Long-horizon equity return predictability: some new evidence for the United Kingdom

*Anne Vila Wetherilt**

and

*Simon Wells ***

Working Paper no. 244

* Monetary Instruments and Markets Division, Monetary Analysis, Bank of England, Threadneedle Street, London, EC2R 8AH.
E-mail: anne.vila-wetherilt@bankofengland.co.uk
** Sterling Markets Division, Bank of England.
E-mail: simon.wells@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We are very grateful to Mark Salmon for his advice on this research. We would further like to thank Roger Clews, Simon Price, Peter Westaway, two referees and seminar participants at the Bank of England and City University Business School for their helpful comments.

Copies of working papers may be obtained from Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone 020 7601 4030, fax 020 7601 3298, e-mail mapublications@bankofengland.co.uk.

Working papers are also available at www.bankofengland.co.uk/wp/index.html.

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

ISSN 1368-5562

Contents

Abstract

This paper revisits the issue of long-horizon equity return predictability for the United Kingdom in the context of the dynamic dividend discount model of Campbell and Shiller. This model attributes predictable variation in equity prices to predictable variation in expected returns. The model is supported by the theoretical asset pricing literature, which shows how the variation in expected returns can be related to investors' time-varying preferences for risk. The paper considers various empirical specifications that are consistent with the Campbell and Shiller model and finds that they are supported by UK equity data. In particular, there is weak evidence that the dividend yield has predictive ability for long-horizon excess returns. The paper also examines some of the econometric issues brought up by recent research, in particular the small-sample bias, and applies appropriate statistical corrections. It further shows that the model's predictive ability depends greatly on the sample period over which the model is estimated.

Summary

In this paper, we revisit the issue of long-horizon equity return predictability for the United Kingdom in the context of the dynamic dividend discount model of Campbell and Shiller. This model attributes predictable variation in equity prices to variations in expected returns. The model is supported by the theoretical asset pricing literature, which shows how the variation in expected returns can be related to investors' time-varying preferences for risk.

In the past, this model has received ample support from the data. In particular, the dividend yield appeared to do a reasonably good job at predicting long-horizon excess returns. Moreover, predictability was found to increase with the return horizon. But more recent research has questioned the statistical validity of these claims. In particular, incorrect econometric treatment may have led to overrejection of the null hypothesis of no predictability. Researchers have also found that simple predictability models may be unstable. In some papers, it appears that simply extending the sample period by a few years, or altering the forecast horizons, can alter both the sign of the regression coefficients and their statistical significance, and that over some periods US dividend yields do not forecast long-horizon equity returns.

Using quarterly data for the United Kingdom for the period 1965 Q1-2002 Q4, we first estimate a simple model of return predictability that relates observed excess returns to the dividend yield. Second, we focus on the small-sample issue and consider a range of statistical corrections. Third, we address the issue of robustness by estimating the dividend yield model across a range of sample periods and forecasting horizons. Although the paper does not formally address the all-important issue of model selection, we briefly discuss the forecasting performance of the earnings yield, as an alternative to the dividend yield.

We find evidence that standard valuation ratios such as the dividend and earnings yield help to forecast UK long-horizon equity returns. This result is not stable across subsample periods. In particular, we find that predictability declined significantly during the period of rapidly rising returns of the late 1990s. But as returns started falling in late 2000, the significance of the regressions was restored. The research also confirms that the relationship between the dividend yield and excess returns is highly sensitive to both the chosen return horizon and the sample period.

1. Introduction

Policymakers are interested in asset prices for a number of reasons. First, asset prices form part of the transmission mechanism and therefore enter forecasts of future growth and inflation. Second, asset prices embody forecasts of market participants about future states of the world, and can – at least in principle – be used to obtain information about market expectations. Related, unexpected changes in asset prices can be attributed to unexpected changes in the key factors that determine asset prices, which may in turn be of interest to policymakers. In order to carry out this type of analysis, policymakers need a model of asset price determination. In this paper, we review the leading empirical model of equity price determination and discuss the efficacy of this model in providing long-horizon equity price forecasts.

Up to the early 1980s, most financial economists believed that long-horizon equity returns could not be forecasted. But, in the late 1980s, people like Campbell and Shiller (1988a) and (1988b) and Fama and French (1989) showed that long-horizon returns are highly predictable. Much of the empirical predictability was attributed to either the dividend yield or the earnings yield. Consequently, when in the late 1990s dividend yields in the United States and United Kingdom fell to unprecedented lows, many feared that equity prices would have to fall by very large amounts in order to bring the dividend yield back to its historical mean.(1)

Recent developments in the theoretical asset pricing literature have deepened our understanding of the main drivers of equity prices. Currently, many economists agree that equity prices display predictable variation over time and that this is more a reflection of predictable variation in expected returns than in expected cash flows. The theoretical models further show how this variation in expected returns can be related to investors' time-varying preferences for risk.

Yet, in spite of substantial theoretical and empirical support for long-run predictability (see eg Campbell, Lo and MacKinlay (1997)), the empirical literature has recently become aware of serious statistical problems that affect tests of long-run predictability. First, some authors have highlighted problems of model selection. Apart from the dividend yield, empirical studies have examined a range of explanatory variables, both financial and macroeconomic. They have found that the explanatory power of these variables can differ markedly, thereby affecting the price forecasts (see eg Lamont (1998)).

(1) See eg Campbell and Shiller (2001). See also Vila Wetherilt and Weeken (2002) for a discussion of the behaviour of the dividend yield in the United Kingdom.

Researchers have also found predictability to vary over time. For example, Ang and Bekaert (2003) show that simply extending the sample by a few years, or altering the forecast horizons, can alter both the sign of the regression coefficients and their statistical significance. Goyal and Welch (2002) claim that the predictive power of the US dividend yield was present only in two years, 1974 and 1975. In related work, Pesaran and Timmerman (1995, 2002) have argued that the forecasting model is likely to change over time, as market participants themselves learn and develop better forecasting models.

Second, in a recent paper, Ang and Bekaert (2003) argue that inaccurate econometric techniques may have led researchers to be overly optimistic about long-horizon predictability. In particular, they demonstrate that the small samples that are typically used in long-horizon regressions may have lead to over rejection of the null hypothesis of no-predictability in many papers. Using more robust estimation techniques, they find that over the period 1975-99, US dividend yields no longer forecast long-horizon equity returns. Researchers have also pointed out that many of the explanatory variables used in predictability regressions may be non-stationary, meaning that estimated relations may be spurious.

In this paper, we revisit the issue of long-horizon return predictability for the United Kingdom. We first estimate the popular dividend yield model and show that predictability does indeed vary over time. Second, we examine the small sample issue and use the statistical corrections proposed by Ang and Bekaert (2003). Third, we address the issue of robustness by estimating the dividend yield model across a range of sample periods and forecasting horizons. Although the paper does not formally address the all-important issue of model selection, we briefly discuss the forecasting performance of the earnings yield, as an alternative to the dividend yield.

Using quarterly data for the period 1965 Q1-2002 Q4, we find evidence that standard valuation ratios such as the dividend and earnings yield help to forecast UK long-horizon equity returns. But this result is not stable across subsample periods. In particular, we find that predictability declined significantly during the period of rapidly rising returns of the late 1990s. As returns started falling in late 2000, long-run predictability returned. The results further suggest that, while the present value model provides us with a forecast model that is consistent with mainstream asset pricing theory and receives support from the UK data, we nevertheless face difficult practical issues that mainly stem from the lack of stability.

2. Understanding return predictability

Predictability of equity returns can arise in a range of theoretical models. In this section, we rely on a simple accounting model to explain the relationship between dividend yields and expected returns. This is the most popular approach in the finance literature. We start with the dynamic dividend discount model of Campbell and Shiller (1988a). In this model, equity prices are determined by the present discounted value of future dividends and future expected returns (serving as discount factors).[2] That expression is linearised to obtain

$$p_t = E_t \sum_{j=0}^{\infty} \rho^j [(1-\rho) d_{t+j+1} - r_{t+j+1}] + \kappa \qquad \textbf{(1)}$$

where p is the log price, d the log dividend, r the rate of return and ρ and κ are constant parameters coming from the linear approximation used to obtain **(1)**.[3] [4] [5] ρ is bounded between 0 and 1 (but is not a 'discount factor' as such), and ensures the sum is finite. It follows that, in this model, equity prices are high when future dividends are expected to be high and/or future equity returns are expected to be low.

Equation **(1)** is often rearranged in terms of the price-dividend ratio (the inverse of the dividend yield).

$$p_t - d_t = E_t \sum_{j=0}^{\infty} \rho^j [\Delta d_{t+j+1} - r_{t+j+1}] + \kappa \qquad \textbf{(2)}$$

Equation **(2)** shows that any variation in the dividend yield must imply changing expectations of future dividend growth and/or future equity returns. Or put differently, the dividend yield $(d_t - p_t)$ forecasts future expected equity returns ($E_t\, r_{t+j}$) in excess of dividend growth (Δd_{t+j}).

To see more clearly why the dividend yield can forecast expected returns, it is useful to consider a simple example developed by Cochrane (2001). In this example, he chooses

---

[2] Expected returns can be thought of as the risk-free rate plus the equity risk premium.

[3] Equation **(1)** is obtained in a number of steps: first, write the one-period equity return (the dividend yield plus capital gain). Second, iterate this identity forward to obtain a multi-period relationship between prices, expected returns and expected dividends. Third, linearise this equation using a Taylor expansion and take logs. See Cochrane (2001) for more details.

[4] If in addition one imposes the assumption of a constant discount rate r and a constant dividend growth rate, then equation **(1)** reduces to the Gordon dividend discount model.

[5] The simple present value model is challenged in the literature that deals with bubbles (eg Froot and Obstfeld (1991), or dividend smoothing (eg Ackert and Hunter (2000)). The resulting models are no longer linear, and the respective authors claim that they are better at reconciling observed equity prices and fundamentals.

specific functional forms for the right-hand side variables in equation **(2)**. First, he assumes that (de-meaned) dividend growth, Δd, is a white noise process and therefore not predictable:

$$\Delta d_{t+1} = \varepsilon_{d,t+1} \tag{3}$$

Second, he assumes that expected returns, $E_t r_{t+j}$, are slow moving:

$$r_{t+1} = z_t + \varepsilon_{r,t+1} \tag{4}$$

$$E_t r_{t+1} = z_t = bz_{t-1} + \delta_t \quad b > 0_t \tag{5}$$

where z_t is an unobservable state variable that drives expected returns, δ_t, $\varepsilon_{d,t}$ and $\varepsilon_{r,t}$ are shocks to expected returns, dividend growth and realised (or *ex-post*) returns respectively. Substituting equation **(5)** into **(2)** yields:

$$p_t - d_t = \frac{z_t}{1-\rho b} \tag{6}$$

$$p_{t+1} - d_{t+1} = b(p_t - d_t) + \frac{\delta_{t+1}}{1-\rho b} \tag{7}$$

Using **(7)**, realised returns and prices (using the basic definition of returns and a Taylor expansion) can be written as:

$$r_{t+1} = (1-\rho b)(d_t - p_t) + \left(\varepsilon_{d,t+1} - \frac{\rho}{1-\rho b}\delta_{t+1} \right) \tag{8}$$

$$\Delta p_{t+1} = (1-b)(d_t - p_t) + \left(\varepsilon_{d,t+1} - \frac{1}{1-\rho b}\delta_{t+1} \right) \tag{9}$$

We can now explain the role of the dividend yield. Suppose the dividend yield is below its mean, so that prices are high compared to their current dividend. From **(2)** we know that a low dividend yield (or a high price dividend ratio) means that investors must expect future stock returns to be low (recall we assumed they could not forecast dividend growth). In order for these low stock returns to materialise, prices must depreciate from the current low level.[6] This mechanism is summarised in equations **(8)** and **(9)**, where we see that the dividend yield predicts positive equity returns.

We need some persistence (b>0) in the dividend yield, otherwise we would have a one-to-one relationship between *ex-post* (log) returns, the (log) dividend yield, and in turn, price growth.

---

[6] This is the argument used in Campbell (2001) and Campbell and Shiller (2001).

Cochrane (2001) shows that this would require implausibly large price changes for the dividend yield to revert to its mean. The Cochrane model also illustrates why past returns are often found to be poor predictors of future returns. This may seem puzzling at first, given the high predictability of the expected return component. Equation **(8)** shows, however, that returns include both dividend ($\varepsilon_{d,t}$) and expected returns news (δ_t), and as such are a poor proxy for expected returns alone. Therefore, lagged returns cannot be expected to have the same predictive power as the dividend yield.

The present value model, as formulated in equation **(2)**, provides the motivation for setting up a forecasting equation that relates equity returns to the dividend yield:

$$r_{t+j} = \alpha + \beta(d_t - p_t) + u_{t+j} \tag{10}$$

In equation **(10)**, predictability of equity returns stems from the dividend yield only. In the next section, we describe how equation **(10)** has performed in empirical applications. We will also briefly comment on alternative predictors.

To summarise the discussion so far, we have shown how in the standard present value model, return predictability arises if one assumes persistence in expected returns. This assumption creates mean reversion in the dividend yield, providing a rationale for the latter's observed role in predicting *ex-post* equity returns. Moreover, the slow-moving, time-varying nature of expected returns that arises in the present value model is entirely consistent with more general asset pricing models that generate time-varying risk premia. For example, in the habit model of Campbell and Cochrane (1999), risk aversion moves countercyclically, producing in turn a countercyclical pattern in expected returns: investors demand low risk premia (or expected returns) at the peak of a business cycle when their risk aversion is low, and require high equity risk premia at the bottom of a recession, when they are more risk-averse.

In a related model, Lettau and Ludvigson (2001) show that variations in aggregate consumption, asset wealth and labour income can explain time variation in expected returns. For example, when equity returns are expected to be lower in the future, consumers in their model will reduce their consumption out of current asset wealth and labour income. As a result, deviations of consumption from its long-run relationship to asset wealth and labour income can predict future expected returns.[7]

[7] More precisely, in their model, deviations of consumption from its shared trend with asset wealth and labour income predict future expected returns.

The dividend yield regression (equation **(10)**) relies on a second, crucial assumption, namely non-predictability of dividend growth. Recently, however, Lettau and Ludvigson (2003) have challenged the assumption that dividend growth is non-predictable. Going back to the present value equation **(2)**, when dividend growth is forecastable, changes in the dividend yield may reflect changes in expected dividend growth and/or expected returns. This makes it more difficult to asses the exact forecasting contribution of the dividend yield. In addition, equation **(1)** may be misspecified.

Ang and Bekaert (2003) formally test whether the dividend yield predicts dividend growth, but fail to find strong evidence to support this hypothesis. Lettau and Ludvigson (2003) compare the present value equation **(10)** with an alternative consumption based present value model that relates expected returns to wealth,[8] consumption growth and dividend growth. They find the dividend yield to have little forecasting power for long-horizon returns (one to six years). In contrast, dividend growth significantly contributes to forecasting returns over horizons greater than one year. [9]

Finally, it is worth noting that the previous discussion did not consider the possibility that return predictability may result from market inefficiency. Instead, the arguments presented rely on rational investor behaviour in response to time variation in expected returns. Market prices are assumed to fully reflect this behaviour. This is the dominant view on return predictability in the finance profession at present.

3. A brief discussion of model selection issues

The finance profession appears in broad agreement on the use of the present value model as a good starting point for an empirical model of return predictability. Specifically, equation **(10)** is routinely used as the main tool to test predictability of excess returns at long horizons. The choice of equation **(10)** rests on the view that expected returns are time varying and are well proxied by the dividend yield. This was explained in some detail in the previous section. But the use of the dividend yield model also receives support from the data. For example, it is well documented that the dividend yield is highly correlated with variables which are thought

(8) More precisely, deviations of consumption from its shared trend with asset wealth and labour income (as in the earlier cited Lettau and Ludvigson (2001)).
(9) Menzly *et al* (2003) provide a theoretical model to highlight the relationship between expected returns, the dividend yield and dividend growth. Specifically, they show that when both risk preferences and dividend growth are time varying, the simple, linear relationship between expected returns and the dividend yield (as in equation **(10)**) no longer holds. As in Lettau and Ludvigson (2003), this result depends on the predictability of dividend growth.

to share the expected return's positive co-movement with the business cycle, such as credit and term spreads (Fama and French (1989)).

Financial economists have also explored alternative explanatory variables when estimating equation **(10)**. A popular approach is to replace the dividend yield by the earnings yield and the payout ratio in equation **(10)**.[10] This approach is used by Lamont (1998) and Nelson (1999) and avoids two problems related to dividends: first, changes in dividend policy are not necessarily captured by the dividend yield alone. This is an important concern in the light of recent evidence that corporations are relying more on share purchases at the expense of dividend payouts.[11] Second, many companies (especially young ones) do not pay any dividends at all. In estimating an earnings model, researchers also follow market practice more closely, since market participants tend to look at earnings more than dividends. A related approach is considered in Sharpe (2002) who uses survey-based expectations for earnings growth rather than actual earnings. In what follows, we will estimate both dividend yield and earnings yield models.

The literature has also considered more complex models that include variables that move closely together with the business cycle. In doing so, researchers aim to capture time variation in expected returns. These models include financial variables that are known to track the business cycle, such as short-term interest rates, the term spread and the default spread (see eg Fama and French (1989)). They also include macroeconomic variables, such as money supply, inflation, GDP, industrial production. But the list of candidate macro and financial variables and their possible combinations is very large, giving rise to a model selection problem. A further problem is highlighted by Pesaran and Timmerman (1995, 2002) is that the set of key forecasting variables is likely to change over time, as market participants themselves learn and develop better forecasting models.

Formal selection criteria are available in the literature (see eg Pesaran and Timmerman (1995, 2002) or Hoover and Perez (1999) and Hendry and Krolzig (1999)). They do not form part of the present research, as our objective is to focus on econometric issues other than model selection. These issues (to be outlined in Section 5) will be treated within the context of the

---

(10) One can easily rewrite equations **(1)** and **(2)** using the identity D/P = (E/P)*(D/E).

(11) See eg Nelson (1999) and Wadhwani (1999).

simple dividend yield model (equation **(10)**), and a few simple variants, including the earnings yield model.

4. Data

Equation **(10)** forms the basis of our empirical work and is estimated for FTSE All-Share excess returns. We obtained quarterly data from Datastream for FTSE All-Share prices, dividends and earnings, from 1963 Q1 to 2002 Q4. We also use a longer set of annual data for the period 1926-2002. For the earlier part of that sample, the data were obtained from Global Financial Data.

To test whether long-run predictability is affected by the return horizon (as in Ang and Bekaert (2003)), we estimate our models for three investment horizons: one, two and four years. The nominal (log) return on equity over a given period is calculated by subtracting the log of the initial price from the log of the terminal price plus dividends paid through the period. We next construct excess returns by subtracting the return on a portfolio of UK Treasury bills from the equity return index. To calculate real dividend and earnings growth rates, we deflate dividends and earnings by the RPI. As we will explain below, we also consider models with a short-term interest rate. For this purpose, we use the three-month Treasury Bill rate.[12] Table A presents summary statistics for our key series, while Charts 1 and 2 plot the dividend yield and real one-year dividend growth rates, respectively.

Charts 1: FTSE All-Share dividend yield

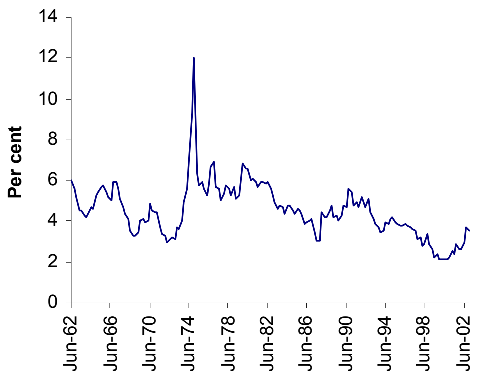

Chart 2: One-year real dividend growth

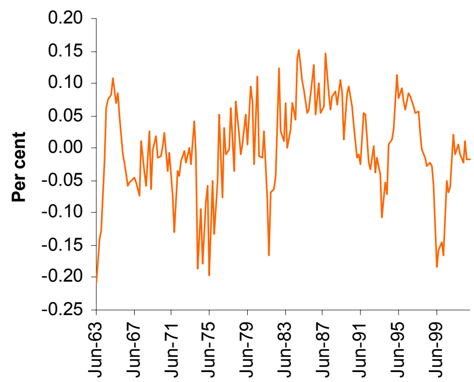

[12] Much of the empirical literature uses a detrended short rate (eg Campbell (1991) and Lamont (1998)). For this purpose, we used the three-month Treasury bill rate relative to its twelve-month moving average. We found, however, that this did not affect the regression results. For this reason, we will only report the results for the raw series.

Table A: Univariate summary statistics (1963:01 – 2002:04)

	Mean	Maximum	Minimum	St.Dev.	Persistence
One-year excess return	0.04	0.82	0.88	0.22	0.75
Two-year excess return	0.09	0.74	-1.17	0.28	0.86
Four-year excess return	0.18	1.01	-0.67	0.30	0.89
Real dividend growth	0.002	0.15	-0.20	0.07	0.66
Real earnings growth	0.005	0.26	-0.43	0.14	0.83
Dividend yield	-3.14	-2.12	-3.86	0.29	0.93
Earnings yield	-2.61	-1.39	-3.35	0.40	0.95
Pay-out ratio	-0.55	-0.14	-1.01	0.17	0.91
Risk-free rate	2.04	2.79	1.30	0.38	0.93

Note: Some observations are lost when constructing the returns series. All variables are in logs. Persistence is measured as the first-order autocorrrelation.

While Chart 1 clearly illustrates the slow-moving nature of the dividend yield, Chart 2 shows that real dividend growth is a more variable and rapidly mean-reverting series. Table A confirms this observation, namely that real dividend growth has lower persistence (0.66) than the dividend yield (0.93). This has two implications. First, the data confirm the assumption made in Section 2 of slow mean reversion in the dividend yield (the b coefficient).

Second, the relatively high persistence of the dividend growth series appears at odds with equation **(3)**, which assumed that dividend growth is white noise, and hence exhibits no persistence. Recall that the model in Section 2 showed that the time-varying dividend yield could explain variations in either expected returns or expected dividend growth, or both. By ruling out the latter possibility, the model allowed us to write down a simple regression equation associating expected returns and the dividend yield only. So how does the apparent persistence in the dividend growth series affect this regression model? Campbell, Lo and MacKinlay (1997) argue that even if there is some small predictable component in dividend growth, it is likely that the variance in expected dividend growth is substantially less than the variance in the expected returns. In this case, it is reasonable to proceed under the assumption that all the time variation in the dividend yield reflects changes in expected returns, thereby validating the simple dividend-yield regression. Cochrane (2001) supports this view by providing strong evidence that the dividend yield has no predictive power for dividend growth. A similar result is found for UK data.[13]

Table A further shows that other candidate regressors, such as the earnings yield, the pay-out ratio, or even the risk-free rate are all highly persistent. How this affects the statistical

(13) This is done by regressing the dividend growth on the dividend yield. Running this regression on non-overlapping data for 1925-2002 yielded an R^2 of 0.02, comparable to Cochrane's 0.06. The full set of results is available from the authors.

inference will be discussed in the next section. Table B presents the results of unit-root tests for all variables of interest, using the familiar augmented Dickey-Fuller (ADF) procedure. In the case of both dividend and earnings growth, we can reject the null hypothesis of a unit root with 90% confidence. Stationary dividend (earnings) growth suggests that the levels of dividends (earnings) are a unit root process.

Table B: Unit root tests (1963:01 – 2002:04)

	ADF t-statistic	ADF p-value	Lags used
One-year excess return	-2.71	0.07	12
Two-year excess return	-3.20	0.02	11
Four-year excess return	-3.07	0.03	3
Real dividend growth	-2.70	0.08	8
Real earnings growth	-3.82	0.003	12
Real dividends	-1.04	0.74	5
Real earnings	-1.15	0.69	5
Dividend yield	-2.30	0.17	0
Earnings yield	-2.00	0.29	4
Pay-out ratio	-2.35	0.16	6
Risk-free rate	-2.54	0.11	1

Note: All ADF tests use a constant. Lags are selected with the AIC criterion.

The table also shows that there is little evidence for stationarity of the dividend yield (and earnings yield) series. This poses an econometric problem, as the estimation of equation **(10)** using OLS requires the stationarity of both the regressors (eg the dividend yield) and the regressands (the return index). The fact that the variables we are interested in may be unit root or near unit root processes suggests the possibility of spurious regressions. But the evidence is not conclusive. In particular our sample size is relatively small, and ADF tests are well known to suffer from low power in small samples. Hence it is possible that the series is indeed stationary, but that this cannot be detected in a relatively short sample. Using the longer data sample (1926-2002), we find strong evidence of stationarity using ADF tests. But even these tests could be misleading, as the quality of the data may deteriorate as we go back through time. Bearing these warnings in mind, we proceed on the assumption that the dividend yield (and earnings yield) series are stationary. In Section 5.2, we discuss in more detail the problem of spurious regressions as one of many econometric issues that we may encounter.

5. Econometric issues

Before estimating our empirical model, we discuss a number of econometric problems that may affect the results and their interpretation. More specifically, we discuss three inter-related problems that frequently arise in financial regressions, and are summarised in Chart 3.[14] We start (in Section 5.1) by discussing the issues that arise when running regression models containing long-horizon returns. We next consider the potential pitfalls of including persistent regressors in the model (Section 5.2). Related to this, we then discuss the problem of stochastic regressors (Section 5.3).

Chart 3: Common econometric problems encountered with financial regressions

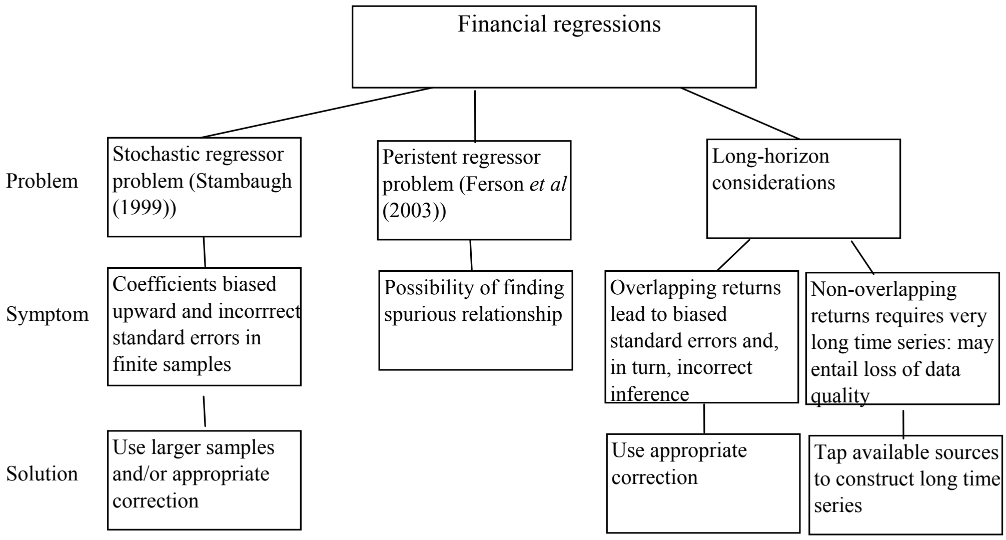

[14] Note that this list of problems is not meant to be exhaustive. Rather, it reflects the issues that we consider most important for the regressions estimated in this paper.

5.1 Long-horizon regressions

As mentioned in Section 1, finance textbooks suggest that return predictability is easier to detect at long horizons, which may be congruent with predictability reflecting slow-moving changes in investors' risk preferences. One possible interpretation of this finding is that it reflects nothing more than a compounding over time of a very small predictable component of one-period returns. We can demonstrate this formally in the simple model of Section 2. If the true data generating process is well approximated by equations **(3)** and **(4)**, then using the fact that the *k*-period return is equal to the sum of *k* one-period returns, we can use repeated substitution to show that the *k*-period return is given by:

$$r_{t,t+k} = r_{t,t+1} + r_{t+1,t+2} + ... + r_{t+k-1,t+k} = (1 + b + b^2 + ... + b^k)z_t + e_t \qquad \textbf{(12)}$$

So if the state variable, z_t, is highly persistent (ie the persistence parameter b is close to 1), then from equation **(12)** we can see that the estimated coefficient on the dividend yield should increase with the return horizon, k. A similar argument can be presented for the R^2 of the regression.

But recently some authors, notably Ang and Bekaert (2003), have challenged the result that equity returns are predictable, whatever the return horizon. They propose that when appropriate small-sample adjustments are made, the significance of the dividend yield is greatly reduced. More importantly, they find that the size, and in some cases the *sign,* of the estimated coefficient is dependent on the return horizon.

Long-horizon regressions present the researcher with an awkward dilemma: either work with a smaller data set or with overlapping returns. The former approach (taken by Fama and French (1988) among others) requires a long run of data and precludes the possibility of studying returns of periods longer than one year.(15) Furthermore this approach may require the splicing together of data from different sources and therefore lead to a reduction in quality. The more common approach is to work with overlapping returns, which complicates inference based on OLS regression results.

These problems were addressed by Hansen and Hodrick (1980) who showed that using *k*-period overlapping returns in OLS regressions results in residuals that are autocorrelated up to order (*k*-1) even under the null hypothesis of 'no predictability'. Formally,

(15) For example, 200 years of data would be required to obtain a series of 100 non-overlapping two-year returns.

$$E(e_{t+k,k}e_{t+k-j,k}) \neq 0, \forall |j| < k$$

where $e_{t+k,k}$ is the k-period forecast error at time $t+k$.

In addition to serial correlation in the residuals, it is generally thought that the volatility of asset returns is also serially correlated (see discussion in Campbell, Lo and MacKinlay (1997)). This makes an assumption of homoscedasticity inappropriate. As is well known, OLS standard errors are biased in the presence of autocorrelated and heteroscedastic residuals. This means that all standard errors need to be corrected and, importantly, the appropriate correction needs to take to account of the small sample size.

We follow Ang and Bekaert (2003) and report test statistics, based on four alternative methods for calculating standard errors. We first compute the usual OLS standard errors. Although inappropriate, we include them simply to provide a benchmark for the set of alternative standard errors, all of which correct for heteroscedasticity as well as for the autocorrelation induced by the use of overlapping returns. These include a modified version of the errors proposed in Hansen and Hodrick (1980) (MHH errors), the familiar Newey-West (NW) estimate, and a measure proposed by Hodrick (1992) (Hodrick errors). All are described in detail in the appendix.

Briefly, the MHH and Hodrick standard errors are valid under the restrictive null hypothesis that equity returns have a constant conditional mean – ie the dividend yield (or any other candidate predictor) contains no information about future expected returns. As mentioned previously, under the null, the residuals will be autocorrelated to order (k-1). This information is used to ensure that both the MHH and Hodrick errors are appropriate under the null hypothesis. Under the alternative hypothesis, where equity returns may have a variable conditional mean, the regressors in the model may not completely capture the predictable component of returns. In this case, there may be serial correlation present in the residuals of order greater than (k-1) and our estimated standard errors do not correct for this. Consequently, we can only be certain of their validity under the null.

Hodrick (1992) errors are constructed by imposing more of the null hypothesis (namely that if returns are not predictable, then the k-period error should equal the sum of k one-period errors) and exploiting the properties of covariance-stationary time series to remove the overlapping structure of the regression residuals. The advantage of this approach is that it avoids the need to sum a large number of estimated covariance matrices, which can cause poor small-sample performance of the estimated standard error.

Ang and Bekaert (2003) argue that much of the literature has relied on standard errors (OLS or GMM) that are upwardly biased in small samples. This has lead to frequent over rejection of the null hypothesis of zero predictability. They show that when using the Hodrick (1992) standard errors, the bias is reduced and evidence of long-run predictability is significantly weakened. To illustrate this, Table C presents the coefficient estimates, t-statistics and corresponding p-values for a simple regression of the one-year excess equity return on the dividend yield.

Table C: Comparison of t-statistics and p-values

(Results for one-year excess returns, 1963-2002)

	Constant	Dividend Yield
Coefficient	**1.90**	**0.58**
t-OLS	**8.02**	**7.56**
p-OLS	0.00	0.00
t-NW	**3.53**	**3.16**
p-NW	0.00	0.00
t-MHH	**3.02**	**2.70**
p-MHH	0.00	0.01
t-Hodrick	**2.12**	**2.02**
p-Hodrick	0.04	0.05
R^2	0.30	

Although the null hypothesis can be confidently rejected in all cases, the table clearly shows how relying on the Newey-West correction could lead to over rejection of the null hypothesis of zero predictability, confirming Ang and Bekaert's (2003) results. Specifically, the low Hodrick t-statistic suggests that the dividend yield is only just significant at the 95% confidence level. Note also that the t-statistic resulting from the Newey-West standard errors is higher than from the MHH errors. This suggests that the Newey-West correction places too low a weight on higher order autocorrelations which are known to exist (see the appendix for details).

Hodrick (1992) and Ang and Bekaert (2003) present Monte Carlo evidence that, in small samples, the Hodrick errors provide test statistics with the best *size* (ie one that minimises the probability of rejecting a true null hypothesis). But, without a specific alternative hypothesis, it is not possible to compare the *power* of the tests (probability of not rejecting a false null hypothesis).

5.2 Spurious regressions

Regardless of whether or not overlapping returns are used, we may encounter other econometric problems. In particular, the model for understanding predictability outlined in Section 2 relies on the persistence of the state variable, z, and its relationship with the dividend yield. But if the model is incorrect, and our regressor (the dividend yield) is in fact unrelated to the state variable, it is still possible that we can estimate a (seemingly) significant relationship between them, if the regressor and the state variable are similarly persistent.

This phenomenon – a spurious regression – has been well documented since Granger and Newbold (1974), and most commonly occurs when regressing non-stationary time series. In a more recent paper, Ferson *et al* (2003) show that spurious regressions may arise even if the dependent variable is not a highly persistent series, for example an index of returns on a financial asset. More specifically, they use Monte Carlo simulations to show that regression models such as **(10)**, in which a persistent lagged variable is used to predict stock returns, can produce spurious results if actual returns are driven by a persistent expected return plus a random shock term. Further, their work is based on non-overlapping returns. This suggests that even after correcting standard errors for the autocorrelation in the estimated residuals, induced by overlapping returns, we still run the risk of estimating a spurious regression. But seeing as there is no way of testing for this, we do not aim to address the issue. Rather, we rely on the fact that we have outlined a theoretical framework that formally links the dividend yield and the state variable, and acknowledge the spurious regression problem as a potential caveat.

5.3 Stochastic regressors

A related problem arises when including regressors that have a stochastic component. Even if one could observe the state variable, there may be a finite sample bias if there is correlation between regressor and the error term (recall that a maintained assumption of the classical regression model is that $E[u_t|X_t]=0$). To understand how this relates to our aim, note that the dividend yield at time (t-1) depends upon the price level, p_{t-1}. Observing the dividend yield at t and t-1 therefore contains information about the evolution of the stock price through period t. In turn, this means we have some information about the change in price from t-1 to t – ie we have some information about the return. As a consequence, $E[u_t|X_t,X_{t-1}]$ may be non-zero.

Stambaugh (1999) shows that, in finite samples, this violation of the classical assumption may introduce significant upward bias to the estimated coefficients. And this bias is increasing in the persistence of the regressor, and can be significant (up to half a standard deviation of the estimated coefficient) even for relatively large sample sizes. Although Stambaugh (1999) provides analytical expressions for the finite-sample moments of the estimated coefficient, these are only valid for non-overlapping returns. Since the focus of the current paper is on models with overlapping returns, we simply acknowledge the stochastic regressor problem as a potential caveat to our results and instead focus on making appropriate corrections to the OLS standard errors.

6. Empirical results

In this section, we present the estimation results. We estimate both the simple dividend yield model of equation **(10)**, labelled model 1, and the related earnings yield model, where the dividend yield is replaced by the earnings yield and the pay-out ratio (referred to as model 2). We also estimate variants of these models, first including the risk-free rate (models 1b and 2b), and second adding the lagged dependent variable (models 1c and 2c). Both variants are frequently encountered in empirical work of this type.

In light of the previous discussion, we first estimate the models with non-overlapping returns for the period 1926-2002. Because using these non-overlapping returns significantly reduces the number of observations, we can only do this estimation with one-year returns. Table D below shows the dividend yield to be significant. Yet the model's explanatory power is limited, as indicated by the relatively low R^2. Similar results are obtained for the remainder models.

Table D: Regression results with non-overlapping one-year returns (1926-2002)

Model	Const	Div Yield	Earns Yield	Pay-out Ratio	Short Rate	Lagged Return	R-squared
1	**1.01**	**0.32**					**0.17**
t-stat	4.14	3.96					
1b	**1.08**	**0.33**			**-0.004**		**0.18**
t-stat	4.18	4.03			-0.81		
1c	**1.07**	**0.33**				**0.07**	**0.18**
t-stat	4.13	3.94				0.66	
2	**1.0**		**0.30**	**0.35**			**0.14**
t-stat	3.30		3.31	2.34			
2b	**1.07**		**0.31**	**0.38**	**-0.004**		**0.14**
t-stat	3.29		3.35	2.40	-0.61		
2c	**1.26**		**0.38**	**0.44**		**-0.26**	**0.19**
t-stat	3.95		3.97	2.88		-2.17	

t-statistics are unadjusted.

To interpret the results reported in Table D, note further that (taking model 1 as an example) a coefficient of 0.32 on the log dividend yield translates into a coefficient of around 7 on the level of the dividend yield.[16] So for each percentage point that the level of the dividend yield is above its mean, the expected excess equity return increases by around 7 percentage points. This suggests that excess returns are highly responsive to small changes in the dividend yield. But note that the coefficient we estimate is only slightly higher than those found by 'typical' dividend yield regressions run on US data.

A richer set of results is obtained using overlapping returns, and they form the basis of the estimations in the remainder of the paper. Table E below presents the results for one-year excess returns. Both MHH and Hodrick t-statistics are reported.

Table E: Predictability regressions for one-year excess returns (1963-2002)

Model	Const	Div Yield	Earns Yield	Pay-out Ratio	Short Rate	Lagged Return	R-squared
1	**1.13**	**0.35**					**0.22**
t-MHH	4.60	4.32					
t-Hod	2.26	2.23					
1b	**1.55**	**0.44**			**-0.02**		**0.25**
t-MHH	3.72	3.87			-1.57		
t-Hod	2.55	2.49			-1.56		
1c	**1.11**	**0.34**				**-0.03**	**0.22**
t-MHH	4.24	3.84				-0.19	
t-Hod	2.57	2.53				-0.12	
2	**1.35**		**0.39**	**0.58**			**0.24**
t-MHH	4.65		4.40	3.76			
t-Hod	2.94		2.62	2.88			
2b	**1.59**		**0.44**	**0.54**	**-0.01**		**0.26**
t-MHH	3.82		3.84	3.56	-1.16		
t-Hod	2.78		2.55	2.67	-1.04		
2c	**1.33**		**0.38**	**0.57**		**-0.04**	**0.24**
t-MHH	4.61		4.17	3.77		-0.27	
t-Hod	3.13		2.86	2.82		-0.17	

The results for model 1 and its variants show that the coefficient on the dividend yield is positive and significant, whether using the MHH or Hodrick t-statistic. These results are consistent with the earlier discussed role of the dividend yield as a proxy for expected returns. Replacing the dividend yield by the earnings yield and pay-out ratio (model 2 and variants)

[16] This follows from taking a Taylor expansion about the mean of the log dividend yield - 4.79% - between 1926 and 2002.

has a small impact on the predictive power of the regression (R^2 increases from 0.20 to 0.24 in model 2). The earnings yield coefficients are all positive and significant.

Turning now to the other regressors, we find the coefficient for the risk-free rate to be negative, but never significant. This result is in line with the empirical literature (eg Lamont (1998) and Campbell (1991)), but contradicts Ang and Bekaert (2003), who find the short rate to be the most significant variable in their predictability regressions. We further find that the lagged dependent variable (models 1c and 2c) is never significant. Recall that equation **(8)** showed actual returns to include both dividend and expected returns news. This implies that lagged returns cannot be expected to have the same predictive power as the dividend yield. The results for models 1c and 2c confirm this prediction from the model.

To understand how the above results change as the return horizon is increased, Tables F and G show the estimation results for two and four-year excess returns. At the two-year horizon, the yield variables continue to be positive and significant. At the four-year horizon, however, the Hodrick t-statistic for the dividend yield and earnings yield are no longer significant at the 95% confidence level, even though the magnitude of the estimated β coefficient is larger. The risk-free rate and the lagged dependent variable continue to be insignificant.

Table F: Predictability regressions for two-year excess returns (1963-2002)

Model	Const	Div Yield	Earns Yield	Pay-out Ratio	Short Rate	Lagged Return	R-squared
1	**1.93**	**0.59**					**0.36**
t-MHH	4.39	3.97					
t-Hod	2.47	2.38					
1b	**2.27**	**0.67**			**-0.01**		**0.38**
t-MHH	4.28	4.12			-1.15		
t-Hod	2.59	2.53			-0.75		
1c	**1.84**	**0.56**				**-0.09**	**0.38**
t-MHH	4.35	3.88				-0.85	
t-Hod	2.21	2.01				-0.23	
2	**2.49**		**0.69**	**1.13**			**0.43**
t-MHH	4.64		4.17	4.13			
t-Hod	2.89		2.69	2.44			
2b	**2.50**		**0.70**	**1.13**	**-0.001**		**0.43**
t-MHH	4.41		4.07	4.05	-0.06		
t-Hod	2.77		2.77	2.63	-0.04		
2c	**2.39**		**0.66**	**1.09**		**-0.10**	**0.44**
t-MHH	5.04		4.45	4.32		-1.04	
t-Hod	2.49		2.20	2.20		-0.25	

Table G: Predictability regressions for four-year excess returns (1965-2001)

Model	Const	Div Yield	Earns Yield	Pay-out Ratio	Short Rate	Lagged Return	R-squared
1	**2.77**	**0.84**					**0.43**
t-MHH	5.23	4.54					
t-Hod	1.58	1.44					
1b	**2.85**	**0.86**			**-0.003**		**0.43**
t-MHH	6.81	5.96			-0.19		
t-Hod	1.46	1.39			-0.09		
1c	**2.78**	**0.84**				**-0.03**	**0.44**
t-MHH	3.96	3.40				-0.25	
t-Hod	1.30	1.17				-0.06	
2	**3.33**		**0.96**	**1.30**			**0.46**
t-MHH	7.36		6.17	6.95			
t-Hod	1.40		1.36	1.05			
2b	**3.20**		**0.93**	**1.34**	**0.01**		**0.47**
t-MHH	7.12		6.43	10.12	0.66		
t-Hod	1.39		1.35	1.04	0.27		
2c	**4.01**		**1.15**	**1.70**		**0.11**	**0.46**
t-NW*	7.27		6.63	6.43		1.35	
t-Hod	1.02		0.98	0.81		0.17	

* MHH errors not available. See the appendix for details.

Taken together, Tables E-G, suggest that both the estimated coefficients on the dividend yield (and the earnings yield) and the R^2 increase with the return horizon. At the same time, the Hodrick t-statistics decline, thereby weakening the contribution of the dividend yield at the longer horizon. These results are not entirely consistent with those predicted by the simple model outlined in Section 2 (see also equation **(12)**). It is interesting to compare these results with the US regressions run by Ang and Bekaert (2003) for the period 1952-2001. In their univariate regressions, they fail to find significance for the dividend yield or the earnings yield. Furthermore, they observe that the estimated dividend (or earnings) yield coefficient decreases with the investment horizon until around four years and increases gradually thereafter. Finally, they find strong evidence of predictive ability of the short interest rate, both in univariate and in bivariate regressions.(17)

7. Stability checks

Our results so far suggest that long-horizon returns contain a predictable component. At the same time, we find that this predictability is sensitive to the chosen return horizon, and that it

(17) Ang and Bekaert (2003) also run regressions on monthly UK data for the period 1975-2001. They find no evidence of long-run predictability in the dividend yield model.

sometimes differs from that reported elsewhere in the literature. In this section, we present some evidence on the time variation in predictability, described in Section 1. This is done in two ways: first by estimating the predictability model over a fixed, but rolling sample period (Section 7.1), and second by employing an expanding sample (Section 7.2).

7.1 Sensitivity to sample period

We first wish to explain whether the significance of the yield variables and the R^2 of the regression are dependent on the period of estimation. We do so by estimating our models over a rolling window.[18] Chart 4 shows the t-statistics that result from estimating model 1 repeatedly over a 24-year rolling window.

Chart 4: Model 1 t-statistics from rolling regressions (24-year window)

4.A One-year returns

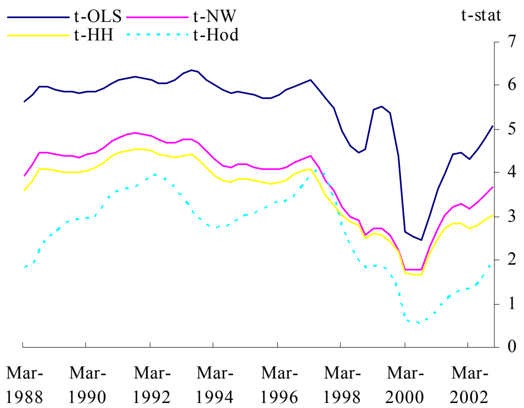

4.B Two-year returns

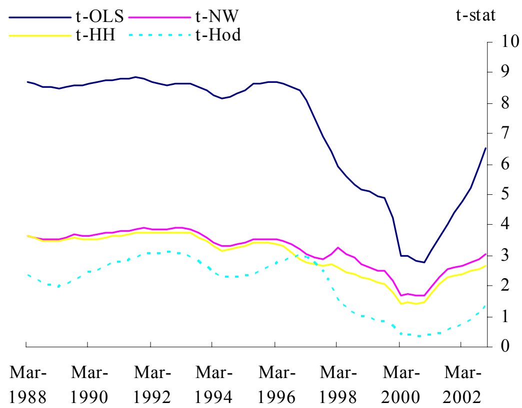

4.C Four-year returns

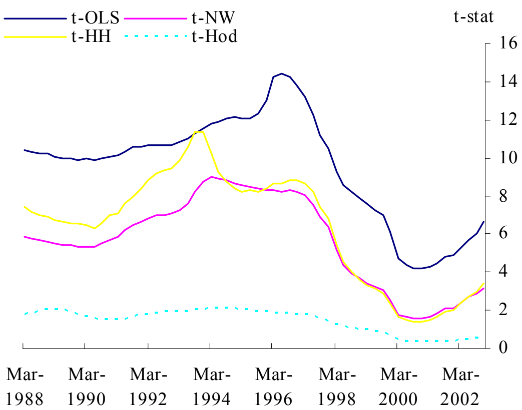

(18) In this section, we focus exclusively on our model 1 specification. Similar patterns were found when estimating other models recursively and over a rolling window.

Charts 4A and 4B suggest that predictability at the one-year and two-year horizon is highly sensitive to the sample period considered. Most strikingly, we observe that significance starts falling in late 1997, to reach an all-time low at the end of 1999. Thereafter, the t-statistics pick up, so that by the end of the sample period, they are back to the levels observed in late 1998. A second observation is that, whereas our earlier reported one-year model reported a Hodrick t-statistic of 2.23 (Table E), a model estimated over the period 1976:1 – 2000:1 yields a Hodrick t-statistic of just 0.91.[19]

Taken together, Charts 4A and 4B suggest that much of the loss in predictability occurred during a period of rapidly rising returns, during which our regression results yielded unusually large standard errors. As explained in the appendix, both the MHH and Hodrick t-statistics employ estimated standard errors. So it is easy to see that a sequence of very high errors will have a great impact on these test statistics. This effect can persist for a long time, even though the estimation window moves forward, because of autocorrelation in the residuals. This persistence explains why the rise in predictability after 2000 (seen in Charts 4A and 4B) is gradual. A similar line of argument may explain why predictability rises significantly between 1988 and 1992. The rise in the t-statistics shown in Charts 4A and 4B corresponds to the large swings in returns of the mid-1970s dropping out of the window.

The results for the rolling regressions may also help to explain why we lose significance of the yield variables in the four-year excess returns models. Chart 4C shows that predictability both falls and recovers at a slower rate than in Charts 4A and 4B. At the end of the sample period, the Hodrick t-statistics have only just regained their end-1999 values.

Charts 5 and 6 below show plots of the estimated dividend yield coefficient (β) and R^2 from the model 1 rolling regressions. Consistent with the results in Charts 4A-C, we see that the size of the coefficient and the fit of the regression fall sharply around 1997, reach bottom in late 2000 and recover slowly thereafter. Chart 5 may also shed light on Ang and Bekaert's result that the size of the coefficient on the dividend yield initially falls with the return horizon up to around two years, only to start rising again at longer horizons. Chart 5 suggests that this result is quite unique to the sample period they consider. The norm is for the size of the coefficient to increase with horizon.

[19] Similar results hold for the other horizons.

Chart 5: Model 1 β coefficients

(24-year rolling window)

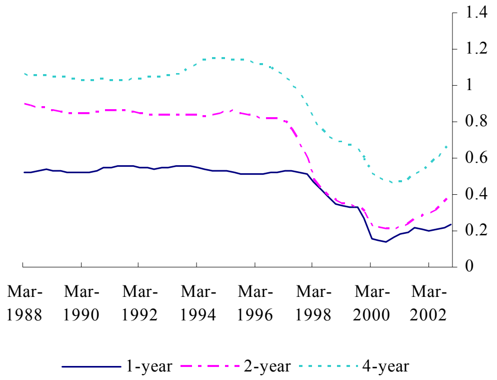

Chart 6: Model 1 R^2

(24-year rolling window)

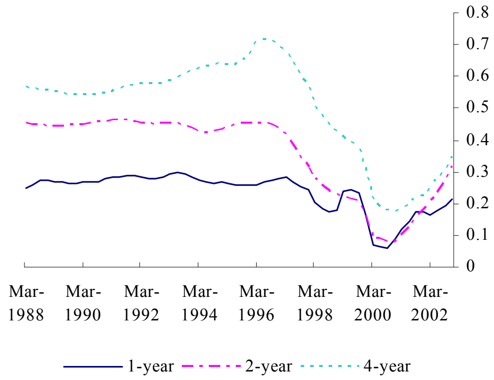

7.2 Sensitivity to sample size

The results so far illustrate substantial variation in predictability, most clearly in the 1990s. The question arises whether this time variation could have been caused by either the large run-up in equity prices through the late 1990s entering the estimation window, or the turbulence of the early 1970s dropping out, or both. To help determine which, we next estimate model 1 for varying sample sizes. Chart 7 shows the t-statistics from model 1 (dividend yield) at the one-year return horizon estimated recursively with the starting point held fixed at 1966:1. Charts 8 and 9 show the estimated coefficients and R^2 when using these same expanding windows.

Chart 7: Model 1 t-statistics from recursive regressions (expanding windows)

7A One-year returns

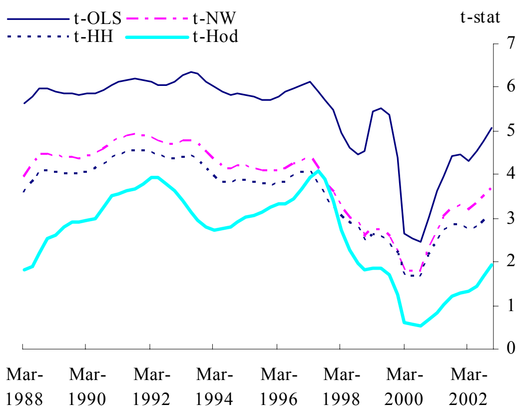

7B Two-year returns

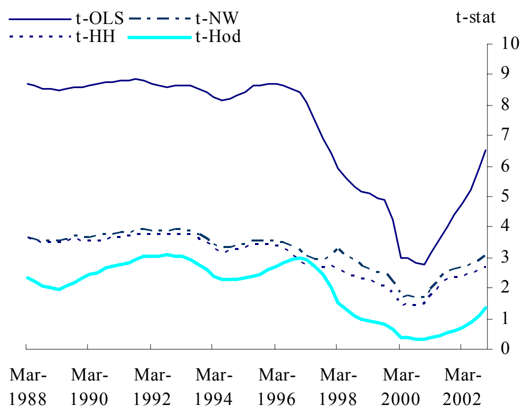

7C Four-year returns

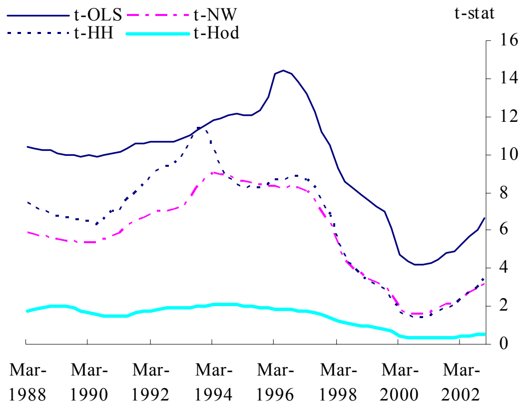

The t-statistics from recursive regressions are clearly less volatile than for the rolling regressions. In particular, the evidence based on the Hodrick t-statistic is remarkably stable. In the case of one-year returns, the Hodrick test statistic for the dividend yield hovers around 2 for most of the period. It lies below 1.96 between 1997 and the first half of 2001. It is only in the last five quarters of our sample that it recovers above 2. A similar story can be told for the two-year returns.

Chart 8: Model 1 β coefficients (expanding window)

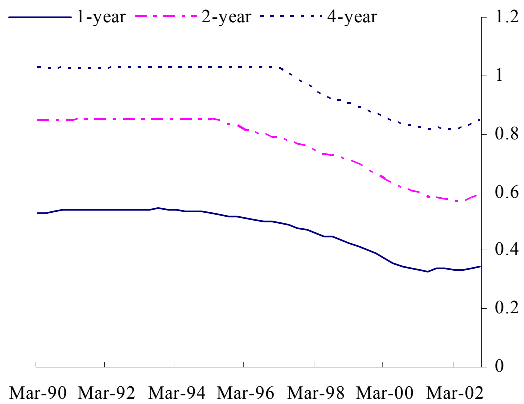

Chart 9: Model 1 R^2 (expanding window)

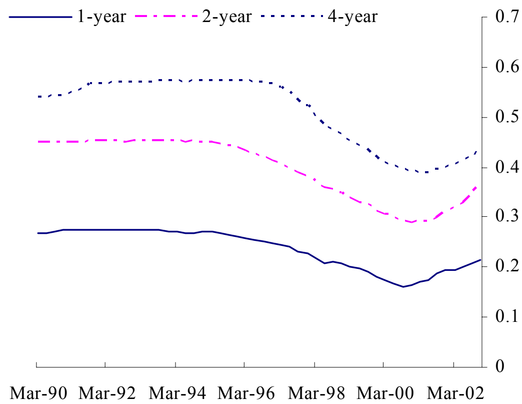

Comparing the rolling and recursive regressions, we can learn the following. First, the charts suggest that not including the less volatile period 1963-1973, when the dividend yield behaved more like it did in the 1980s (see Chart 1), leads one to conclude that the dividend yield has no predictive power for UK returns. Second, the run-up in equity prices through the 1990s and accompanying falls in the dividend yield affect the size of the coefficient and the fit of the regression; both have fallen gradually since around 1997. The subsequent decline in equity prices has led to a restoration of predictability, albeit at a slow rate. This becomes clear when one compares the rate of recovery of the t-statistics in Charts 4 (rolling windows) and 7 (expanding windows).

Taken together, the results of the rolling and the recursive regressions suggest that the relationship between the dividend yield and excess returns is highly sensitive to both the chosen return horizon and the sample period. This result is comparable to Ang and Bekaert (2003), who find evidence of return predictability only in samples that exclude the 1990s. Moreover, in these samples, the contribution of the dividend yield was dominated by that of the risk-free rate. This result does not appear to hold for the UK sample.[20]

8. Conclusions

In this paper, we estimate empirical models of equity return predictability for the United Kingdom that are motivated by the dynamic dividend discount model of Campbell and Shiller (1988a). We offer evidence that in the United Kingdom, long-horizon excess equity returns are predictable. In line with existing US studies, we confirm that both the dividend yield and the earnings yield help forecast such returns. At the same time, we highlight that these results could be model-dependent and that alternative models may dominate. We provide evidence that the size of the regression coefficients and R^2 of the regressions increase with the forecast horizon and that these results depend on the chosen sample period. We further demonstrate that the strength of the empirical dividend yield – equity return relationship depends crucially on the chosen sample period, with little apparent evidence of long-horizon equity predictability in the early 1970s and late 1990s. We do not confirm the predictive role of the risk-free rate reported in Ang and Bekaert (2003).

Ang and Bekaert (2003) conclude that the weak performance of the dividend yield model demonstrates the need for a different approach to the dynamic dividend discount model in empirical applications. In particular, they suggest that most predictability may stem from the risk-free rate, the second component of the time-varying discount rate in the theoretical model, and not from the expected return (for which the dividend yield is a proxy). As indicated in Section 2, some researchers are currently redirecting their focus towards the predictable component in cash flows (dividend and earnings growth).

Others consider the possibility of discrete shifts in the empirical dividend yield – equity return relationship. For example, research by Carlson, Pelz and Wohar (2001) argues that even if

[20] Results available from the authors.

the dividend yield is mean-reverting, its mean value may have changed over the past decades.(21) Using structural break point tests, they find evidence of a small number of breaks in the US dividend yield and earnings yield series, with the latest regime producing a mean well below the total sample mean.(22) In related work, Vila Wetherilt and Weeken (2002) discuss factors that might have contributed to a structural shift in the mean dividend (or earnings) yield, which in turn would have affected their relationship with expected returns.

Finally, much work is being done in the area of model selection, as indicated in Section 3. Also worth noting is work by Pesaran and Timmerman (2002), who demonstrate that equity return predictability is significantly improved when using time-varying parameters. Other papers suggest that equity returns might exhibit long memory (Henry and Zaffaroni (2001)). This means that even though they continue to exhibit mean reversion, unanticipated shocks have very long-lasting effects. Another interesting issue is that of rational bubbles. Modelling rational bubbles would require non-linear functions of fundamentals. The research presented in the present paper may highlight that the linear predictability model is not stable.

(21) Recall that we could not reject non-stationarity of the valuation ratios, a result, which may suggest a non-constant mean.
(22) They find strong evidence of breaks in the US dividend yield in 1955 and 1982 using annual data from 1872. They also find tentative evidence of a break in the dividend yield and price earnings ratio in 1992 Q4, using quarterly data from 1945.

Appendix: Long-horizon regressions

As discussed in the main text, estimating OLS regressions of long-horizon returns with a small data set provides an awkward dilemma for the researcher: either work with an even smaller sample of non-overlapping returns, or use overlapping returns which entail econometric complications.

Regressions involving overlapping returns can bias inference. This is because the data are sampled more finely than the return interval, so the estimated residuals will be serially correlated up to the order of return interval (as discussed in Hansen and Hodrick (1980)). This means that the OLS standard errors will be biased unless an appropriate adjustment is made.

Some authors overcome the problem by using non-overlapping observations (see, among others, Fama and French (1988)). This is not a viable option using our data set because a non-overlapping series of annual returns would leave just 36 observations. And, given that an aim of the current paper is to study the predictive power of data as the return horizon is increased we choose to work with 'overlapping residuals' and investigate appropriate adjustments to the standard error.

A further complication is that, in general, the volatility of asset returns appears to be serially correlated (see discussion in Campbell, Lo and MacKinlay (1997)). This makes the assumption of homoscedasticity inappropriate. Below, we describe three estimates of standard errors that allow for autocorrelation in the residuals and heteroscedasticity of unknown form.

Consider the generalised linear regression model with normally distributed errors

$$\begin{aligned} y_{t+k,k} &= x'_t\,\beta + \varepsilon_{t+k,k} \\ \varepsilon &\sim N(0,\sigma^2\Omega) \end{aligned} \qquad \textbf{(A1)}$$

where $y_{t+k,k}$ denotes the k-period excess return to $t+k$ and epsilon is an error term which belongs to the $t+k$ information set. x_t contains the candidate predictors so that in the simplest case (model 1), x_t' = [constant dividend yield]'.

Mapping the regression model into GMM provides an expression for the variance of the OLS estimate of beta, $\hat{\beta}_T$, which is

$$\text{var}(\hat{\beta}_T) = \frac{1}{T} E(x_t x_t')^{-1}.S.E(x_t x_t')^{-1} \quad \textbf{(A2)}$$

where

$$S = \sum_{j=-\infty}^{\infty} E(\varepsilon_{t+k,k} x_t x'_{t-j}\ \varepsilon_{t+k-j,k}) = \sum_{j=-\infty}^{\infty} E(w_{t+k} w'_{t+k-j}) \quad \textbf{(A3)}$$

Equation **(A3)** shows that we face the common problem of estimating *S* from a finite sample. Simply replacing the true autocovariances of w_{t+k} with their sample autocovariances can be problematic. To understand this, note that when $j = -(T-1)$ the estimate is based on just one observation and for $j < -(T-1)$ no estimate is available. Further, estimates based on sample autocovariances are, in general, not guaranteed to be positive definite. But tractable estimates of *S* are available. We follow Ang and Bekaert (2003) and consider three alternatives to OLS standard errors.

(Modified) Hansen-Hodrick errors

First, we consider standard errors similar to those proposed by Hansen and Hodrick (1980). As mentioned above, Hansen and Hodrick (1980) show that even under the null hypothesis that returns are not predictable, the error terms will be autocorrelated up to the order of the return horizon, *k*. Formally,

$$E(w_{t+k} w_{t+k-j}) \neq 0, \forall |j| < k$$
$$= 0, \text{ otherwise.}$$

Incorporating this information into **(A2)** and **(A3)** and using the estimated residuals provides a heteroscedasticty consistent version of the Hodrick and Hansen (1980) standard errors. We denote these Modified-Hansen-Hodrick (MHH) errors.

$$\text{var}(\hat{\beta}_T^{MHH}) = \frac{1}{T} E(x_t x_t')^{-1} \left[\sum_{j=-k}^{k} E(\varepsilon_{t+k,k} x_t x'_{t-j}\ \varepsilon_{t+k-j,k}) \right] E(x_t x_t')^{-1} \quad \textbf{(A4)}$$

There are two potential problems with this estimator. First, the null may not be correct. If so then any variation in expected returns not captured by changes in the explanatory variables may cause the error term to be autocorrelated beyond the return horizon considered. In this case, the MHH standard errors would be inconsistent, so we can only be sure of their validity under the null. Second, there is nothing that guarantees the estimate of *S* will be positive definite.

Newey-West estimator

One estimator that does guarantee a positive definite estimate of S is the familiar Newey-West (NW) estimator. Unlike the MHH estimator, the NW estimator downweights higher order autocorrelations. The NW standard errors are given by

$$\mathrm{var}(\hat{\beta}_T^{NW}) = \frac{1}{T} E(x_t x_t')^{-1} \left[\sum_{j=-k}^{k} \left(\frac{q-|j|}{q} \right) E(\varepsilon_{t+k,k} x_t x'_{t-j} \, \varepsilon_{t+k-j,k}) \right] E(x_t x_t')^{-1} \qquad \textbf{(A5)}$$

As in the MHH case, setting $q = (k+1)$ would be sufficient to remove the error autocorrelation under then null hypothesis. Although this estimator is guaranteed positive definite, when using it we run the risk of excessively downweighting autocorrelations which we know to be non-zero.

Hodrick errors

Another estimator that guarantees positive definiteness of the estimated S was proposed by Hodrick (1992). Under the null hypothesis, the k-horizon error is simply the sum of k one-period errors,

$$\varepsilon_{t+k,k} = (u_{t+1} + \ldots + u_{t+k})$$

where u_{t+1} is the (serially uncorrelated) one-step-ahead forecast error. Hodrick (1992) shows that, by substituting this into **(A3)** and maintaining the assumption that autocovariances of a higher order then the return horizon k are zero, gives

$$S = \sum_{j=-k}^{k} E\left[\left(\sum_{i=1}^{k-j} u_{t+i}^2 \right) x_t x'_{t-j} \right] \qquad \textbf{(A6)}$$

Notice that **(A6)** contains the expectation of two stationary series. Since the expectation between two (covariance) stationary time series depends not on the particular points in time, but only the lag between them, we can rewrite the expectation in **(A6)** as

$$E\left[\left(\sum_{i=1}^{k-j} u_{t+i}^2 \right) x_t x'_{t-j} \right] = E\left[u_{t+1}^2 \left(\sum_{i=0}^{k-j-1} x_t x'_{t-j-i} \right) \right] \qquad \textbf{(A7)}$$

The right-hand side of **(A7)** sums the regressors into the past instead of summing the error autocovariances forward, but keeps the interval between the two terms unchanged. Applying this logic to each expectation in **(A6),** we get

$$S = E\left[u_{t+1}^2\left(\sum_{i=0}^{k-1} x_{t-i}\right)\left(\sum_{i=0}^{k-1} x_{t-i}\right)\right]$$

where u_t can be estimated by a simple regression of the one-period return on a constant. Finally, we get an expression for the Hodrick errors.

$$\mathrm{var}(\hat{\beta}_T^H) = \frac{1}{T} E(x_t x_t')^{-1}\left[\sum_{j=k}^{T} wk_t wk'_t\right] E(x_t x_t')^{-1} \qquad \textbf{(A8)}$$

where

$$wk_t = u_{t+1} \sum_{i=0}^{k-1} x_{t-i}$$

This method also guarantees positive definiteness of the estimated *S*, and avoids the summation of autocorrelation matrices (as in **(A4)** and **(A5)**). Hodrick (1992) suggests that the latter property is potentially important, as in small samples this summation leads to poor properties of standard errors. Hodrick (1992) and Ang and Bekaert (2003) find that in Monte Carlo simulation, test statistics based on the Hodrick errors have the best *size*. That is, they produce a test with a lower probability of rejecting a true null hypothesis, relative to tests based on NW or MHH errors. This result is intuitive. The Hodrick and MHH errors are only valid under the null hypothesis, but the Hodrick errors impose more information from the null hypothesis. One might expect that imposing more about the null would reduce the probability of rejecting it, if it is in fact true.(23)

(23) See discussion in chapter 11 of Cochrane (2001).

References

Ackert and Hunter (2000), ‘An empirical examination of the price-dividend relation with dividend management’, Federal Reserve Bank of Chicago, *Working Paper no. 00-22*.

Ang and Bekaert (2003), ‘Stock return predictability: is it there’, *mimeo*, Columbia University, July.

Campbell (1991), ‘A variance decomposition for stock returns’, *Economic Journal*, Vol. 101, pages 157-79.

Campbell (2001), ‘Valuation ratios and the long-run stock market outlook: an update’, *NBER Working Paper no. 8221*, April.

Campbell and Cochrane (1999), ‘By force of habit: a consumption based explanation of aggregate stock market behaviour’, *Journal of Political Economy,* pages 205-51.

Campbell, Lo and MacKinlay (1997), *The econometrics of financial markets*, Princeton University Press.

Campbell and Shiller (1988a), ‘Stock prices, earnings and expected dividends’, *Journal of Finance*, Vol. 43, pages 661-76.

Campbell and Shiller (1988b), ‘The dividend-price ratio and expectations of future dividends and discounts factors’, *Review of Financial Studies*, Vol. 1, pages 195-228.

Campbell and Shiller (2001), ‘Valuation ratios and the long-run stock market outlook’, *NBER Working Paper no. 8221*, April.

Carlson, Pelz and Wohar (2001), ‘Will the valuation ratios revert to their historical means? Some evidence from breakpoint tests’, Federal Reserve Bank of Cleveland, *Working Paper no. 01-13*.

Cochrane (2001), *Asset pricing*, Princeton University Press.

Fama and French (1988), ‘Dividend yields and expected stock returns’, *Journal of Financial Economics*, Vol. 22, pages 3-25.

Fama and French (1989), ‘Business conditions and expected returns on stocks and bonds’, *Journal of Financial Economics*, Vol. 25, pages 23-49.

Ferson, Simin and Sarkissian (2003), ‘Spurious regressions in financial economics?’, *Journal of Finance*, Vol. 58, pages 1,393-414.

Froot and Obstfeld (1991), ‘Intrinsic bubbles: the case of stock prices’, *American Economic Review*, Vol. 81, pages 189-214.

Goyal and Welch (2002), ‘Predicting the equity premium with dividend ratios’, *NBER Working Paper no. 8788.*

Granger and Newbold (1974), ‘Spurious regressions in economics’, *Journal of Econometrics*, Vol. 2, pages 111-20.

Hansen and Hodrick (1980), 'Forward exchange rates as optimal predictors of future spot rates: an econometric analysis', *Journal of Political Economy*, Vol. 88, pages 829-53.

Hendry and Krolzig (1999), 'Improving on 'data mining reconsidered' by K.D. Hoover and S.J. Perez', *Econometrics Journal*, Vol. 2, pages 202-19.

Henry and Zaffaroni (2001), 'The long range dependence paradigm for macroeconomics and finance', *mimeo*.

Hodrick (1992), 'Dividend yields and expected stock returns: alternative procedures for inference and measurement', *Review of Financial Studies*, Vol. 5, pages 357-86.

Hoover and Perez (1999), 'Data mining reconsidered: encompassing and the general-to-specific approach to specification search', *Econometrics Journal*, Vol. 2, pages 167-91.

Lamont (1998), 'Earnings and expected returns', *Journal of Finance*, Vol. 53, pages 1,563-87.

Lettau and Ludvigson (2001), 'Consumption, aggregate wealth, and expected stock returns', *Journal of Finance*, Vol. 56, pages 815-49.

Lettau and Ludvigson (2003), 'Expected returns and expected dividend growth', *NBER Working Paper no. 9605*.

Menzly, Santos and Veronesi (2004), 'Understanding predictability', *Journal of Political Economy*, Vol. 112, pages 1-47.

Neeley and Weller (2000), 'Predictability in international asset returns: a re-examination', *Journal of Financial and Quantitative Analysis*, Vol. 35, pages 601-20.

Nelson (1999), 'The aggregate change in shares and the level of stock prices', Federal Reserve Board, *Working Paper no. 1999-08*.

Pesaran and Timmerman (1995), 'Predictability of stock returns: robustness and economic significance', *Journal of Finance*, Vol. 50, pages 1,201-08.

Pesaran and Timmerman (2002). 'Market timing and return prediction under model instability', *Journal of Empirical Finance*, Vol. 9, pages 495-510.

Sharpe (2002), 'Reexamining stock valuation and inflation: the implication of analysts' earnings forecasts', *Review of Economics and Statistics*, Vol. 84, pages 632-48.

Stambaugh (1999), 'Predictive regressions', *Journal of Financial Economics*, Vol. 54, pages 375-421.

Vila Wetherilt and Weeken (2002), 'Equity valuation measures: what can they tell us?', *Bank of England Quarterly Bulletin*, Winter, pages 391-403.

Wadhwani, S (1999), 'The US stock market and the global economic crisis', *National Institute Economic Review*, pages 86-105.