

Using Cash Flow Dynamics to Price Thinly Traded Assets

Walter I. Boudry
School of Hotel Administration
Cornell University

Crocker H. Liu
School of Hotel Administration
Cornell University

Tobias Mühlhofer*
University of Miami

Walter N. Torous
MIT

October 26, 2015

Abstract

Are cash flows informative and predictive in valuing thinly traded assets? We investigate the extent to which cash-flow and discount-factor information plays a role in pricing thinly traded assets. We focus on pricing the various traded tranches in commercial mortgage-backed securities (CMBS) by developing an adaptation of the Campbell-Shiller dynamic Gordon growth model, which we term a *Self-Propagating Rolling-Window VAR*. We apply this to cash flows and actual bond prices. In contrast to stocks, we find that cash flows are informative in valuing thinly traded assets. Our predicted cash flow yields closely resemble ex-post realized transaction yields, and these predicted yields even outperform yields based on matrix prices especially for subordinated tranches. We also find that discount-factor information, while important is not as informative as cash flows in this setting, except after the financial crisis where the impact of discount-factor information increases somewhat. Our results provide a good representation of CMBS yields; investors can readily apply this algorithm to infer values of other types of thinly traded assets where cash flows are observable.

Keywords: Thinly traded assets, asset pricing, panel vector autoregression, Commercial Mortgage-Backed Securities.

JEL Codes: G12, R33

*The emails of the authors are: wb242@cornell.edu, chl62@cornell.edu, tmuhlhofer@bus.miami.edu, wtorous@mit.edu respectively. Contact author: Tobias Mühlhofer, School of Business Administration, University of Miami, Department of Finance, 514-S Jenkins Building, Coral Gables, FL 33131. Telephone: +1-305-284-9490.

1 Introduction

A debate exists on the relative importance of cash-flow and discount-factor information in pricing financial assets. At the heart of this question lies the discounted cash flow (DCF) pricing framework, which suggests a natural division among these two categories of information. Although the cash flows generated by a financial asset should play an important part in this asset's price formation process, in the stock market the variation in the discount factor seems to play a much more important role.

Our study contributes to this literature by examining the question regarding the informativeness of cash-flow- relative to discount-factor variability in a financial market other than the stock market, in which assets are thinly traded, and therefore prices are observed infrequently. On the other hand, in the setting we examine, cash-flow information is observable and readily available. Specifically, we examine the market for Commercial Mortgage-Backed Securities (CMBS) bonds which exhibits these characteristics. In addition to offering us this type of setting, the market for Mortgage-Backed Securities (MBS) bonds has been the subject of considerable interest in recent years, and our study helps shed light on this market's drivers of volatility, by modeling its price-formation process.

Thin trading can distort the measurement of return and risk and thus lead to an incorrect conclusion 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 fewer transactions exist in thin markets, these assets are less liquid with prices having higher volatility and typically larger bid-ask spreads. Even though a thin asset market exists, the underlying price formation process is not necessarily ad-hoc in nature if the asset in question generates regular observable 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.

In order to model the CMBS market's price-formation process, we model the the cash-flow yields that this market generates. To do this, we use as a basis the 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, by playing a role in determining a yield. The challenge in applying this framework lies in the thin trading of the market we examine, in that yields only exist when an asset trades. To overcome this, we develop a rolling-window panel VAR technique to estimate cash flow yields for untraded bonds, 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 75,000 bond 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 thus develop a methodology to infer yields of CMBS bonds in between trades. Having done this, we can then analyze the drivers of volatility in these yields.

We find that our VAR procedure generates yields that are statistically sound. Further, when compared to transactions where they do occur, our predicted yields capture a substantial portion of variation in transaction yields. Using state variables that capture both cash-flow and discount-factor information, we predict, out of sample, nearly 29% of variation in transaction yields. Omitting discount-factor information, we still explain nearly 20% of transaction yields out of sample. This indicates that cash-flow information plays an important role in this market's price formation process. For the time period after the 2008 financial crisis, the fraction of variation we predict using only cash flows drops to 15% suggesting a more pronounced role for discount-factor information in this time period. In sub-samples of bonds, we find that the role played by cash flows is even more pronounced. We further find that our model performs well both during and after the 2008 financial crisis, which indicates that even in highly volatile times, this market's yield dynamics are successfully modeled in our paradigm.

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 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.

Campbell and Ammer (1993) first examine the question of the relative importance of cash-flow versus discount-factor information in the setting of bond markets.¹ While Campbell and Ammer study liquid zero-coupon Treasury bonds of various maturities we concentrate on risky coupon bonds of various maturities that trade in thin markets. In contrast to Campbell and Ammer who investigate the relationship between bond and stock markets at an aggregate level, we focus on decomposing the returns on individual CMBS bonds. While we do not directly explore the link between bonds and stocks, we do include returns on real estate investment trust (REIT) stocks as one of our determinants of bond prices, in order to proxy for a real estate specific discount factor, if one exists.² Campbell and Ammer (1993) explain the low correlation between the returns on stocks and bonds by demonstrating that stock returns are driven mainly by news about future stock returns while bond returns are predominantly driven by news about future inflation rates.

¹Later, Campbell and Taksler (2004) also consider this question. However, the focus of that paper is explore the extent to which equity volatility acts as a suitable proxy for the variation in corporate credit spreads.

²This is also in line with previous papers, which find that low-grade bonds are more sensitive to stock returns (see e.g., Fama and French (1989, 1993)) and Cornell and Green (1991)).

This suggests that the news about future discount rates vis-a-vis future inflation rates impact bond returns.

Our study, on the other hand, addresses the question of what type of news drives the prices of risky bonds in illiquid, thinly traded markets where cash flows are observable. A distinguishing feature of our paper is our finding that cash flows are informative with respect to pricing CMBS tranches. The biggest improvement over matrix prices are for bonds in non-senior CMBS tranches of deals, which have more volatile (and therefore more informative) cash flows than senior tranches, but which are also much more difficult to price, in that they are much less liquid and show more price volatility. Consistent with the prior literature on stocks, we also find that expected returns influence pricing especially in the post-crisis period although cash flows exert a relatively stronger impact on price. In contrast to the stock market, most of the variation in the cash flow-to-price ratio for CMBS does not come from expected returns. For the CMBS market, the cash flow-to-price ratio does forecast growth in cash flows. We further explore whether a distinction in the price formation process exists between the AAA tranches which are the most liquid bonds 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 asset class. Another contribution of our study is that we develop a technique to infer cash flow yields for assets that are not continuously traded but for which continuous cash flow data exists. We observe little to no deterioration in performance of our proposed methodology during the financial crisis, during which liquidity dried up, making an accurate estimation of cash flow-to-price ratios especially important. This cash flow yield can then be applied to cash flow levels for inferring prices directly.

Although we apply this methodology to CMBS 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 commercial real estate, 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 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

In order to investigate the relative importance of cash flow and discount-rate information in our asset market's price formation process, we focus on modeling the dynamics of the cash-flow yields of CMBS bonds. Given a series of cash flows, an asset market determines a value for this series of cash flows; this value forms the basis for the price that will be paid in the market, subject to this market's competitiveness. This value (or price), conditional upon the underlying series of cash flows, can be expressed as the ratio known as the cash-flow yield. Modeling the dynamics of a market's cash flow yields based on exogenous sources of information, then reveals which of these information sources enter the market's price formation process and to what extent.

Specifically, like any investment asset, CMBS bonds are priced by an investor as the present discounted value of all future cash flows the asset entitles him to. While CMBS bonds have a finite life, a reduced form expression such as the Gordon Growth Model (which implies an indefinite life span for an investment) still offers an illustrative starting point for the dynamics of cash flow yields:³

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

Here, P_t is the price of the asset at time t , CF_{t+1} is the income cash flow produced by the bond 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 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 value to an asset. The observed yield (the denominator of Equation 1) represents the ex-post result of this price formation process, and the discount factor and the cash flow growth rate constitute the two broad categories of information that enter it.

³This is, in part, due to the fact that once a bond repays its face value, the investor will likely need to roll this capital over into a new investment.

To capture these two categories of information and make variance attributions to them, 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.⁴

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). For the purpose of our study, the parameter matrix of this VAR would show exactly how a particular submarket uses past cash flow and discount-factor information in order to construct yields and therefore prices. The VAR coefficient matrix becomes a quantitative representation of the market’s price formation process.

In our setting of a thinly traded market, we face an additional challenge in using this empirical framework. In order to estimate a VAR coefficient matrix to model the yield, we would require a full set of time series observations for the yield itself. However, since a yield requires a price, and a price requires a transaction, the result of working in a thinly-traded market is that we cannot construct a full time series of the yield itself. We overcome this problem by taking advantage of this empirical framework in the following additional way: once we have a VAR coefficient matrix estimated (and therefore modeled the market’s price formation process), we apply this coefficient matrix to existing information, in order to infer the yield of a bond in between trades as an out-of-sample VAR prediction. This becomes our *self-propagating* rolling VAR, whose mechanics are explained in detail in Section 2.3.

Given an observable series of underlying cash flows, the task that must be accomplished in

⁴In practice, however, other sources of exogenous discount-factor information are often necessary, and modeling this segment of the market’s information set empirically is difficult.

order to analyze the drivers of asset price volatility, as well as simultaneously infer the price of an untraded asset is to model the cash flow 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 model 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. While this methodology was developed for modeling equity prices, Campbell and Ammer (1993) show that red it also applies to the prices of finite-lived assets such as bonds. In fact, they argue that while this methodology constitutes an approximation for equity prices, it holds exactly for bond prices. The above 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, especially in a thinly traded market), is omitted.⁵ 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 dynamic Gordon growth model process. Such an observation justifies the use of this methodology to model inferred cash flow yields.

⁵Since the risk premium is not modeled directly, it becomes part of the VAR residuals. In theory, this should be the primary source of variation of these residuals.

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

$$z_t = Az_{t-1} + v_t \tag{4}$$

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 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

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

This prediction exercise fulfills a dual purpose in the context of our study. First, it allows us to infer yields for bonds between trades. Second, it serves to cleanly isolate sources of information: in fact, the predicted yields generated from this VAR will be affected exclusively by the state variables included in the VAR system. The amount of actual yield variation captured by the variation in predicted yields indicates how important a role the state variables in the VAR system play in the market’s price formation process.

2.2 In-Sample Estimation

To initially establish the feasibility of our approach, we conduct this procedure in-sample, using Matrix prices to construct yields for bonds that do not trade at a particular time period.

Since the main focus of our study is to infer the drivers of yield dynamics of individual bonds (and simultaneously to infer values of untraded assets), our empirical analysis focuses on a bond-by-bond level. This means we are dealing with panel data consisting of a cross-section of bonds over time, and so we modify the VAR approach for this setting⁶. In the panel setting, we stack

⁶VARs 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).

time series observations for individual bonds. 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 the yield for bond i at time t , through lagged values of the yield for the same bond, the common long-term interest rate (which is identical for all bonds at a particular time), and the lagged cash flow growth for that bond. Economically, the coefficient matrix estimated in this way characterizes the common price formation process applied to all bonds 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 across the entire panel. In the final implementation, we also estimate separate VAR systems for senior- and subordinated bonds, to allow the coefficient matrix to vary between market segments. All data is demeaned.

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, such studies often use fixed effects and clustering of standard errors at various levels of panel observation.⁷ 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

⁷Love 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.

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 force us to exclude a large part of the early portion of the data sample.

To assess the quality of predicted yields in this setting, the literature looks at two measures. The first is a ratio of the standard deviation of the series of predicted yields, divided by that of the actual realized yields. This shows the fraction of total volatility captured in the VAR; if the VAR system resembles the market's actual price formation process, this should be high (in theory close to 1). For the in-sample estimation, this suffices to assess the quality of the estimated yields. The second measure that is commonly used is the correlation coefficient between predicted and actual yields. While we report these in our out-of-sample estimation, it can be shown that for one-period predictions from in-sample estimation of a demeaned VAR system, the ratio of standard deviations will be equal to this correlation coefficient, and so we only report the former.

At this stage, for δ (which we call *yield* in the tables) we use the natural log of the ratio of bond-level coupon cash flow (interest plus any penalty) divided by the matrix price of the bond (or transaction price in months in which the bond trades). For r (which we call *lt.rate* in the tables) we use the long-term interest rate, and for Δd (which we call *cf.growth* in the tables) we use the difference in log of coupon cash flow per dollar of outstanding face value.⁸

2.3 Self-Propagating Rolling-Window VAR Procedure

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, as described above. 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 model yields and infer prices for thinly assets, we then expand this procedure to generate out-of-sample estimates of yields from a rolling-window VAR.

⁸Defining the cash flow measure in this way prevents unwanted effects inadvertently introduced by the bond's amortization according to a pre-determined schedule, which is also part of the cash flows paid to investors. By contrast, using total cash flows when a bond amortizes would otherwise initially raise the cash flow (as principal is being amortized in addition to interest payments), and then lower it (since the same percentage coupon is paid on less principal outstanding), even though, economically, the bond is still paying its regularly-scheduled interest service at the pre-determined coupon rate.

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 full-sample run described above, we initially use matrix prices to generate the yields that are input into the VAR estimation, to generate the coefficients. Subsequently, we then move away from such a heavy reliance on matrix prices, to a more transaction-based approach. The CMBS bond market is too thinly traded to estimate a VAR such as Equation 6 solely using yields based on transaction prices. Instead, we use a rolling-window VAR to populate the data with VAR-predicted yields where a transaction yield does not exist. The primary difficulty in this context becomes estimating the VAR system which is then used to predict yields, 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 bonds, $i \in 1, 2, 3, \dots, I$, 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 yield* for each bond, i at time t , which we initially populate with a matched transaction yield. That means, for each bond, i , in whose market segment, m , one or more transactions occur at time t , we set $mixed.yield_{i,t}$ equal to the average of transaction yields in market segment m at that time. We define a bond’s matched market segment m , by the interaction of its vintage (i.e. the year in which the deal was issued), and its initial credit rating (which proxies for the bond’s level of subordination within its deal). We believe that a matched market yield defined in this way constitutes a reasonable proxy for a bond’s own yield in the absence of trades. Applying this procedure fills in yields for some panel observations but still leaves the vast majority of yield observations blank. In practice, out of 17,290 month-vintage-rating combinations, 4,077, or 23.6% of matched segments 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 bonds present in the dataset from $t = 1$ to w , we have no choice but to fill in the missing yields using matrix prices as bond 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 yields (i.e. yields for properties in whose markets no matched yield 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 yield 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 are used in the estimation. The one-period forward predictions from this new VAR run are then used to populate the previously-unpopulated yield observations for $t = w + 2$. These 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 yields for $w + 3$, and so forth. It is thus clear 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 yield for each individual bond simultaneously, rather than an overall market yield. To see how this works, consider a panel of four bonds ($i = 1, \dots, 4$), for which we have estimated the coefficient matrix, as illustrated in Equation 6. The bond-specific predictions of the state variables are then calculated by multiplying the time- t matrix of bond-by-bond values of state-variable realizations with the transcript of the coefficient matrix⁹, 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 bond. As described above, we then take these yields and fill them in for bonds at time $t + 1$ where this field is unpopulated (i.e.

⁹In 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.

no matched yield exists). Before doing this, of course, they must be demeaned and exponentiated, as the VAR is run with demeaned 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, or otherwise to slowly decay to zero, like an impulse-response function. Our results show, however, that this is not the case. There are two reasons for this. First, in the panel, some new yield data is used, since matched transactions and some trades 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 bonds in this study. We choose to include bonds that, for a given number of lags L included in the VAR, have at least $2L + 1$ time-series observations. This restriction helps ensure that enough of a time series exists to estimate reasonable autoregressive coefficients, even within bonds and not just across them. Further, for bonds that first appear in the dataset at a time beyond the initial window from 1 to w (where we automatically filled unpopulated *mixed.yield* observations with matrix-price based yields), we populate any empty *mixed.yield* observations in the bond's first $2L + 1$ months of existence with matrix-price based yields. There are two points to be considered with regards to this rule. First, if we did not do this, we could only use bonds for our estimation which in their first $2L + 1$ quarters had matched transactions in their market. This would be a non-randomly drawn sample and would therefore introduce selection bias. 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 74,708 yields. To be parsimonious, and to use as few matrix prices as possible, we try to use a short window length, w , of 36 months. 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 substantial part of a market cycle. 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 two subsets of this market separately. We divide CMBS bonds into bonds that are rated AAA at issuance, and bonds that are not. This allows us to distinguish between bonds of senior tranches and those of subordinate tranches. Economically, it is conceivable that the price formation process (and therefore the VAR coefficients) for the different market segments might differ amongst each other.

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 cash flow 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 our case this inclusion should be beneficial. For the same purpose, in our full model we also include total returns to the REIT industry. This series should serve as an additional proxy for a real-estate specific discount factor, if this exists.

We further attempted our estimation by including other variables that might proxy for discount-factor news. These included credit spreads (in line with Campbell and Ammer (1993)), covariance with CMBS-market aggregate cash flows (as a type of β), as well as delinquency rates in a particular market segment. In all cases, each variable added noise to our estimation, resulting in predictive performance that was worse than what we attained without that variable's inclusion.

We construct the following statistics to examine the predictive power of our self-propagating VAR. As before, we first use the ratio of standard deviation of the yields we predict over the standard deviation of a realized (or matrix-price based) yield series. Since we predict out of sample, however, if we were to produce VARs that were extremely noisy, but whose estimated yields in no way resembled ex-post realized yields, the previously discussed measure would still be high (and could even be above one), but this would 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 realized (or matrix-based) yields.

We begin by comparing our predicted bond yields to yields implied by matrix prices. This allows us to gauge whether our VAR produces results that are statistically reasonable. While matrix prices themselves are not always indicative of transaction prices (and we therefore do not

strive to fully imitate this price series in order to make true variance attributions), this comparison is still a useful one, also because it allows us to assess the results over the entire sample of yields produced. In contrast, the comparisons that follow use only yields produced around bond trades.

Next, we assess our self-propagating VAR's ability to predict eventual transaction yields (when the bond is sold). To do this, we first produce the two statistics mentioned above, for predicted yields versus transaction yields. Since these are out-of-sample predictions of a value that is realized ex-post (as opposed to a matrix price which is also sort of a prediction), we then also introduce the additional measure of the 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 yield at time $t + 1$, δ'_{t+1} is the predicted yield and $\bar{\delta}_t$ is the historical average yield 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 yield. The more of an improvement the prediction offers over the historical mean, the closer to one 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.

Next, we compare our *mixed.yield* with transaction-based yields. Since mixed yields are augmented by transactions that occur in other market segments at a particular time, we compare, for every trade at time t , the mixed yield the month before (i.e. time $t - 1$), to ensure that we are truly examining a prediction, and not a value that potentially contains the information of the subject trade itself. Here, too, we report the ratio of standard deviations, the correlation coefficient, and the out-of-sample R-squared.

Lastly, we benchmark our predicted yields against matrix-price based yields. To do this, we report the same statistics as for *mixed.yield* $_{t-1}$, for the matrix price at time $t - 1$. Once again, we use the price the previous period since, if we were to report the matrix price at time t , this could potentially contain information from the subject trade, and therefore not be a prediction.

3 Data

3.1 CMBS Deal Structure

Our sample is made up of bonds from the Commercial Mortgage Backed Securities (CMBS) market. The CMBS market is the second largest source of mortgage funding for the commercial real estate market.¹⁰ A typical CMBS deal is based on a senior/subordinate cash flow waterfall. Debt service from the underlying pool of mortgages is passed through to the bonds in sequential order.¹¹ The principal component of debt service is paid first to the most senior bond class until it is fully paid off, then to the next most senior class until it is paid off, and so on. There is typically no overcollateralization in the market, so the face value of the bonds equals the face value of the mortgage pool. The sizing of the senior bond classes is set up in such a way as to create bonds that have targeted maturity dates of 3, 5, 7, and 10 years. By volume, the majority of bonds are 10-year bonds because the mortgages in the pool typically have a 10-year term.¹²

Interest payments from the mortgage pool are used to pay coupon payments on the bonds. This is once again done in a sequential fashion, with the AAA bonds typically all paid together, and the more junior bonds then paid in sequential order. Most deals have one or more excess interest only classes of bonds (IO bonds). These bonds receive the excess between the weighted average rate paid by the mortgages and the weighted average coupon paid to the bonds.

Losses in a CMBS pool are assigned in reverse order to principal. The most junior bond receives losses first, and once it is wiped out, the next most junior bond starts receiving losses, and so on. This mechanism combined with the sequential ordering of principal payments is what creates differing levels of credit risk in the bonds. The senior bonds get paid first and receive losses last, while the junior bonds get paid last and receive losses first.

Note that mortgage delinquencies do not necessarily lead to instantaneous cash flow changes to

¹⁰The mortgages in the CMBS market are senior liens and tend to be made only to stabilized properties. The underlying property types and credit quality is quite diverse and runs the spectrum of the commercial real estate market.

¹¹The underlying mortgages in the pool pay debt service on a monthly basis and this is passed through to the bondholders on a monthly basis also. The date of payment is known as the distribution date for the bonds and it is also the date on which remittance reports are released to the market.

¹²This also means that the majority of the bonds look like coupon bonds. Any amortization is paid to the most senior bonds, and all the 10-year bonds are paid back by the balloon payments on the mortgages.

the bonds, even to those bonds most exposed to credit risk. When a mortgage becomes delinquent or is in imminent default it is transferred to a special servicer. Under the servicing standard, the role of the special servicer is to maximize the net present value (NPV) of the loan. This obviously given them considerable leeway in dealing with delinquent borrowers. For delinquent loans, the special servicer will advance the debt service the loan should have made, to the extent that they believe it is recoverable. These servicer advances are designed to reduce cash flow volatility in the bonds.

Unlike in the RMBS market, prepayment plays less of a role in CMBS cash flow dynamics because of the typical prepayment restrictions placed on the underlying mortgages. Most commercial mortgages in the market will either be locked out from prepaying, require defeasance, or require yield maintenance. Of these mechanisms, only yield maintenance is going to lead to cash flow volatility in the waterfall. This is not to say there is no prepayment risk in the market, just that it is less significant than in the RMBS market.

3.2 CMBS Data

Our CMBS sample consists of bonds traded by insurance companies during the period 2003 to 2013. Unlike the corporate bond market where transactions are observable through the TRACE system, the CMBS market is a very opaque broker/dealer market. To our knowledge, the only investor group that systematically reports their CMBS trades are insurance companies. These trades are reported in Schedule D Parts 1, 3, and 4 of their regulatory filings. We obtain these filing from the NAIC and SNL Financial. In terms of magnitude, insurance companies are a key purchaser of CMBS, holding approximately 20-30% of the market's outstanding bonds at any given time. The one bias that using this sample does present, is that insurance companies don't hold the B piece of CMBS transactions, so the sample really represents the investment grade portion of the CMBS market.¹³

¹³The most junior portion of the CMBS deal, which is known as the "B piece" is bought by specialist B piece investors. This usually includes all the bonds up to a BBB- credit rating, and accounts for approximately 5% of the deal. The B piece is bought through a private transaction at origination and not usually traded. The B piece buyer is given more time and more information (they will have all the information the loan originator has) to do due diligence in making this initial purchase decision.

We calculate the traded price of the bond as the trade size weighted average of the reported dollar prices for that bond on a given day. Because the insurance company filing don't time stamp a trade, it is impossible to identify the last trade of the day. This is why we take the average for that day. The only restriction we place on the price is that the dollar price must be less than 140.¹⁴

We obtain cash flow information for the bonds from Trepp. Trepp is the leading provider of CMBS market information and the standard source used by market participants to obtain such information. Trepp collects and disseminates information contained in remittance reports. The remittance report contains information on the cash flows flowing into the pool from the mortgages and how this was allocated to each bond. It also contains information on mortgage delinquencies and also any potential commentary made by the special servicer regarding the delinquent mortgages. Remittance reports are made at the same time that cash flows are distributed to bond holders, so for any given distribution date, the information in the report and the cash flow on the bonds are known to the market at the same time.

From CMAAlert's Deal Database we obtain deal and bond level characteristics (bond type, bond credit rating, issuer, subordination level etc.) for each CMBS deal. Using this information we filter our sample to include only fixed rate bonds from fixed rate US CMBS deals.¹⁵ We include only P&I bonds (principal and interest bonds) and exclude all IO bonds.

Our final source of CMBS data is the Thomas Reuters Eikon database. This database provides us with proprietary matrix prices for each bond. Availability of matrix prices is the main limit on our data, as these prices are only available back to 2003.

To create cash flow yields for each bond on a monthly basis, we use the following procedure. For each distribution date, we collect bond cash flow information from Trepp. Because some bonds may be amortizing, we define our cash flow variables as a fraction of the outstanding principal balance. In this sense we don't treat lower coupon payments due to amortization as being informative. For any given month, we include any coupon payment and any prepayment penalties the bond receives as its cash flow for that month. We match the bond cash flow data from Trepp to the most recent

¹⁴By convention a dollar price of 100 is par value for the bond. An examination of the data suggests that these are very likely to be data errors and our conversations with CMBS traders also indicate that bonds don't trade at these prices in practice.

¹⁵WAC bonds are considered fixed rate for our analysis.

matrix price available in the Eikon database, and to the most recent transactions price available from the insurance company filings. For both the matrix price and the transaction price to be valid, we require that it occurred since the last remittance report (i.e. it occurred in the last month.) This ensures that the prices contain information about the latest distribution date.

In total our sample contains 3980 bonds from 496 CMBS deals and 283,319 bond/month observations. These bonds have an aggregate face value of \$652.5 Billion at origination. The earliest deal in the sample is from 1995, while the latest deal is from 2013. Of the 3980 bonds, 2063 have a AAA credit rating at origination and an aggregate face value of \$593.2 Billion, 1795 have a non-AAA investment grade rating and an aggregate face value of \$56.3 Billion, and 125 have a non-investment grade rating and an aggregate face value of \$3 Billion.

4 Results

We begin by examining the coefficient matrix from our full-sample panel VAR shown in Table 1. The state variables used are *yield*, the log of the capitalization rate, *lt.rate*, the log of the long-term interest rate, and *cf.growth*, the change in quarterly log cash flow. The table shows the coefficients for the VAR using one lag. 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.

Observe in the first line of the coefficient matrix that a very large portion of the variation in the yield (88 percent) is accounted for by the combination of a one period lag in the yield, long term interest rate, and cash flow growth.¹⁶ The t-statistics are high in this case, in part due to the large number of observations in our data, so we do not place great emphasis on these hypothesis tests. We highlight, rather, the signs of the coefficients. These are positive for the lagged yield (i.e. the yield is highly persistent), positive for the risk-free rate, and negative for cash flow growth. Since the interest rate and cash flow growth represent the denominator ($r - g$) in the Gordon growth

¹⁶In untabulated results, we find that a two-lag VAR model does not offer much additional information over a one lag VAR model in terms of each equation's R^2 . This seems to suggest that the market is fairly quick at integrating this information into prices. One caveat here, of course, is that the majority of yields comes from matrix prices, and so this statement may apply more to the algorithm used to construct these, rather than to the market's actual price formation process.

model, this suggests that our VAR does a good job in simulating this framework. Recall that both r and g are linear since we are using logs, which is the idea underlying the Campbell-Shiller method. More generally, it appears that prior information embedded in each of our variables is useful for prediction purposes.

Table 2 presents statistics comparing the predicted yields from our single panel VAR, with realized yields where available, and otherwise matrix-price based yields. Specifically, the table presents, for each VAR specification, the ratio of the standard deviations between the predicted and the realized/matrix series of yields. All predicted yields are fitted values from a set of full-sample panel VARs, with one, two, three, and four lags. The first line shows yields from a VAR system that contains only the Campbell-Shiller variables (yield, the risk-free rate, and cash flow growth). The second line shows yields from a system that, in addition, contains the two discount-rate variables (squared cash-flow growth and REIT returns).

It is apparent from Table 2 that, in sample, our VAR systems capture a substantial amount of the variation of CMBS yields. In fact, all ratios of standard deviations are at or above .94. The first line of the table indicates that cash flow information by itself already captures a large portion of the variation of CMBS yields as defined in this context. Comparing these values with those in the second line leads to the conclusion that adding discount-rate information to the VAR gives only very little improvement in the amount of variation captured. The same can be said about increasing the number of lags in the system: little more variation of yields is captured by the VAR in this way. It is important to view this result with caution, as in this setup the majority of the yields used consists of matrix-price based yields. However, these results should present evidence that overall the dynamic Gordon growth model seems to present a useful paradigm in which to model the price dynamics of CMBS bonds.

We now turn to the results illustrating the performance of our full estimation technique, the *Self-Propagating* Rolling-Window Panel VAR estimation of yields for CMBS bonds. Table 3 presents these results. Panel A shows the results for the full time series. This means that, given that our data starts in 2003 accounting for the lags needed in the VAR, and the length of our estimation window, the first prediction is for the beginning of 2006. The first section of the panel shows

how our model performs for all bonds, while the subsequent sections show results for estimations conducted over only the set of bonds rated AAA at origination (i.e. senior securities), and only the set of bonds rated Non-AAA (i.e. subordinated securities).¹⁷

The first column ($Pred_t$ vs Mat_t) compares our predicted yields to matrix prices at the same time. We report the ratio of the standard deviation of our predicted series over that of the matrix-based 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 matrix prices, we also report the correlation between the two series and a t-test, testing the null hypothesis of zero correlation. For all bonds, we find that we capture .372 of the variation, with a positive significant correlation coefficient of .1536. It should be remembered that matrix prices are not transaction prices, but are appraisals. Therefore, matching these as closely as possible is not a primary goal for our predicted yields. The main point of 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) is therefore unfounded. We can be confident of the reasonable performance of our procedure. Performing this comparison with matrix prices for this purpose has the added advantage that we have observations for nearly the entire panel, while, as previously stated, transactions tend to be scarce.

The second column ($Pred_t$ vs $Trans_t$) compares our predicted yields to actual realized transaction yields. The ratio of standard deviations is .7734 and the correlation coefficient is positive and significant at .5733. The Out-of-Sample R-squared (R_{OOS}^2) is .2881. All this indicates that our VAR estimation procedure based on a dynamic Gordon growth model captures a substantial amount of the variation of actual CMBS-bond transaction yields.¹⁸ This should indicate that the market's valuation process resembles the one we model. It should also be remembered that for much of our VAR estimation we do not have actual yields, but use the predicted yields from our VAR instead, in our *self-propagating* procedure. All this should indicate that, in our setting of

¹⁷Once again, we conduct separate estimation runs for each of the three segments, in order to allow the VAR coefficient matrices to vary among market segments.

¹⁸ R_{OOS}^2 below 10% are very common in the stock-market predictability literature.

a thinly traded market where, conversely, good cash flow information is available, the dynamic Gordon growth model is an effective way to characterize the market's price formation process.

The third and fourth columns ($Mixed_{t-1}$ vs $Trans_t$ and Mat_{t-1} vs $Trans_t$) respectively compare the *mixed.yield* the period before a transaction, as well as the matrix-price yield the period before a transaction, to the actual transaction yield. Viewed jointly, these compare the quality of our VAR predictions, with matrix prices. The third column of the table shows a correlation coefficient that is similar to that in the second column (.55 versus .57), with a higher ratio of standard deviations (.89 versus .77), and a lower R^2_{OOS} (.19 versus .28). This indicates that in this context our estimates are overall noisier than the ones reported in column 2, which makes sense as the estimate comes from one month before the trade. However, we still capture a high fraction of yield variation, especially when compared to matrix prices. For these (in column 4), there is a positive significant correlation (of .246), with a ratio of standard deviations of 1.89 and an R^2_{OOS} of -2.58 . This indicates that these matrix prices are extremely noisy, while our VAR produces much more stable estimates.¹⁹

The second section of Panel A shows results for senior bonds only (i.e. bonds that were rated AAA at origination). The first column indicates that here, too, we produce statistically reasonable estimates. The second column shows that here we predict a slightly lower fraction of the variation of transaction yields with the R^2_{OOS} , as the most telling statistic, at .22 instead of .29 (the ratios of standard deviations as well as the correlation coefficient are in line with this). This is despite the fact that overall variation in senior-bond yields is lower, and that therefore these should be easier to price. Since, in the section below for subordinated bonds, we also find that we predict yields less accurately than for the whole sample, this is consistent with the idea that subordinated bonds provide useful information for pricing senior bonds, and not having these in the estimation sample is detrimental to predictions.

The third column for senior bonds, however, reveals that, according to the R^2_{OOS} our mixed yields the period before a trade, predict transaction yields better in this subsample than they do in

¹⁹In further investigation that we do not tabulate, we find that matrix prices become especially problematic when bonds are in unusual situations, such as a deal's nearing the end of its life, or a bond becoming late on interest payments and then catching up.

the entire sample. Not only this, but mixed yields the period before a trade are actually a better predictor of transaction yields than pure predictions, the period of the trade. This is likely due to the matched yields that we use to form mixed yields. In the senior-bond segment, transactions are plentiful and so are matched yields. Further, there should be less heterogeneity among bonds in this universe than in the entire sample, and so the matched yields that are substituted are probably a closer proxy for the yield of the bond in question. The fourth column, once again, show that in this subset, matrix prices, though slightly better than before, are still noisy, with a ratio of standard deviations of 1.23 and an R_{OOS}^2 of $-.55$. Here, too, our VAR predictions seem to be more stable.

The final section of Panel A shows the results for subordinated bonds. While the first column shows that our estimates are statistically reasonable, the subsequent columns show that we predict the yields in this subsample less accurately (although we still capture a substantial amount of variation). While in the second column, the correlation coefficient is $.55$, the combination of the ratio of standard deviations ($.91$) and the R_{OOS}^2 ($.17$) indicate that in this subsample our predictions are slightly noisier and therefore less accurate. This is congruous with economic intuition, as variation in yields in this segment is higher, and transactions are more scarce and more heterogeneous. The third column shows that the mixed yield the period before the transaction suffers from more noise, with a higher ratio of standard deviations ($.97$) and a lower R_{OOS}^2 ($.11$). This also indicates that the argument about matched yields described for senior bonds does not apply in this segment, which, again is consistent with intuition. Lastly, column 4 indicates that in this segment the noise problem with matrix prices is especially severe, with a ratio of standard deviations of nearly 2 and an R_{OOS}^2 of -3.12 . This highlights how difficult it is to price bonds in this market segment. However, our estimation is still able to produce reasonable yields, even in this challenging environment.

Panel B shows a subsample of results from 2009 (i.e. post financial crisis) going forward. Overall, our results are qualitatively maintained, with primarily a slight increase in prediction accuracy from our mixed yields (column 3), for all bonds (R_{OOS}^2 of $.24$ instead of $.19$). For senior bonds, the accuracy of the mixed yields (column 3 again) decreases and is now lower than that of pure predicted yields. It is possible that in this subperiod there was more heterogeneity in yields within this universe, and that therefore the advantage of the matched yield was lost, in favor of

the disadvantage of making the prediction one month earlier. The predictions for subordinated bonds in this time period are of similar quality, with only slightly more noise for pure predictions (higher ratio of standard deviations, lower R^2_{OOS} and slightly less noise for matched yields). The noise problem with matrix prices is more severe in this subsample with higher ratios of standard deviations and more negative R^2_{OOS} . Overall, Table 3 shows that our *Self-Propagating Rolling Panel VAR* estimation procedure seems to constitute a good way to quantitatively model the price-formation process for the CMBS market.

Having established this, we can now address the question of the relative importance of cash-flow and discount-factor information in the CMBS market's valuation process. To accomplish this, we re-run our *self-propagating* VAR estimation, omitting discount-factor variables (squared cash-flow growth and REIT returns) and using only cash-flow variables (yield, the risk-free rate, and cash-flow growth). Table 4 shows the results from this estimation. We compare these results to those from the full system in Table 3, to assess the extent to which this reduced system captures less variation in yields when only cash-flow information is included.

The entire first column of Table 4, Panel A, looks qualitatively similar to that in the previous table. This indicates that once again our VAR procedure produces reasonable estimates. In columns 2 and 3 for all bonds, we find that the R^2_{OOS} is somewhat reduced by a (.29 to .20 for pure predictions and .20 to .15 for mixed yields). However, of more note is the level of predictive power that remains when only cash flow information is used. This suggests that in this thinly traded market where transactions cash flow information is good, cash flows play an important role in price formation. This point is made even more forcefully in the second and third sections, which show that the predictive power of our VAR setup is virtually unaffected in the subsamples of only senior- or only subordinated bonds. In these subsamples, virtually the entire predictive power of our VAR comes from our cash-flow variables, suggesting that these play a crucial role in the market's price formation process.

We conjecture that the reason for why we only see a drop in predictive power for all bonds, once again has to do with the information content that is provided across the two subsets. If a subordinated bond has a cash-flow shortfall, this affects its cash-flow volatility estimate, and

indicates to holders of more senior bonds that the mortgage pool underlying a certain deal is experiencing problems. This presents useful discount-factor information to holders of senior bond tranches (even though their cash flows are not affected, and most likely never will be). For the subordinated bond itself, on the contrary, this information would come too late, as the bond would likely soon be in default and cease to exist. For this reason, it would be consistent with economic intuition that discount-factor information received across the two sub-universes would be useful primarily for valuing bonds over the entire sample. It is in this sense that news about the future discount rate has an impact on prices.

As before, Panel B of Table 4 shows results that are qualitatively maintained. As before, even with the (small) reduction in explanatory power that comes from omitting discount factor information, the predictive ability of our VAR exceeds that of matrix prices.

Overall, these results show that in the CMBS market, which is characterized by thin trading, but good availability of cash flow information, a dynamic Gordon growth model constitutes an accurate way to characterize the market's price formation process. We further show that within such a setting, cash flow information plays an important role in valuing assets.

5 Conclusion

We model the price formation process of the CMBS market. This is an asset class characterized by thin trading and therefore poor pricing information, but good cash flow information. We compare the relative importance of cash-flow and discount-factor information in this price formation process. In order to be able to do so, we devise a *Self-Propagating* Rolling-Window Panel VAR technique, based on a dynamic Gordon growth model, with which we predict yields of bonds in between trades. Having done this, we then find that in this setting, while discount-factor information helps explain some variation in CMBS bond yields, cash-flow information plays a substantial role as a driver of volatility in this market. We thus shed light on the price formation process in the CMBS bond market, while at the same time exploring the relative importance of the two major drivers of asset price volatility in a setting of thin trading and good cash flow information.

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.
- Campbell, J. and Taksler, G.: 2004, Equity volatility and corporate bond yields, *Journal of Finance* .
- Campbell, J. Y.: 1991, A variance decomposition for stock returns, *The Economic Journal* **101**(45), 157–179.
- Campbell, J. Y. and Ammer, J.: 1993, What moves the stock and bond markets? a variance decomposition for long-term asset returns, *The Journal of Finance* **48**(1), 3–37.
- 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.
- 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.
- Cornell, B. and Green, K.: 1991, The investment performance of low-grade bond funds, *The Journal of Finance* **46**(1), 29–48.
- 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.
- Fama, E. F. and French, K. R.: 1989, Business conditions and expected returns on stocks and bonds, *Journal of financial economics* **25**(1), 23–49.
- Fama, E. F. and French, K. R.: 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* **33**, 3–56.
- 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.
- Ismail, B. and Kim, M.: 1989, On the association of cash flow variables with market risk: further evidence, *Accounting Review* pp. 125–136.

- 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.
- Petersen, M. A.: 2009, Estimating standard errors in finance panel data sets: Comparing approaches, *Review of financial studies* **22**(1), 435–480.
- Scholes, M. and Williams, J.: 1977, Estimating betas from nonsynchronous data, *Journal of financial economics* **5**(3), 309–327.
- 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: 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 cash-flow yield implied by matrix prices, *lt.rate*, the log of the long-term interest rate, and *cf.growth*, the change in quarterly log of bond cash flow. The table shows the coefficients for the VAR using one lag. 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>cf.growth</i> _{<i>t</i>-1}	$\overline{R^2}$	<i>F</i>
<i>yield</i> _{<i>t</i>}	0.9546 (1355.46)***	0.0153 (6.75)***	-0.4299 (-222.19)***	0.8738	612592
<i>lt.rate</i> _{<i>t</i>}	0.0064 (32.62)***	0.9415 (1492.07)***	-0.0044 (-8.19)***	0.8936	742909
<i>cf.growth</i> _{<i>t</i>}	-0.0376 (-59.76)***	0.0099 (4.92)***	-0.4357 (-252.01)***	0.2175	24598

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

Table 2: Predictive Power of Panel Vector Autoregressions.

This table presents statistics comparing the predicted yields 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. Both sections present this for the predictions (i.e. fitted values) from a full-sample panel VAR. The variable $pred.yield_{i,t}$, is the fitted value from the VAR for this variable. The realized (matrix-price based) yield is $yield_{i,t}$. 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 $cf.growth_{i,t}$, the quarterly difference of bond cash flows. In addition, the second system contains $cf.growth_{i,t}^2$, the log of the squared difference of bond cash flow, and $reit.ret_t$, the return to the REIT market. All variables are de-meanded.

Lags:	1	2	3	4
Full-Sample Panel VAR System:				
$[yield_{i,t}, lt.rate_t, cf.growth_{i,t}]'$				
$\sigma(pred.yield_{i,t})/\sigma(yield_{i,t})$	0.9348	0.9417	0.9461	0.9483
Full-Sample Panel VAR System:				
$[yield_{i,t}, lt.rate_t, cf.growth_{i,t}, cf.growth_{i,t}^2, reit.ret_t]'$				
$\sigma(pred.yield_{i,t,t+1})/\sigma(yield_{i,t+1})$	0.9359	0.9432	0.9476	0.9496

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

Table 3: Predictive Power of Self-Propagating Transaction-Based Panel Vector Autoregressions, for CMBS Bonds, Full Specification.

This table presents statistics comparing the predicted CMBS-bond yields from our Self-Propagating Transaction-Based Rolling Panel VARs, with ex-post realized yields. The table first presents statistics that compare our predicted yields with matrix-price-based yields. Then, the table shows how our predicted yields compare with ex-post realized transaction-based yields in the period of each transaction. We then show how our *Mixed* yields one month before the transaction compare to transaction yields, and for benchmarking, how the matrix-price yield the period before the transaction compares to the transaction yield. The statistics we present are ratios of standard deviations of the predicted over the realized series, 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 yields, an out-of-sample R^2 (R_{OOS}^2). The set of state variables for the VAR consists of $mixed.yield_{i,t}$, $lt.rate_t$, the log of the long-term interest rate, $cf.growth_{i,t}$, the difference in log of bond cash flows, $cf.growth_{i,t}^2$, the log of the squared difference of bond cash flow, and $reit.ret_t$, the return to the REIT market. Mixed yield consists of transaction yields (where they occur), matched yields, where transactions occur for a bond of the same vintage and initial rating, and predicted yield from the rolling VAR system where they do not. The window size for the rolling VAR is 36 months. Panel A presents results from the entire available time period (2006-2013), while the subsequent panel shows a later subset, as labeled. The frequency is monthly.

Panel A: Full Sample				
Measure	Pred_t vs. Mat_t	Pred_t vs. Trans_t	Mixed_{t-1} vs. Trans_t	Mat_{t-1} vs. Trans_t
All Bonds				
Ratio of σ	0.372	0.7734	0.8878	1.8882
Correlation	0.1536	0.5733	0.5502	0.246
t-statistic	(68.89)***	(102.29)***	(95.45)***	(36.77)***
R_{OOS}^2		0.2881	0.1948	-2.5804
AAA at Origination				
Ratio of σ	0.3954	0.4188	0.5429	1.2316
Correlation	0.4834	0.4765	0.48	0.3832
t-statistic	(174.23)***	(73.31)***	(73.34)***	(55.61)***
R_{OOS}^2		0.2221	0.2321	-0.5508
Non-AAA at Origination				
Ratio of σ	0.3323	0.9096	0.9661	1.9992
Correlation	0.1464	0.5497	0.5361	0.2027
t-statistic	(44.77)***	(35.33)***	(33.77)***	(11.01)***
R_{OOS}^2		0.1724	0.1109	-3.1212

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

Panel B: Post-crisis (2009 forward)

Measure	Pred _t vs. Mat _t	Pred _t vs. Trans _t	Mixed _{t-1} vs. Trans _t	Mat _{t-1} vs. Trans _t
All Bonds				
Ratio of σ	0.3074	0.8305	0.8466	2.0648
Correlation	0.1654	0.5785	0.5254	0.2086
t-statistic	(63.83) ^{***}	(90.87) ^{***}	(78.43) ^{***}	(27.1) ^{***}
R^2_{OOS}		0.271	0.2397	-3.2738
AAA at Origination				
Ratio of σ	0.3686	0.4048	0.515	1.4085
Correlation	0.4835	0.4648	0.3926	0.3172
t-statistic	(154.08) ^{***}	(62.59) ^{***}	(50.45) ^{***}	(39.53) ^{***}
R^2_{OOS}		0.2113	0.1766	-1.0388
Non-AAA at Origination				
Ratio of σ	0.3167	0.9217	0.9001	2.1583
Correlation	0.1308	0.5436	0.486	0.1669
t-statistic	(33.9) ^{***}	(30.14) ^{***}	(25.7) ^{***}	(7.82) ^{***}
R^2_{OOS}		0.1525	0.1342	-3.8132

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

Table 4: Predictive Power of Self-Propagating Transaction-Based Panel Vector Autoregressions, for CMBS Bonds, Cash-Flow-Only Specification.

This table presents statistics comparing the predicted CMBS-bond yields from our Self-Propagating Transaction-Based Rolling Panel VARs, with ex-post realized yields. The table first presents statistics that compare our predicted yields with matrix-price-based yields. Then, the table shows how our predicted yields compare with ex-post realized transaction-based yields in the period of each transaction. We then show how our *Mixed* yields one month before the transaction compare to transaction yields, and for benchmarking, how the matrix-price yield the period before the transaction compares to the transaction yield. The statistics we present are ratios of standard deviations of the predicted over the realized series, 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 yields, an out-of-sample R^2 (R_{OOS}^2). The set of state variables for this VAR consists of only $mixed.yield_{i,t}$, $lt.rate_t$, the log of the long-term interest rate, and $cf.growth_{i,t}$. Mixed yield consists of logs of transaction yields (where they occur), matched yields, where transactions occur for a bond of the same vintage and initial rating, and predicted yield from the rolling VAR system where they do not. The window size for the rolling VAR is 36 months. Panel A presents results from the entire available time period (2006-2013), while the subsequent panel shows a later subset, as labeled. The frequency is monthly.

Panel A: Full Sample

Measure	Pred _t vs. Mat _t	Pred _t vs. Trans _t	Mixed _{t-1} vs. Trans _t	Mat _{t-1} vs. Trans _t
All Bonds				
Ratio of σ	0.3406	0.9112	0.9599	1.8882
Correlation	0.1788	0.5634	0.5548	0.246
t-statistic	(80.54)***	(99.69)***	(96.62)***	(36.77)***
R_{OOS}^2		0.1964	0.1515	-2.5804
AAA at Origination				
Ratio of σ	0.3955	0.4538	0.543	1.2316
Correlation	0.4834	0.4753	0.48	0.3832
t-statistic	(174.22)***	(73.07)***	(73.35)***	(55.61)***
R_{OOS}^2		0.2254	0.2322	-0.5508
Non-AAA at Origination				
Ratio of σ	0.3335	0.9278	0.9589	1.9992
Correlation	0.1457	0.5505	0.5401	0.2027
t-statistic	(44.53)***	(35.4)***	(34.13)***	(11.01)***
R_{OOS}^2		0.1607	0.1246	-3.1212

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

Panel B: Post-crisis (2009 forward)

Measure	Pred _t vs. Mat _t	Pred _t vs. Trans _t	Mixed _{t-1} vs. Trans _t	Mat _{t-1} vs. Trans _t
All Bonds				
Ratio of σ	0.3285	0.9781	0.9345	2.0648
Correlation	0.1671	0.5679	0.5287	0.2086
t-statistic	(64.49) ^{***}	(88.4) ^{***}	(79.11) ^{***}	(27.1) ^{***}
R^2_{OOS}		0.1539	0.1776	-3.2738
AAA at Origination				
Ratio of σ	0.3687	0.4347	0.5153	1.4085
Correlation	0.4835	0.4625	0.3927	0.3172
t-statistic	(154.07) ^{***}	(62.21) ^{***}	(50.47) ^{***}	(39.53) ^{***}
R^2_{OOS}		0.2129	0.1767	-1.0388
Non-AAA at Origination				
Ratio of σ	0.3172	0.9221	0.8872	2.1583
Correlation	0.1296	0.5453	0.4885	0.1669
t-statistic	(33.57) ^{***}	(30.27) ^{***}	(25.87) ^{***}	(7.82) ^{***}
R^2_{OOS}		0.1547	0.1488	-3.8132

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