for Commercial Property, we achieve values between .62 and .93. Thus, the explanatory power of our procedure for stocks is dramatically lower than for commercial property. The relatively high ratios of standard deviations, together with the low R^2_{OOS} values indicate that our procedure here generates noisy and much less accurate estimates of dividend yields. Together with the cash-flow informativeness idea that we began to establish for Commercial Property, these results lend strong support to the hypothesis that the effectiveness of this procedure hinges upon the informativeness of cash flows: in a market in which cash flows are documented to contain less information for prices, our procedure does significantly worse, than in a market in which cash flows should be more informative.

5 Conclusion

We address the measurement of return and risk in a market where assets trade infrequently by focusing on observable cash flow dynamics in modeling the market's price formation process. The underlying premise is that cash flow fundamentals are the most important driver for pricing an asset and its risk. To evaluate the extent to which we can use observable cash flows to infer property price, we construct a *Self-Propagating* Rolling-Window Panel VAR technique, based on a Dynamic Gordon Growth Model. Through this, we infer cash flow yields (cap rates) for untraded commercial properties as out-of-sample predictions from our empirical procedure. Our results show that our predicted yields explain between 75% and 93% of variation in ex-post realized transaction yields, while appraisal-based yields, by contrast, explain only half to two thirds of this variation. These findings are consistent with the underlying premise that cash flow fundamentals are the most important driver for pricing an asset and its risk. The algorithmic nature of our technique makes our procedure readily applicable to investors wishing to infer cash-flow yields and thereby asset

values for investment assets that are thinly traded.

References

- Beaver, W., Kettler, P. and Scholes, M.: 1970, The association between market determined and accounting determined risk measures, *The Accounting Review* **45**(4), 654–682.
- Beaver, W. and Manegold, J.: 1975, The association between market-determined and accounting-determined measures of systematic risk: Some further evidence., *Journal of Financial and Quantitative Analysis* **10**(2), 231 – 284.
- Boudoukh, J., Michaely, R., Richardson, M. and Roberts, M. R.: 2007, On the importance of measuring payout yield: Implications for empirical asset pricing, *Journal of Finance* **62**(2), 877–915.
- Campbell, J. Y.: 1991, A variance decomposition for stock returns, *The Economic Journal* **101**(45), 157–179.
- Campbell, J. Y. and Shiller, R. J.: 1988a, The dividend price ratio and expectations of future dividends and discount factors, *The Review of Financial Studies* **1**(3), 195–228.
- Campbell, J. Y. and Shiller, R. J.: 1988b, Stock prices, earnings and expected dividends, *Journal of Finance* **43**(3), 661–676.
- Chen, L.: 2009, On the reversal of return and dividend growth predictability: A tale of two periods, *Journal of Financial Economics* **92**, 128–151.
- Cohen, K. J., Hawawini, G. A., Maier, S. F., Schwartz, R. A. and Whitcomb, D. K.: 1983, Friction in the trading process and the estimation of systematic risk., *Journal of Financial Economics* **12**(2), 263 – 278.
- Da, Z. and Warachka, M. C.: 2009, Cashflow risk, systematic earnings revisions, and the cross-section of stock returns., *Journal of Financial Economics* **94**(3), 448 – 468.
- Dimson, E.: 1979, Risk measurement when shares are subject to infrequent trading., *Journal of Financial Economics* **7**(2), 197 – 226.
- Driessen, J., Lin, T. and Phalippou, L.: 2011, A new method to estimate risk and return of non-traded assets from cash flows: the case of private equity funds, *Journal of Financial and Quantitative Analysis* . Forthcoming.
- Gonedes, N.: 1975, A note on accounting-based and market-based estimates of systematic risk, *Journal of Financial and Quantitative Analysis* **10**(02), 355–365.
- Hochberg, Y. V. and Mühlhofer, T.: 2011, Market timing and investment selection: Evidence from real estate investors. Working paper, available at: <http://tobias.muhlhofer.com>.
- Ismail, B. and Kim, M.: 1989, On the association of cash flow variables with market risk: further evidence, *Accounting Review* pp. 125–136.
- Kallberg, J. G., Liu, C. H. and Srinivasan, A.: 2003, Dividend pricing models and reits, *Real Estate Economics* **31**(3), 435–450.

- LeRoy, S. and Porter, R.: 1981, The present-value relation: Tests based on implied variance bounds, *Econometrica: Journal of the Econometric Society* pp. 555–574.
- Love, I. and Zicchino, L.: 2006, Financial development and dynamic investment behavior: Evidence from panel var, *The Quarterly Review of Economics and Finance* **46**(2), 190–210.
- Mandelker, G. and Rhee, S.: 1984, The impact of the degrees of operating and financial leverage on systematic risk of common stock, *Journal of Financial and Quantitative Analysis* **19**(1), 45–57.
- Mühlhofer, T. and Ukhov, A. D.: 2011, Do stock prices move too much to be justified by changes in dividends? evidence from real estate investment trusts. Working Paper, Available at <http://tobias.muhlhofer.com>.
- Petersen, M. A.: 2009, Estimating standard errors in finance panel data sets: Comparing approaches, *Review of financial studies* **22**(1), 435–480.
- Plazzi, A., Torous, W. and Valkanov, R.: 2010, Expected returns and expected growth in rents of commercial real estate, *Review of Financial Studies* **23**(9), 3469.
- Scholes, M. and Williams, J.: 1977, Estimating betas from nonsynchronous data, *Journal of financial economics* **5**(3), 309–327.
- Shiller, R. J.: 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *The American Economic Review* **71**(3), 421–436.
- Shiller, R. J.: 1992, *Market Volatility*, The MIT Press.
- Welch, I. and Goyal, A.: 2008, A Comprehensive Look at The Empirical Performance of Equity Premium Prediction, *The Review of Financial Studies* **21**(4), 1455–1508.
- Wooldridge, J. M.: 2010, *Econometric analysis of cross section and panel data*, The MIT press.

Table 1: Summary Statistics, NCREIF

This table shows summary statistics for the NCREIF property database. We show time-series distributional statistics of the number of properties in the dataset, as well as time-series distributions subdivided by region or type. Lastly, we show statistics on cross-sectional distributions of Net Operating Income (NOI) per square foot, subdivided by property type, at three points in the dataset: first, the end of 2011 (the last full year of data we have), and then ten and twenty years earlier.

	Mean	StDev	1st Quart.	Median	3rd Quart.
Number of Properties					
Total	3,426	2,797	1,274	2,244	5,607
Apartment	566.9	575.5	44.25	355	919.2
Hotel	52.2	60.66	13	29.5	91.75
Industrial	1,405	1,260	625.8	811.5	1,918
Office	889.9	692.9	380.5	531	1,711
Retail	512.2	348.9	211.8	484	684
East	709.2	618.8	211.2	461.5	1,192
Midwest	539.6	351.8	281.8	419	822.8
South	962.3	864	340	544	1,440
West	1,215	972.4	439.8	800	2,168
Property Sizes (Sqm)					
Total	234,903	356,281	75,000	149,330	287,049
Apartment	225,457	240,027	85,932	214,994	316,385
Hotel	116,040	212,879	0	79,428	125,860
Industrial	234,485	411,583	70,788	134,791	272,548
Office	241,949	309,172	86,158	155,058	287,023
Retail	249,489	364,252	73,600	129,605	264,973
East	245,534	325,372	81,027	160,611	299,180
Midwest	261,667	469,610	80,158	157,852	306,931
South	230,254	334,977	76,732	151,200	288,544
West	220,133	331,144	66,817	135,065	267,967
2011 NOI per Square Foot					
Apartment	2.759	9.266	1.421	2.008	2.9
Hotel	3.467	2.934	1.748	2.907	4.654
Industrial	0.9854	0.9986	0.5401	0.9283	1.306
Office	3.265	3.21	1.549	2.883	4.46
Retail	3.736	4.315	1.918	2.957	4.282
2001 NOI per Square Foot					
Apartment	1.69	0.9911	1.196	1.493	1.879
Hotel	5.497	6.233	0.9893	2.533	10.27
Industrial	1.518	1.936	0.796	1.114	1.627
Office	3.807	2.699	2.398	3.432	4.829
Retail	3.026	4.922	1.624	2.592	3.524
1991 NOI per Square Foot					
Apartment	1.151	0.4893	0.8427	1.09	1.429
Hotel	1.583	1.128	0.898	1.201	1.86
Industrial	0.8656	0.7627	0.4726	0.7434	1.083
Office	1.847	1.924	0.789	1.542	2.522
Retail	2.627	3.45	1.035	1.722	2.668

Predicted and Actual Log Cap Rates,
LA Office

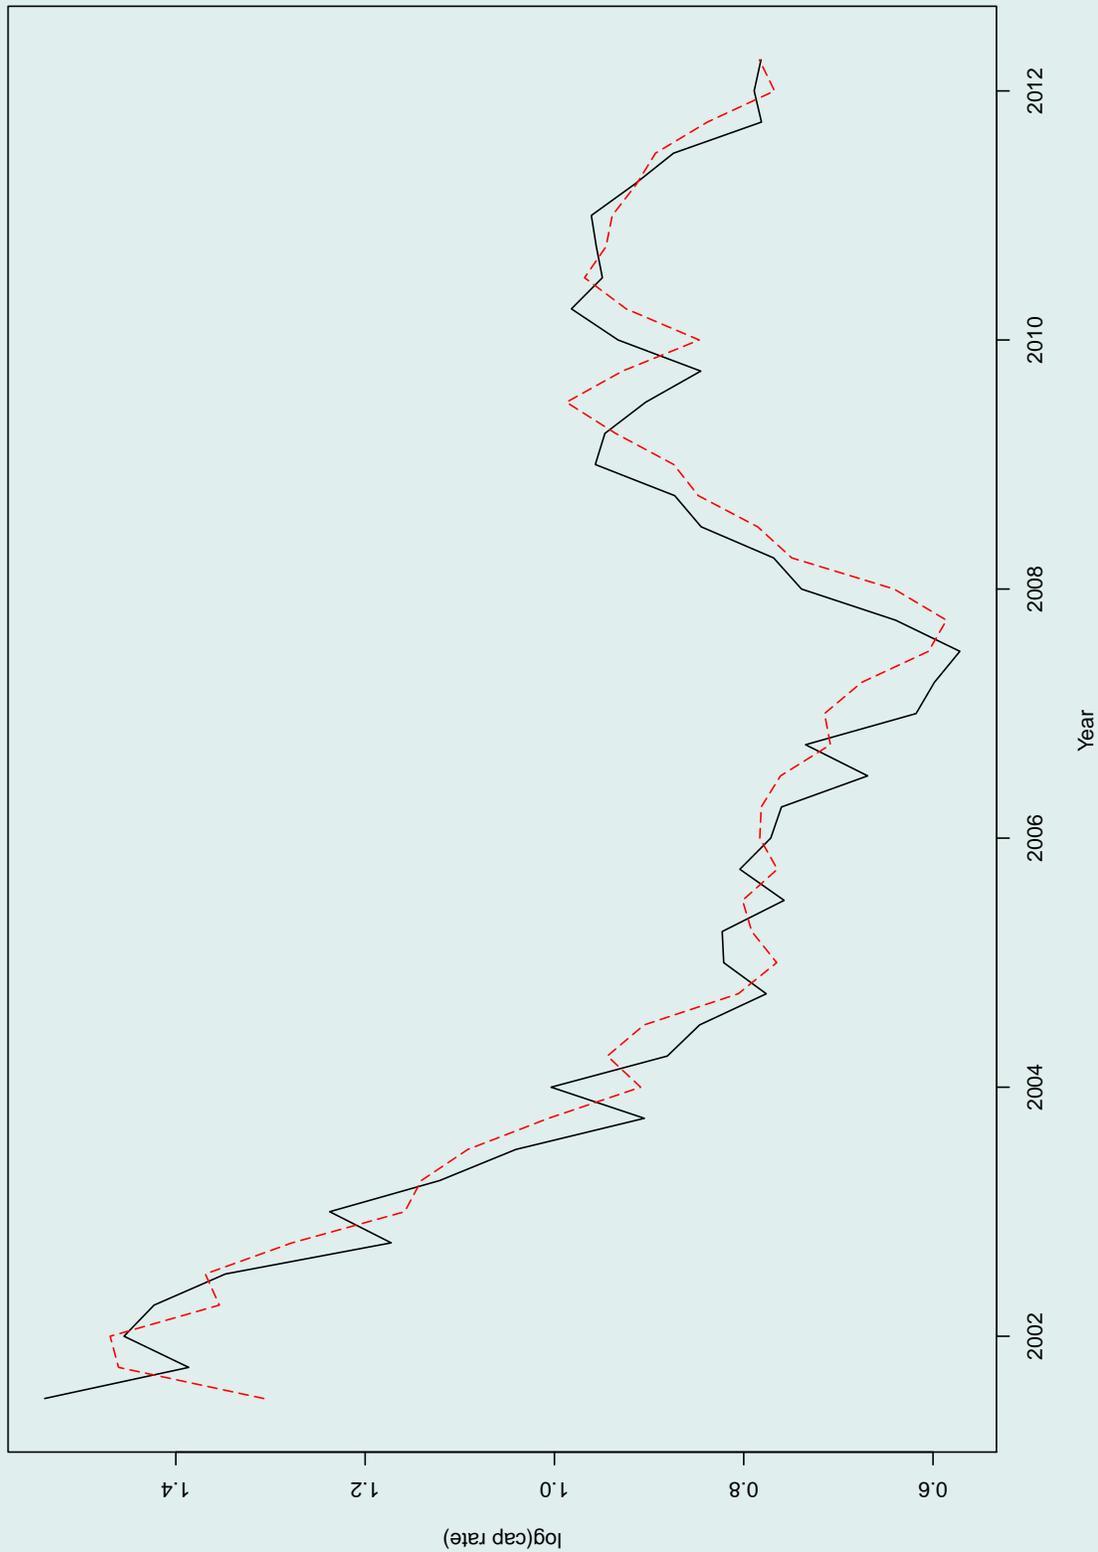


Figure 1: This figure shows, in the black solid line, the actual log capitalization rate for Los Angeles office properties, and in the red dashed line, fitted values for cap rates from our VAR for the same value, containing log of lagged cap rate, log of lagged long-term interest rate, and log of quarterly NOI growth. In this VAR, we employ two lags.

Table 2: Full-Sample Panel VAR Coefficients

This table shows VAR coefficients for our full-sample panel VAR. The state variables used are *yield*, the log of the capitalization rate, *lt.rate*, the log of the long-term interest rate, and *noi.growth*, the change in quarterly log NOI per square foot. The first panel shows the coefficients for the VAR using one lag and the second for the VAR using two lags. While we run these VARs for up to four lags, we omit tabulating the coefficients for our higher-lag VARs, to save space. The VARs are estimated through equation-by-equation OLS.

Dependent	<i>yield</i> _{<i>t</i>-1}	<i>lt.rate</i> _{<i>t</i>-1}	<i>noi.growth</i> _{<i>t</i>-1}	<i>yield</i> _{<i>t</i>-2}	<i>lt.rate</i> _{<i>t</i>-2}	<i>noi.growth</i> _{<i>t</i>-2}	$\overline{R^2}$	<i>F</i>
One Lag								
<i>yield</i> _{<i>t</i>}	0.593725 (383.23)***	0.097079 (44.06)***	-0.230434 (-138.05)***				0.3083	53069
<i>lt.rate</i> _{<i>t</i>}	0.013188 (26.87)***	0.908111 (1300.89)***	-0.007392 (-13.98)***				0.8306	583848
<i>noi.growth</i> _{<i>t</i>}	-0.271511 (-175.9)***	0.059176 (26.96)***	-0.215326 (-129.48)***				0.1846	26960
Two Lags								
<i>yield</i> _{<i>t</i>}	0.507721 (193.5)***	0.063548 (11.89)***	-0.17697 (-67.16)***	0.131734 (49.3)***	0.020304 (3.8)***	-0.113281 (-64.39)***	0.3145	25747
<i>lt.rate</i> _{<i>t</i>}	0.009432 (11.26)***	0.776312 (455.17)***	-0.005663 (-6.73)***	0.00334 (3.92)***	0.144518 (84.75)***	-0.004593 (-8.18)***	0.8337	281377
<i>noi.growth</i> _{<i>t</i>}	-0.313955 (-119.75)***	0.029017 (5.44)***	-0.198268 (-75.31)***	0.076831 (28.78)***	0.017996 (3.37)***	-0.093577 (-53.24)***	0.1929	13418

° : significance level < 10%. * : significance level < 5%. ** : significance level < 1%. *** : significance level < 0.1%.

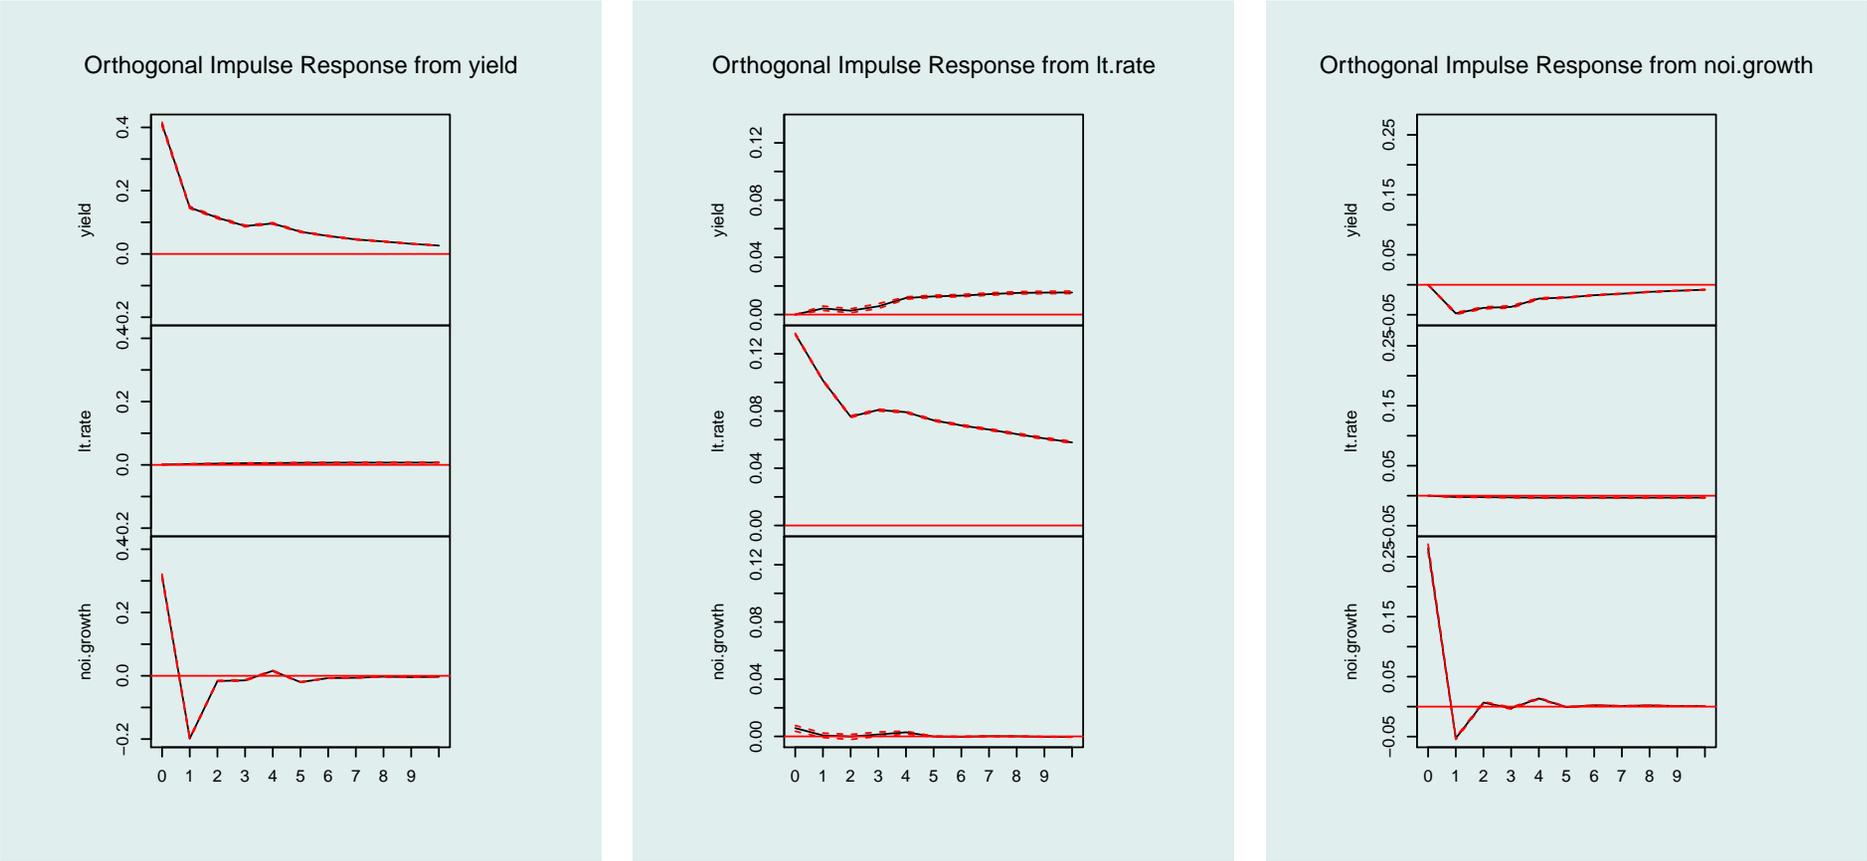


Figure 2: This figure shows impulse response functions from the full-panel VAR with four lags. The first panel shows the response on all state variables obtained by shocking *yield*, the second the response obtained by shocking *lt.rate*, and the third the response obtained by shocking *noi.growth*.

Table 3: Predictive Power of Panel Vector Autoregressions.

This table presents statistics comparing the predicted yields (log cap rates) from our VARs, with actual ex-post realized yields. Specifically, the table presents, for each VAR specification, the ratio of the standard deviations between the predicted and the realized series of yields. The first section presents this for the predictions (i.e. fitted values) from a set of full-sample panel VARs. The second section presents out-of-sample predictions for a set of rolling-window panel VARs. The predicted yields are for quarter $t+1$ and are generated by a VAR using 40 quarters' worth of data, ending at t . The variable $pred.yield_{i,t}$, in the first section is the fitted value from the VAR for this variable. The realized yield at quarter, for the same property t is $yield_{i,t}$. In the second section, for the rolling VAR $pred.yield_{i,t,t+1}$ is the prediction constructed in quarter t , for the yield for property i in quarter $t+1$, while $yield_{i,t+1}$ is the ex-post realized yield for the same property at that time. For the out-of-sample predictions we present both the ratio of standard deviations of the two series, as well as their correlation coefficient. In parentheses there is the value of a t-statistic testing the hypothesis that the actual correlation between the two series is 0. The set of state variables for the VAR consists of $yield_{i,t}$, $lt.rate_t$, the log of the long-term interest rate, and $noi.growth_{i,t}$, the quarterly difference of log NOIs. All variables are de-meanned.

Lags:	1	2	3	4
Full-Sample Panel VAR System: $[yield_{i,t}, lt.rate_t, noi.growth_{i,t}]'$				
$\sigma(pred.yield_{i,t})/\sigma(yield_{i,t})$	0.5552	0.5608	0.5673	0.5654
Rolling Panel VAR System: $[yield_{i,t}, lt.rate_{i,t}, noi.growth_{i,t}]'$				
$\sigma(pred.yield_{i,t,t+1})/\sigma(yield_{i,t+1})$	0.5299	0.5353	0.5397	0.5261
$cor(pred.yield_{i,t,t+1}, yield_{i,t+1})$	0.6177	0.6155	0.6175	0.6102
	(428.17)***	(398.13)***	(375.84)***	(345.71)***

° : significance level < 10%. * : significance level < 5%. ** : significance level < 1%. *** : significance level < 0.1%.

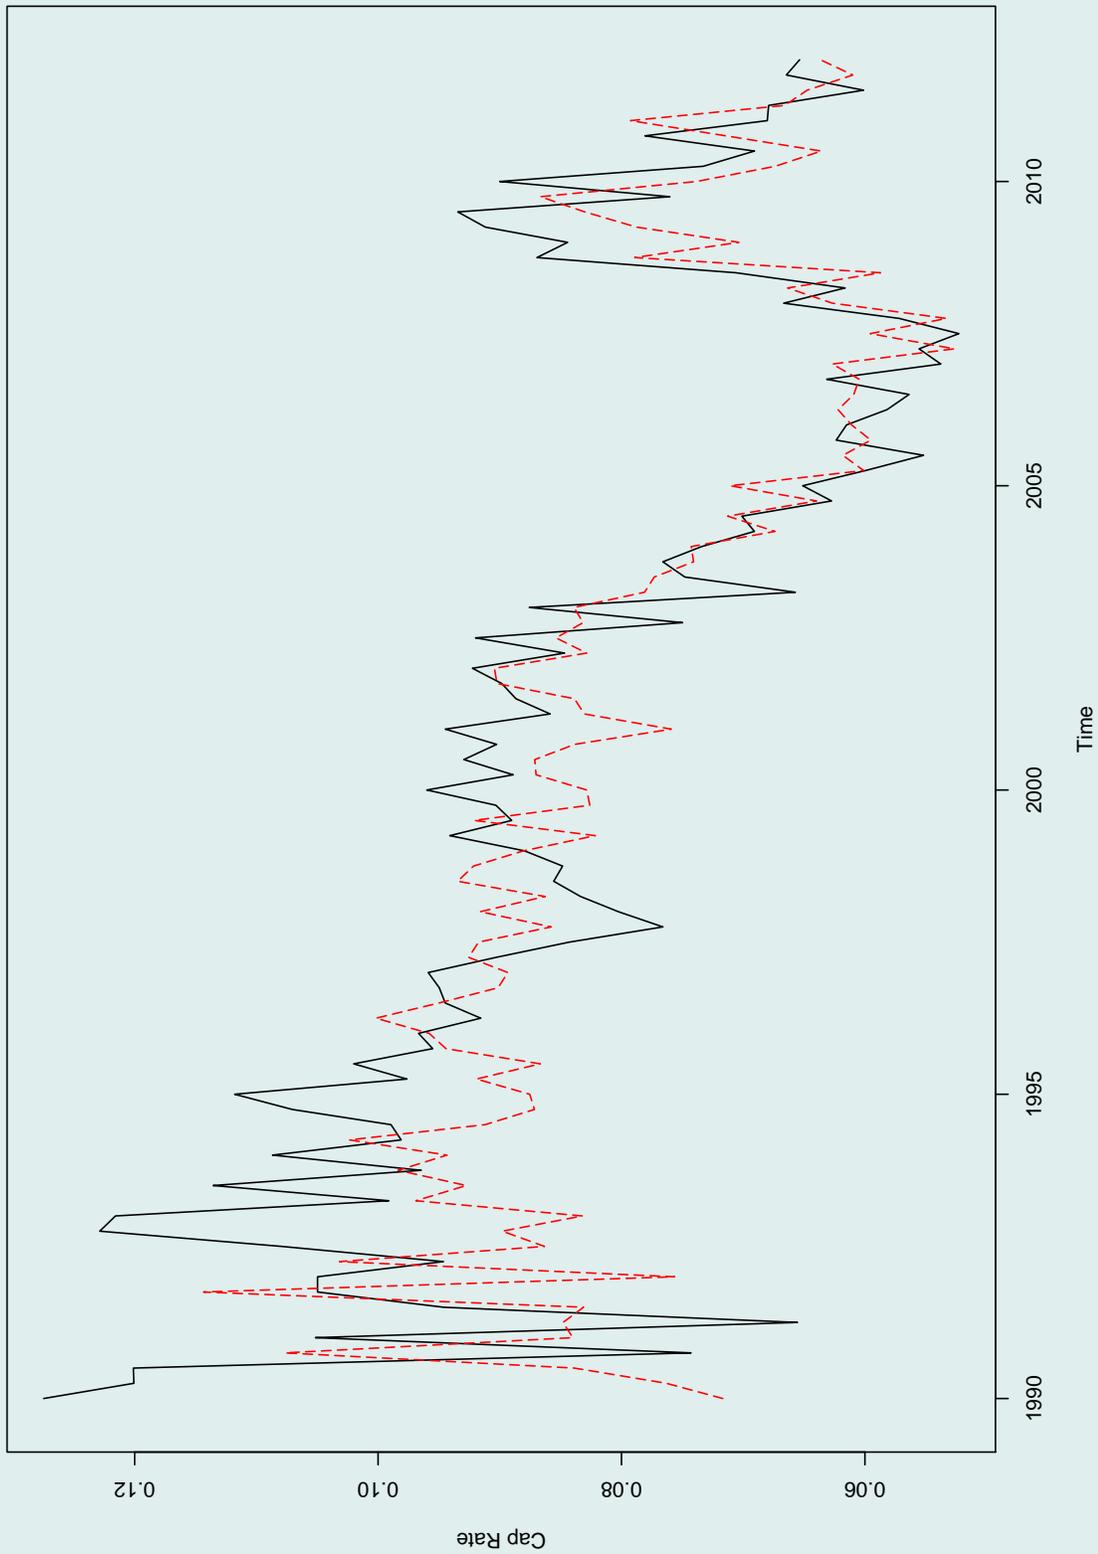


Figure 3: This figure shows, in the black solid line, the actual time series of cross-sectional average realized capitalization rates for all transactions in our data, and in the red dashed line, the time series of cross-sectional average predicted cap rates from our Self-Propagating Rolling Panel VAR.

Plot of Average CBSA Transaction Cap Rate, minus Average CBSA Predicted Cap Rate

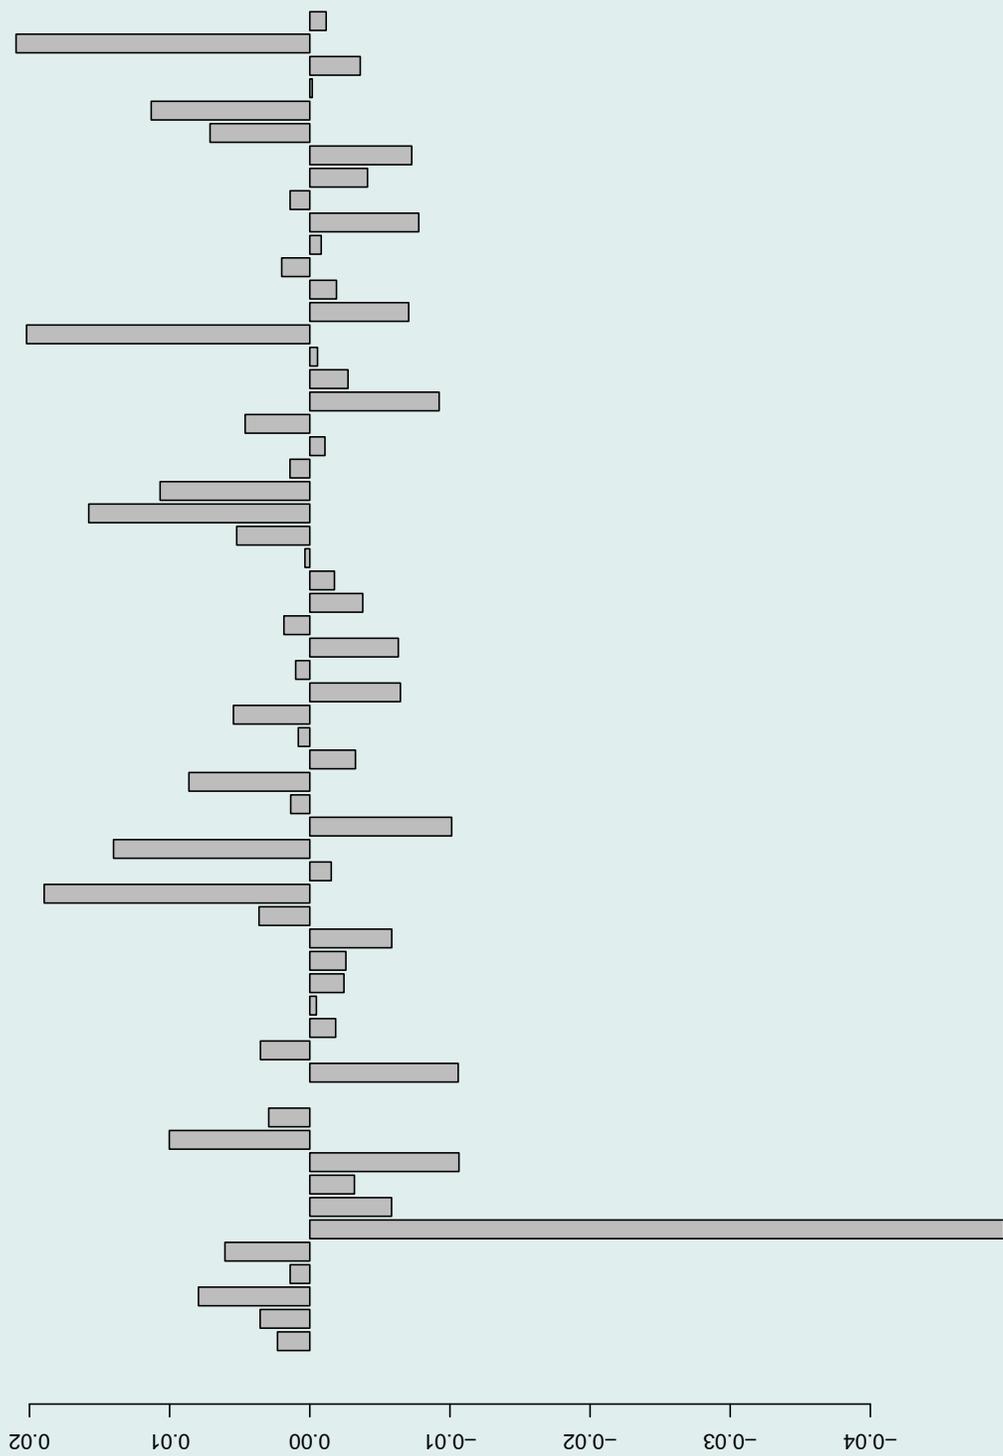


Figure 4: This figure shows, for each CBSA, the difference between the time-series average of actual transaction cap rate minus the time-series average of predicted cap rates from our Self-Propagating Rolling Panel VAR.

Table 4: Predictive Power of Self-Propagating Transaction-Based Panel Vector Autoregressions.

This table presents statistics comparing the predicted yields (cap rates) from our Self-Propagating Transaction-Based Rolling Panel VARs, with actual ex-post realized yields. The table first presents statistics that compare our predicted cap rates with appraisal-based cap rates. Then, the table shows how our predicted cap rates compare with ex-post realized transaction-based cap rates and for comparison also shows how appraisal-based cap rates compare with transaction-based cap rates. The statistics we present are ratios of standard deviations (*predicted/appraisal*, *appraisal/transaction*, and *predicted/transaction*, respectively), correlation coefficients between each pair of series, with t-statistics for the null hypothesis that the actual correlation between the two series is 0 in parentheses, and lastly, for comparison with actual transaction-based cap rates, an out-of-sample R^2 (R_{OOS}^2). The set of state variables for the VAR consists of *mixed.cap.rate_{i,t}*, *lt.rate_t*, the log of the long-term interest rate, *noi.growth_{i,t}*, the difference in log of NOI, and *noi.growth_{i,t}²*, the log of the squared difference of NOI. Mixed cap rate consists of local average cap rate, where transaction occur in the local market, and predicted cap rate from the rolling VAR system where they do not. The window size for the rolling VAR is 20 quarters for the combined dataset (*All Types*) and 30 quarters for individual property types. Panel A presents results from the entire available time period (1985-2012) for *All Types* and (1988-2012) for individual property types, while the subsequent panels show later subsets, as labeled. The frequency is quarterly.

Panel A: Full Sample

Measure	Predicted vs. Appraisal	Appraisal vs. Transaction	Predicted vs. Transaction
All Types			
Ratio of σ	0.823	0.9324	0.7
Correlation	0.3514	0.6756	0.3944
t-statistic	(117.67)***	(45.49)***	(21.35)***
R_{OOS}^2		0.438	0.7511
Apartments			
Ratio of σ	0.9217	0.9823	0.8922
Correlation	0.6239	0.7785	0.5699
t-statistic	(116.73)***	(33.58)***	(18.79)***
R_{OOS}^2		0.5713	0.844
Industrial			
Ratio of σ	0.8496	0.8995	0.6452
Correlation	0.2831	0.5887	0.271
t-statistic	(65.7)***	(23.71)***	(9.18)***
R_{OOS}^2		0.3168	0.6857
Office			
Ratio of σ	0.7504	1.0061	0.6834
Correlation	0.2444	0.7481	0.234
t-statistic	(35.67)***	(27.22)***	(5.84)***
R_{OOS}^2		0.5371	0.7361
Retail			
Ratio of σ	0.8166	0.9492	0.6764
Correlation	0.3549	0.7116	0.3288
t-statistic	(39.68)***	(17.25)***	(5.96)***
R_{OOS}^2		0.4832	0.8495

° : significance level < 10%. * : significance level < 5%. ** : significance level < 1%. *** : significance level < 0.1%.

Panel B: 1995 and After

Measure	Predicted vs. Appraisal	Appraisal vs. Transaction	Predicted vs. Transaction
All Types			
Ratio of σ	0.8171	0.9604	0.7246
Correlation	0.3586	0.7162	0.4033
t-statistic	(113.58) ^{***}	(48.71) ^{***}	(20.96) ^{***}
R^2_{OOS}		0.4932	0.7423
Apartments			
Ratio of σ	0.9006	1.0054	0.9008
Correlation	0.6168	0.7847	0.5609
t-statistic	(113.09) ^{***}	(33.92) ^{***}	(18.17) ^{***}
R^2_{OOS}		0.571	0.834
Industrial			
Ratio of σ	0.8345	0.9337	0.6611
Correlation	0.2911	0.626	0.2826
t-statistic	(63.67) ^{***}	(24.41) ^{***}	(8.97) ^{***}
R^2_{OOS}		0.3585	0.6712
Office			
Ratio of σ	0.7763	0.9936	0.7042
Correlation	0.2539	0.7682	0.2496
t-statistic	(35.21) ^{***}	(27.86) ^{***}	(6.01) ^{***}
R^2_{OOS}		0.581	0.7253
Retail			
Ratio of σ	0.8227	0.9185	0.6656
Correlation	0.3425	0.7426	0.338
t-statistic	(33.91) ^{***}	(17.81) ^{***}	(5.8) ^{***}
R^2_{OOS}		0.5505	0.8485

° : significance level < 10%. * : significance level < 5%. ** : significance level < 1%. *** : significance level < 0.1%.

Panel C: 2005 and After

Measure	Predicted vs. Appraisal	Appraisal vs. Transaction	Predicted vs. Transaction
All Types			
Ratio of σ	0.8056	0.9511	0.7232
Correlation	0.2568	0.7119	0.294
t-statistic	(65.37)***	(35.44)***	(10.77)***
R^2_{OOS}		0.492	0.6944
Apartments			
Ratio of σ	0.9323	0.923	0.8699
Correlation	0.3206	0.6919	0.3345
t-statistic	(39.2)***	(19.33)***	(7.17)***
R^2_{OOS}		0.4652	0.6749
Industrial			
Ratio of σ	0.8167	0.9109	0.6586
Correlation	0.2024	0.6346	0.1219
t-statistic	(36.98)***	(18.25)***	(2.74)**
R^2_{OOS}		0.3708	0.6471
Office			
Ratio of σ	0.7352	1.0193	0.7315
Correlation	0.1333	0.7947	0.2322
t-statistic	(14.7)***	(22.6)***	(4.13)***
R^2_{OOS}		0.6133	0.691
Retail			
Ratio of σ	0.8979	0.8497	0.6379
Correlation	0.0675	0.6954	0.1814
t-statistic	(4.68)***	(8.65)***	(1.66)
R^2_{OOS}		0.5694	0.8997

° : significance level < 10%. * : significance level < 5%. ** : significance level < 1%. *** : significance level < 0.1%.

Panel D: 2007Q3 and After

Measure	Predicted vs. Appraisal	Appraisal vs. Transaction	Predicted vs. Transaction
All Types			
Ratio of σ	0.7934	0.9302	0.7027
Correlation	0.2559	0.7045	0.2994
t-statistic	(56.16)***	(24.87)***	(7.88)***
R^2_{OOS}		0.5032	0.7314
Apartments			
Ratio of σ	0.9215	0.9504	0.8643
Correlation	0.3248	0.7336	0.3143
t-statistic	(33.8)***	(15.97)***	(4.9)***
R^2_{OOS}		0.5719	0.7361
Industrial			
Ratio of σ	0.7987	0.8918	0.6184
Correlation	0.1918	0.5713	0.1038
t-statistic	(30.56)***	(11.44)***	(1.72) ^o
R^2_{OOS}		0.2726	0.6263
Office			
Ratio of σ	0.7257	0.9784	0.7743
Correlation	0.1303	0.8565	0.1856
t-statistic	(11.96)***	(17.56)***	(2)*
R^2_{OOS}		0.7577	0.803
Retail			
Ratio of σ	0.9533	1.0115	0.629
Correlation	0.0695	0.7958	0.2424
t-statistic	(4.14)***	(8.52)***	(1.64)
R^2_{OOS}		0.6339	0.929

^o : significance level < 10%. * : significance level < 5%. ** : significance level < 1%. *** : significance level < 0.1%.

Table 5: Predictive Power of Rolling Panel VAR for Stocks.

This table presents statistics comparing the predicted log-dividend-yields from a Rolling Panel VAR, akin to the one run for the whole property sample in Table 3, run for the entire CRSP stock panel, from 1980 to 2012. The table shows how our predicted lag-dividend-yields compare with ex-post realized log-dividend-yields. The statistics we present are ratios of standard deviations, correlation coefficients between each pair of series, with t-statistics for the null hypothesis that the actual correlation between the two series is 0 in parentheses, and lastly, for comparison with actual log dividend-yields, an out-of-sample R^2 (R_{OOS}^2). The set of state variables for the VAR consists of the log-dividend-yield, the log of the long-term interest rate, and the log dividend growth. The frequency is quarterly.

Measure	Full Sample	From 1995Q1	From 2005Q1	From 2007Q3
Ratio of σ	0.5682	0.6009	0.6246	0.6385
Correlation	0.5075	0.5442	0.5575	0.5719
t-statistic	(481.68)***	(438.7)***	(293.19)***	(249.31)***
R_{OOS}^2	0.2538	0.2929	0.3062	0.3226

° : significance level < 10%. * : significance level < 5%. ** : significance level < 1%. *** : significance level < 0.1%.