



BANK OF ENGLAND

Working Paper No. 520

A forecast evaluation of expected equity return measures

Michael Chin and Christopher Polk

January 2015

Working papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee or Financial Policy Committee.



BANK OF ENGLAND

Working Paper No. 520

A forecast evaluation of expected equity return measures

Michael Chin⁽¹⁾ and Christopher Polk⁽²⁾

Abstract

Recent studies find evidence in favour of return predictability, and argue that their positive findings result from their ability to capture expected returns. We assess the forecasting performance of two popular approaches to estimating expected equity returns, a dividend discount model (DDM) commonly used to estimate ‘implied cost of capital’, and a vector autoregression (VAR) model commonly used to decompose equity returns. In line with recent evidence, in-sample tests show that both estimates generate substantially lower forecast errors compared to traditional predictor variables such as price-earnings ratios and dividend yields. Out-of-sample, the VAR and DDM estimates generate economically and statistically significant forecast improvements relative to a historical average benchmark. Our results tentatively suggest that the VAR approach better captures expected returns compared to the DDM.

Key words: Expected returns, implied cost of capital, dividend discount model, return predictability, forecasting.

JEL classification: G10, G11, G12, G17.

(1) Macro Financial Analysis Division, Bank of England. Email: michael.chin@bankofengland.co.uk

(2) Department of Finance, London School of Economics. Email: c.polk@lse.ac.uk

The views expressed in this paper are those of the authors, and do not necessarily reflect those of the Bank of England, the Monetary/Financial Policy Committees or the PRA Board. We would like to thank Katie Farrant, Andrew Meldrum, Simon Price, Chris Young and an anonymous referee for extensive and valuable comments. This paper was finalised on 18 December 2014.

The Bank of England’s working paper series is externally refereed.

Information on the Bank’s working paper series can be found at www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

Summary

Central banks pay close attention to movements in equity prices, as they can have important implications for the real economy. For example, increases in equity prices may increase the value of individuals' investments and pensions, and can lower firms' cost of financing new projects, and these effects can lead to increases in aggregate consumption and investment. Understanding why equity prices have changed is also important. An investor who holds a firm's equity is entitled to a share of the profits that the firm generates in the immediate and distant future. In theory, this investor is therefore willing to pay a price for that equity that reflects expectations of these future cash flows (which are often dividend payments). At the same time, however, the equity price will not exactly equal the total expected future sum of dividends, since in general investors dislike having to wait for dividend payments, and also need to be compensated for the possibility that dividends rise and fall depending on the future state of the economy. As a result, when evaluating the price of an equity, investors discount expected dividends by an additional amount reflecting the return on a safe asset and a premium for risk, which is the expected return.

At the aggregate level, expected returns are important equity market indicators, as they summarise the attitudes toward risk of a range of investors. But they cannot be observed directly, and one needs a model to estimate them. In theory, these expected return estimates should be able to forecast actual future returns on the stock market to an extent. Academics and practitioners have devoted substantial effort to testing whether equity returns are indeed predictable, often using variables such as dividend-price and price-earnings ratios. Successes in this area have, in general, been attributed to the ability of a forecasting measure to capture expected returns. Put another way, the better the measure of expected return, the better that estimate should be at forecasting future returns.

Our paper focuses on two competing measures of expected returns, and examines their ability to forecast returns on the equity market. The first is estimated using Campbell's 1991 vector autoregression (VAR) model, a simple statistical model that describes the relationship between short-term returns and a range of other variables. These short-term dynamics can be used to make predictions about returns over longer horizons, and we use these as one measure of expected returns. The second measure is estimated using an adaptation of Gordon's 1962 dividend discount model (DDM). This approach directly models investor dividend expectations using analyst survey measures, and then solves for the expected return as the difference between these expectations and the equity price. We examine the ability of these measures to forecast future returns in a range of tests. In addition to the two models, we also consider a selection of popular variables used to forecast returns, and ask whether the expected return estimates forecast better compared to these predictors.

We use UK and US data to examine the forecasting performance of the range of variables over short and long forecast horizons (from three months to three years). We initially consider each forecasting variable individually, and run thorough statistical tests of their predictive power. Here, we find that the DDM and VAR expected return measures perform well, whereas the more traditional predictors do not. In a related test, we compare the relative accuracy of the

expected return measures to the traditional predictors, and find that the former generate much smaller errors when predicting returns. Finally, we consider how these predictors would have done in 'real-time', where we imagine we had to generate forecasts on a month-by-month basis using only information available to us at the time of making each forecast. Previous work has found that forecasting performance is often very weak under these test conditions. In fact, in many cases, it has been shown that one can achieve better forecast performance by simply using the past average return on the market to predict its future direction. In contrast, we show that our expected returns estimates perform better compared to this simple rule, in particular when forecasting returns at longer horizons, which is strong evidence in favour of the approaches we consider.

1 Introduction

Many models have been used to understand movements in equity prices, and a large proportion of these rely on a present-value framework, where the equity price is expressed as the expected sum of its future cash flows discounted at an appropriate rate. This discount rate comprises a risk-free element and an equity risk premium, the compensation an investor requires to hold an equity, and is also known as the expected return on equity. The expected return on the aggregate equity market is an important concept in asset pricing and macroeconomics, as it reflects perceptions of aggregate risk and risk-aversion, and is directly related to firms' cost of capital and household wealth. Within the finance literature, expected returns have been the focus of much theoretical and empirical work, where there is a lot of evidence to suggest that they vary substantially over time and play a large role in stock market movements (see Cochrane (2011)). Expected returns are unobservable, however, and need to be estimated.

A consequence of time-varying expected returns is that they allow for the presence of predictability within the equity market, and there is a large related literature examining aggregate return predictability. The existence of return predictability has been a somewhat controversial topic, where earlier positive findings have been challenged on the basis of issues regarding statistical inference in regression-based tests, and the (lack of) out-of-sample performance of forecasting models (Welch and Goyal (2008)). Recent studies, however, have provided evidence in favour of return predictability. Li, Ng and Swaminathan (2013) show that an expected return estimate, generated from an adaptation of the Dividend Discount Model (DDM) framework of Gordon (1962), out-performs a range of traditional forecast variables within in- and out-of-sample forecast tests. In addition, Rapach, Strauss and Zhou (2010) show substantial improvements in forecasting performance can be achieved through combining information contained in a range of individual predictor variables. The aforementioned studies argue that the improvements in forecast performance arise due to the fact that their forecasting models better capture *expected returns*.

This paper considers two popular approaches to modelling expected equity returns. The first approach also uses the Gordon (1962) DDM framework, in a similar formulation to that proposed in Li, Ng and Swaminathan (2013) and Chen, Da and Zhao (2013). This approach uses analyst forecasts of future dividends in a present-value formula, and solves for the expected return used to discount these cash flows by residual. The second approach models the short-run dynamics of equity returns in a Vector Autoregression (VAR), originally proposed in Campbell (1991), which is more commonly used to decompose equity returns

into cash flow and discount rate news. The discount rate news in these studies simply reflects changes to the expected return forecasts made by the VAR, and this expected return estimate is the focus of our study. The VAR approach can be thought of as combining information in a range of individual forecast variables. We use a standard VAR specification to generate expected return estimates, which closely resemble the models used in recent studies (e.g. Campbell and Vuolteenaho (2004), Campbell, Giglio and Polk (2013), Campbell, Giglio, Polk and Turley (2014)).

We test whether the expected return measures are able to forecast realised returns, both in- and out-of-sample, and compare their performance to a range of traditional predictor variables such as the price-earnings ratio and term spread, using UK and US data. The expected return measures perform favourably in simple forecast regression tests, where they significantly predict realised returns at a range of horizons, and where the traditional predictors generally do not. In-sample, the VAR and DDM estimates also generate substantially lower forecast errors compared to the alternative predictors. Out-of-sample, we compare the range of forecast variables to a historical average benchmark forecast, and find that the VAR and DDM offer economically and statistically significant forecast improvements. In line with recent findings, we provide evidence that direct measures of expected returns perform well in forecast tests, and our paper generally provides evidence in favour of return predictability.

We also explore the relative performance of the VAR and DDM approaches in measuring expected returns. The VAR and DDM approaches are both housed within the present value framework, but the models differ in their estimation, where one chooses to model either cash flows or discount rates as a residual. Since the DDM expected return is estimated by residual, it is important that the analyst forecast inputs, usually taken from I/B/E/S surveys, are well-measured. There are, however, well-known measurement issues and biases in I/B/E/S earnings forecasts, which have the potential to distort the expected return estimates. In particular, the surveys can be slow to update in response to news, and the forecasts can also be subject to large data corrections (Ljungqvist, Malloy and Marston (2009)). There are also potential limitations to the VAR approach, where this method may struggle to capture long-horizon expectations, and where outputs from this approach clearly depend on the variables included within the model. Within our in- and out-of-sample forecast evaluations, we attempt to explore the extent to which these issues are detrimental to the forecasting ability of these approaches. In-sample, we find that the VAR-based estimates generate substantially lower forecast errors compared to the DDM. Within the out-of-sample tests, however, the performances of the VAR model and DDM are harder to differentiate.

Few studies have assessed the relative merits of the DDM and VAR approaches, and as



well as providing evidence on their relative forecasting performance, we also explore the return decompositions from both models. Chen, Da and Zhao (2013), noting the limitations of the VAR and DDM approaches, use both approaches to decompose equity returns at the firm and aggregate levels, and find a much larger role for expected cash flows in return variation, relative to the preceding VAR and predictive regression-based evidence (e.g. Campbell and Ammer (1993), Ammer and Mei (1996), Vuolteenaho (2002)). We find that the DDM and VAR find similar relative roles for cash flow and discount rate news, though we show that they often differ in their interpretations of stock market movements at given points in time.

We also contribute to the return predictability literature by demonstrating that an alternative DDM expected return estimate also successfully predicts returns. While our DDM estimate closely resembles that used in Li, Ng and Swaminathan (2013), there are some differences compared to our model. First, the previous versions jointly model earnings forecasts and ‘plowback’ rates – the proportion of earnings paid out as dividends in each period. We use I/B/E/S dividend forecasts directly, which removes the need to make any assumptions about plowback rate expectations, since these will already be captured in the dividend forecasts. Second, the previous versions use an expanding moving average of GDP growth rates as their long-term growth assumption, whereas we assume a constant long-run growth rate. While the differences are not large, we demonstrate that the forecasting performance is robust to changes in specification and a different forecast sample period. Finally, our application to UK data also adds to a relatively small literature decomposing UK equity returns (e.g. Cuthbertson, Hayes and Nitzsche (1999), Bredin, Hyde, Nitzsche and O’Reilly (2007)) and assessing return predictability in the UK (e.g. Wells and Wetherilt (2004), Price and Schleicher (2005)).

The paper proceeds as follows: Section 2 briefly outlines the theoretical foundations of the two approaches, and provides details of their empirical implementation. Section 3 examines the outputs and diagnostics of the models, and draws a comparison between the different approaches. Section 4 tests the ability of the models to predict future returns in and out of sample. Section 5 concludes.

2 Modelling Expected Equity Returns

2.1 VAR Model: Campbell (1991)

The VAR approach is associated with the return decomposition originally proposed in Campbell (1991). A log-linear approximation expresses unexpected returns in terms of revisions

to both expected future dividend growth and expected returns.¹

$$r_{t+1} - E_t r_{t+1} \approx (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{t+1+j} \quad (1)$$

Here, r_{t+1} is the log total return from holding a stock from period t to period $t + 1$, d_{t+1} is the log dividend paid during period t , ρ is a by-product of the log-linearization of returns, defined as the ratio of steady-state prices to prices plus dividends, which empirically is a number a little less than one. E_t is the conditional expectation formed at time t , and Δ denotes a one-period backward difference operator. We refer to the revisions to future dividends and expected returns as ‘cash flow news’ and ‘discount rate news’ respectively, and rewrite the identity using a tilde to denote revisions to each of the components:

$$\tilde{e}_{t+1} = \tilde{e}_{CF,t+1} - \tilde{e}_{DR,t+1} \quad (2)$$

2.1.1 VAR Implementation

The empirical implementation of this framework requires a model for generating forecasts of dividends and returns. We follow a common approach and avoid modelling dividends, modelling only the expected return components in a VAR, and calculating revisions to expected dividends as the residual component of the identity in (2) (e.g. Campbell (1991), Campbell and Vuolteenaho (2004), Bernanke and Kuttner (2005), Campbell, Giglio and Polk (2013)).² We define an m -variable state vector z_{t+1} , where the first element is the stock return r_{t+1} . We assume that z_{t+1} follows a first-order VAR process:

$$z_{t+1} = a + \Gamma z_t + u_{t+1} \quad (3)$$

z_{t+1} is the m -by-1 state vector, Γ is the m -by- m matrix of VAR coefficients, and u_{t+1} is an i.i.d. m -by-1 shock vector. The VAR model can be used to generate forecasts of single period returns at any point in the future, where the h -period ahead return forecast is given by:

$$E(z_{t+h}) = (I - \Gamma)(I - \Gamma^h)a + \Gamma^h z_t \quad (4)$$

Expected returns are generated for a range of horizons, where the h -period expected return is the sum of the single-period log returns over the forecast horizon. These expected

¹See Campbell and Shiller (1988) and Campbell (1991) for detailed derivations.

²This is motivated by the complications that arise when modelling the dividend growth process, which has a near unit root, exhibits strong seasonal variation, and is subject to changes in corporate pay-out policies.

return estimates are the subject of the forecast evaluation tests later on. One can use revisions to these forecasts from one period to the next as measures of the components of equation (2). Specifically, the left-hand side $r_{t+1} - E_t r_{t+1}$ can be written as:

$$\tilde{e}_{t+1} = \mathbf{e}\mathbf{1}'u_{t+1} \quad (5)$$

where $\mathbf{e}\mathbf{1}$ is the first column of a m -by- m identity matrix. Revisions to expected returns and dividends can be written as:

$$\tilde{e}_{DR,t+1} = \mathbf{e}\mathbf{1}'\lambda u_{t+1} \quad (6)$$

$$\tilde{e}_{CF,t+1} = (\mathbf{e}\mathbf{1}' + \mathbf{e}\mathbf{1}'\lambda)u_{t+1} \quad (7)$$

where $\lambda = \rho\Gamma(I - \rho\Gamma)^{-1}$. In addition to the market return, we include variables in the VAR that have been shown to forecast equity returns, and that are commonly used in applications of this framework (e.g. Campbell and Vuolteenaho (2004), Bernanke and Kuttner (2005), Campbell, Giglio and Polk (2013) and Campbell, Giglio, Polk and Turley (2014)). We include a smoothed dividend-price ratio, defined as the index price at time t scaled by the trailing average of 10 year earnings (Campbell and Shiller (1988)). This ratio, or a close variant of it, is one of the most frequently used predictor variables within studies that implement the Campbell (1991) framework.³ We include the term spread, defined as the 10-year constant maturity government bond yield less the yield on the 3-month Treasury bill (Keim and Stambaugh (1986)), also commonly used in related studies. In the US VAR, we include a default spread defined as the difference between BAA and AAA corporate bond yields. For the UK model, due to the short sample of available data for the UK default spread, we include a pay-out ratio, defined as the ratio of aggregate dividends paid to aggregate earnings over the past twelve months, another commonly used predictive variable (Lamont (1998)). This model ought to provide a good reflection of the VAR specifications that have been used in previous studies, though the model more closely resembles the specifications in Campbell and Vuolteenaho (2004), Campbell, Giglio and Polk (2013) and Campbell, Giglio, Polk and Turley (2014).

Recent evidence suggests that the inclusion of the price-dividend ratio and short-term interest rate capture a significant amount of the relevant information for modelling expected returns. Maio and Philip (2014) test the sensitivity of the return decomposition to the inclusion of a large amount of information extracted from alternative macroeconomic and

³As described in Campbell and Ammer (1993) and Engsted, Pederson and Tanggaard (2012), it is crucial that a stock price variable is included within the VAR, since this is an underlying assumption used to derive the Campbell (1991) framework. The inclusion of the price-earnings ratio ensures that this condition is met.

financial variables. They find that there is little incremental information in the extra series, over and above the dividend yield and short-term interest rate. To the extent that our VAR specification includes information contained in both of these predictor variables, we might take some reassurance that our model provides a reasonable representation of the market's information set, given that one can only ever hope to capture a subset of this information.

2.2 Dividend Discount Model

The common starting point of the DDM approach is the approximation of the standard present value equation. The standard two-period equation is as follows:

$$P_t = E_t \left[\frac{P_{t+1} + D_{t+1}}{1 + r_{e,t}} \right] \quad (8)$$

Here P_t is the index price in month t , $E_t[D_{t+1}]$ is the dividend forecast for year $t + 1$, and r_e is the cost of capital. This can be extended to an infinite horizon model, which rules out the existence of bubbles:

$$P_t = E_t \sum_{k=1}^{\infty} \frac{D_{t+k}}{(1 + r_{e,t+k})^k} \quad (9)$$

The DDM approach is then a special case of the infinite horizon present value equation, which makes the further assumption that expected returns are a deterministic constant, and allows the expectations operator to be moved inside the summation sign – applying only to future dividends:

$$P_t = \sum_{k=1}^{\infty} \frac{E_t(D_{t+k})}{(1 + r_e)^k} \quad (10)$$

There are many ways to implement this version of the model empirically. The majority of applications, however, tend to explicitly model the path of expected dividends over two to three years from period t . We use the following adaptation:⁴

$$P_t = \sum_{k=1}^5 \frac{E_t(D_{t+k})}{(1 + r_e)^k} + \frac{E_t(D_{t+6})}{(1 + r_e)^5(r_e - g)} \quad (11)$$

Dividends from $t + 1$ to $t + 4$ are measured using I/B/E/S survey data. Specifically, we use aggregate I/B/E/S growth rate forecasts for years 1 to 3 to calculate the level of expected dividends over these three years.⁵ We then assume that year 4 earnings grow at

⁴This is a similar model to that used in Inkinen, Stringa and Voutsinou (2010).

⁵This differs slightly to Li, Ng and Swaminathan (2013) and Chen, Da and Zhao (2013), who use individual

the average of the year 3 growth rate and the I/B/E/S medium term growth rate. Year 6 growth is set equal to a long-run growth assumption of 5%, which allows calculation of a perpetuity starting at this point in time.⁶ The final Year 6 cash flow is valued using the constant growth perpetuity formula, hence the appearance of $(r_e - g)$ in the denominator of the final cash flow. By assuming that the discount rate is constant at each maturity, we can solve for the implied discount rate numerically.

Our model closely resembles that used in Chen, Da and Zhao (2013) and Li, Ng and Swaminathan (2013), which in turn are based on the formulation of Pastor, Sinha and Swaminathan (2008), though there are some differences compared to our model. First, the other versions jointly model earnings forecasts and ‘plowback’ rates – the proportion of earnings paid out as dividends in each period. We use dividend forecasts directly, which removes the need to make any assumptions about plowback rate expectations, since these will already be captured in the I/B/E/S dividend forecasts. Second, whereas our long-term growth rate is constant, the previous versions use an expanding moving average of GDP growth rates. In practice, this difference does not materially affect the DDM estimate. Since the expanding moving average is based on growth rates over a long history (since 1930), the long-term growth rate over our sample (1998–2014) changes little. As a result, changes in the discount rates are almost perfectly correlated (0.99). There is a small difference in the levels of the constant and average long-term growth rates, which generates a small level difference between the discount rate estimates under the two different assumptions.

2.2.1 Dividend Discount Model: Cash flow and Discount Rate News

Following Chen, Da and Zhao (2013), we estimate the impact of revisions to cash flow and discount rates on equity returns, allowing a comparison of ‘news’ series with the outputs from the VAR decomposition. The j -period return from holding an equity, excluding dividends, is equal to the change in price over this period, divided by the price in period t :

$$r_{t+j} = \left(\frac{P_{t+j} - P_t}{P_t} \right) \quad (12)$$

Within the present value framework, the equity price can be expressed in terms of ex-

firm forecasts in the present value formula to estimate firm level expected returns, and calculate the aggregate expected return from these values. We use aggregate forecasts in the present value formula to estimate aggregate expected returns directly.

⁶The expected return estimate is only sensitive to the long-run growth assumption in the level. Roughly speaking, a 100 basis point increase in the long-run assumption increases the estimate by around 80 basis points. Estimates under different growth assumptions are almost perfectly correlated, however (0.99 or greater when varying the assumption from 1-10%).

pected dividends and expected returns at time t and $t + j$:

$$P_t = \sum_{k=1}^{\infty} \frac{E_t(D_{t+k})}{(1 + r_{e,t})^k} = f(c_t, q_t) \quad (13)$$

$$P_{t+j} = \sum_{k=1}^{\infty} \frac{E_{t+j}(D_{t+j+k})}{(1 + r_{e,t+j})^k} = f(c_{t+j}, q_{t+j}) \quad (14)$$

$$r_{t+j} = \left(\frac{f(c_{t+j}, q_{t+j}) - f(c_t, q_t)}{P_t} \right) \quad (15)$$

where c_t represents the stream of expected future dividends, and q_t represents the future path of expected returns. Returns can be decomposed into cash flow and discount rate news as follows:

$$NCF_{t+j} = 0.5 \left(\frac{f(c_{t+j}, q_{t+j}) - f(c_t, q_{t+j})}{P_t} + \frac{f(c_{t+j}, q_t) - f(c_t, q_t)}{P_t} \right) \quad (16)$$

$$NDR_{t+j} = 0.5 \left(\frac{f(c_t, q_{t+j}) - f(c_t, q_t)}{P_t} + \frac{f(c_{t+j}, q_{t+j}) - f(c_{t+j}, q_t)}{P_t} \right) \quad (17)$$

$$r_{t+j} = NCF_{t+j} + NDR_{t+j} \quad (18)$$

These ‘news’ series do not correspond exactly to those obtained through the VAR decomposition. The Campbell (1991) framework is a decomposition of *unexpected return*, which also includes dividend payments. In contrast, the DDM decomposition decomposes *price changes* i.e. excluding dividends, and is not relative to an expectation. These differences should not materially affect a comparison between the two models, for two reasons. First, there is a lot of noise in short-term return variation, implying that returns are difficult to forecast, and that most of return variation is unexpected. Second, as pointed out in Chen, Da and Zhao (2013), dividends play a minor role in total return volatility. Over the period 1965 to 2013, the standard deviation of monthly total returns on indices in the UK and US are 5.55% and 4.51%, respectively. These numbers are little affected when excluding dividends, where the UK and US standard deviations are 5.52% and 4.50%, respectively.⁷

2.3 Data

The data that enter the VAR are obtained from Global Financial Data (GFD). The data are monthly frequency, and values are taken at the end of the month. Returns are calculated

⁷And the series are near perfectly correlated, with correlation coefficients of 0.97 and 0.99 for the UK and US respectively.

from total return equity indices for the FTSE All-Share and the S&P 500, where returns on these indices include reinvested dividend payments from constituent firms. Within the VAR, and also within the forecast tests, we use Price-Earnings (PE) ratios, Dividend-Price (DP) ratios and 10-year and 3-month Treasury bill rates for the UK and US, provided by GFD. PE and DP ratios are based on 12 month trailing earnings and dividends respectively, and these ratios are used to calculate the smoothed Price-Earnings ratios and payout ratio. We use I/B/E/S analyst forecasts in the DDM, which are obtained from Datastream. Due to the short sample of I/B/E/S data, our forecast tests are limited to data from August 1998 to December 2013, while the variables that enter the VAR estimations and other predictive regressions are available further back in time. For the UK, our sample starts in January 1965 and ends in December 2013. For the US, we use data from January 1978 to December 2013.

3 Expected Return and News Estimates

The results from the UK and US VAR models are provided in Table A1 in the appendix. The coefficient estimates in the returns equation are in line with what one might expect, and with past evidence. For both the UK and US, there is a small degree of momentum in short-term returns, shown by the positive coefficient on the lagged return. The price-earnings ratios and default spread negatively predict returns, and the pay-out ratio has a positive coefficient.⁸

3.1 Expected Returns

As outlined in the previous section, the VAR model and DDM can be used to generate estimates of expected returns. The DDM assumes that expected returns are constant across horizons, while the VAR allows for variation in expected returns at different horizons. Figures A1 and A2 show the expected return estimates from the two models for the UK and US. It is not immediately obvious which horizon estimate from the VAR should be compared against the DDM estimate, though since the DDM can be thought of as averaging across short to very long horizons, we opt for the 10-year expected return estimate from the VAR. Visually, the two measures are comparable. Broadly speaking, the series increase around 2002 and 2008, consistent with the equity market falls and economic slowdowns around these periods. Table A2 shows that the estimates co-move strongly across models and across countries. When comparing the two estimates for the same country, the correlation between

⁸For the lagged return, see Conrad and Kaul (1998), for the price-earnings ratio, see Campbell and Shiller (1988), for the payout ratio see Lamont (1998)

the DDM and VAR estimates is 0.75 for the UK, and 0.74 for the US. When comparing across countries, we observe a high degree of international co-movement. The correlation between UK and US estimates is in excess of 0.7, regardless of which model we use. While the levels of the UK and US VAR estimates are similar, the level of the UK DDM estimate is around 2 percentage points higher than the US DDM estimate on average.⁹

3.2 Cash flow and discount rate news

Using the methods described earlier, we decompose monthly returns into their cash flow and discount rate components, and investigate the role of each component in explaining the variance of returns. For both models, the variance decomposition is obtained using the following identity, which expresses return variance as a function of the cash flow and discount rate news series:

$$\begin{aligned} \text{var}(r_t) &= \text{cov}(NCF_t, r_t) + \text{cov}(NDR_t, r_t) \\ 1 &= \left(\frac{\text{cov}(NCF_t, r_t)}{\text{var}(r_t)} \right) + \left(\frac{\text{cov}(NDR_t, r_t)}{\text{var}(r_t)} \right) \end{aligned} \quad (19)$$

The relative proportions can therefore be obtained from the OLS regression coefficients of equity returns regressed on the DDM news series, and the regression coefficients of unexpected returns regressed on the news components from the VAR model. Table A3 reports the return variance decomposition from the two models. The implications of the two models are broadly consistent: discount rate variation tends to account for a larger proportion of return variance, in line with past evidence. The distinction between discount rate and cash flow news is less marked for the US VAR model, however. The results are consistent with Chen, Da and Zhao (2013), who find a qualitatively similar role for cash flow and discount news at the 12-month horizon between the two model types.

The return decompositions are shown in Figures A3 and A4. Each series has been smoothed using a 12-month moving average, and the discount rate news series has been inverted, such that negative contributions correspond to higher discount rates. Focusing on the stock market downturn over 2000-02, we observe slightly different interpretations of market movements. The VAR decompositions attribute a large part of the stock market declines around 2002 to downward revisions to expectations of future cash flows, although

⁹This may reflect differences in payout policies between UK and US equity markets, where firms in the US market have a greater tendency to return profits to shareholders through share repurchases, rather than dividends. All else equal, this would result in a lower level of forecasted dividends in the US market, which reduces the estimated expected return.

they also assign some role to higher discount rates. While the US DDM decomposition tells a similar story, showing a large negative contribution from cash flow news in 2001, the UK DDM decomposition points to discount rates as the main driver (perhaps intuitive given the smaller economic slowdown that the UK experienced over this period). The UK DDM decomposition is more consistent with the results of Campbell, Giglio and Polk (2013), who use the VAR framework, and also attribute most of the downturn in the early 2000s to higher discount rates. The other decompositions are consistent with Garrett and Priestly (2012), however, who explicitly model revisions to cash flow forecasts in the VAR framework, and find significant negative cash flow news in the early 2000s stock market decline. At the height of the 2007-09 financial crisis, all models attribute the stock market declines to a combination of negative cash flow and discount rate news, consistent with both Campbell, Giglio and Polk (2013) and Garrett and Priestly (2012). The DDM decompositions, while attributing declines to cash flow news eventually, interpret the large initial falls as reflecting large upward revisions to discount rates.

One of the primary concerns regarding the I/B/E/S data is the potential for relatively long lags between relevant information being released and the forecasts being updated. Figure A5 shows a comparison between the cash flow news estimates from the VAR and DDM models, and realised dividend growth 12 months ahead. While the two cash flow news series cannot be considered direct forecasts of future dividend growth, one would still expect to see correlation between these measures and realised future dividends. The broad trends are similar, however there are clear differences in timing of large movements and of higher frequency variation. In particular, the fall in actual dividends over 2008 was a little better predicted by the cash flow news series estimated from the VAR, suggesting that the VAR cash flow news series may contain more timely information about future dividends. This observation may confirm the issue of lags in I/B/E/S survey updates, and motivates a more systematic evaluation of the ability of these models to capture expectations, which is the focus of the next section.

4 Forecast Tests

In this section, we evaluate the forecasting performance of the DDM and VAR model, focusing on the expected return estimates the models generate. We do this by testing the ability of the estimates to forecast realised future returns, using common forecast evaluation frameworks (e.g. Welch and Goyal (2008) and Campbell and Thompson (2008)). We also consider the forecasting performance of each model relative to a range of alternative



predictor variables that are commonly assessed in the return predictability literature, including: the dividend-price ratio (DP), three price-earnings measures (a 12 month trailing ratio ‘PE’, a cyclically adjusted ratio ‘PES’ and a forward looking ratio ‘PEF’), and default and term spreads (‘DEF’ and ‘TERM’), further details for which can be found in Table A4 in the appendix. In general, we refer to this set of forecast variables as the ‘alternative’ or ‘traditional’ predictors. Summary statistics are provided in Table A5.

Early evidence on return predictability suggested that each of the traditional predictor variables were able to forecast future returns (see Rozeff (1984), Fama and French (1988), Fama and French (1989)). These findings were later challenged, however, on the basis of issues regarding statistical inference in regression-based tests, that led to results that were biased toward finding predictability (Stambaugh (1999)). An additional challenge came from Welch and Goyal (2008), who showed that these forecasting models performed poorly out-of-sample. In theory, price-dividend and price-earnings ratio variables ought to be able to forecast returns to the extent that they capture expected returns. These ratio variables are not clean measures of expected returns, however, as expected dividend growth is also a potentially important component that has distortionary effects (Lettau and Ludvigson (2005)). One would prefer to consider an equity price relative to its future earnings, though the denominators of these ratios are usually backward looking. In this sense, one might expect the DDM-implied expected return to provide a better estimate of expected returns, and therefore better forecast future returns. Indeed, Li, Ng and Swaminathan (2013) show that a range of DDM expected return estimates outperform standard ratio and business cycle variables using US data. Their results complement studies emphasizing the importance of modelling expected cash flow growth components of price ratio variables, and showing enhanced forecasting performance once this is accounted for (e.g. Lacerda and Santa-Clara (2010) and Golez (2014)). The VAR methodology may also provide an improvement upon single predictor variables, to the extent that the model combines information from a range of predictor variables in a parsimonious manner to generate return forecasts. In some cases it has been shown that poor performance of individual forecasting variables can be improved upon by forming combination forecasts based on the individual predictors or separate forecasts, for example in stock return prediction (Rapach, Strauss and Zhou (2010)) and in studies forecasting output and inflation (Stock and Watson (2003)). The VAR does not combine forecasts in the same way as these studies, which often use weighting methodologies that are based on recent forecasting performance, but incorporates a handful of predictor variables nonetheless.

In the remainder of the paper, we first test whether the expected return estimates and



other predictors can forecast future returns using in-sample regressions. We then proceed to compare the VAR and DDM with each other and the alternative predictors. Finally, we explore the out-of-sample forecastability of the range of predictors. Within the out-of-sample analysis, we test the forecasts from the DDM and VAR model against one another, and against a historical average benchmark.

4.1 In-Sample Prediction

We explore the in-sample relationship between the set of predictor variables and realised future returns using a standard framework. We estimate forecast regressions with average k -period returns as the dependent variable regressed on individual predictors at time t :

$$\frac{r_{t,t+k}}{k} = a_k + b_k X_t + e_{t+k} \quad (20)$$

where k is the forecast horizon in months, and X_t is the predictor variable. The k -period return is calculated as the cumulative sum of the monthly total log returns over the forecast horizon. The regressions are estimated for $k = 3, 12, 24$ and 36-month horizons.

An issue when considering multi-period returns is that overlapping observations in the regression model induce serial correlation in the residual term. To correct for this issue in inference tests, we use Newey-West standard errors with $k+1$ lags. In addition, a well-known issue with this predictive regression set-up is that in-sample tests are biased in small samples with persistent predictors, and when the predictor and return innovations are correlated (Stambaugh (1999)). To account for this bias, we follow a bootstrapping procedure similar to the one used in Rapach and Wohar (2005) and Neely, Rapach, Tu and Zhou (2014). This approach involves estimating a data generating process for single-period returns, r_{t+1} , under the null hypothesis of no predictability.¹⁰ We also estimate an autoregressive process for the predictor variable, X_t , and combine the return and predictor processes to generate a range of pseudo-samples.¹¹ These samples are generated by randomly drawing (with replacement) from the sets of residuals from the two processes. We generate 2000 samples, and construct empirical distributions of Newey-West t -statistics by running predictive regressions over the relevant horizon for each pseudo-sample. We summarise results using p -values based on a comparison of the Newey-West statistic from the regression based on the original data compared to the estimated empirical distribution.¹² The significance of the coefficient estimates

¹⁰In applying this procedure and the null of no predictability, we abstract from the presence of predictability in the VAR-based expected return estimate.

¹¹The AR lag-length is selected using AIC, where we consider up to a maximum of four lags.

¹²We conduct a one-sided test, where the p -value for $\beta_k > 0$ ($\beta_k < 0$) reflects the proportion of bootstrapped statistics that are greater (less) than the statistic based on the original data.

gives a simple indication of the ability of each variable to predict future returns.

The results for the UK and US are provided in Tables 1a and 1b. For each predictor, the coefficient estimates are in the second column, p -values in the third column, and the adjusted R^2 in the fourth column. In general, across the range of predictors and forecast horizons there are relatively few examples of statistically significant relationships. Popular forecast variables, such as the simple and forward looking price-earnings ratios show little predictive power over this period, in both the UK and US. There are some examples of significant relationships (at the 10% level) for the UK and US term spreads at the 24- and 36-month horizons, and the US default spread and smoothed price-earnings ratio also at longer horizons (consistent with the predictive power of these variables shown in Campbell and Vuolteenaho (2004), Campbell, Giglio and Polk (2013), and Campbell, Giglio, Polk and Turley (2014)).

Turning to the model-based estimates, we find some evidence that the VAR and DDM significantly predict realised returns over this period. The UK and US VAR estimates are the only forecast variables that significantly predict 3-month returns at the 10% significance level, and the predictive power appears to also extend to longer forecast horizons, where the 12- to 36-month forecasts are significant in both the UK and US. The DDM does not display the same forecasting ability at short horizons, where the 3-month relationship is insignificant and predicts with the incorrect sign. On the other hand, however, the DDM significantly predicts 24- and 36-month returns at or below the 5% level in both the UK and the US, and 12-month returns in the US at the 10% level. The results are in line with Li, Ng and Swaminathan (2013), who run similar forecast regressions over an earlier sample using US data. They show that a similar DDM measure can significantly forecast returns, and that a range of ratio variables cannot except for a few cases at long horizons. The results suggest that the VAR and DDM differ in their forecasting ability depending on the forecast horizon. The DDM appears to perform better at long horizons rather than short, while the VAR appears to outperform others at shorter horizons. This difference in performance could arise naturally from the underlying assumptions of both models. When assuming constant expected returns in the DDM, it appears that the resulting estimate better captures long-horizon expected returns. To the extent that the assumption implies that the constant return estimate averages across expected returns over short to long (in fact, infinite) horizons, it is intuitive that the average maturity of this estimate is longer rather than shorter. Since the VAR is a model of the short-run dynamics of returns, one might expect it to perform well in forecasting shorter horizon returns, though our results suggest it also performs well at longer horizons.

Table 1a: Forecast Regressions: UK

The k -period horizon return is regressed on the predictor variable, using OLS: $\frac{r_{t,t+k}}{k} = a_k + b_k X_t + e_{t+k}$. Sample period from August 1998 to December 2013. p -values based on Newey-West t -statistics (with lag length $k + 1$), compared with bootstrapped empirical distributions.

| VAR | k | b | $p - val$ | R^2 | Term | k | b | $p - val$ | R^2 |
|-----|-----|-------|-----------|-------|---------|-----|-------|-----------|-------|
| | 3 | 0.57 | 0.10 | 0.03 | | 3 | 0.22 | 0.24 | 0.01 |
| | 12 | 0.70 | 0.11 | 0.12 | | 12 | 0.30 | 0.18 | 0.08 |
| | 24 | 1.12 | 0.04 | 0.42 | | 24 | 0.30 | 0.15 | 0.16 |
| | 36 | 1.25 | 0.06 | 0.70 | | 36 | 0.31 | 0.08 | 0.26 |
| DDM | k | b | $p - val$ | R^2 | Def | k | b | $p - val$ | R^2 |
| | 3 | -0.41 | 0.58 | 0.01 | | 3 | 0.05 | 0.43 | 0.00 |
| | 12 | 0.44 | 0.15 | 0.05 | | 12 | 0.16 | 0.15 | 0.06 |
| | 24 | 0.77 | 0.03 | 0.30 | | 24 | 0.18 | 0.13 | 0.16 |
| | 36 | 0.19 | 0.03 | 0.55 | | 36 | 0.17 | 0.13 | 0.24 |
| DP | k | b | $p - val$ | R^2 | PES | k | b | $p - val$ | R^2 |
| | 3 | -0.09 | 0.84 | 0.00 | | 3 | -0.04 | 0.29 | 0.00 |
| | 12 | 0.80 | 0.08 | 0.12 | | 12 | -0.10 | 0.21 | 0.11 |
| | 24 | 0.93 | 0.10 | 0.33 | | 24 | -0.13 | 0.14 | 0.34 |
| | 36 | 0.77 | 0.29 | 0.41 | | 36 | -0.14 | 0.14 | 0.59 |
| PE | k | b | $p - val$ | R^2 | Fwd. PE | k | b | $p - val$ | R^2 |
| | 3 | -0.03 | 0.31 | 0.00 | | 3 | -0.02 | 0.42 | 0.00 |
| | 12 | -0.11 | 0.29 | 0.10 | | 12 | -0.11 | 0.27 | 0.07 |
| | 24 | -0.11 | 0.19 | 0.22 | | 24 | -0.15 | 0.14 | 0.24 |
| | 36 | -0.11 | 0.17 | 0.34 | | 36 | -0.15 | 0.10 | 0.39 |

Table 1b: Forecast Regressions: US

The k -period horizon return is regressed on the predictor variable, using OLS: $\frac{r_{t,t+k}}{k} = a_k + b_k X_t + e_{t+k}$. Sample period from August 1998 to December 2013. p -values based on Newey-West t -statistics (with lag length $k + 1$), compared with bootstrapped empirical distributions.

| VAR | k | b | $p - val$ | R^2 | Term | k | b | $p - val$ | R^2 |
|-----|-----|------|-----------|-------|------|-----|-------|-----------|-------|
| | 3 | 2.57 | 0.10 | 0.10 | | 3 | -0.07 | 0.34 | 0.00 |
| | 12 | 1.64 | 0.10 | 0.10 | | 12 | 0.21 | 0.27 | 0.03 |
| | 24 | 2.30 | 0.08 | 0.35 | | 24 | 0.55 | 0.02 | 0.36 |
| | 36 | 2.60 | 0.05 | 0.63 | | 36 | 0.51 | 0.00 | 0.62 |

| DDM | k | b | $p - val$ | R^2 | Def | k | b | $p - val$ | R^2 |
|-----|-----|-------|-----------|-------|-----|-----|-------|-----------|-------|
| | 3 | -0.52 | 0.86 | 0.00 | | 3 | -0.10 | 0.29 | 0.00 |
| | 12 | 1.64 | 0.06 | 0.13 | | 12 | 0.15 | 0.34 | 0.02 |
| | 24 | 1.70 | 0.03 | 0.29 | | 24 | 0.42 | 0.05 | 0.33 |
| | 36 | 1.28 | 0.05 | 0.32 | | 36 | 0.42 | 0.01 | 0.62 |

| DP | k | b | $p - val$ | R^2 | PES | k | b | $p - val$ | R^2 |
|----|-----|------|-----------|-------|-----|-----|-------|-----------|-------|
| | 3 | 0.34 | 0.17 | 0.00 | | 3 | -0.00 | 0.39 | 0.00 |
| | 12 | 1.55 | 0.14 | 0.15 | | 12 | -0.07 | 0.25 | 0.10 |
| | 24 | 1.60 | 0.13 | 0.31 | | 24 | -0.10 | 0.11 | 0.35 |
| | 36 | 1.23 | 0.23 | 0.34 | | 36 | -0.10 | 0.06 | 0.58 |

| PE | k | b | $p - val$ | R^2 | Fwd.PE | k | b | $p - val$ | R^2 |
|----|-----|-------|-----------|-------|--------|-----|-------|-----------|-------|
| | 3 | -0.04 | 0.19 | 0.00 | | 3 | -0.03 | 0.40 | 0.00 |
| | 12 | -0.09 | 0.31 | 0.08 | | 12 | -0.11 | 0.27 | 0.08 |
| | 24 | -0.09 | 0.33 | 0.16 | | 24 | -0.13 | 0.13 | 0.23 |
| | 36 | -0.08 | 0.44 | 0.20 | | 36 | -0.12 | 0.17 | 0.30 |

The results also suggest differences in performance between the model-based estimates and the range of the alternative predictors, where commonly used ratio and business cycle variables struggle to forecast returns. We explore the performance of the VAR and DDM relative to each other, and to the alternative predictors on a pair-by-pair basis. We summarise the relative forecast errors from each model with a standard R^2 statistic (e.g. Welch and Goyal (2008), Campbell and Thompson (2008), Rapach, Strauss and Zhou (2010)), that

gives the ratio of forecast errors from two competing forecasting regressions:

$$R_{IS}^2 = 1 - \left(\frac{\sum_{k=1}^q (r_{m+k} - \hat{r}_{m+k})^2}{\sum_{k=1}^q (r_{m+k} - \tilde{r}_{m+k})^2} \right) \quad (21)$$

\hat{r} and \tilde{r} are the fitted return projections from two competing forecasting models. When $R_{IS}^2 > 0$, the \hat{r} model outperforms the benchmark \tilde{r} model. Statistical significance is assessed with the Diebold-Mariano (1995) test. Table 2a shows the R_{IS}^2 values for the UK. In Panel A the VAR is set as the benchmark model. $R_{IS}^2 > 0$ implies that the VAR underperforms the alternative predictor. In Panel B the DDM is the benchmark model (\tilde{r}), where $R_{IS}^2 > 0$ implies that the DDM underperforms.

Table 2a: In-sample Relative Forecasting Performance – UK

Panel A shows the relative forecast errors with the VAR as the benchmark. Panel B shows the relative performance where the DDM is used as the benchmark. p-values are based on the Diebold-Mariano (1995) test, adjusted for serial correlation using a $k - 1$ lag parameter.

| Panel A | $k = 3$ | | 12 | | 24 | | 36 | |
|---------|------------|-------------|------------|-------------|------------|-------------|------------|-------------|
| | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> |
| DDM | -0.02 | 0.34 | -0.07 | 0.06 | -0.22 | 0.03 | -0.47 | 0.01 |
| DP | -0.03 | 0.14 | 0.01 | 0.47 | -0.17 | 0.15 | -0.94 | 0.00 |
| PE | -0.03 | 0.08 | -0.02 | 0.43 | -0.36 | 0.13 | -1.16 | 0.03 |
| Term | -0.02 | 0.20 | -0.04 | 0.26 | -0.46 | 0.01 | -1.43 | 0.01 |
| Def | -0.03 | 0.07 | -0.06 | 0.12 | -0.47 | 0.02 | -1.49 | 0.03 |
| PES | -0.03 | 0.06 | -0.01 | 0.45 | -0.14 | 0.10 | -0.33 | 0.03 |
| PEF | -0.03 | 0.08 | -0.05 | 0.28 | -0.31 | 0.14 | -0.99 | 0.03 |

| Panel B | $k = 3$ | | 12 | | 24 | | 36 | |
|---------|------------|-------------|------------|-------------|------------|-------------|------------|-------------|
| | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> |
| VAR | 0.02 | 0.34 | 0.07 | 0.06 | 0.18 | 0.03 | 0.32 | 0.01 |
| DP | -0.01 | 0.23 | 0.07 | 0.11 | 0.04 | 0.36 | -0.33 | 0.04 |
| PE | -0.01 | 0.33 | 0.05 | 0.24 | -0.11 | 0.31 | -0.48 | 0.16 |
| Term | 0.00 | 0.50 | 0.03 | 0.36 | -0.20 | 0.08 | -0.66 | 0.07 |
| Def | -0.01 | 0.32 | 0.01 | 0.42 | -0.20 | 0.13 | -0.70 | 0.09 |
| PES | -0.01 | 0.37 | 0.06 | 0.09 | 0.06 | 0.27 | 0.09 | 0.34 |
| PEF | -0.01 | 0.30 | 0.03 | 0.32 | -0.08 | 0.35 | -0.36 | 0.19 |

For the UK, the negative R_{IS}^2 values in Panel A indicate that the VAR generates smaller forecast errors relative to all the other predictors considered, and across almost every forecast horizon. This outperformance is statistically significant at the 10% level in many cases, and is particularly marked at the 36-month horizon. The first rows of results in both panels compare the performance of the VAR model and DDM. The results show that the VAR estimates perform significantly better compared to the DDM at the 12- to 36-month horizons. The DDM also generates smaller forecast errors in many cases, though statistically significant outperformance is less widespread compared to the VAR, and is concentrated at longer horizons, consistent with the simple regression results. Table 2b shows the results for the US.

Table 2b: In-sample Relative Forecasting Performance – US

Panel A shows the relative forecast errors with the VAR as the benchmark. Panel B shows the relative performance where the DDM is used as the benchmark. p-values are based on the Diebold-Mariano (1995) test, adjusted for serial correlation using a $k - 1$ lag parameter.

| Panel A | $k = 3$ | | 12 | | 24 | | 36 | |
|---------|------------|-------------|------------|-------------|------------|-------------|------------|-------------|
| | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> |
| DDM | -0.11 | 0.02 | 0.03 | 0.38 | -0.09 | 0.26 | -0.85 | 0.00 |
| DP | -0.11 | 0.02 | 0.05 | 0.31 | -0.07 | 0.37 | -0.80 | 0.00 |
| PE | -0.11 | 0.01 | -0.02 | 0.42 | -0.30 | 0.13 | -1.18 | 0.00 |
| Term | -0.11 | 0.02 | -0.08 | 0.07 | 0.01 | 0.47 | -0.03 | 0.46 |
| Def | -0.11 | 0.02 | -0.09 | 0.06 | -0.02 | 0.46 | -0.04 | 0.45 |
| PES | -0.11 | 0.01 | -0.00 | 0.50 | -0.00 | 0.49 | -0.16 | 0.14 |
| PEF | -0.11 | 0.01 | -0.03 | 0.38 | -0.19 | 0.21 | -0.90 | 0.01 |

| Panel B | $k = 3$ | | 12 | | 24 | | 36 | |
|---------|------------|-------------|------------|-------------|------------|-------------|------------|-------------|
| | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> | R_{IS}^2 | <i>pval</i> |
| VAR | 0.10 | 0.02 | -0.03 | 0.38 | 0.08 | 0.26 | 0.46 | 0.00 |
| DP | -0.00 | 0.47 | 0.02 | 0.35 | 0.03 | 0.43 | 0.03 | 0.44 |
| PE | -0.00 | 0.50 | -0.05 | 0.23 | -0.19 | 0.19 | -0.18 | 0.21 |
| Term | -0.00 | 0.40 | -0.12 | 0.11 | 0.10 | 0.31 | 0.44 | 0.00 |
| Def | -0.00 | 0.45 | -0.13 | 0.09 | 0.06 | 0.38 | 0.44 | 0.00 |
| PES | -0.00 | 0.40 | -0.06 | 0.18 | 0.06 | 0.32 | 0.40 | 0.02 |
| PEF | -0.00 | 0.45 | -0.06 | 0.17 | -0.09 | 0.32 | -0.03 | 0.45 |

In the US, the VAR model again performs relatively strongly, where the majority of the

R_{IS}^2 values are negative, and statistically significant in many cases. The DDM also outperforms the alternative predictors in many cases, though the DDM significantly underperforms a handful of alternative predictors at the 36-month horizon. Overall, the assessment of relative forecasting performance is less indicative of the VAR and DDM performing differently based on the forecast horizon, where the VAR generates significantly lower forecast errors at both short and long horizons.

To further explore the relative forecasting ability of the VAR and DDM estimates. We consider a multiple regression that includes only the model-based predictors, where we regress k -period returns on the VAR and DDM estimates. The results are shown in Table 3. For the UK, both the VAR and DDM significantly predict 3-month returns. At 3-, 12- and 24-month horizons, the DDM forecasts have an incorrect sign. For the US, the VAR significantly predicts 3-month returns, while the DDM forecast is insignificant at this horizon and predicts with the incorrect sign. At the 24-month horizon, the DDM predicts significantly at the 10% level. These results mildly suggest that the VAR and DDM estimates contain information relevant for future returns at different horizons.

Table 3: VAR vs. DDM In-sample Relative Forecasting Performance

Parameter estimates for in-sample forecast regressions for the UK (top panel) and the US (bottom panel), estimated using OLS. The k -period horizon return is regressed on the VAR and DDM expected returns: $\frac{r_{t+k}}{k} = a + b_{1,k}VAR_t^k + b_{2,k}DDM_t + e_{t+k}$. The sample period is from August 1998 to December 2013. p -values based on Newey-West standard errors, with lag length set to $k + 1$, compared with bootstrapped empirical distributions.

| UK | k | b_1 | $p - val$ | b_2 | $p - val$ | R^2 |
|----|-----|-------|-----------|-------|-----------|-------|
| | 3 | 2.44 | 0.00 | -2.50 | 0.00 | 0.20 |
| | 12 | 1.17 | 0.16 | -0.53 | 0.26 | 0.12 |
| | 24 | 1.28 | 0.13 | -0.15 | 0.44 | 0.42 |
| | 36 | 1.16 | 0.14 | 0.07 | 0.37 | 0.69 |
| US | k | b_1 | $p - val$ | b_2 | $p - val$ | R^2 |
| | 3 | 2.85 | 0.05 | -1.24 | 0.51 | 0.11 |
| | 12 | 1.04 | 0.25 | 1.25 | 0.16 | 0.15 |
| | 24 | 1.65 | 0.15 | 0.98 | 0.08 | 0.41 |
| | 36 | 2.37 | 0.09 | 0.24 | 0.27 | 0.64 |

Finally, we consider a multiple regression that relaxes the VAR specification, and tests the forecasting ability of the variables that are used to generate VAR expected return es-

timates. The multiple regression results for the UK and US are shown in Tables A6a and A6b in the appendix. For the UK, the price-earnings and payout ratio significantly predict returns, without imposing the parametric assumptions of the VAR. For the US, all variables are generally unable to predict returns, which suggests, given the results presented earlier, that modelling the short-run dynamics of these variables in a VAR provides incremental forecasting power.

Overall, the in-sample evidence suggests that the model-based expected return estimates are able to forecast returns, and in many cases outperform commonly used forecast variables. Given that the DDM can forecast returns and performs better than many other predictors, it appears that any biases or lags in the I/B/E/S analyst forecasts are not entirely detrimental to the DDM estimate. The in-sample results favour the VAR estimates over the DDM, however, where the VAR generates significantly lower forecast errors. It could be the case that I/B/E/S data imperfections underlie the underperformance of the DDM relative to the VAR.

4.2 Out-of-sample Prediction

The in-sample tests suggest that the VAR and DDM expected return estimates can predict future returns in many cases, as can some traditional predictors at certain horizons. The parameters of the in-sample forecast regressions, however, do not take into account whether information is available contemporaneously, and incorporate data ahead of time. This is a well-known limitation of in-sample tests. In this section, we proceed to run out-of-sample tests, which indicate whether the range of forecast variables would have been able to predict returns in real-time. We follow a standard out-of-sample testing procedure that evaluates forecasting performance where only contemporaneously available information is incorporated (e.g. Welch and Goyal (2008), Campbell and Thompson (2008), Rapach, Strauss and Zhou (2010)). Using a similar predictive regression framework to the in-sample analysis:

$$r_{t+k} = a_k + b_k X_t + e_{t+k} \quad (22)$$

we now estimate initial regressions based on data up to August 1998 (the first m observations), so that the first out-of-sample forecast is given by:

$$\hat{r}_{m+k} = \hat{a}_{m,k} + \hat{b}_{m,k} X_m \quad (23)$$

following which the data set is expanded by a single observation, and the model parameters re-estimated, such that the next out-of-sample forecast is given by:

$$\hat{r}_{m+1+k} = \hat{a}_{m+1,k} + \hat{b}_{m+1,k}X_{m+1} \quad (24)$$

This process is repeated to generate a set of out-of-sample forecasts at various horizons, from August 1998 to December 2013. To compare forecasting performance, we once again use the relative R^2 measure, which has the same interpretation as before, though is denoted as R_{OS}^2 to identify it as the out-of-sample statistic:

$$R_{OS}^2 = 1 - \left(\frac{\sum_{k=1}^q (r_{m+k} - \hat{r}_{m+k})^2}{\sum_{k=1}^q (r_{m+k} - \tilde{r}_{m+k})^2} \right) \quad (25)$$

It is also the case in the in-sample exercise that the VAR expected return estimates are generated using information not available at the time, and so we need to re-estimate the model to account for this. Following the same methodology as outlined above, we expand the sample period of the VAR one period at a time and re-estimate the VAR model parameters and projections at each period. The DDM return estimates do not need to be re-estimated at this stage, since all the inputs to the model (in particular the survey forecasts) are available contemporaneously.

We first compare the out-of-sample performance of the DDM relative to the VAR model, where the R_{OS}^2 values (with the VAR estimate as the benchmark) are shown in Table 4. The negative R_{OS}^2 values suggest that the VAR forecasts again generate lower forecast errors in most cases, across the UK and US. In contrast with the in-sample results, however, the difference in predictive ability is not statistically significant. This suggests that the stronger performance of the VAR in-sample partly resulted from the look-ahead bias in its return projections.

Table 4: Out-of-sample forecasting performance: VAR vs. DDM

p-values based on the Diebold-Mariano (1995) test

| UK | | | US | | |
|---------|------------|-----------|---------|------------|-----------|
| Horizon | R_{OS}^2 | $p - val$ | Horizon | R_{OS}^2 | $p - val$ |
| 3 | -0.00 | 0.46 | 3 | 0.00 | 0.50 |
| 12 | -0.09 | 0.32 | 12 | 0.02 | 0.41 |
| 24 | -0.33 | 0.20 | 24 | -0.01 | 0.46 |
| 36 | -0.67 | 0.11 | 36 | -0.07 | 0.32 |

To explore the out-of-sample forecasts further, we compare the forecasting performance of the range of predictors against a historical average benchmark $\bar{r}_{t+k} = \sum_{j=1}^t r_j$. As summarised in Welch and Goyal (2008) and Campbell and Thompson (2008), this benchmark has been surprisingly difficult to beat in predicting returns. We again rely on the R_{OS}^2 to summarise relative forecast errors, with the historical average as the benchmark. We formally test the reduction in the mean squared prediction error using the Clark and West (2007) test, where the test statistic is defined as:

$$f_{t+1} = (r_{m+k} - \tilde{r}_{m+k})^2 - [(r_{m+k} - \hat{r}_{m+k})^2 - (\tilde{r}_{m+k} - \hat{r}_{m+k})^2]$$

\hat{r}_{m+k} is the forecast from the prediction model, and \tilde{r}_{m+k} is the benchmark average forecast. The test statistic, f_{t+1} , is regressed on a constant, and the Newey-West t -statistic is used to calculate the associated p -value using a standard normal distribution. Comparisons against the historical average are based on the Clark and West (2007) test since the historical average model is nested within the predictive regression models. We continue to use the Diebold and Mariano (1995) test for the VAR and DDM, however, as we use the expected return estimates directly from these models rather than using the regression framework.

The comparison of each predictor against the historical average benchmark is shown in Tables 5 and 6 for the UK and US respectively. Here, a positive R_{OS}^2 value indicates that the predictor outperforms the historical average.

Table 5: Out-of-sample forecasting performance against historical average benchmark: UK

| UK | $k = 3$ | | 12 | | 24 | | 36 | |
|------|------------|--------|------------|--------|------------|--------|------------|--------|
| | R_{OS}^2 | $pval$ | R_{OS}^2 | $pval$ | R_{OS}^2 | $pval$ | R_{OS}^2 | $pval$ |
| VAR | 0.05 | 0.24 | 0.24 | 0.18 | 0.49 | 0.11 | 0.66 | 0.06 |
| DDM | 0.04 | 0.04 | 0.17 | 0.06 | 0.31 | 0.05 | 0.43 | 0.03 |
| DP | 0.04 | 0.02 | 0.25 | 0.01 | 0.48 | 0.01 | 0.64 | 0.01 |
| PE | 0.01 | 0.22 | 0.15 | 0.03 | 0.32 | 0.07 | 0.44 | 0.10 |
| Term | -0.01 | | 0.04 | 0.07 | 0.03 | 0.06 | 0.03 | 0.03 |
| Def | 0.01 | 0.17 | 0.03 | 0.09 | -0.01 | | 0.03 | 0.08 |
| PES | -0.00 | | 0.10 | 0.02 | 0.30 | 0.03 | 0.52 | 0.04 |
| PEF | -0.00 | | 0.06 | 0.16 | -0.01 | | -0.12 | |

Table 6: Out-of-sample forecasting performance against historical average benchmark: US

| US | $k = 3$ | | 12 | | 24 | | 36 | |
|------|------------|--------|------------|--------|------------|--------|------------|--------|
| | R_{OS}^2 | $pval$ | R_{OS}^2 | $pval$ | R_{OS}^2 | $pval$ | R_{OS}^2 | $pval$ |
| VAR | 0.03 | 0.30 | 0.12 | 0.25 | 0.31 | 0.14 | 0.48 | 0.05 |
| DDM | 0.03 | 0.12 | 0.14 | 0.10 | 0.30 | 0.06 | 0.44 | 0.03 |
| DP | 0.01 | 0.23 | 0.06 | 0.08 | 0.05 | 0.12 | 0.08 | 0.14 |
| PE | 0.01 | 0.24 | 0.08 | 0.05 | 0.09 | 0.00 | 0.05 | 0.02 |
| Term | -0.05 | | -0.06 | | 0.02 | 0.05 | 0.02 | 0.11 |
| Def | -0.03 | | -0.01 | | 0.07 | 0.02 | 0.10 | 0.01 |
| PES | -0.03 | | -0.07 | | -0.21 | | -0.32 | |
| PEF | -0.03 | | -0.04 | | -0.16 | | -0.39 | |

For both the UK and US, the VAR and DDM estimates perform well relative to the historical average, where the R_{OS}^2 values are positive and economically significant at most forecast horizons. The DP and PE predictors also generate smaller forecast errors relative to the benchmark across all forecast horizons. For the UK, the magnitude of the reduction in forecast errors is comparable across the predictors, though only the DDM and dividend-price ratio generate statistically significant reductions at the 5% level for the majority of cases. For the US, the improvements relative to the historical average are most marked for the VAR and DDM, though the level of significance varies a little. The relatively weaker performance of the US ratio and business cycle variables is in line with past evidence (e.g. Welch and Goyal (2008), Li, Ng and Swaminathan (2013)). The relatively strong performance of the DDM is again consistent with Li, Ng and Swaminathan (2013), who show similar findings for the US at the 1-month horizon, using a 2-year moving average of their DDM estimate. The additional contribution of our findings is to show that the out-of-sample predictability extends far beyond the 1-month horizon, and remains when using the contemporaneous DDM estimate rather than the moving average.

As outlined earlier, the DDM and VAR may offer improvements upon traditional predictor variables to the extent that they better account for cash flow expectations, and combine information contained in a range of model inputs. The larger reduction in forecast errors of the two models in the US suggests that this extra sophistication does indeed lead to forecast improvements. This is less the case in the UK, however, where dividend-price and price-earnings ratios perform similarly well against the historical average. Overall, the out-of-sample analysis provides evidence in favour of return predictability, and provides ad-

ditional support to recent studies showing enhanced forecast performance using implied cost of capital measures and combined forecasts.

5 Conclusion

We test whether two expected equity return estimates can forecast realised returns. Expected returns are generated using two popular equity modelling approaches, which have been used in a variety of applications. The first approach is the VAR model associated with Campbell (1991), and the second the Dividend Discount Model (DDM) adaptation of Gordon (1962), which has recently been used in Li, Ng and Swaminathan (2013) and Chen, Da and Zhao (2013). We examine the predictive power of these estimates in a range of in- and out-of-sample tests, and also consider their forecasting performance against commonly used predictors such as the dividend-price ratio and term spread. In-sample, we find evidence that both the expected return measures can significantly forecast returns, whereas the traditional predictors generally can to a lesser extent. When exploring the relative performance further, we find that the VAR-based estimate generates substantially lower forecast errors compared to the DDM estimate and the traditional predictors. Within the out-of-sample tests, there is further evidence that the model-based estimates are useful for forecasting, though the performances of the VAR model and DDM are harder to differentiate. We compare the range of forecast variables to a historical average benchmark forecast, and find that the VAR and DDM offer economically and statistically significant forecast improvements.

6 Appendix

Table A1: UK and US VAR Results

VAR parameter estimates for the UK and US, estimated using OLS. The model is a first order VAR with a constant, that includes the market return r_t , the smoothed price-earnings ratio PE_t , the term spread y_t , and the default spread DEF_t for the US and the pay-out ratio PAY_t for the UK. Standard errors are shown in parentheses. The sample periods are January 1970 to December 2013 for the UK VAR and January 1977 to December 2013 for the US VAR.

| UK | c | r_t | PE_t | y_t | PAY_t | R^2 |
|-------------|-----------------|-----------------|-----------------|-----------------|-----------------|-------|
| r_{t+1} | 1.76 (1.46) | 0.11 (0.04) | -0.20 (0.06) | -0.03 (0.15) | 0.05 (0.03) | 0.03 |
| PE_{t+1} | -0.05 (0.22) | 0.01 (0.01) | 0.97 (0.01) | 0.00 (0.02) | 0.01 (0.00) | 0.96 |
| y_{t+1} | 0.09 (0.11) | -0.00 (0.00) | -0.01 (0.01) | 0.97 (0.01) | 0.00 (0.00) | 0.94 |
| PAY_{t+1} | 0.96 (0.40) | 0.03 (0.01) | -0.03 (0.02) | -0.08 (0.04) | 0.99 (0.01) | 0.98 |
| US | c | r_t | PE_t | y_t | DEF_t | R^2 |
| r_{t+1} | 5.99 (0.03) | 0.08 (0.10) | -1.52 (0.05) | 0.03 (0.87) | -0.66 (0.27) | 0.02 |
| PE_{t+1} | 0.04 (0.16) | 0.00 (0.10) | 0.99 (0.00) | 0.00 (0.15) | -0.01 (0.34) | 0.99 |
| y_{t+1} | -0.14 (0.57) | -0.00 (0.35) | 0.04 (0.53) | 0.95 (0.00) | 0.10 (0.06) | 0.91 |
| DEF_{t+1} | 0.25 (0.00) | -0.01 (0.00) | -0.05 (0.01) | -0.01 (0.02) | 0.94 (0.00) | 0.94 |

Table A2: Expected return correlations

| | UK DDM | UK VAR | US DDM | US VAR |
|--------|--------|--------|--------|--------|
| UK DDM | 1 | | | |
| UK VAR | 0.76 | 1 | | |
| US DDM | 0.80 | 0.74 | 1 | |
| US VAR | 0.75 | 0.87 | 0.74 | 1 |

Table A3: Return variance decomposition using VAR and DDM

| | | \tilde{e}_{CF} | \tilde{e}_{DR} |
|-----|----|------------------|------------------|
| VAR | UK | 0.34 | 0.66 |
| | US | 0.43 | 0.57 |
| | | NCF_t | NDR_t |
| DDM | UK | 0.26 | 0.74 |
| | US | 0.38 | 0.62 |

Table A4: Alternative Predictor Variable Descriptions

| Predictor | | Description |
|------------------------------------|------------|---|
| 12 month Dividend-Price | <i>DP</i> | <i>Dividends paid over the previous 12 months, divided by today's price</i> |
| 12 month Price-Earnings | <i>PE</i> | <i>Today's price divided by earnings paid over the previous 12 months</i> |
| Cyclically Adjusted Price-Earnings | <i>PES</i> | <i>Today's price divided by a 10 year trailing moving average of earnings</i> |
| Forward Price-Earnings | <i>PEF</i> | <i>Today's price divided by I/B/E/S analyst forecasts of earnings over the next 12 months</i> |
| Term Spread | <i>TS</i> | <i>Yield on 10 year government bond less rate on 3 month Treasury Bill</i> |
| Default Spread | <i>DEF</i> | <i>Yield on BAA-rated corporate bond basket less yield on AAA-rated basket</i> |

Table A5a: Forecast Variable Summary Statistics: UK

| Predictor | Mean | Standard Deviation |
|------------------------------------|-------------|---------------------------|
| VAR (3 month) | 6.87 | 9.96 |
| VAR (6 month) | 7.24 | 9.60 |
| VAR (12 month) | 7.87 | 8.75 |
| VAR (24 month) | 8.95 | 7.25 |
| VAR (36 month) | 9.81 | 6.09 |
| DDM | 9.20 | 0.71 |
| 12 month Dividend-Price | 3.25 | 0.61 |
| 12 month Price-Earnings | 15.28 | 4.24 |
| Cyclically Adjusted Price-Earnings | 16.99 | 4.55 |
| Forward Price-Earnings | 13.48 | 3.38 |
| Term Spread | 0.76 | 1.35 |
| Default Spread | 2.19 | 2.18 |

Table A5b: Forecast Variable Summary Statistics: US

| Predictor | Mean | Standard Deviation |
|------------------------------------|-------------|---------------------------|
| VAR (3 month) | 8.00 | 4.38 |
| VAR (6 month) | 8.11 | 4.06 |
| VAR (12 month) | 8.27 | 3.78 |
| VAR (24 month) | 8.46 | 3.39 |
| VAR (36 month) | 8.55 | 3.07 |
| DDM | 7.43 | 0.34 |
| 12 month Dividend-Price | 1.74 | 0.40 |
| 12 month Price-Earnings | 20.35 | 4.90 |
| Cyclically Adjusted Price-Earnings | 23.41 | 6.46 |
| Forward Price-Earnings | 16.50 | 3.92 |
| Term Spread | 1.86 | 1.21 |
| Default Spread | 3.45 | 1.48 |

Table A6a: Multiple regression: UK VAR variables

Parameter estimates for in-sample forecast regressions estimated using OLS. The k -period horizon return is regressed on the variables that enter into the UK VAR:

$$\frac{r_{t+k}}{k} = a_k + b_{1,k}r_t + b_{2,k}PE_t + b_{3,k}y_t + b_{4,k}PAY + e_{t+k}$$

The sample period is from August 1965 to December 2013. p -values based on Newey-West t -statistics, based on lag length $k + 1$, compared with bootstrapped empirical distributions.

| k | b_1 | $pval$ | b_2 | $pval$ | b_3 | $pval$ | b_4 | $pval$ | R^2 |
|-----|-------|--------|-------|--------|-------|--------|-------|--------|-------|
| 3 | 0.35 | 0.20 | -0.13 | 0.05 | 0.02 | 0.41 | 1.59 | 0.11 | 0.33 |
| 12 | 0.10 | 0.23 | -0.18 | 0.01 | 0.01 | 0.44 | 1.73 | 0.10 | 0.23 |
| 24 | 0.05 | 0.25 | -0.19 | 0.03 | -0.05 | 0.40 | 1.78 | 0.06 | 0.42 |
| 36 | 0.03 | 0.03 | -0.17 | 0.02 | -0.01 | 0.48 | 1.40 | 0.09 | 0.56 |

Table A6b: Multiple regression: US VAR variables

Parameter estimates for in-sample forecast regressions estimated using OLS. The k -period horizon return is regressed on the variables that enter into the US VAR:

$$\frac{r_{t+k}}{k} = a_k + b_{1,k}r_t + b_{2,k}PE_t + b_{3,k}y_t + b_{4,k}DEF + e_{t+k}$$

The sample period is from January 1977 to December 2013. p -values based on Newey-West t -statistics, based on lag length $k + 1$, compared with bootstrapped empirical distributions.

| k | b_1 | $pval$ | b_2 | $pval$ | b_3 | $pval$ | b_4 | $pval$ | R^2 |
|-----|-------|--------|-------|--------|-------|--------|-------|--------|-------|
| 3 | 0.35 | 0.59 | -0.75 | 0.28 | 0.10 | 0.46 | -2.61 | 0.32 | 0.34 |
| 12 | 0.10 | 0.41 | -1.05 | 0.40 | 0.16 | 0.42 | -2.54 | 0.33 | 0.16 |
| 24 | 0.05 | 0.52 | -1.18 | 0.37 | 0.25 | 0.31 | -4.66 | 0.19 | 0.25 |
| 36 | 0.04 | 0.09 | -1.22 | 0.27 | 0.36 | 0.29 | -5.73 | 0.15 | 0.39 |

Figure A1: Model-based Measures of UK Expected Returns




Figure A2: Model-based Measures of US Expected Returns

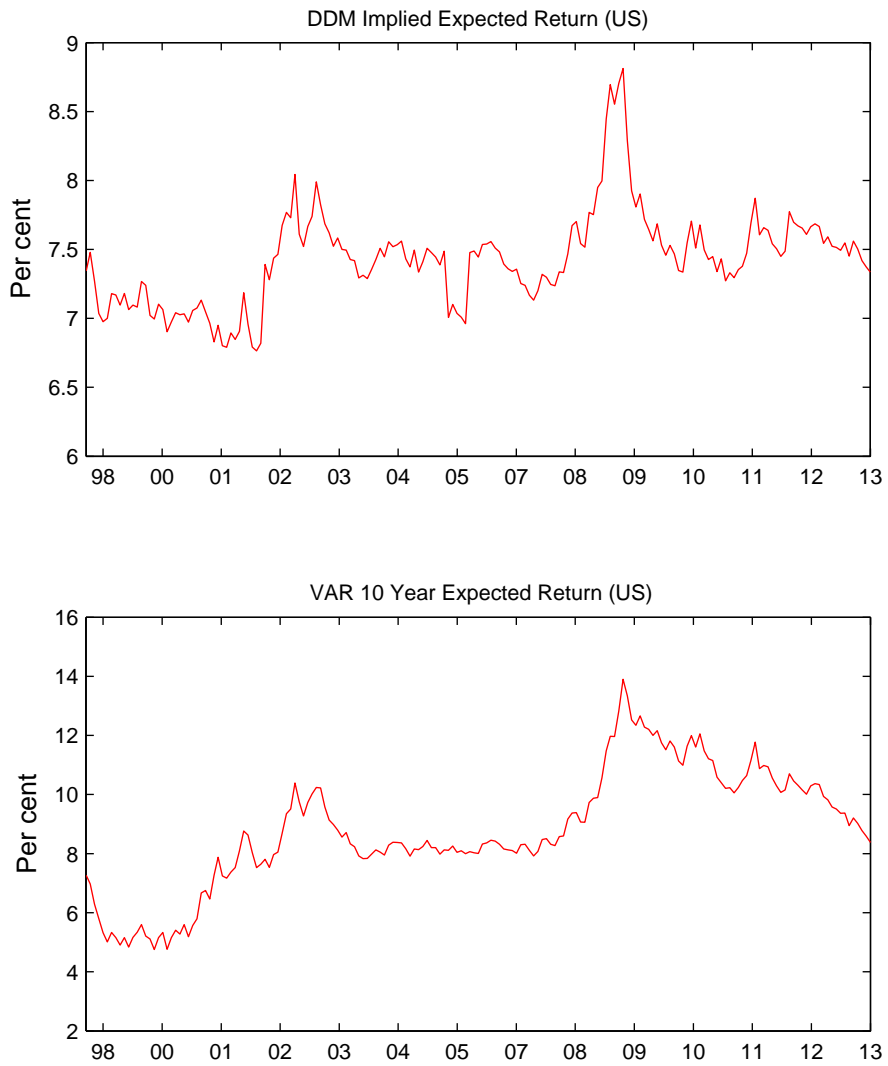


Figure A3: UK Equity Return Decomposition


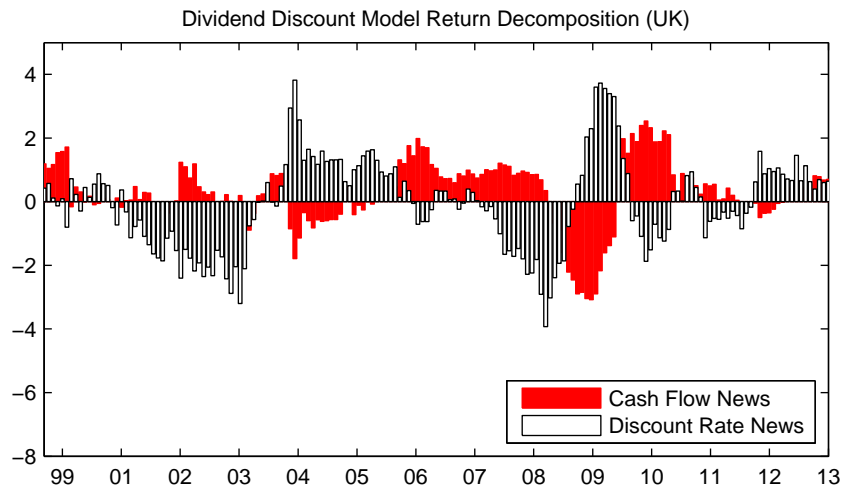


Figure A4: US Equity Return Decomposition



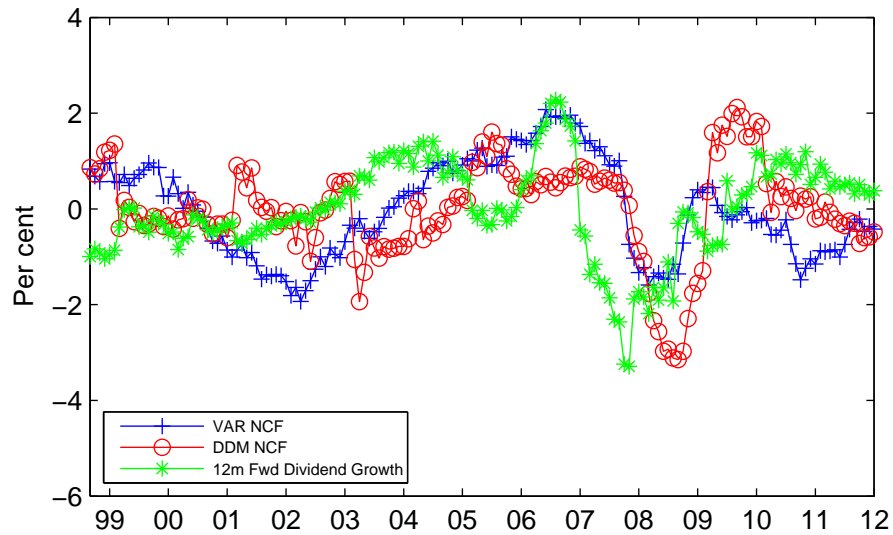
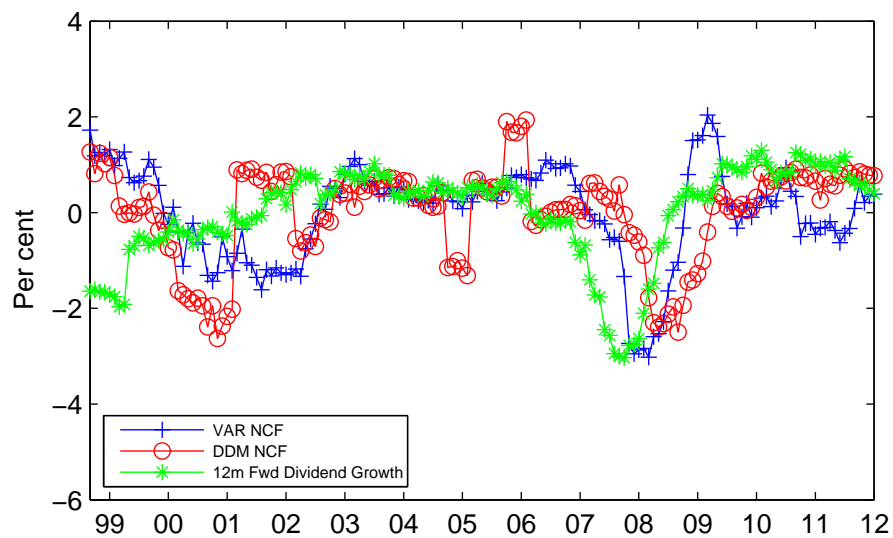


Figure A5: Cash Flow News Against Realised Dividends (normalised)

a. UK



b. US



References

- Ammer, J. and Mei, J. (1996). “Measuring International Economic Linkages with Stock Market Data”, *Journal of Finance*, 51(5), pp. 1743–1763
- Bernanke, B. S. and Kuttner, K.N. (2005). “What Explains the Stock Market’s Reaction to Federal Reserve Policy?”, *Journal of Finance*, 60(3), pp. 1221–125
- Bredin, D., Hyde, S. J., Nitzsche, D. and O’Reilly, G. (2007). “UK Stock Returns and the Impact of Domestic Monetary Policy Shocks”, *Journal of Business, Finance & Accounting*, 34(5&6), pp. 872–888
- Campbell, J. Y. (1991). “A Variance Decomposition for Stock Returns”, *Economic Journal*, 101(405), pp. 157–179
- Campbell, J. Y. and Ammer, J. (1993). “What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns”, *Journal of Finance*, 48(1), pp. 3–37
- Campbell J. Y, Giglio, S. and Polk, C. (2013). “Hard Times”, *Review of Asset Pricing Studies*, 3(1), pp. 95–132
- Campbell, J. Y., Giglio, S., Polk, C. and Turley, R. (2014). “An Intertemporal CAPM with Stochastic Volatility”, *NBER Working Paper* No. 18411
- Campbell, J. Y., Polk, C. and Vuolteenaho, T. (2010). “Growth or Glamour? Fundamentals and Systematic Risk in Stock Returns”, *Review of Financial Studies*, 23(1), pp. 305–344
- Campbell, J. Y., and Shiller, R. J. (1988). “Stock Prices, Earnings, and Expected Dividends”, *Journal of Finance*, 43(3), pp. 661–676
- Campbell, J. Y., and Shiller, R. J. (1989). “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors”, *Review of Financial Studies*, 1(3), pp. 195–228
- Campbell, J. Y., and Thompson, S. B. (2008). “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?”, *Review of Financial Studies*, 21(4), pp. 1509–1531
- Campbell, J. Y. and Vuolteenaho, T. (2004). “Bad Beta, Good Beta”, *American Economic Review*, 94(5), pp. 1249–1275



Chen, L., Da, Z. and Zhao, X. (2013). “What Drives Stock Price Movements?”, *Review of Financial Studies*, 26(4), pp. 841–876

Clark, T. E. and West, K. D. (2007). “Approximately Normal Tests for Equal Predictive Accuracy in Nested Models”, *Journal of Econometrics*, 138(1), pp. 291–311

Cochrane, J. H. (2011). “Discount Rates”, *Journal of Finance*, 66(4), pp. 1047–1108

Conrad, J. and Kaul, G. (1998). “An Anatomy of Trading Strategies”, *Review of Financial Studies*, 11(3), pp. 489–515

Cuthbertson, K., Hayes, S. and Nitzsche, D. (1999). “Explaining Movements in UK Stock Prices”, *Quarterly Review of Economics and Finance*, 39(1), pp.1–20

Diebold, F. X. and Mariano, R. (1995). “Comparing Predictive Accuracy”, *Journal of Business and Economic Statistics*, 13(3), pp. 253–265

Engsted, T., Pederson, T. Q. and Tanggaard, C. (2012). “Pitfalls in VAR-based Return Decompositions: A Clarification”, *Journal of Banking and Finance*, 36(5), pp. 1255–1265

Fama, E. F., and French, K. R. (1988). “Dividend Yields and Expected Stock Returns”, *Journal of Financial Economics*, 22(1), pp. 3–25

Fama, E. F., and French, K. R. (1989). “Business Conditions and Expected Returns on Stocks and Bonds”, *Journal of Financial Economics*, 25(1), pp. 23–49

Garrett, I. and Priestly, R. (2012). “Dividend Growth, Cash Flow and Discount Rate News”, *Journal of Financial and Quantitative Analysis*, 47(5), pp. 1003–1028

Golez, B. (2014). “Expected Returns and Dividend Growth Rates Implied by Derivative Markets”, *Review of Financial Studies*, 27(1), pp. 790–822

Gordon, M. J. (1962). “The Investment, Financing, and Valuation of the Corporation”, Homewood, IL, R. D. Irwin

Inkinen, M., Stringa, M. and Voutsinou, K. (2010). “Interpreting Equity Price Movements since the start of the Financial Crisis”, *Bank of England Quarterly Bulletin*, 2010Q1

Keim, D. B., Stambaugh, R. F. (1986). “Predicting Returns in the Stock and Bond Markets”, *Journal of Financial Economics*, 17(2), pp. 357–390

Lacerda, F. and Santa-Clara, P. (2010). “Forecasting Dividend Growth to Better Predict Returns”, *University of Chicago Working Paper*

Lamont, O. (1998). “Earnings and Expected Returns”, *Journal of Finance*, 53(5), pp. 1563–1587

Lettau, M. and Ludvigson, S. (2005). “Expected Returns and Expected Dividend Growth”, *Journal of Financial Economics*, 76(3), pp. 583–626

Li, Y., Ng, D. T. and Swaminathan, B. (2013). “Predicting Market Returns Using Aggregate Implied Cost of Capital”, *Journal of Financial Economics*, 110(2), pp. 419–436

Ljungqvist, A., Malloy, C. and Marston, F. (2009). “Rewriting History”, *Journal of Finance*, 64(4), pp. 1935–60

Maio, P. and Philip, D. (2014). “Macro Variables and the Components of Stock Returns”, Working Paper

Neely, C.J., Rapach, D.E., Tu, J. and Zhou, G. (2014). “Forecasting the Equity Risk Premium: The Role of Technical Indicators”, *Management Science*, 60(7), pp. 1772–1791

Pastor, L., Sinha, M. and Swaminathan, B. (2008). “Estimating the Intertemporal Risk-return Tradeoff Using the Implied Cost of Capital”, *Journal of Finance*, 63(6), pp. 2859–97

Price, S. and Schleicher, C. (2005). “Returns to Equity, Investment and Q: Evidence from the UK”, *The Manchester School*, 43(s1), pp. 32–57

Rapach, D. E., Strauss, J. K. and Zhou, G. (2010). “Out-of-sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy”, *Review of Financial Studies*, 23(2), pp. 821–862

Rapach, D. E. and Wohar, M. E. (2005). “Valuation Ratios and Long-horizon Stock Price Predictability”, *Journal of Applied Econometrics*, 20(3), pp. 327–344

Rozeff, M. S. (1984). “Dividend Yields are Equity Risk Premiums”, *Journal of Portfolio Management*, 11(1), pp. 68–75

Stambaugh, R. F. (1999). “Predictive Regressions”, *Journal of Financial Economics*, 54(3), pp. 375–421

Stock, J. H., and Watson, M. W. (2003). “Forecasting Output and Inflation: The Role of Asset Prices”, *Journal of Economic Literature*, 41(3), pp. 788–829

Vuolteenaho, T. (2002). “What Drives Firm-Level Stock Returns?”, *Journal of Finance*, 57(1), pp. 233–64

Welch, I. and Goyal, A. (2008). “A Comprehensive Look at the Empirical Performance of Equity Premium Prediction”, *Review of Financial Studies*, 21(4), pp. 1455–508

Wells, S. and Wetherilt, A. (2004). “Long-horizon Equity Return Predictability: Some New Evidence for the UK”, *Bank of England Working Paper 244*

