Forecasting the Market

Elroy Dimson, Paul Marsh, and Mike Staunton
London Business School, Regents Park, London NW14SA, UK

Abstract: *Within markets, there has been a tendency for higher-yielding stocks to enjoy greater returns. In this paper we ask whether this type of relationship extends to markets as a whole, across countries and over time. We investigate whether buying equities when the aggregate dividend yield of the market is high, and selling when yields are low, gives rise to superior long-run returns. In order to judge the informativeness of aggregate dividends about subsequent stock market returns, we study real annual returns in 16 different equity markets over the last 104 years. We find that dividend-based forecasts of the market underperform naïve extrapolation of returns in three-quarters of the 16 markets.*

Within markets, there has been a tendency for higher yielding stocks to enjoy greater subsequent returns. In this paper we ask whether this type of relationship extends to markets as a whole, across countries and over time. We examine whether buying equities when the aggregate dividend yield of the market is high, and selling when yields are low, gives rise to superior long-run returns. Since the earliest days of security analysis, investors have been aware of the potential benefits from buying when stock prices are low relative to fundamentals, and selling when prices are high. Dividends are probably the oldest yardstick. Dow (1920) stressed the relevance of the aggregate dividend yield on the overall market, and Graham and Dodd (1934) popularised an investment focus on fundamental value. In more recent—and widely cited—research, Fama and French (1988) report that the aggregate dividend yield of the US market has been a reliable predictor of subsequent stock returns over intervals of one to four years ahead. Other research on this topic includes Rozeff (1984) and Campbell and Shiller (1988), and the work described in Cochrane (1997) and Campbell and Viceira (2002).

These intriguing findings have been re-examined by researchers and practitioners. At least one stock market—the United Kingdom—has exhibited an even stronger relationship between dividend yields and subsequent stock market returns than is the case in the United States. Goetzmann and Jorion (1995) attribute this to the major reversal from the 1974 collapse of the UK equity market to the 1975 recovery. On the other hand, a number of researchers find that the dividend-return relationship documented by Fama and French is specific to their 1927-86 sample period, and is not robust to the choice of research design; see, for example, Goetzmann and Jorion (1995), Cremers (2002) and Goyal and Welch (2003). The objective of this paper is to take a new look at market predictability around the world. Apart from some studies that have looked at non-US/UK markets over intervals of under 25 years, notably Ang and Bekaert (2001), we are unaware of any long-term evidence on predicting international returns from aggregate dividends.

To examine market predictability globally, we interrogate the 1664 annual equity returns for the 16 markets in the database that underpins the research in Dimson, Marsh, and Staunton (2002, 2003, 2004). We focus throughout on inflation-adjusted total returns. Our objective is to investigate the informativeness of aggregate dividend yields about subsequent stock market returns. In the next section we review the link between dividend yields and subsequent market returns over time and across countries. The subsequent section examines the performance of the equity market following years with an extreme dividend yield. We then ask when dividend-based models can be expected to perform well as predictors of subsequent equity returns. This informs development of a dividend model, the performance of which is then tested on an out-of-sample basis. We find that dividend-based forecasts of the market underperform naïve extrapolation of equity returns in three-quarters of the markets we study. We end the paper with a brief summary and conclusion.

Dividend yields and stock returns

The modern literature on the linkage between dividend yields and expected stock returns dates back to the 1980s. The intuition is that the level of the stock market is low when expected returns are high; similarly, the level of the stock market is high when expected returns are low. We measure whether the market is low or high by examining its level relative to fundamental value—in this case, relative to dividends. The dividend-to-price ratio, hereafter referred to as the dividend yield, is therefore an indicator of the rate at which subsequent cash flows are discounted to the present. Many researchers share the view that dividend yields impound a forecast of expected returns. That is, even in a completely efficient market, dividend yields should have the power to predict subsequent returns.

This view is sometimes referred to as the discount-rate effect: holding cash flows constant, anything that increases the rate at which cash flows are discounted (such as changes in liquidity or risk) lowers market value and increases both expected returns and financial ratios such as the dividend yield. Fama and French (1988) test this hypothesis using data on aggregate dividend yields for the US stock market for the period 1927-1986. They examine dividend-based return forecasts over investment horizons of one to four years (and also over one and three months, but with little success in developing a return-forecasting model). They confirm that dividend yields forecast subsequent market returns over a horizon of a year or more; and typically, the relationship has statistical significance, as is apparent in the table below which is extracted from the Fama-French study.

Table 1 : Regression of annualised real equity returns on prior dividend yield

Period	Return horizon	Observations	Slope	*t*-value	R-squared %
1927-86	1 year	60	3.35	1.72	3
	2 years	59	8.77	2.59	15
	3 years	58	11.53	2.93	21
	4 years	57	14.43	3.25	31
1941-86	1 year	46	4.82	2.38	16
	2 years	45	10.26	3.15	21
	3 years	44	13.10	3.53	21
	4 years	43	15.71	3.31	22

Source: Fama and French (1988)

In Table 1, the slope coefficient measures the sensitivity of returns to yields. For example, if the slope is 3.35, then an additional 10% yield would give rise to an additional 0.335 percent annualised return (i.e., 3.35 multiplied by 10%). The *t*-values indicate statistical significance of most coefficients at conventional levels (a value close to or above 2). The economic and statistical significance of dividend yields, as a predictor of annualised returns, is at least as strong over longer intervals as over single years. The strength of the relationship, indicated by the R-squared of the regression, increases with the forecasting horizon. Furthermore, more recent intervals such as 1941-86 (see table) or 1957-86 (not shown) revealed a strengthened relationship. The yield-return relationship was evident whether the dividend yield was defined as the ratio of annual dividend to the recent or lagged (twelve-month-old) index level; whether returns were in nominal or real terms; and whether the index was value- or equally-weighted. The observation that dividend yields appear to predict subsequent stock returns is said by Goyal and Welch (2003) to "*rank among the most important findings of academic finance*" (page 639). This phenomenon underpins a variety of approaches to dynamic asset allocation, such as those described in Campbell and Viceira (2002).

The robustness of the dividend-return relationship has been questioned because of its reliance on long-horizon regressions using overlapping data over a small number of decades. Some of the criticisms are: [1] the dangers of in-sample validation (Goetzmann and Jorion, 1993; Nelson and Kim, 1993); [2] biased coefficient estimates when the predictor is highly persistent (Stambaugh, 1999); [3] inappropriate long-horizon statistical testing and the inadequacy of R-squared (Valkanov, 2003); [4] the impact of an extreme outlier (Kothari and Shanken, 1997); and [5] uncertainty about coefficient magnitude (Cremers, 2002). Even more fundamentally, there is a recurrent concern about the impact of decades of data-mining by researchers examining the same information for the same country over the same time period. As Lo and MacKinlay (1990) and Bossaerts and Hillion (1999) observe, it is to be expected that repeated searching comes up with a model for return that works well historically, since we usually have prior knowledge of what is likely to perform convincingly, even in supposedly out-of-sample tests.

The literature is split between two schools of thought. On the one hand, some consider the association between dividend yields and expected returns to be a consequence of the discount-rate effect. For them, the Fama-French findings summarised above are consistent with market efficiency. The alternative is a behavioural interpretation: investors tend to overreact and to bid prices up to an excessive level or down to too low a level. From this (inefficient markets) point of view, ratios such as the dividend yield can reveal intervals of non-rational undervaluation and overvaluation. For example, when stocks sell for an irrationally low price, yields will be high and expected returns will be large. In time, the mispricing will be corrected, and investors can expect to earn superior returns as the stock market reverts toward its equilibrium level. Barberis and Thaler (2003) provide a recent guide to the behavioural literature.

Time-varying expected returns

If the behavioural view—that market prices are bid up or dragged down to inappropriate levels—is valid, then we might construct an even simpler model of time-varying expected returns. In particular, we could look solely at the sequence of annual returns, to examine whether prices get driven too high or low. Chart 1 presents a graphical overview of evidence drawn from our annual database.

Chart 1 : Annual returns in the UK (top) and US (bottom) stock markets, 1900–2003

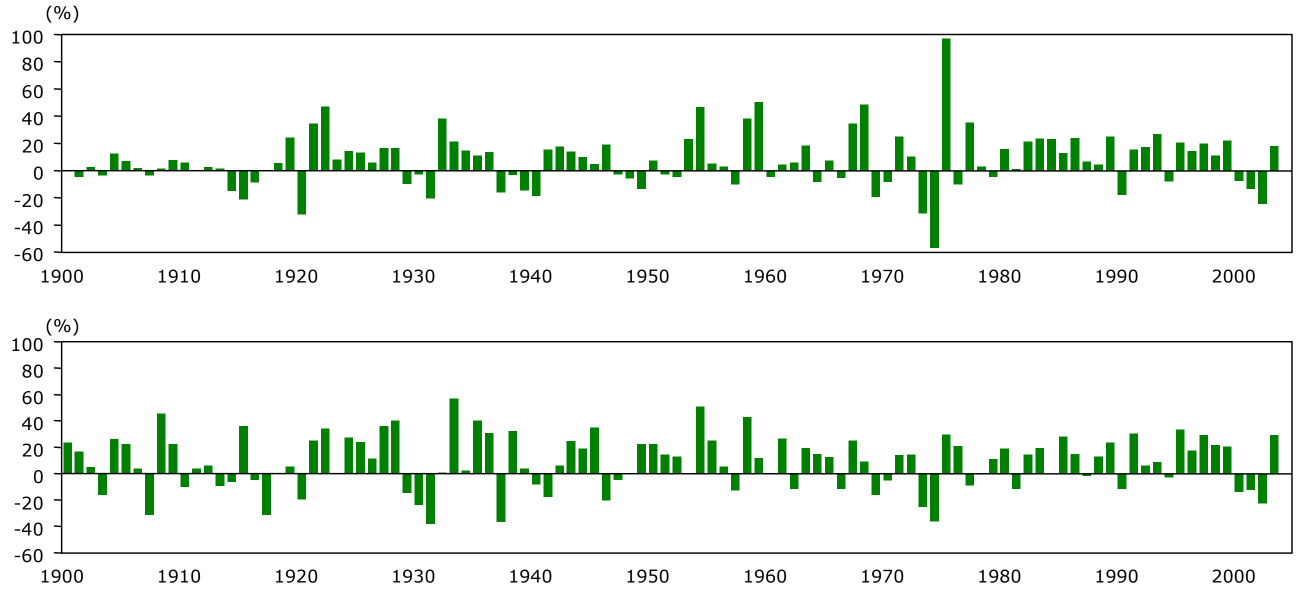

Chart 1 shows the most recent 104 annual returns on the UK and US equity markets. If prices deviate from fundamentals, we should expect to find a pattern of reversals in the sequence of returns depicted in the diagram. There should be a tendency for markedly favourable returns to be followed by a negative reaction, and for markedly unfavourable returns to be followed by a positive recovery. We therefore search for predictable patterns of behaviour in the return series portrayed above, analysing data over intervals of single or multiple years. There is no evidence of predictability in Chart 1. The autocorrelation of continuously-compounded, annual real returns for these two countries is -0.03 and +0.02 respectively. In addition, we have estimated the autocorrelation of the biannual real returns (commencing in 1900 or 1901) and of the triannual real returns (commencing in 1900 or 1901 or 1902), none of which is significantly different from zero. There is no sign of significant reversals in annual, biannual or triannual returns.

Table 2 reports the autocorrelations of the 16-country world index, the 16 individual countries, and the average of the individual countries. We also calculate autocorrelations using investment horizons of one, two and three years; and based on all data or on post-1950 data. The autocorrelation coefficients are on average equal to +0.08, -0.03 and -0.08 for the full period, and +0.02, -0.09 and -0.07 for the more recent interval. Coefficients that are significant at the 95 percent level are typeset in bold; only five out of 96 coefficients are significant at the 95 percent level—exactly the proportion that would be expected by chance alone. Even if they were highly stable—and the reality, of course, is that these coefficients are notoriously difficult to predict out of sample—such levels of autocorrelation would defy profitable investment.

Many researchers have therefore focussed on return-forecasting models that consider stock prices relative to fundamentals. Dividend yield is not the only financial ratio, and alternative proxies for fundamental value include the book-to-market ratio, Q ratio, and earnings-to-price ratio. Dividends, however, are the most widely used yardstick at the aggregate market level, perhaps because they are more readily compared across national markets than valuation ratios that involve accounting conventions in estimating book values or earnings. In the next section we turn to dividend yields and their time series properties.

Table 2: Serial correlation of stock market returns, 1900-2003

Country	**Returns of varying length, 1900-2003**			**Returns of varying length, 1950-2003**		
	1 year	2 years*	3 years*	1 year	2 years*	3 years*
World	0.14	-0.07	-0.13	0.06	-0.08	0.08
UK	-0.03	-0.10	-0.09	-0.11	-0.22	-0.13
US	0.02	-0.20	-0.18	-0.02	-0.13	0.13
Japan	-0.21	-0.11	-0.01	0.06	-0.10	-0.09
Germany	**0.34**	0.13	0.04	0.05	-0.01	0.19
Netherlands	0.10	0.04	0.07	0.21	-0.02	-0.01
France	0.16	0.10	0.00	-0.02	0.02	-0.11
Italy	0.04	-0.06	-0.10	**0.30**	0.00	-0.08
Switzerland	0.20	-0.02	-0.07	0.10	-0.14	-0.24
Australia	-0.01	-0.11	-0.07	-0.09	-0.17	-0.13
Canada	0.16	-0.11	-0.27	-0.07	-0.28	-0.09
Sweden	0.12	-0.04	-0.11	0.00	-0.16	-0.12
Denmark	-0.15	-0.13	-0.21	-0.26	-0.21	-0.24
Spain	**0.34**	0.22	0.03	**0.35**	0.31	0.10
Belgium	**0.22**	0.01	-0.01	0.03	0.05	0.05
Ireland	-0.06	0.01	0.00	-0.14	-0.11	-0.14
South Africa	0.00	-0.18	-0.24	-0.05	-0.31	-0.24
Average	0.08	-0.03	-0.08	0.02	-0.09	-0.07

** For each country, we show the mean of 2 or of 3 coefficients based on non-overlapping data starting in 1900, 1901 and 1902. Standard errors for 1-, 2- and 3-year returns are 0.10, 0.14 and 0.17 over 104 years; and 0.14, 0.20 and 0.24 over 54 year. The data source is Dimson, Marsh, and Staunton (2004).*

Dividend yields around the world

If dividend yields are a proxy for fundamental values, it is important to look at the time series of dividend yields. Two features of the series are important. First, dividend yields are close to stationary; and second, they have had a slight downward trend over the very long term. This is illustrated in Chart 2, which shows the dividend yield for UK and US equities from 1900 to the beginning of 2004. In the UK, the dividend yield has fluctuated in the range 2.4–7.9 percent (and a somewhat wider band in the US). Dividend yields declined over the total period spanned by our database, but there was considerable variation over the 104 years, with sub-periods that coincided with transitory rises in yields. Examination of the entire time series reveals that successive yields are close to a random walk. The level of the yield has a high serial correlation (0.71 for the UK, 0.87 for the US), so that dividend yield levels predict subsequent yield levels with considerable accuracy. Yield changes are essentially unpredictable; this is similar to stock market returns, where changes in levels are not serially correlated, and cannot be used to predict subsequent market movements.

Chart 2 : Dividend yields in the US and UK, 1900–2003

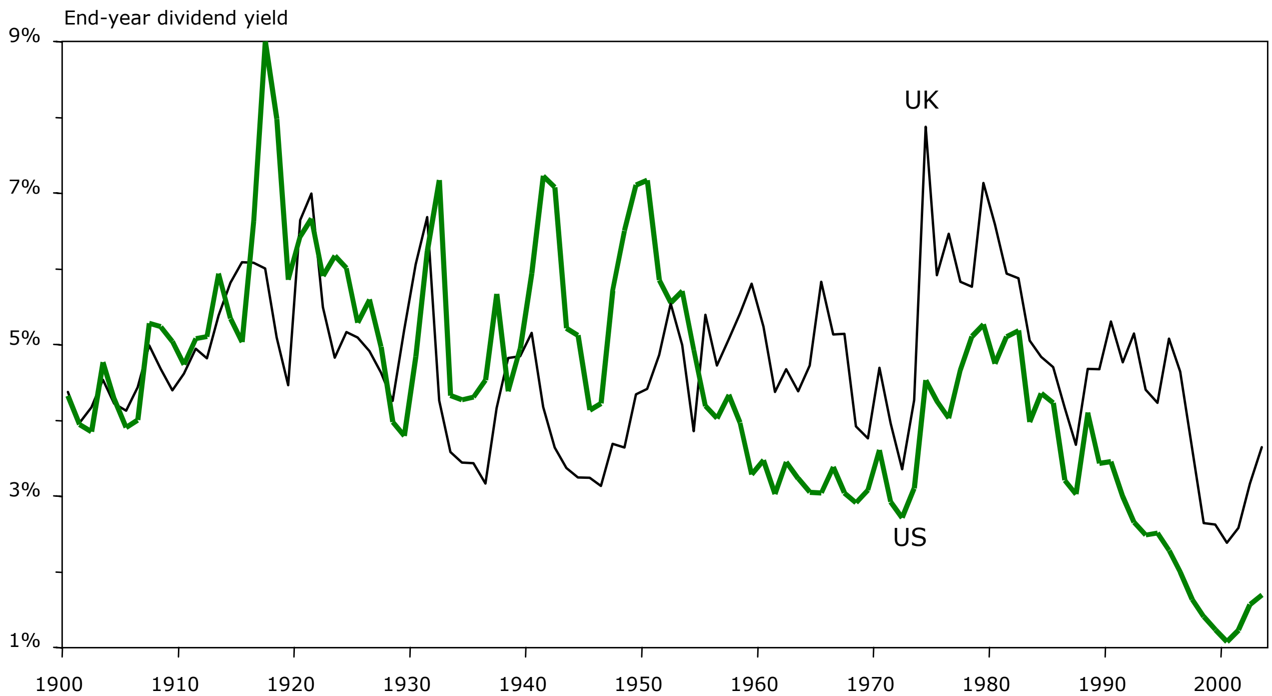

Chart 3 displays the level of dividend yields in each country at (the end of) 1900, and at the start of 1950 and 2004. The yields are in all cases calculated as the aggregate dividends paid during the year, divided by the end-of-year index level. Yields have varied greatly over time. US equities gave a yield of 4.3 percent in 1900, 7.2 percent in 1950, and (after a low of 1.1 percent in 2000) 1.7 percent at the beginning of 2004. Although Japan's current yield of 1.2 percent is the lowest of all sixteen markets, in both 1900 and 1950 Japanese equities had a dividend yield that ranked among the highest two out of the sixteen countries. Over the twentieth century, Danish yields fell from being the highest, at 7 percent, to one of the lowest, at today's 2 percent.

Beside changing over time, yields also vary across countries. Despite greater integration of world equity markets, the cross-sectional dispersion of yields at the start of 2004 remains very wide. Yields ranged from 4.3 percent in Spain, down to 3.6 percent in the UK, 1.7 percent in the US, and just 1.2 percent in Japan. There are many reasons for this dispersion, including national approaches to dividend payouts, differing tax regimes, and cross-country differences in growth opportunities. In addition, in any particular year countries will be at varying stages in their economic cycles, and this may be associated with yield differentials.

Across all sixteen countries, the equally weighted average dividend yield was around 5 percent in 1900, and this was still the case in 1950. By end-2000, the average yield had fallen to 1.9 percent, while the yield on the capitalization-weighted 16-country world equity index was even lower at 1.4 percent. At the start of 2004, the average was 2.8 percent, and the yield on the world index was 2.0 percent. The long-term fall in yields may be attributed to a decline in the discount rate and/or to an increase in equity growth projections, plus the impact of a reduced average payout ratio.

It is important that these secular trends in yield are incorporated within a dividend-return forecasting model. If the model fails to take account of gradual changes over time in levels of yield, there is the danger that the model is uninformative about contemporary markets. Suppose, for example, the model indicates that dividend yields below "X" percent are a signal of potentially lower expected returns. In a scenario where the market has consistently traded at a yield well below X percent for very many years, then the model will be unable to indicate when prospective returns are low or high. Most previous research has ignored long-term trends in financial ratios. For this reason, we develop a model that seeks to take account of the differing yields across markets and the changing levels of yield over time.

Chart 3 : Dividend yields around the world: 1900, 1950 and 2004

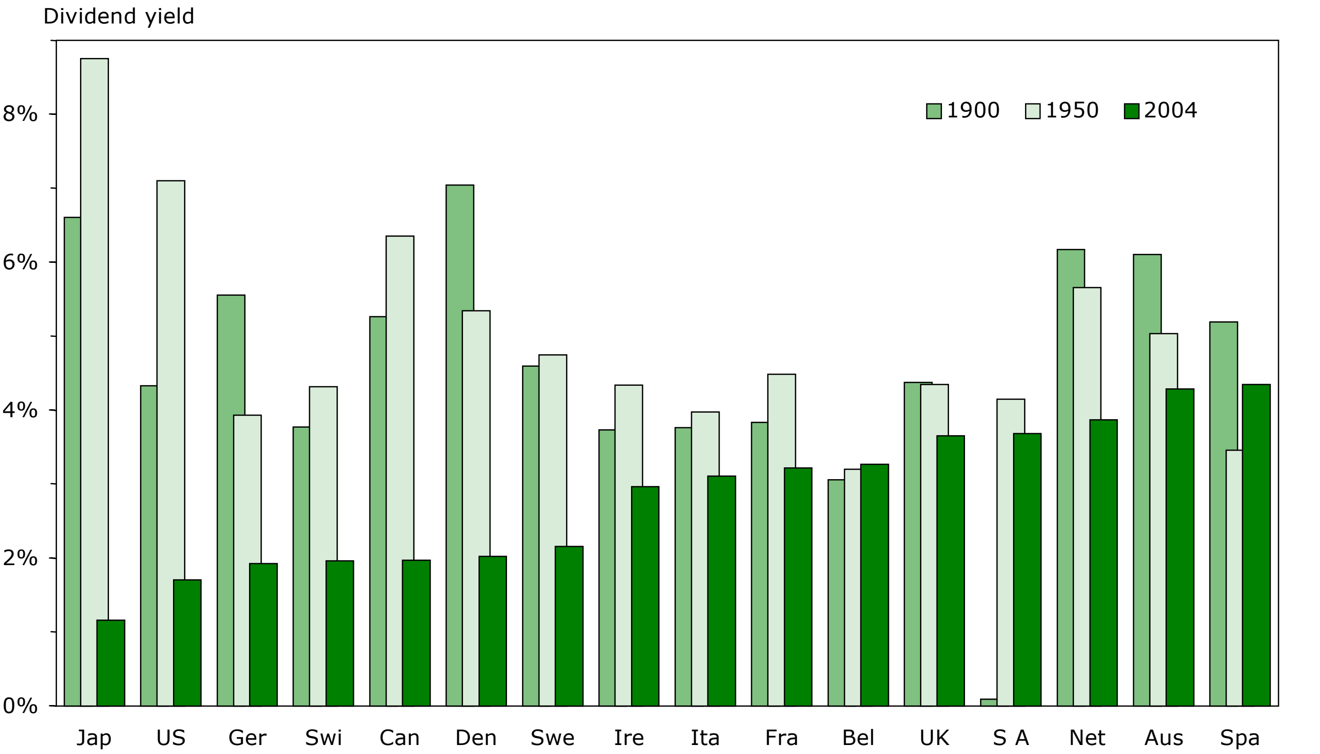

Performance after an extreme yield

At a time of low dividend yields, we must be concerned with the level of performance that can reasonably be anticipated for the future. When stock prices in most countries are high relative to fundamentals, should we project future returns that fall below the historical average? Campbell and Shiller (2004) propose an analogy with climbing: "*When one is mountaineering, one can enjoy the exhilarating view from high up on a mountain, and may look forward to the possibility of discovering a way up to a much higher level. But one will reflect that, realistically, at a random date years from now, one will probably be down at ground level*". It sounds plausible that valuation ratios revert toward a long-term mean, so there is a prima facie case for looking at the long-term record of dividend yields as a predictor of subsequent stock returns.

Sensible mountain climbers embark with knowledge of the height of their mountain, and a means of judging their elevation. An issue raised by Campbell and Shiller's analogy is that investors do not know the maximum level for a valuation measure. They cannot tell where they are located between the maximal and minimal values for a financial ratio. With foresight about the range of possible dividend yields, investors would know when yields are at their peak, and would know that at some future date yields would indeed be lower. That knowledge could be profitable. Though the strategy implicitly requires foresight about the future, its potential profitability is still revealing about the impact of dividend yields on returns. This idea is illustrated in Chart 4.

In this chart, we estimate dividend yields for the UK annually from 1900 to 1999, and identify the ten year-ends with the highest yield; the ten year-ends with the next highest yield; and so on to the ten lowest yields of all. We then calculate the one-year and four-year holding-period return subsequent to the date of each yield. We find that after the highest dividend yields, one- and four-year real returns were on average stunning. Conversely, after the lowest yields, real returns were on average negative. In many other countries, there are also dramatic differences between buying when yields are high and when they are low.

Some authors interpret evidence like Chart 4 (though presented at a more sophisticated level) as indicating large predictable return differentials. This interpretation is invalid. To persist with the Campbell-Shiller analogy, investors are in the position of individuals dropped by helicopter onto a cloudy mountain. They are unaware of the height of the mountain and of their elevation; they cannot judge whether they are at a high or a low level. They do not know whether dividend yields are destined to fall lower (which might lead to today's low yield eventually being categorised into, say, group 5 in the previous diagram); or whether today's yield is at a level that is destined to be an all-time low (which would indicate poor expected returns, as on the right-hand side of the previous chart).

Chart 4 : UK equity real returns after dates with dividend yields ranked from high to low, 1900-2003

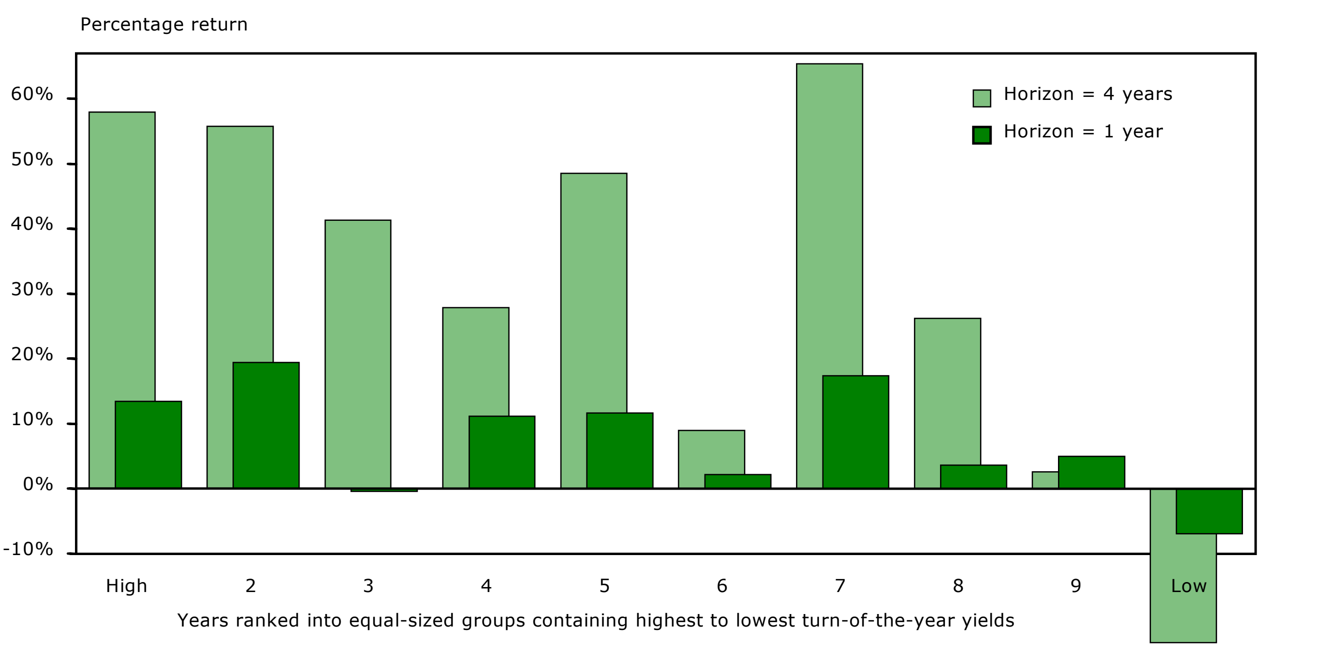

We therefore conduct a more careful experiment, in which there is no look-ahead bias. At each year-end, we designate the dividend yield as falling in the highest to the lowest of ten groups, estimated from the historical distribution of dividend yields *prior to* that date. For instance, in the UK the first eight end-year dividend yields of the twentieth century were 4.37, 3.96, 4.17, 4.54, 4.22, 4.13, 4.44 and 4.99 percent respectively. The historical distribution had included a yield that was above 4.54 but no more than 4.99 percent; there had never been a twentieth-century yield of 5 percent or more. The end-1908 dividend yield was actually 4.68 percent, so the 1908 yield fell in the one-from-highest yield group.

This generates varying numbers of observations in each group. There were 14 yields classified as group 1 (highest yields); 16 in group 2; then 9, 7, 12, 6, 7, 4 and 11 in groups 3 to 9 respectively; and 17 in group 10 (lowest yields). We calculate the average return from buying equities after yields were high (groups 1–3), and the average return from buying equities after yields were low (groups 8–10). Chart 5 reports the average annualised return from buying after yields were high, and from buying after yields were low, estimated over holding periods of 1 to 15 years.

In years after yields were high, returns greatly exceeded what was achieved in years after yields were low. Holding-period returns over long horizons diverged. Further, although annualised returns converged over time, holding-period returns were larger over longer intervals. The average annual return following a date when the market's yield is high is 13.2 percent, compared with the return after a low yield of 2.3 percent. The average biannual return following a high yield was an annualised 10.5 percent, compared to an annualised 0.4 percent after a low yield. Over five years, the corresponding returns were an annualised 8.5 percent and 0.6 percent respectively.

We emphasise market-wide return differences following high-yield and low-yield dates. This is quite different from the cross-sectional concept of an HML ("High minus Low") factor. First, HML typically refers to the difference in contemporaneous returns between portfolios with a high or low book-to-market ratio, whereas we focus here on portfolios that differ in their dividend-to-price ratios (Dimson, Nagel and Quigley, 2003, use the term IMC—income minus capital-gain—to refer to return differences based on dividend yield). Second, and this is crucial, in contrast to the cross-sectional interpretation of HML, the returns we consider are calculated over intervals that are dispersed through time. Both return differences may, however, be regarded as a reward for risk, in either a time-series or a cross-sectional context.

Chart 5 : UK real returns after high and after low dividend yields, over horizons of 1–15 years, 1900–2003

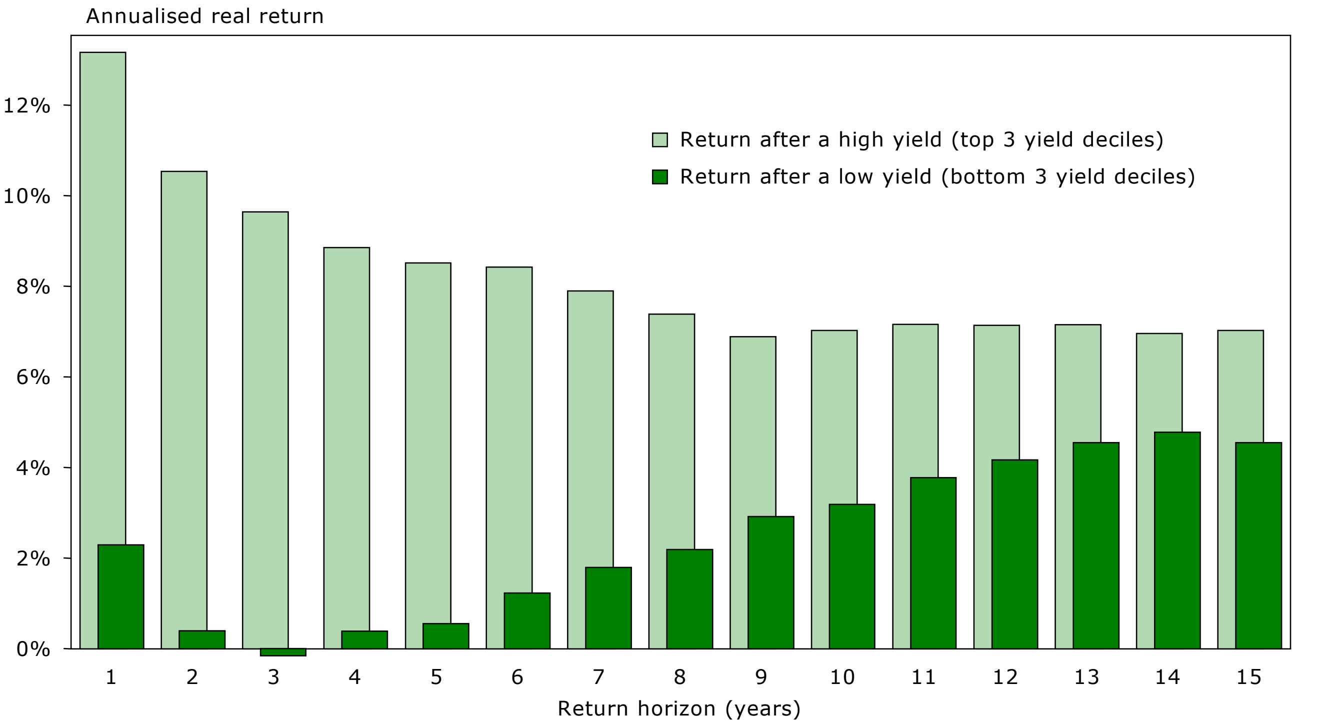

In Chart 5 the return differences following a high or low yield in the UK (namely, the gap between the two bars for each return horizon) are 11 percent over one year; 10 percent over two years; 8 percent over five years; and 4 percent over 10 years. The return difference between buying when yields are high and when they are low looks dramatic for the UK, but what about other markets? We shall see that in most countries, higher dividend yields signalled higher expected returns.

Chart 6 portrays these return differences between investing after high- and low-yield dates, over horizons of two and ten years. The UK is toward the middle of the sample of 16 countries. The 16-country average return difference has been 6.3 percent over two-year periods, and 3.4 percent over ten-year periods. We can also estimate these return differences omitting the two countries—Switzerland and Sweden—that maintained neutrality during the last century, and the average return difference was 7.6 and 4.3 percent over intervals of two and ten years respectively.

Let us now review the historical stock market performance after the aggregate dividend yield was at an extreme level. If an investor had purchased equities when dividend yields were at a historical high, he would subsequently have experienced a relatively large mean return. If he had purchased equities when dividend yields were low, he would subsequently have experienced a relatively small mean return. Most countries shared this pattern. Note, however, that the advantage favouring high-yield investing varied greatly across countries. Within each country, the mean return, after a given level of yield, can mask a wide range of returns across the individual years that are being averaged. These differences may reflect a considerable degree of estimation error because of the relatively small number of observations that underpin our charts. Note that we have omitted statistical tests of significance from this preliminary report on our research.

We classify countries according to their broad levels of performance after high-and low-yield dates. Since the first year after observing an extreme dividend yield can have anomalous performance, we rank countries by their annualised two–year performance statistics. We use the same investment horizons as before: intervals that range from 1 year to 15 years. For each country/horizon combination, we then plot the arithmetic difference between annualised real returns after years with a high and with a low dividend yield.

Chart 6 : Global estimates of return after high yields *minus* return after low yields, 1900–2003

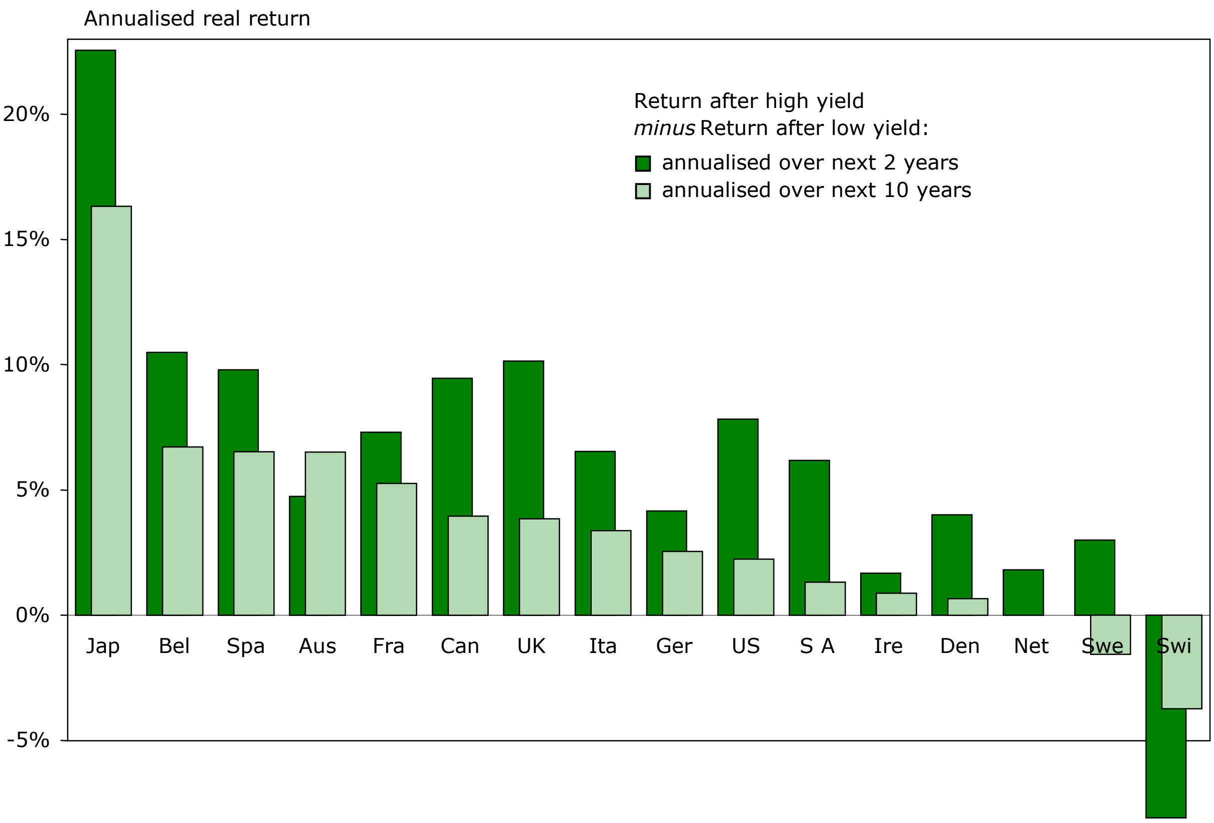

The combined cross-sectional and time-series evidence is summarised in Chart 7, which has two pronounced features. First, most countries experienced higher returns after the market stood on a high yield, rather than on a low yield. The exceptions were Switzerland and Sweden. Second, most countries experienced persistence in this pattern. Ignoring the first year after an extreme yield, the return premium from investing subsequent to a high, rather than a low, yield appears to be highly persistent.

There is no look-ahead bias in the data used in these calculations. All the investment decisions embodied in this analysis are based on information that was available prior to the start of each year. There was no presumption that the investor would know the future. Decision rules did not incorporate information on yields or returns that were destined to happen in the remainder of the twentieth century.

Two aspects of the analysis should be pointed out:

1. Apart from the one-year returns, the performance statistics span overlapping time-periods. For instance, buying Japanese equities when their yield was high meant buying in the early part of and in the middle of the last century (which turned out to be excellent timing), and avoiding intervals when yields were low, such as the period from the 1980s to date (at first disastrous, then a wonderful decision). The multi-year results reflect the continuing Japanese bear market of recent times, and statistical interpretation of the multi-year returns requires particular care.

2. Even though we have followed an investment process that avoids look-ahead bias, it still generates investment strategies that are likely to vary the investor's market exposure over time. The resulting investment-date clustering imposes risks of underperformance in calendar time. We look in more detail at several countries, among them Japan, later on.

Chart 7 : Return after high yields *minus* return after low yields, 16 countries and 15 horizons, 1900–2003

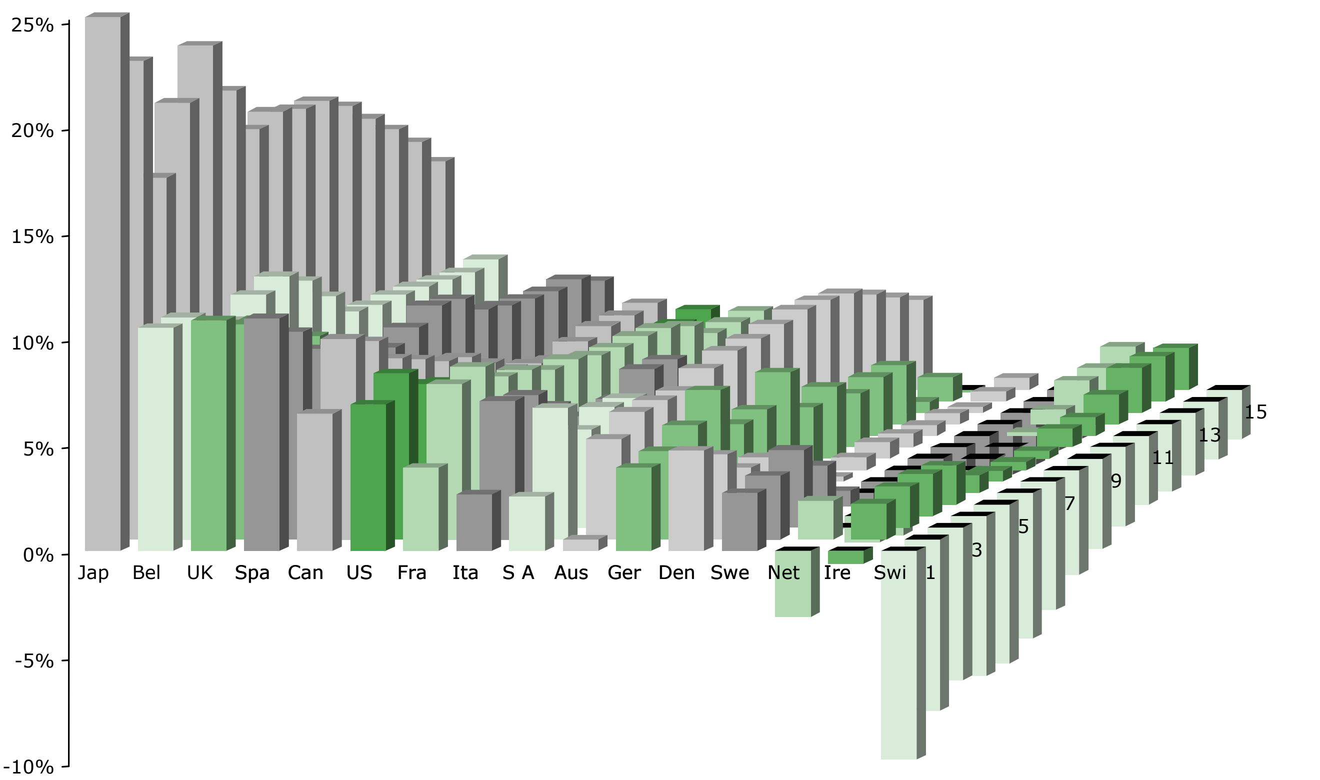

When should dividend models outperform?

In his review article, Cochrane (1997) explains that when dividend yields are low, one of three things—or a combination of them—must be true: [1] Investors expect dividends to rise in the future. Or [2] Investors expect returns to be low in the future (since future cash flows are discounted at a low rate, giving rise to a high price in relation to dividends). Or [3] Investors expect security prices to keep rising, in a bubble that is disconnected from fundamentals. This statement is not a theory or a hypothesis. As Cochrane notes, it is an accounting identity: it *must* be true. When yields are low, either dividends must grow, or prices must fall, or yields must continually diverge from their historical average.

We do not consider further the notion that stock market levels can be permanently decoupled from dividend fundamentals. Statistically, dividend yields are, in fact, close to a random walk. Year-to-year fluctuations in the dividend yield are unlikely to reflect bubbles, and are more likely to reflect the changing perceptions of prospective dividend growth or of time-varying expected returns. Which of these two is more likely to be an important factor? Aggregate dividends change gradually, and annual dividend growth rates fall within a narrower range than fluctuations in the level of the market. One would expect—and empirical research confirms—that changing dividend yields are more closely linked with changes in expected stock returns.

Why might there be a positive association between dividend yields and expected returns? Yields are typically higher when markets have fallen. When markets have fallen, equity investors are poorer; and if they are poorer, they are likely to be more risk averse. There could also be other factors at work. For instance, market falls may be associated with greater volatility and hence risk; and when equity markets have fallen, most companies' leverage rises (in market value terms) which makes equities still more volatile. This, too, means investors will require a higher reward for risk.

Chart 8 reports average stock market performance *before* each turn of the year, when the yield was high or low. We rank years by their start-year dividend yields, and identify the tenth of all years with the highest start-year yield, and the tenth with the lowest start-year yield. We then calculate the mean four-year return prior to the date of the yield estimate, and compute the difference between returns prior to a high- and a low-yield date. At times of high yield, investors have typically experienced a loss—and so may exhibit increased risk aversion, and demand higher required rates of return.

Chart 8 : Real return before a high-yield *minus* return before a low-yield year, 1900-2003

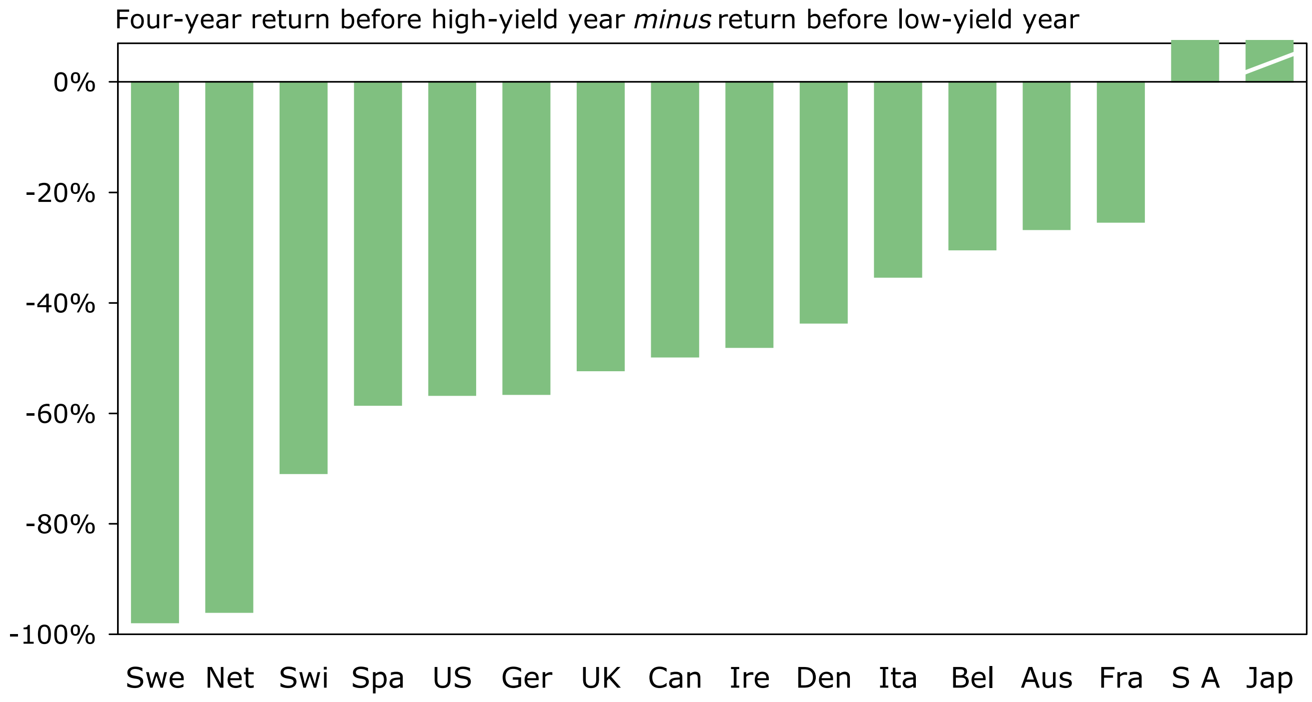

In 14 out of 16 countries, high-yield years had prior performance that was inferior to the prior performance of low-yield years. Investors who bought stocks in high-yield years tended to be buying when equity investors as a whole were poorer, and had lost money over the previous four years, while precisely the opposite was true for investors who bought stocks in low-yield years. Compared to their position a few years beforehand, the median wealth of stock market investors at a time of high dividend yields was about one-half that of investors at a time of low dividend yields. This is consistent with the view that, at times of high yields, investors are more risk averse and on average required a higher expected return to compensate them for equity market risk. We can now move toward developing a model that links the aggregate dividend yield with expected returns on the equity market.

Towards a dividend model

Dividend yields have the potential to indicate higher expected returns from equities. When were yields high, and in what years was there a higher expected return from equities? Chart 9 presents a histogram of UK dividend yields. The lowest yields (of below 3 percent) occurred in the four years 1998-2001; these years are in the left tail of the distribution. The highest yields (above 7 percent) occurred in 1974 and 1979; these years are in the extreme right tail of the distribution.

In the histogram, we have highlighted the most recent two decades, 1984-2003, in dark green, and the preceding decade, 1974-83, in gray. The remaining years, 1900-73, are shaded in light green. There were no extra-high yields during the two decades from 1984-2003 (shaded in dark green), and there were no extra-low yields in the previous decade from 1974-83 (in gray). There is overrepresentation of low yields in recent years, and of higher yields in the 1970s. Similar patterns can be seen in equivalent charts for other markets. It is an indication that, despite variation over the years, dividend yields have drifted downward over time.

Chart 9 : Frequency distribution of UK dividend yields at year-ends, 1900-2003

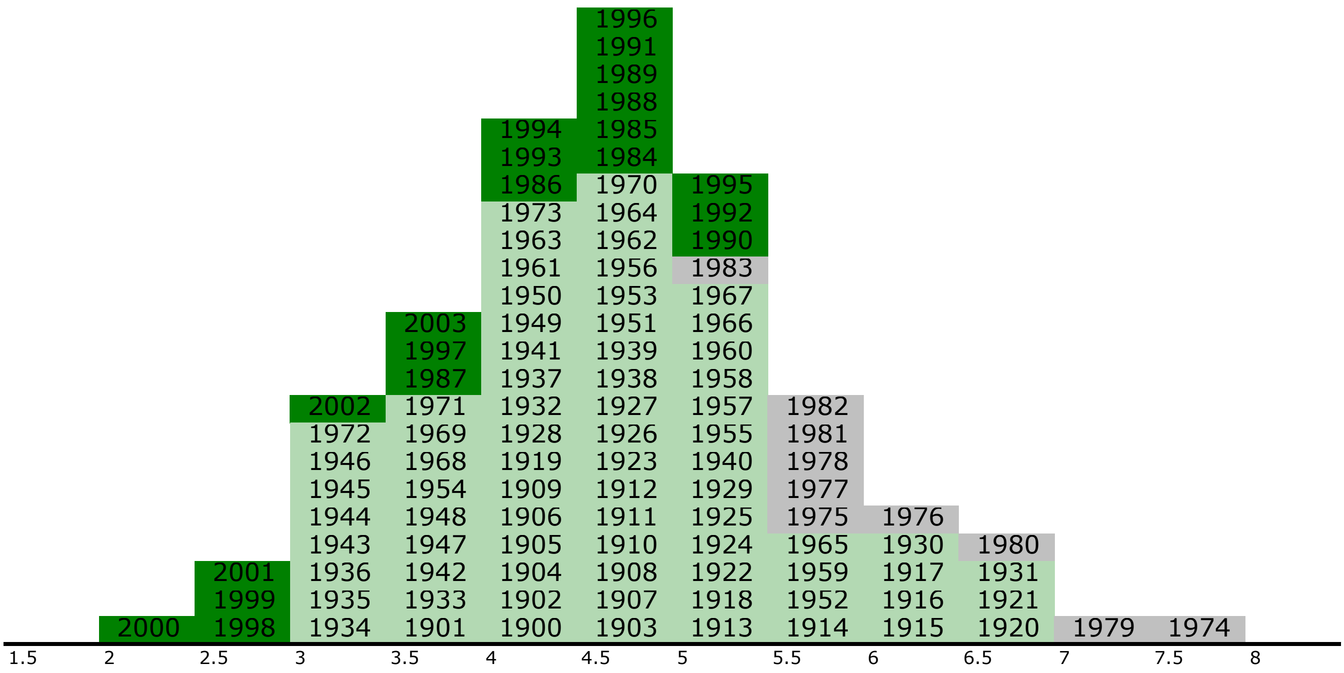

In fact, as Chart 10 shows, in all 16 countries, dividend yields have been lower over the last two decades than over the preceding long-term history in our database. Any model that is to be usable in today's market conditions needs to take account of drift in the average level of dividend yields. Apart from a small number of papers—notably Ang and Bekaert (2001), Cremers (2002) and Goyal and Welch (2003)—most researchers have ignored this phenomenon. This is despite the fact that, to be implemented, it is important that a dividend model does not simply tell you how you should have invested long ago, but also what to do today that is relevant to the future.

Chart 10: Average 1984-2003 dividend yield *minus* average 1900-1983 dividend yield

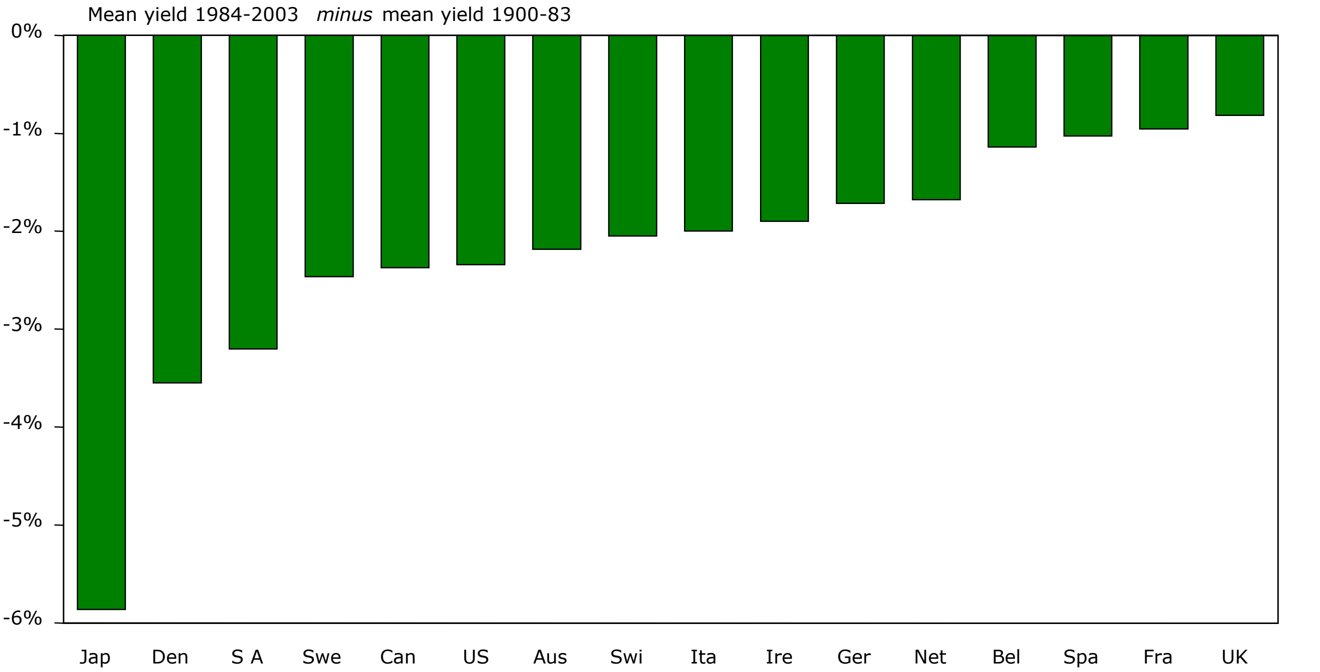

The simplest model—one that has the greatest statistical precision as a summary of past stock market behaviour—uses all available data. We therefore commence with an examination of the relationship between dividend yields and subsequent returns, using the entire dataset from 1900 through 2003. The data for an illustrative country are portrayed in Chart 11. This scatter diagram plots the continuously-compounded annual returns on the UK equity market, on the vertical axis, against dividend yield, plotted on the horizontal axis using a logarithmic scale. For the UK, there is a clear relationship between the dividend yield and subsequent stock market return. The relationship, which we estimate using OLS regression, can be used to infer the expected return implied by a particular level of the dividend yield.

The regression model, linking yields to returns, is as follows (*t*-values in brackets):

Continuously compounded return = 0.9134 + 0.2803 x *Logarithmic dividend yield*
(3.6311) (3.4354)

The regression slope coefficient, which in this case is 0.28, has a highly significant *t*-value. It is usual for the R-squared to be small in regressions like this, but in this case there appears to be a fair degree of explanatory power for the regression: the model has an R-squared of 10.6 percent. However, the R-squared (and the slope coefficient) would be smaller if it were not for the stock market reversal between 1974 and 1975 (see the upper panel of Chart 1 above). In Chart 11 below, 1974 is the outlier that lies beneath the plotting area, while 1975 is the observation in the north-east corner of the plotting area.

To interpret this regression, consider a year that opens with a low dividend yield of 3 percent (the log of 0.03 is a logarithmic dividend yield of –3.5066). When yields are this low, the model indicates an expected return of 0.9134 – 0.2803 x 3.5066 = -0.0694 = –6.9 percent on a continuously-compounded basis. This equates to a predicted loss in value of 6.7 percent over a year. Now consider a year that opens with a high dividend yield of 6 percent (a logarithmic dividend yield of –2.8134). When yields are this high, the model indicates an expected return of 0.9134 – 0.2803 x 2.8134 = 0.1249 = +12.5 percent on a continuously-compounded basis. This equates to a predicted gain in value of +13.3 percent over a year.

The model we have estimated here embraces some peculiarities of UK financial market history that many consider unlikely to be repeated. In the case of the 1975 stock market recovery, for example, not only was the magnitude large, but the reversal happened to be centred on the turn of the calendar year. Two questions present themselves: What would the model have looked like prior to 1974? And is the British experience typical of other countries? We address these questions next.

Chart 11 : Scatterplot of annual UK equity market returns against start-year dividend yields, 1900-2003

Table 3 provides summary statistics on the dividend-return regressions for all 16 countries. The UK is the first line in the table. The R-squared estimates run from a high of 10.6 percent (for the UK) down to just above zero. All but one of the estimated slopes are positive, indicating that, virtually everywhere, higher dividend yields tend to be associated with higher expected real returns. The *t*-statistics for the slope coefficient are significant at the 95 percent confidence level for the first six countries. For those searching for a reliable and profitable investment decision rule, these levels of explanatory power are, on balance, disappointing. Given that the Fama and French (1988) study also looked at multi-year intervals, one could try out regressions with a time interval that is longer than a year. However, longer horizons introduce statistical problems, especially from the use of overlapping observations. Indeed, Goyal and Welch (2003) note that annual horizons have "*the best or close-to-best performance*" (page 645). Rather than experiment with alternative model specifications, we dig deeper into the properties of one-year-ahead forecasts of return, and delay results based on four-year-ahead forecasts until later.

Such models might in principle guide future investment strategy, especially in the countries that have significant *t*-values in the regressions. However, we have so far looked only at the in-sample performance of the regressions; we do not know how reliably such models might predict market returns on an out-of-sample basis. That is because our models have been estimated using data for 1900-2003, and statistics that indicate model validity (such as the R-squared and the *t*-value) are computed for the same countries over the same period 1900-2003, and are not based on an independent dataset. To judge the robustness of the dividend-based models, we therefore explore how they would have performed in the past, not on an in-sample basis but on an out-of-sample basis. We focus on an observer who might have re-estimated the model annually, using all data that was available up to the time of the computations, but without foreknowledge of data that was unavailable at that time.

Table 3 : Regression of annual real return on prior dividend yield, 1900-2003

Country	Intercept	*t*-value	Slope	*t*-value	R-squared
United Kingdom	0.91	3.63	0.28	3.44	10.6%
France	0.37	2.71	0.10	2.50	5.9%
Italy	0.39	2.20	0.11	2.11	4.3%
Ireland	0.43	2.31	0.13	2.10	4.2%
Spain	0.30	2.27	0.08	2.07	4.1%
United States	0.37	2.31	0.10	1.95	3.7%
Canada	0.31	2.26	0.08	1.85	3.3%
Japan	0.20	1.73	0.05	1.47	2.1%
Sweden	0.32	1.81	0.08	1.43	2.0%
Belgium	0.20	1.34	0.05	1.21	1.4%
South Africa	0.16	1.91	0.03	1.17	1.3%
Australia	0.31	1.51	0.08	1.16	1.3%
Netherlands	0.23	1.32	0.06	1.05	1.1%
Denmark	0.09	0.79	0.01	0.40	0.2%
Germany	0.10	0.35	0.02	0.25	0.1%
Switzerland	-0.01	-0.06	-0.01	-0.35	0.1%

A time-varying model

In this section we develop a time-varying dividend-return model, which uses only data that was available at each year-end over the last century. This methodology overcomes the data-mining challenge of researching and then testing model validity. Every annual prediction is computed using what was then historical data. It is then compared to the subsequent outcome on a strictly out-of-sample basis. We compare predictions with outcomes over a long timescale and over all sixteen countries. The key to the model is the slope coefficient presented earlier. The slope coefficient tells us how much weight to give to the dividend yield in order to forecast stock market returns.

Chart 12 presents the results of an annual estimation procedure for the slope coefficient. The bold black line records successive estimates of the slope for the UK; the dark green line is for the US. We start at the end of 1903, estimating the regression of 1901, 1902 and 1903 returns on dividend yields for end-1900, end-1901 and end-1902. This gives us a regression slope estimated as at end-1903. We repeat the process, using four observations up to 1904, five to 1905, and so on until we use 103 observations ending in 2003. This gives us a sequence of 100 estimates of the regression slope. These recursive estimates of the slope coefficient fluctuate wildly until 1920, so the values are not plotted until then.

The gray shaded bars indicate the standard error of the estimated slope coefficient for the UK; the light green area is the standard error for the US coefficients. Until about 1920, the UK slope coefficient was smaller than its standard error, and that situation recurred during the middle third of the century. The UK slope coefficient did not achieve statistical significance (i.e., around twice its standard error) until the 1974-5 stock market reversal. For most of the period represented in the diagram, a UK observer would have asserted that the empirical relationship was not significant.

The pattern was slightly different in the US (see the dark green line and the light green area). The slope coefficient exceeded two standard errors for an interval that mapped closely onto the Fama-French research period, but it is now insignificant.

Chart 12 : Recursive estimates of slope coefficient for regression of return on prior yield, 1900-2003

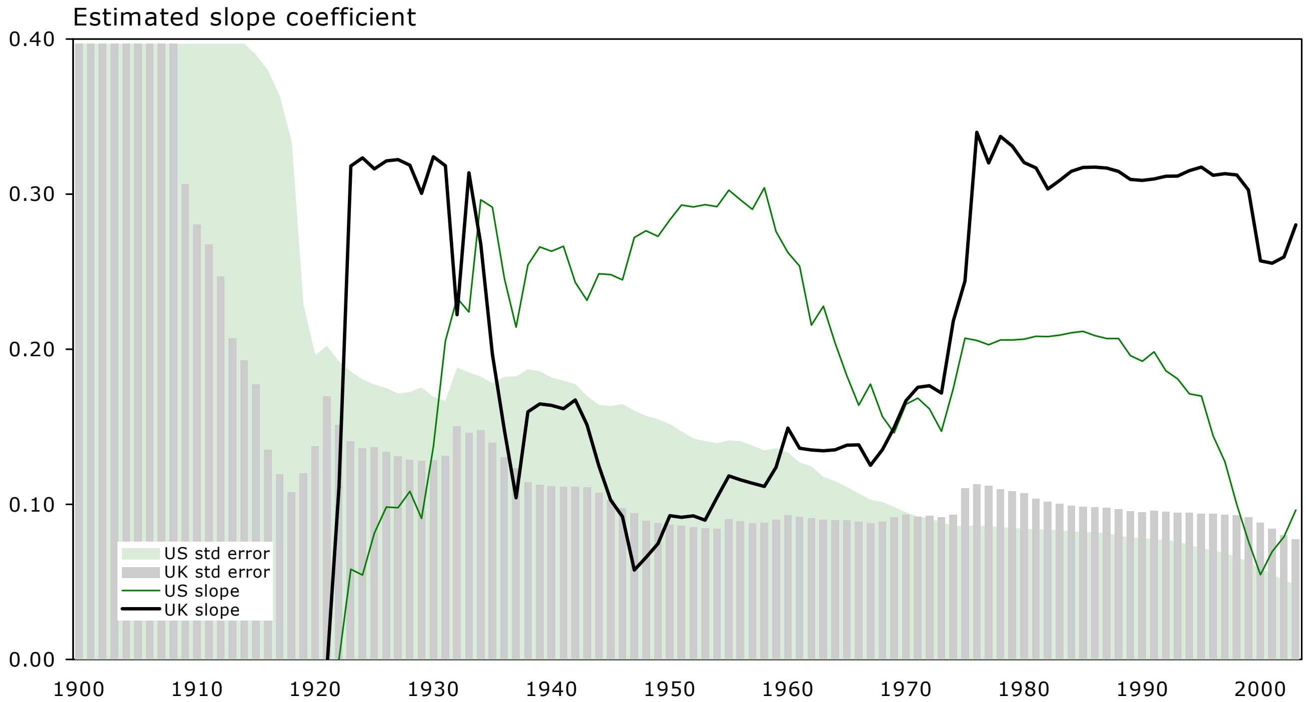

Diagnosing the model's out-of-sample performance

Goyal and Welch (2003) introduce a graphical diagnostic method for appraising the relative performance of a forecasting model. The diagnostic distinguishes between two alternatives:

1. Generating forecasts by extrapolating returns by projecting the historical average into the future

2. Generating forecasts by means of the dividend-based model.

The former is an unconditional forecast; the latter, a conditional forecast. The diagnostic procedure involves calculating the cumulative sum-squared error from the unconditional model, and deducting the cumulative sum-squared error from the dividend-based model. The diagnostic is plotted for the US, using a base date of 1909, in Chart 13.

In the diagnostic diagram, a positive slope indicates that the dividend-based model has lower forecasting error in a specific year. A cumulative value that is positive at a particular date indicates that the dividend-based model has, by that date, outperformed the unconditional model. From 1909 to date, the dividend-based model performed poorly. Its forecasting performance was much worse than simply extrapolating expected returns from the historical record that could have been observed at each date.

Closer inspection reveals that from the date of the 1929 Wall Street Crash to 1997, every start date gave rise to cumulative underperformance in the US from the dividend-based model. Like Fama and French (1988), Goyal and Welch used a dataset that began in 1926, so neither could produce plots like these until after 1929, though their chosen base dates for out-of-sample testing are in fact 1946 (Goyal-Welch) and 1967 (Fama-French). As can be seen from the diagnostic, for every possible base date between 1929-1997 the diagnostic has a cumulative value that is negative by the end of 2003.

The US experience reveals that the apparently significant dividend-return model achieves significance only when tested on an in-sample basis. Dividend-based forecasts are not successful out of sample. This raises the question of how well such models might perform in other countries. We will therefore compare the United Kingdom with other major markets—the United States, Japan and France.

Chart 13 : Cumulative out-of-sample prediction errors for the US dividend-return model, 1900-2003

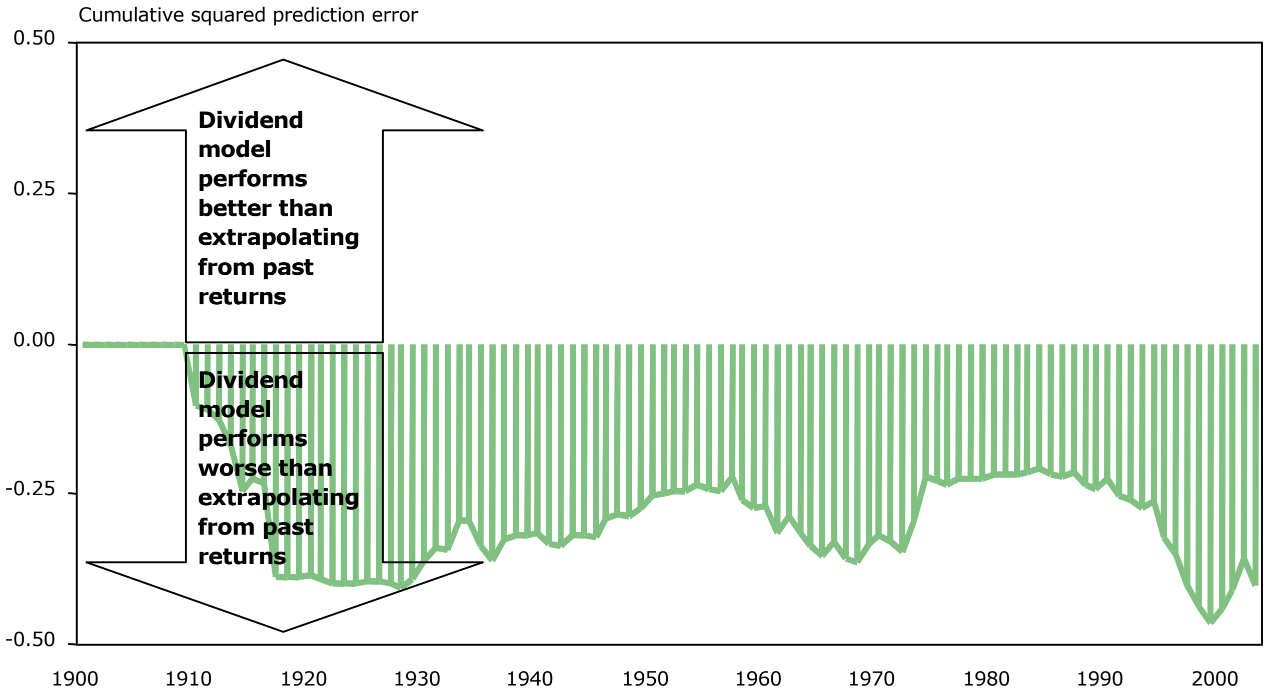

The four-country comparison is shown in Chart 14. Turning first to the UK, the cumulative performance of the dividend model appears reasonably good. The diagnostic ends up in positive territory (more so, if the base date is close to the middle third of the last century). In the UK, the dividend-based model did better than extrapolating from past market returns. As noted earlier, superior performance reflected two factors. First, the model indicated higher expected returns after the 1973-4 oil price crisis, when the return on the market was –71 percent in real terms in just two years; the market in fact recovered in 1975 with a real return of 97 percent. Second, it indicated lower expected returns at the end of the 1990s tech boom; the UK market in fact fell in the three years after 1999 by 41 percent. Yet history could have been different, and 1975 could have witnessed further collapse rather than recovery. In the 1990s the dividend-based signal came too early for some, and those who prematurely exited the market, underperformed. A model that makes two correct "calls" in a century cannot be useful for investors who measure their horizons in years or decades.

Japan's experience of dividend-based forecasting is more of a roller coaster, as the third panel of the chart reveals. The explanation is linked partly to the high yield at which Japanese equities traded after WW2. The dividend model signalled Japan's impending post-war economic recovery. Once that had happened, Japanese equities traded at a low yield. The dividend model erroneously indicated low expected returns, and by the late 1980s had relinquished all of the model's outperformance. Only in recent years has the low yield of Japanese equities correctly predicted poor returns. So while two "calls" were successful—the post-WW2 recovery and the recession of the 1990s—overall, the dividend-return model was uninformative. France was another roller-coaster experience. The dividend model worked well for a decade or so in the middle of the twentieth century, from the end of WW2 until around the mid-50s. But it performed poorly early in the century, and especially poorly during the period since 1990.

Did most countries have an experience like the UK, in which the dividend model outperformed (albeit in a spotty way)? Or were they like Japan or France, with no support for using dividend-based forecasting? Summarising all 16 countries, the dividend model outperformed simple extrapolation of next year's real return in just one market: the UK. A further market (South Africa) was on the borderline. For 14 markets, dividends were uninformative about the following year's market returns. We also looked at using the same dividend model to predict multiple years into the future. The results were similar. For example, out of 16 countries, the dividend model outperformed simple extrapolation of the next *four* years' real return in just one market: once again, the UK. Two additional markets (Spain and Ireland) were on the borderline. For 13 markets, dividends remained uninformative about longer-term stock market returns.

Chart 14 : Cumulative out-of-sample prediction errors for major global markets, 1900–2003

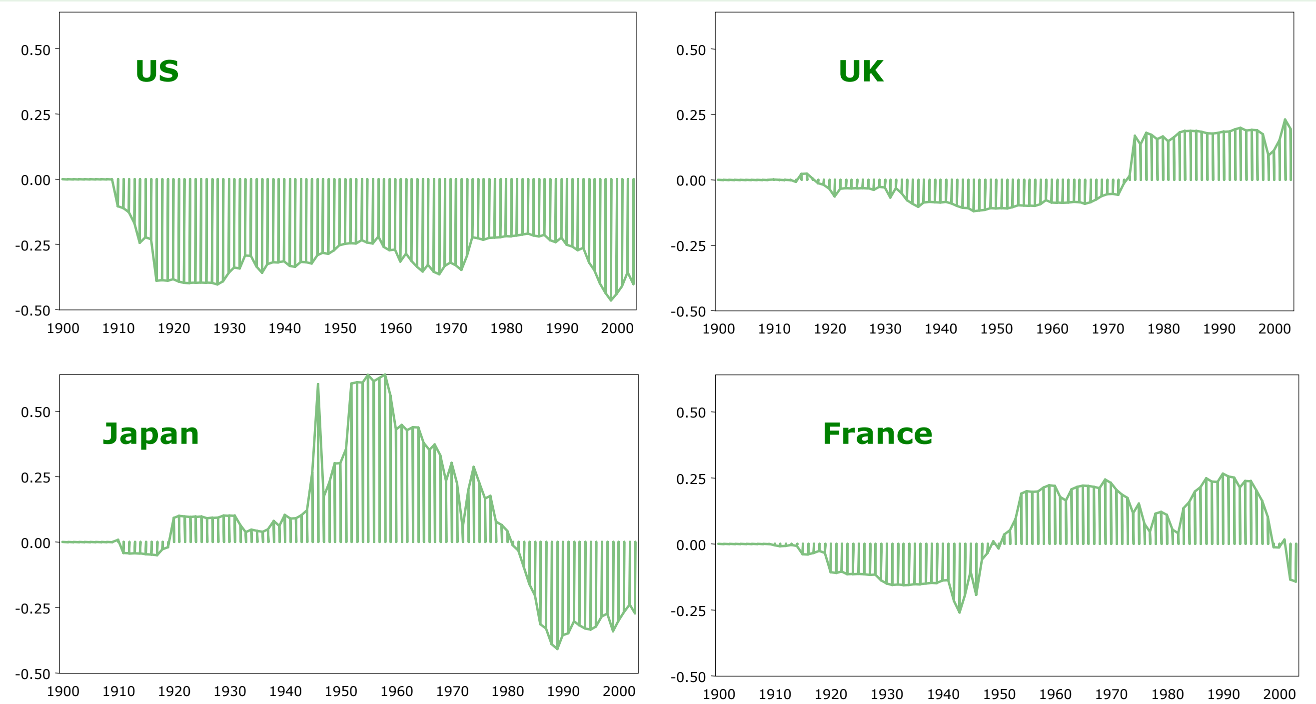

Conclusion

Buying equities when dividend yields are high should, in theory, lead to exposure to the market when other investors are most risk-averse and expected returns are at their highest. While this is not a recipe for beating the herd—since it takes investors into equities at precisely the time they are most risk averse—we do find evidence that at times of high yield, expected returns are larger than they otherwise would be. Unfortunately, it is hard to create a model that reliably captures such facets of the market, and which works successfully on an out-of-sample basis. For most countries, the diagnostic reveals that dividend yields forecast stock market returns less accurately than assuming that the historical record will be repeated.

Of course, there are limitations to our research that we plan to address in the future. We have chain-linked market indices to create a 104-year history for each country, and we wish to check the robustness of our methods for constructing these series. We have focused on inflation-adjusted total returns, and though we also find similar patterns when the research is repeated using risk premia instead of real returns, other performance measures are possible. Finally, we wish to examine longer-run predictions in more detail, with appropriate statistical testing.

Dividends are a crucial component of long-run returns. Over the last 104 years, Dimson, Marsh, and Staunton (2004) report that a portfolio of UK equities, with dividends reinvested, would have grown to almost 119 times the value it would have attained if dividends had been spent. Aggregate dividends also underpin forecasts of the overall return on the stock market. In this paper, we have researched how to use the dividend yield to predict subsequent stock market movements. Investors who bought equities at times when yields were high achieved larger returns than buying at times when yields were low. However, this does not translate into a straightforward decision rule for predicting stock market returns. We find that, when making forecasts one or several years into the future, dividend-based return predictions are in general inferior to extrapolation from the historical record of equity returns. Given our current state of knowledge, the long-run historical record should be the starting point for forecasts of long-run stock market returns.

References

Ang, A. and G. Bekaert, 2001, *Stock return predictability: is it there?*, NBER and Columbia University.

Barberis, N. and R. Thaler, 2003, A survey of behavioural finance, *Handbook of the Economics of Finance*, Edited by G. M. Constantinides, M. Harris and R. Stultz, North Holland.

Bossaerts, P. and P. Hillion, 1999, Implementing statistical criteria to select return forecasting models: what do we learn?, *Review of Financial Studies* 55(1) 225-264.

Campbell, J. Y. and R. J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount rates, *Review of Financial Studies* 1(3) 195-228.

Campbell, J. Y. and R. J. Shiller, 2004, Valuation ratios and the long-run stock market outlook: an update, *Advances in Behavioral Finance—Volume II*, Edited by N. Barberis and R. Thaler, Russell Sage Foundation.

Campbell, J. Y., and L. M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.

Cochrane, J. H., 1997, Where is the market going? Uncertain facts and novel theories, *Federal Reserve Bank of Chicago—Economic Perspectives* 21(6) 3-37.

Cremers, K. J. M., 2002, Stock return predictability: a Bayesian model selection perspective, *Review of Financial Studies* 15(4) 1223-1249.

Dimson, E., P. R. Marsh and M. Staunton, 2004, *Global Investment Returns Yearbook 2004*, ABN AMRO and London Business School.

Dimson, E., P. R. Marsh and M. Staunton, 2003, Global evidence on the equity risk premium, *Journal of Applied Corporate Finance* 15(4) 27-38.

Dimson, E., P. R. Marsh and M. Staunton, 2002, *Triumph of the Optimists: 101 Years of Global Investment Returns*, New Jersey: Princeton University Press.

Dimson, E., S. Nagel and G. Quigley, 2003, Capturing the value premium in the UK, *Financial Analysts Journal* 59(6) 35-45.

Dow, C. H., 1920, Scientific stock speculation, *The Magazine of Wall Street*, New York.

Fama, E. F. and K. R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22(1) 2-25.

Goetzmann, W. N. and P. Jorion, 1993, Testing the predictive power of dividend yields, *Journal of Finance* 48(2) 663-679.

Goetzmann, W. N. and P. Jorion, 1995, A longer look at dividend yields, *Journal of Business* 68(4) 483-508.

Goyal, A. and I Welch, 2003, Predicting the equity premium with dividend ratios, *Management Science* 49(5) 639-654.

Graham, B. and D. L. Dodd, 1934, *Security Analysis*, First edition, McGraw Hill.

Kothari, S. P. and J Shanken, 1997, Book-to-market, dividend yield, and expected market returns: a time series analysis, *Journal of Financial Economics* 44(2) 169-203.

Lo, A. W. and A. C. MacKinley, 1990, Data snooping biases in tests of financial asset pricing models, *Review of Financial Studies* 3(3) 431-468.

Nelson, C. R. and M. J. Kim, 1993, Predictable stock returns: the role of small sample bias, *Journal of Finance* 48(2) 641-661.

Rozeff, M. S., 1984, Dividend yields are equity risk premiums, *Journal of Portfolio Management* 11(1) 68-75.

Stambaugh, R. F., 1999, Predictive regressions, *Journal of Financial Economics* 54(3) 375-421.

Valkanov, R., 2003, Long-horizon regressions: theoretical results and applications, *Journal of Financial Economics* 68(2) 201-232.