1 Market efficiency and forecasting

Wayne E. Ferson

1.1 Introduction

The interest in predicting stock prices or returns is probably as old as the markets themselves, and the literature on the subject is enormous. Fama (1970) reviews early work and provides some organizing principles. This chapter concentrates selectively on developments following Fama's review. In that review, Fama describes increasingly fine information sets in a way that is useful in organizing the discussion. Weak-form predictability uses the information in past stock prices. Semi-strong form predictability uses variables that are obviously publicly available, and strong form uses anything else. While there is a literature characterizing strong-form predictability (e.g. analyzing the profitability of corporate insider's trades), this chapter concentrates on the first two categories of information.

For a while, predicting the future price or value (price plus dividends) of a stock was thought to be easy. Early studies, reviewed by Fama (1970), concluded that a martingale or random walk was a good model for stock prices, values or their logarithms. Thus, the best forecast of the future price was the current price. However, predicting price or value changes, and thus rates of return, is more challenging and controversial. The current financial economics literature reflects two often-competing views about predictability in stock returns. The first argues that any predictability represents exploitable inefficiencies in the way capital markets function. The second view argues that predictability is a natural outcome of an efficient capital market.

The exploitable inefficiencies view of return predictability argues that, in an efficient market, traders would bid up the prices of stocks with predictably high returns, thus lowering their return and removing any predictability at the new price (see, for example, Friedman, 1953; Samuelson, 1965). However, market frictions or human imperfections are assumed to impede such price-correcting, or 'arbitrage' trading. Predictable patterns can thus emerge when there are important market imperfections, like trading costs, taxes, or information costs, or important human imperfections in processing or responding to information, as studied in behavioural finance. These predictable patterns are thought to be exploitable, in the sense that an investor who could avoid the friction or cognitive imperfection could profit from the predictability at the expense of other traders.

The 'efficient markets' view of predictability was described by Fama (1970). According to this view, returns may be predictable if required expected returns vary over time in association with changing interest rates, risk or investors' risk-aversion. If required expected returns vary over time there may be no abnormal trading profits, and thus no incentive to exploit the predictability. Predictability may therefore be expected in an efficient capital market. The return is written as $R = E(R|\Omega) + u$, where Ω is the information at the beginning of the period and u is the unexpected return. Since $E(u|\Omega) = 0$, the unexpected return cannot be predicted ahead of time. Thus, predictability, in the 'efficient markets' view, rests on systematic variation through time in the expected return. Modelling and testing for this variation is the focus of the conditional asset pricing literature (see reviews by Ferson, 1995; Cochrane, 2005).

While this chapter focuses on return predictability, not all of the predictability associated with stock prices involves predicting the levels of returns. A large literature models predictable second moments of returns (e.g. using ARCH and GARCH-type models (see Engle, 2004) or other stochastic volatility models). Predictability studies have also examined the third moments (see, for example, Harvey and Siddique, 2001).

1.2 A modern view of market efficiency and predictability

As described by Fama (1970), any empirical analysis of stock return predictability or the market's informational efficiency involves a 'joint hypothesis'. There must be an hypothesis about the model for equilibrium expected returns, and also an hypothesis about the informational efficiency of the markets. These can be easily described using a modern representation for asset pricing models.

Most of the asset pricing models of financial economics can be described as versions of equation (1.1):

$$E\{m_t(1 + r_t)|\Omega_{t-1}\} = 1 \tag{1.1}$$

where Ω_{t-1} is the information set of economic agents at the beginning of the period and r_t is the rate of return of a financial asset. The scalar random variable m_t is the stochastic discount factor. Different models imply different stochastic discount factors, and the stochastic discount factor should price all of the assets in the model through equation (1.1).

The joint hypothesis of stock return predictability and market efficiency tests may now be described. The assumption of a model of market equilibrium amounts to a specification for m_t. For example, the Capital Asset Pricing Model of Sharpe (1964) implies that m_t is a linear function of the market portfolio return (see, for example, Dybvig and Ingersoll, 1982). Assume that the analyst uses the lagged variables, Z_{t-1}, to predict stock returns. The hypothesis of informational efficiency is simply the statement that Z_{t-1} is contained in Ω_{t-1}. For example, weak-form efficiency says that past stock prices are in Ω_{t-1}, while semi-strong form efficiency says that other publicly available variables are in Ω_{t-1}.

The martingale model for stock values follows as a special case of this modern view. If we assume that m_t is a constant over time (as implied by risk neutral agents with fixed time discounting), then equation (1.1) implies that $E\{r_t|\Omega_{t-1}\}$ is a constant over time. Since $r_t = E(r_t|\Omega_{t-1}\} + u_t$, and u_t is unpredictable, it follows that the returns r_t cannot be predicted by any information in Ω_{t-1}.

1.3 Weak-form predictability

Much of the literature on weak form predictability can be characterized through an autoregression. Let R_t be the continuously compounded rate of return over the shortest measurement interval ending at time t. Let $r(t, t+H) = \Sigma_{j=1,\ldots,H} R_{t+j}$. Then,

$$r(t, t+H) = a_H + \rho_H r(t-H, t) + \varepsilon(t, t+H) \tag{1.2}$$

is the autoregression and H is the return horizon. Studies can be grouped according to the return horizon.[1]

Many studies measure small but statistically significant serial dependence in daily or intra-daily stock return data. Serial dependence in daily returns can arise from end-of-day price quotes that fluctuate between bid and ask (Roll, 1984), or from non-synchronous trading of the stocks in an index (see, for example, Fisher, 1966; Scholes and Williams, 1977). These effects do not represent predictability that can be exploited with any feasible trading strategy. Spurious predictability due to such data problems should clearly not be attributed to time-variation in the expected discount rate for stocks. On the other hand, much of the literature on predictability allows that high frequency serial dependence may reflect changing conditional means. For example, Lo and MacKinlay (1988) and Conrad and Kaul (1988) model expected returns within the month as following an autoregressive process.

Conrad and Kaul (1988, 1989) studied serial dependence in weekly stock returns. They point out that if the expected returns, $E(R|\Omega)$, follow an autoregressive process, the actual returns would be described by the sum of an autoregressive process and a white noise, and thus follow an ARMA process. The autoregressive and moving average coefficients would be expected to have the opposite signs: If current expected returns increase, it may signal that future expected returns are higher, but stock prices may fall in the short run because the future cash flows are discounted at a new, higher rate. The two effects, offset and returns, could have small autocorrelations. Estimating ARMA models, they found that the autoregressive coefficient for weekly returns on stock portfolios are positive, near 0.5, and can explain up to 25% of the variation in the returns on a portfolio of small-firm stocks.

Even with weekly returns, however, some of the measured predictability can reflect non-synchronous trading effects. Lo and MacKinlay (1990) and Muthaswamy (1988) use statistical models that attempt to separate out the various effects in measured portfolio returns. Boudoukh *et al.* (1994) use stock index futures contracts, which are not subject to non-synchronous trading, and find little evidence for predictability at a weekly frequency.

Much of the literature on weak-form predictability studies broad stock market indexes or portfolios of stocks, grouped according to the market capitalization (size) or other characteristics of the firms. However, another significant stream of the literature studies relative predictability. Stocks have relative predictability if the future returns of one group of stocks are predictably higher than the returns of another group. Thus, if a trader could

[1] An alternative to the autoregression is the Variance Ratio statistic, $\mathrm{Var}\{r(t, t+H)\}/HVar(R_t)$, proposed by Working (1949) and studied for stock returns by Lo and MacKinlay (1988, 1989) and others. Cochrane (1988) shows that the variance ratio is a function of the autocorrelation in returns. Kaul (1996) provides an analysis of various statistics that have been used to evaluate weak-form predictability, showing how they can be viewed as combinations of autocorrelations at different lags, with different weights assigned to the lags.

buy the stocks in the high-return group and sell short the stocks in the low-return group, the trader could profit even if both groups were to go up (or down). In a weak-form version of relative predictability, past stock prices or returns are used to form the groups. If past winner (loser) stocks have predictably higher (lower) returns, we have continuation or 'momentum'. If past winner stocks can be predicted to have lower future returns, we have 'reversals'. Relative predictability can be evaluated by viewing equation (1.2) as a cross-sectional regression – an approach taken by Jegadeesh (1990). Lehman (1990) finds some evidence for reversals in the weekly returns of US stocks.

Monthly returns are commonly used in the literature that tests asset pricing models. At this frequency, the evidence on weak-form predictability is relatively sparse. Jegadeesh (1990) finds some evidence for reversals at a monthly return frequency. Ferson *et al.* (2005) make indirect inferences about the time-variation in monthly expected stock returns by comparing the unconditional sample variances of monthly returns with estimates of expected conditional variances. The key is a sum-of-squares decomposition: $Var(R)\} = E\{Var(R|\Omega)\} + Var\{E(R|\Omega)\}$, where $E(.|\Omega)$ and $Var(.|\Omega)$ are the conditional mean and variance, and $Var\{.\}$ and $E\{.\}$, without the conditioning notation, are the unconditional moments. The interesting term is $Var\{E(R|\Omega)\}$; that is, the amount of variation through time in conditionally expected stock returns. This quantity is inferred by subtracting estimates of the expected conditional variance, $E\{Var(R|\Omega)\} = E\{[R - E(R|\Omega)]^2\}$, from estimates of the unconditional variance. The expected conditional variance is estimated following Merton (1980), who showed that while the mean of a stock return is hard to estimate, it is almost irrelevant for estimating the conditional variance, when the time between observations is short. Using high-frequency returns to estimate the conditional variance for each month, then subtracting its average from the monthly unconditional variance, the difference – according to the decomposition – is the variance of the monthly conditional mean.

Ferson *et al.* (2005) find that while historical data prior to 1962 suggests economically significant weak form predictability in monthly stock market returns, there is little evidence of weak-form predictability for monthly returns in modern data. In particular, the evidence for the period after 1992 suggests that any weak-form predictability in the stock market as a whole has vanished. At the same time, a simulation study shows that the indirect tests have the power to detect even modest amounts of predictability.

Jegadeesh and Titman (1993) find that relatively high-past-return stocks tend to repeat their performance over 3- to 12-month horizons. They study US data for 1927–1989, but focus on the 1965–89 period. The magnitude of the effect is striking. The top 20% winner stocks over the last 6 months can outperform the loser stocks by about 1% per month for the next 6 months. This momentum effect has spawned a huge subsequent literature which is largely supportive of the momentum effect, but which has not reached a consensus about its causes. The efficient markets view of predictability suggests that momentum trading strategies should be subject to greater risk exposures which justify their high returns. Most efforts at explaining the effect by risk adjustments have failed.[2]

The momentum effect has inspired a number of behavioural models, suggesting that momentum may occur because markets under-react to news in the pricing of stocks.

[2]There are some partial successes. For example, Ang *et al.* (2001) associate some of the momentum strategy profits with high exposure to 'downside risk' – that is, the covariance with market returns when the market return is negative.

Elsevier UK CH01-H8321 jobcode: FEF 9-4-2007 4:32p.m. Page:5 Trimsize:165×234MM

For example, one argument (Daniel *et al*., 1998) is that traders have 'biased self attribution', meaning that they think their private information is better than it really is. As a result they do not react fully to public news about the value of stocks, so the news takes time to get impounded in market prices, resulting in momentum. In another argument, traders suffer a 'disposition effect', implying that they tend to hold on to their losing stocks longer than they should, which can lead to momentum (Grinblatt and Han, 2003). These arguments suggest that traders who can avoid these cognitive biases may profit from momentum trading strategies. However, Lesmond *et al*. (2004) and Korajczyk and Sadka (2004) measure the trading costs of momentum strategies and conclude that the apparent excess returns to the strategies are consumed by trading costs.

Perhaps the most controversial evidence of weak-form predictability involves long-horizon returns. Fama and French (1988) use autoregressions like equation (1.2) to study predictability in portfolio returns, measured over 1-month to multi-year horizons. They find U-shaped patterns in the autocorrelations as a function of the horizon, with negative serial dependence, or mean reversion, at 4- to 5-year horizons. Mean reversion can be consistent with either view of predictability. If expected returns are stationary (reverting to a constant unconditional mean) but time-varying, mean reversion can occur in an efficient market. Mean reversion would also be expected if stock values depart temporarily from the fundamental, or correct, prices, but are drawn back to that level. The evidence for weak-form predictability in long-horizon returns is subject to a number of criticisms on statistical grounds, as described below.

DeBondt and Thaler (1985) find that past high-return stocks perform poorly over the next 5 years, and *vice versa*, a form of relative predictability. They interpret reversals in long-horizon relative returns as evidence that the market over-reacts to news about stock values, and then eventually corrects the mistake. The reversal effect was shown to occur mainly in the month of January, by Zarowin (1990) and Grinblatt and Moskowitz (2003), which is interpreted as related to 'tax loss selling'. In this story, investors sell loser stocks at the end of the year for tax reasons, thus depressing their prices, and buy them back in the new year, subsequently raising their prices. McLean (2006) finds that reversals are concentrated in stocks with high idiosyncratic risks, which is thought to present a deterrent to arbitrage traders who might otherwise correct temporary errors in the market prices.

Like momentum, behavioural models attempt to explain reversals as the result of cognitive biases. Models of Barberis *et al*. (1998), Daniel *et al*. (1998) and Hong and Stein (1999) argue that both short-run momentum and long-term reversals can reflect biases in under- and over-reacting to news about stock values. Research in this area continues, and it's fair to say that the jury is still out on the issue of weak-form predictability in long-horizon returns.

1.4 Semi-strong form predictability

Studies of semi-strong form predictability can be described with the regression:

$$(t, t+H) = \alpha_H + \beta_{H'} Z_t + \upsilon(t, t+H) \qquad (1.3)$$

where Z_t is a vector of variables that are publicly available by time t. Many predictor variables have been analyzed in published studies, and it is useful to group them into

Basal Fonts:Sabon Margins:Top:36pt Gutter:15mm Font Size:10/12 Text Width:135mm Depth:47 Lines

categories. The first category of predictor variables comprises 'valuation ratios', which are measures of cash flows divided by the stock price. Keim and Stambaugh (1986) use a constant numerator in the ratio and 'detrend' the price. Rozeff (1984), Campbell and Shiller (1988) and Fama and French (1989) use dividend/price ratios, Pontiff and Schall (1998) and Kothari and Shanken (1997) use the book value of equity divided by price. Boudoukh *et al.* (2004) and Lei (2006) add share repurchases and other non-cash payouts, respectively, to the dividend measure. Lettau and Ludvigson (2001) propose a macroeconomic variation on the valuation ratio: aggregate consumption divided by a measure of aggregate wealth. All of these studies find the regression coefficients β_H to be significant.

Malkiel (2004) reviews a valuation ratio approach that he calls the 'Federal Reserve Model.' Here, the market price/earnings ratio is empirically modelled as a function of producer prices, Treasury yields and other variables, and the difference between the model's output and the ratios observed in the market are used to predict the market's direction. (Malkiel finds that the model does not outperform a buy-and-hold strategy.)

Rozeff (1984) and Berk (1995) argue that valuation ratios should generally predict stock returns. Consider the simplest model of a stock price, P, as the discounted value of a fixed flow of expected future cash flows or dividends: $P = c/R$, where c is the expected cash flow and R is the expected rate of return. Then, $R = c/P$, and the dividend price ratio is the expected return of the stock. If predictability is attributed to the expected return, as in the efficient markets view, then a valuation ratio should be a good predictor variable.

Predictability of stock returns with valuation ratios is also related to the expected growth rates of future dividends or cash flows. Consider the Gordon (1962) constant-growth model for a stock price: $P = c/(R - g)$, where g is the future growth rate. Then, $c/P = R - g$. This suggests that if dividend/price ratios vary, either across stocks or over time, then expected returns should vary, and/or expected cash flow growth rates should vary, and the dividend/price ratio should be able to predict one or the other. Campbell and Shiller (1988) show that the intuition from this example holds to a good approximation in more general discounted cash flow models, where the growth rates and expected returns are not held fixed over time. They find that market dividend/price ratios do not significantly predict future cash flow growth rates. Cochrane (2006) uses this result to re-evaluate the empirical evidence for stock return predictability using dividend/price ratios. He essentially argues that if you know that the dividend/price ratio does not forecast future cash flow growth, then it must forecast future stock returns.

Studies of semi-strong form predictability in stock index returns typically report regressions with small R-squares, as the fraction of the variance in returns that can be predicted with the lagged variables is small – say 10–15% or less for monthly to annual return horizons. The R-squares are larger for longer-horizon returns – up to 40% or more for 4- to 5-year horizons. This is interpreted as the result of expected returns that are more persistent than returns themselves, as would be expected if returns are expected returns plus noise. Thus, the variance of the sum of the expected returns accumulates with longer horizons faster than the variance of the sum of the returns, and the R-squares increase with the horizon (see, for example, Fama and French, 1989). However, small R-squares can mask economically important variation in the expected returns.

Stocks are long 'duration' assets, so a small change in the expected return can lead to a large change in the asset value. To illustrate, consider another example using the Gordon model, where the dividend is $c = kE$, E is the earnings and k is the dividend payout ratio. The price/earnings ratio $P/E = 15$, the payout ratio $k = 0.6$, and the expected growth rate $g = 3\%$. The expected return $R = 7\%$. Suppose there is a shock to the expected return, *ceteris paribus*. A change of 1% in R leads to approximately a 20% change in the asset value. Of course, this overstates the effect, to the extent that a shock that changes the required return also changes the expected future cash flows.

As the example suggests, small changes in expected returns can produce economically significant changes in asset values. Consistent with this argument, studies such as Kandel and Stambaugh (1996), Campbell and Viceira (2002) and Fleming *et al*. (2001) show that optimal portfolio decisions can be affected to an economically significant degree by return predictability, even when the amount of predictability, as measured by R-squared, is small.

The second category of semi-strong form predictor variables for stock returns includes calendar and seasonal effects. The list of effects that have been related to stock returns and the list of studies are too long to cite here (see Haugen and Lakonishok (1988) and Schwert (2003) for reviews). Some examples include the season (winter versus summer), the month of the year (especially, high returns in January), the time of the month (first versus last half), holidays, the day of the week (low returns on Mondays), the time of the day, the amount of sunlight (as in seasonal affective disorder) and even the frequency of geomagnetic storms.

The third category of predictor variables in equation (1.3) is a catch-all, 'other' variables. Prominent among these are bond yields and yield spreads. Fama and Schwert (1977) were among the first to observe that the level of short-term Treasury yields predicts returns in equation (1.3) with a negative coefficient. They interpreted the short-term yield as a measure of expected inflation. Ferson (1989) argues that the regressions imply that the systematic risk of stocks that determines the expected returns must vary over time with changes in interest rates. Keim and Stambaugh (1986) study the yield spreads of low-quality over high-quality bonds, and find predictive ability for stock returns, and Campbell (1987) studies a number of yield spreads in shorter-term Treasury securities. Fama and French (1989) assemble a list of variables from studies in the 1980s and describe their relations with US business cycles.

Another interesting set of predictor variables in equation (1.3) includes measures of the conditional variance or volatility of stock returns. Merton (1980) shows that a simple version of his Intertemporal Asset Pricing Model implies that expected returns on the aggregate stock market should be positively related to the conditional variance of the market returns; that is, there should be a positive risk-return tradeoff for the market as a whole. Sharpe's (1964) Capital Asset Pricing Model also makes this prediction. Early studies that tried to predict market returns with predetermined market volatility measures found mixed results – weak or even negative β_H coefficients (e.g. French *et al*., 1987; Breen *et al*., 1993). Scruggs (1998) showed that if additional risk factors as suggested by Merton's (1973) model were included, the partial coefficients became positive as predicted by the theory.

Of course, many other semi-strong form predictor variables have been proposed and more will doubtless be proposed in the future. Some recent variables include the fraction of equity issues in new issues of corporate securities (Baker and Wurgler, 2000), firms' investment plans (Lamont, 2000), the average 'idiosyncratic' or firm-specific component of past return volatility (Goyal and Santa Clara, 2003), the level of corporate cash holdings

(Greenwood, 2004), the aggregate rate of dividend initiation (Baker and Wurgler, 2000), share issuance (Pontiff and Woodgate, 2006), and the political party currently in office (Santa Clara and Valkanov, 2003).

1.5 Methodological issues

Even though the regressions in equations (1.2) and (1.3) seem pretty straightforward, interpreting the predictability evidence for stock returns based on these regressions is not. It can be argued that one of the greatest contributions of the literature on stock return predictability is the methodological lessons it has taught researchers in the field. Perhaps the most difficult issues involve selection bias and data mining. Additional issues that have been addressed in the literature include small sample biases in the coefficients, standard error estimation, multiple comparisons, efficient estimation, regime shifts, spurious regression and the interactions among these effects.

Selection bias and data mining are serious concerns. Data mining refers to sifting through the data in search of predictive or associative patterns. There are two kinds of data mining. Sophisticated data mining accounts for the number of searches undertaken when evaluating the statistical significance of the finding (see, for example, White, 2000). This is important, because if 100 independent variables are examined, we expect to find five that are 'significant' at the 5% level, even if there is no predictive relation. Naive data mining does not account for the number of searches. The big problem, given the strong interest in predicting stock returns among academics and practitioners, and the many studies using the same data, is that it is difficult to account for the number of searches. There probably have been at least as many regressions run using the Center for Research in Security Prices (CRSP) database as there are numbers in the database. Compounding this problem are various selection biases. Perhaps most difficult is the fact that only 'significant' results are circulated and published in academic papers. No one knows how many insignificant regressions were run before those results were found.

A reasonable response to these concerns is to see if the predictive relations hold out-of-sample. This kind of evidence is mixed. Some studies find support for predictability in step-ahead or out-of-sample exercises (for example, Fama and French, 1989; Pesaran and Timmerman, 1995). Semi-strong form variables show some ability to predict returns outside of the US data where they were originally studied (for example, Harvey, 1991; Ferson and Harvey, 1993, 1999; Solnik, 1993). Other studies conclude that predictability using many of the semi-strong form variables does not hold outside of the original samples (for example, Goyal and Welch, 2003, 2004; Simin, 2006). Even this evidence is difficult to interpret, because a variable could have real predictive power yet still fail to outperform a naive benchmark when predicting out of sample (Campbell and Thompson, 2005; Hjalmarsson, 2006).

A large literature has addressed statistical issues in predictive regressions. Boudoukh and Richardson (1994) provide an insightful review. From the perspective of testing the null hypothesis of no predictability, we are interested in whether β_H is zero. The coefficient is proportional to the covariance, $Cov\{r(t, t+H); Z_t\}$. The literature has employed three basic approaches to estimating the covariance. The first is to run the regression in equation (1.3) with overlapping data on $r(t, t+H)$, observed each period t, usually 1 month (see, for example, Fama and French, 1989). Given overlapping data, this approach uses all of the data, as opposed to a sampling scheme, which measures non-overlapping returns every H months, and should therefore be more efficient than the sampling scheme. However, the overlap induces autocorrelation in the error terms in the form of an $H-1$ order moving average process. To conduct inference about β_H, it is necessary to estimate the coefficients and the standard errors without bias in the presence of the moving average error terms. This is complicated by the evidence that stock return data are conditionally heteroscedastic.

When $H > 1$ and several horizons are examined together, the issue of multiple comparisons arises. If 20 independent horizons are examined, it is expected that there will be one 'significant' t-statistic found at the 5% level when the null hypothesis of no predictability is true. The slope coefficients for the different horizons are correlated, which complicates the inference. Richardson (1993) shows this implies that the U-shaped patterns in the autocorrelations across return horizons, observed by Fama and French (1988), are likely to be observed by chance. Boudoukh *et al.* (2005) argue that high correlation of test statistics across the return horizons renders much of the evidence for semi-strong form long-horizon predictability suspect.

Even when $\mathrm{H}=1$ and there is no overlap, correlation in the error terms can lead to finite sample biases in estimates of the slope. Stambaugh (1999) studies one such bias that arises because the regressor is stochastic and its future values are correlated with the error term in the regression. For example, if Z_t is a dividend/price ratio, then shocks to the dividend price ratio at time $t+1$ are related to stock returns at time $t+1$ through the stock price. Stambaugh provides corrections for this bias, which he shows is related to the bias in a sample autocorrelation coefficient, as derived by Kendal (1954). Amihud and Horvitz (2004) explore solutions for this bias in a multiple regression setting.

A second approach to estimating $Cov\{r(t, t+H); Z_t\} = Cov\{\Sigma_{j=1,\dots,H}\ R_{t+j}; Z_t\}$ is to recognize that, if the variables are covariance stationary, then $Cov\{\Sigma_{j=1,\dots,H}\ R_{t+j}; Z_t\} = Cov\{R_t; \Sigma_{j=1,\dots,H}\ Z_{t-j}\}$. This suggests using the horizon $H=1$ on the left-hand side of equation (1.3), and replacing Z_t on the right hand side with $\Sigma_{j=1,\dots,H} Z_{t-j}$. This is the approach taken by Hodrick (1992) and Cochrane (1988). In this approach, the error terms in the regression are not overlapping and there is no induced moving average structure to account for. There may be efficiency gains compared with the first approach, but these depend on the stochastic process that drives the variables.

A third approach to estimating the predictive regression coefficient is to model the single period data $\{R_t,\ Z_{t-1}\}$ using a vector autoregression, and then infer the value of the long-horizon coefficient β_H from the autoregression parameters. This is the approach taken by Kandel and Stambaugh (1990), Campbell (1993), and others. This can be efficient if the vector autoregression is correctly specified, but is subject to error if the autoregression is not correctly specified.

A potential issue with all of the approaches to predictive regressions is spurious regression bias. Spurious regression is studied by Yule (1926) and Granger and Newbold (1974), who warn that empirical relations may be found between the levels of trending time series that are actually independent. For example, given two independent random walks, it is likely that a regression of one on the other will produce a 'significant' slope coefficient, evaluated by the usual t-statistics. In equation (1.3) the dependent variables are stock returns, which are not highly persistent. However, recall that the returns may be considered the sum of the expected return, plus unpredictable noise. If the expected returns are persistent, even if stationary time series, there is still a risk of spurious regression in finite samples. Because the unpredictable noise represents a substantial portion of the variance of stock returns, spurious regression effects, will differ for stock returns, from those in the classical setting of Granger and Newbold. This version of the spurious regression problem has received some attention in recent econometric studies, but more attention to the problem is probably due.

The interaction among the various statistical issues with predictive regressions has received relatively little research to date, and I would expect this to be an active field in the future. Ferson *et al.* (2003) study the interaction between data mining and spurious regression effects. They find that data mining for predictor variables interacts with spurious regression bias in equation (1.3). The two effects reinforce each other, because more highly persistent series are more likely to be found significant in the search for predictor variables. Simulations suggest that many of the regressions in the literature, based on individual predictor variables, may be spurious. Powell *et al.* (2006) extend the analysis of the interaction between data mining and spurious regression to conclude that recent studies of presidential regimes (Santa-Clara and Valkanov, 2003), the 'Halloween Indicator' (Bouman and Jacobsen, 2002) and business cycle effects in momentum (Chordia and Shivakumar, 2002) appear insignificant in view of their combined effects.

1.6 Perspective

Does the current body of evidence lead to the conclusion that there actually is predictability in stock returns? I think there are good reasons to be sceptical of predictability and good reasons to believe in predictability.

Why should we be sceptical? First, the logic of the old random walk = efficient markets literature is compelling to many. In that view, traders would bid up the prices of stocks with predictably high returns, thus lowering their returns and removing any predictability at the new price. Furthermore, data mining and selection bias conspire to make us see predictable patterns where none may exist. There are many choices for the return horizon, H. It need not be the same on both sides of the regression in equation (1.2). Jegadeesh (1990) provides a statistical analysis of the choice of horizons in this context. However, it is the number of choices that leads to scepticism.

For another example, studies of relative predictability, such as momentum, have sliced and diced common stocks into portfolios based on many characteristics of the data, at which point the effect often retreats into subsets of stocks, subperiods of time, phases of the business cycle or other parts of the data. Studies find momentum to be concentrated according to industries, the size of the firm (more momentum in large stocks), the price of the share (more when the price is above $5 per share), etc. The effect does not appear prior to 1940, appears stronger after 1968, and appears stronger during economic expansions than contractions.

The evidence for long-term return reversals has a similar problem. Reversals have been found to be concentrated in small stocks, low-priced stocks, the month of January, high idiosyncratic risk stocks, and to be more pronounced in earlier samples than in more recent data. For each approach to slicing and dicing the data there is a clever story. That is not bad. By digging into subsamples it should be possible in principle, to isolate what is driving an effect. But many of the patterns, especially those documented in the weak-form predictability literature, appear to be sample-specific. Finding results that vary with the sample period stimulates research featuring structural breaks and regime shifts. This reader is left more with concerns about naive data mining, improper multiple comparisons and statistical issues than with an understanding of why and where these weak-form effects occur systematically.

Predictive regressions are subject to a host of statistical issues. There are finite sample biases and problems associated with structural breaks and regime shifts. It is hard to get reliable standard errors for the regressions. There are potential spurious regression problems. And, as we are now beginning to understand, these effects can interact with each other. If semi-strong form predictability is spurious, as a result of statistical bias and naive data mining, we would expect predictor variables to appear in the empirical literature, then fail to work with fresh data. To some extent, the literature has evolved in this way.

There are also many good reasons to believe that stock returns are predictable. First of all, theory suggests that some amount of predictability is likely. If expected returns vary over time with some degree of persistence, predictability is expected. Most people find it easy to believe that expected stock returns and risks might be different coming out of a recession, for example, from going into one, and the predictability evidence tends to confirm such commonsense patterns. Studies find that predictability using lagged variables is largely explained by asset pricing models with multiple risk factors, if they allow the premiums associated with those risks to vary over time (see, for example, Ferson and Harvey, 1991; Ferson and Korajczyk, 1995; Avramov and Chordia, 2006). The behavioural models of predictability are compelling to many, as we see ourselves making the same cognitive errors as made by the agents in those models.

The momentum effect has been found to hold 'out of sample', relative to the original study of Jegadeesh and Titman (1993). Jegadeesh and Titman (2001) find momentum in data for 1990–1998. Momentum is found in stock markets outside of the US by Rouwehorst (1998) and Chui *et al.* (2000), among others.

The evidence of semi-strong form predictability has also survived a number of out-of-sample tests, working in other countries and over different time periods. Many of the variables identified as predictors for stock returns also seem to have some predictive power for other types of securities, such as bonds and futures, and also for the growth rates in 'fundamental' macroeconomic data. The evidence for semi-strong form predictability has survived corrections for a host of statistical and data problems.

Some studies that find semi-strong form stock market predictability, measured directly using lagged variables, has weakened in recent samples. It may be that the predictability was never really there, or that it was 'real' when first publicized, but diminished as traders attempted to exploit it. Ferson *et al.* (2005) examine semi-strong form predictability by regressing individual stocks on firm-specific predictors, then measuring the average covariances of the fitted values. It can be shown that the variance of the expected return on a large portfolio is approximately the average covariance. They find no evidence that predictability, measured in this way, is weaker in recent subperiods. As the firm-specific predictors have not been examined extensively in the literature, they may be less subject to naive data mining biases. Campbell and Thompson (2005), using step-ahead tests, also find that semi-strong form predictability holds up in recent data.

1.7 Conclusion

The issue of predictability in stock returns has important and broad economic implications. For example, it relates to the efficiency of capital markets in allocating resources to their highest valued uses. However, the interpretation of predictability, and the evidence for its very existence, remains controversial. This review of the literature finds the evidence for weak-form predictability (using the information in past stock prices) to be more fragile and less compelling than the evidence for semi-strong form predictability (using publicly available information more generally).

For the field of financial economics and asset pricing in particular, allowing for predictability through time-variation in expected returns, risk measures and volatility has been one of the most significant developments of the past two decades. Such conditional asset pricing models have provided a rich setting for the study of the dynamic behaviour of asset markets. For example, Conditional Performance Evaluation is the application of these models to the problem of evaluating the performance of portfolio managers. Models that allow for time-varying conditional moments produce different inferences about performance than do the traditional measures that do not allow for predictability (e.g. Ferson and Schadt, 1996), and they have influenced both academic views and professional investment practice. Research on predictability has stimulated numerous advances in the statistical and econometric methods of financial economics. The interactions between statistical biases and data mining in stock return studies should be a fruitful area for future research. Research on predictability in asset markets is likely to continue, and remain both useful and controversial for some time.

Acknowledgements

The author acknowledges financial support from the Collins Chair in Finance at Boston College.

References

Amihud, Y. and Horvitz, C. (2004). Predictive regressions: a reduced-bias estimation method, *Journal of Financial and Quantitative Analysis*, 39:813–842.

Ang, A., Chen, J. and Xing, Y. (2001). Downside risk and the momentum effect. NBER Working Paper No. 8643.

Avramov, D. and Chordia, T. (2006). Asset pricing models and financial market anomalies. *Review of Financial Studies*, 19:1001–1040.

Baker, M. and Wurgler, J. (2000). The equity share in new issues and aggregate stock returns. *Journal of Finance*, 55:2219–2257.

Barberis, N., Shleifer, A. and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, 49:307–43.

Berk, J. (1995). A critique of size-related anomalies. *Review of Financial Studies*, 8:275–286.

Boudoukh, J. and Richardson, M. (1994). The statistics of long-horizon regressions revisited. *Mathematical Finance*, 4:103–120.

Boudoukh, J., Richardson, M. and Whitelaw, R. (1994). A tale of three schools: insights on the autocorrelations of short-horizon returns. *Review of Financial Studies*, 7:539–573.

Boudoukh, J., Richardson, M. and Whitelaw, R. (2005). The myth of long-horizon predictability. NBER Working Paper No. 11 841.

Bouman, S. and Jacobsen, B. (2002). The Halloween indicator: sell in May and go away: another puzzle. *American Economic Review*, 92:1618–1635.

Breen, W., Glosten, L. and Jagannathan, R. (1993). On the relation between the volatility and the expected value of the nominal excess returns on stocks. *Journal of Finance*, XLVII:1779–1801.

Campbell, J. Y. (1993). Intertemporal asset pricing without consumption data. *American Economic Review*, 83:487–512.

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

Campbell, J. Y and Thompson, S. (2005). Predicting the equity premium out of sample: can anything beat the historical average? NBER Working Paper No. 11 468.

Campbell, J. Y. and Viceira, L. (2002). Strategic Asset Allocation: Portfolio Choice for Long-term Investors. New York, NY: Oxford University Press.

Chordia, T. and Shivakumar, L. (2002). Momentum, business cycle and time-varying expected returns. *Journal of Finance*, 57:985–1019.

Chui, A., Titman, S. and Wei, K. C. J. (2000). Momentum, ownership structure and financial crises: an analysis of Asian stock markets. Working Paper, University of Texas at Austin.

Cochrane, J. H. (2005). *Asset Pricing*. Princeton, NJ: Princeton University Press.

Cochrane, J. H. (2006). The dog that did not bark: a defense of return predictability. NBER Working Paper No. 12 026.

Cochrane, J. Y. (1988). How big is the random walk in GNP? *Journal of Political Economy*, 96: 893–920.

Conrad, J. and Kaul, G. (1988). Time-variation in expected returns. *Journal of Business*, 61:409–425.

Conrad, J. and Kaul, G. (1989). Mean reversion in short-horizon expected returns. *Review of Financial Studies*, 2:225–240.

Daniel, K., Hirshleifer, D. and Subrahmanyam, A. (1998). Investor psychology and security market under and over-reaction. *Journal of Finance*, 53:1839–1885.

DeBondt, W. and Thaler, R. (1985). Does the stock market overreact? *Journal of Finance*, 40:793–805.

Dybvig, P. H. and Ingersoll, J. (1982). Mean variance theory in complete markets. *Journal of Business*, 55:233–252.

Engle, R. (2004). Risk and volatility: econometric modelling and financial market practice. *American Economic Review*, 94:405–420.

Fama, E. F. (1970). Efficient capital markets: a review of theory and empirical work. *Journal of Finance*, 25:383–417.

Fama, E. and French, K. (1988). Permanent and temporary components of stock prices. *Journal of Political Economy*, 96:246–273.

Fama, E. and French, K. (1989). Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics*, 25:23–49.

Fama, E. F. and Schwert, G. W. (1977). Asset returns and inflation. *Journal of Financial Economics*, 5:115–146.

Ferson, W. E. (1989). Changes in expected security returns, risk and the level of interest rates. *Journal of Finance*, 44:1191–1217.

Ferson, W. E. (1995). Theory and empirical testing of asset pricing models. In: R. Jarrow, V. Maksimovic and B. Ziemba (eds), *Finance, Handbooks in Operations Research and Management Science*. Amsterdam: Elsevier, pp. 145–200.

Ferson, W. E. and Harvey, C. R. (1991). The variation of economic risk premiums. *Journal of Political Economy*, 99:385–415.

Ferson, W. E. and Harvey, C. R. (1993). The risk and predictability of international equity returns. *Review of Financial Studies*, 6:527–566.

Ferson, W. E. and Harvey, C. R. (1999). Conditioning variables and cross-section of stock returns. *Journal of Finance*, 54:1325–1360.

Ferson, W. and Korajczyk, R. (1995). Do arbitrage pricing models explain the predictability of stock returns? *Journal of Business*, 68:309–349.

Ferson, W. E. and Schadt, R. W. (1996). Measuring fund strategy and performance in changing economic conditions. *Journal of Finance*, 51:425–462.

Ferson, W. E., Heuson, A. and Su, T. (2005). Weak and semi-strong form stock return predictability revisited. *Management Science*, 51:1582–1592.

Ferson, W. E., Sarkissian, S. and Simin, T. (2003). Spurious regressions in financial economics? NBER Working Paper No. W9143.

Fisher, L. (1966). Some new stock market indexes. *Journal of Business*, 39:191–225.

Fleming, J., Kirby, C. and Ostdiek, B. (2001). The economic value of volatility timing. *Journal of Finance*, 56:329–352.
Foster, D., Smith, T. and Whaley, R. (1997). Assessing goodness-of-fit of asset pricing models: the distribution of the maximal R-squared. *Journal of Finance*, 52:591–607.
French, K., Schwert, G. W. and Stambaugh, R. (1987). Expected stock returns and volatility. *Journal of Financial Economics*, 19:3–29.
Friedman, M. (1953). The case for flexible exchange rates. In: *Essays in Positive Economics*. Chicago, IL: University of Chicago Press, pp. 157–203.
Gordon, M. J. (1962). *The Investment, Financing and Valuation of the Corporation*. Homewood, IL: Irwin.
Goyal, A. and Santa Clara, P. (2003) Idiosyncratic risk matters. *Journal of Finance*, 43:975–1007.
Goyal, A. and Welch, I. (2003). Predicting the equity premium with dividend ratios. *Management Science*, 49:639–654.
Goyal, A. and Welch, I. (2004). A comprehensive look at the empirical performance of equity premium prediction. NBER Working Paper No. 10 483.
Granger, C. W. K. and Newbold, P. (1974). Spurious regressions in econometrics. *Journal of Econometrics*, 2:111–120.
Greenwood, R. (2004). Aggregate corporate liquidity and stock returns. Working Paper, Harvard University.
Grinblatt, M. and Han, B. (2003). The disposition effect and momentum. NBER Working Paper No. 8734.
Grinblatt, M. and Moskowitz, T. (2003). Predicting price movements from past returns: the role of consistency in tax-loss selling. *Journal of Financial Economics*, 71:541–579.
Harvey, C. R. (1991). The world price of covariance risk. *Journal of Finance*, 46:111–57.
Harvey, C. R. and Siddique, A. (2001). Autoregressive conditional skewness. *Journal of Financial and Quantitative Analysis*, 34:465–487.
Haugen, R. and Lakonishok, J. (1988). *The Incredible January Effect*. Homewood, IL: Dow Jones-Irwin.
Hjalmarsson, E. (2006). Should we expect significant out-of-sample results when predicting stock returns? Working Paper, Federal Reserve Board International Finance Discussion Paper No. 855.
Hodrick, R. (1992). Dividend yields and expected stock returns: alternative preocedures for inference and measurement. *Review of Financial Studies*, 5:257–286.
Hong, H. and Stein, J. (1999). A unified theory of underreaction, momentum trading and overreaction in asset markets. *Journal of Finance*, 45:265–295.
Jegadeesh, N. (1990). Evidence of predictable behavior of security returns. *Journal of Finance*, 45:881–898.
Jegadeesh, N. and Titman, S. (1993). Returns to buying winners and selling losers: implications for stock market efficiency. *Journal of Finance*, 48:65–91.
Jegadeesh, N. and Titman, S. (2001). Profitability of momentum portfolios: an evaluation of alternative explanations. *Journal of Finance*, 56:699–720.
Kandel, S. and Stambaugh, R. (1990). Expectations and volatility of consumption and asset returns. *Review of Financial Studies*, 3:207–232.
Kandel, S. and Stambaugh, R. (1996). On the predictability of stock returns: an asset allocation perspective. *Journal of Finance*, 51:385–424.
Kaul, G. (1996). Predictable components in stock returns. In: G. S. Maddala and C. R. Rao (eds), *Handbook of Statistics 14*. Amsterdam: Elsevier, pp. 269–96.
Keim, D. and Stambaugh, R. (1986). Predicting returns in the stock and bond markets. *Journal of Financial Economics*, 17:357–390.
Kendal, M.G. (1954). A note on bias in the estimation of autocorrelation. *Biometrika*, 41:403–404.
Korajczyk, R. and Sadka, R. (2004). Are momentum profits robust to trading costs? *Journal of Finance*, 59:1039–1082.
Kothari, S. P. and Shanken, J. (1997). Book-to-market, dividend yield and expected market returns: a time-series analysis. *Journal of Financial Economics*, 44:169–203.
Lamont, O. (2000). Investment plans and stock returns. *Journal of Finance*, 40(6):2719–2745.
Lehmann, B. N. (1990). Fads, martingales and market efficiency. *Quarterly Journal of Economics*, 105:1–28.
Lei, Q. (2006). Pricing dynamics of dividend and nondividend payout yields: on the stock return predictability. Working Paper, Southern Methodist University, Dallas.
Lesmond, D., Schill, M. and Zhou, C. (2004). The illusory nature of momentum profits. *Journal of Financial Economics*, 71:349–380.
Lettau, M. and Ludvigson, S. (2001). Consumption, aggregate wealth and expected stock returns. *Journal of Finance*, 56:815–626.

Lo, A. and MacKinlay, A. C. (1988). Stock market prices do not follow random walks: evidence from a simple specification test. *Review of Financial Studies*, 1:41–66.

Lo, A. and MacKinlay, A. C. (1990). An econometric analysis of infrequent trading. *Journal of Econometrics*, 45:181–211.

Malkiel, B. (2004). Models of stock market predictability. *Journal of Financial Research*, 28:449–459.

McLean, D. (2006). Why do momentum and long-term reversals persist? An evaluation of mispricing explanations. Working Paper, University of Alberta.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica*, 41:867–887.

Merton, R. C. (1980). On estimating the expected return on the market: an exploratory investigation. *Journal of Financial Economics*, 8:323–361.

Muthaswamy, J. (1988). Asynchronous closing prices and spurious correlation in portfolio returns. Working Paper, University of Chicago.

Pesaran, M. H. and Timmermann, A. (1995). Predictability of stock returns: robustness and economic significance. *Journal of Finance*, 50:1201–1228.

Pontiff, J. and Schall, L. (1998). Book-to-market as a predictor of market returns. *Journal of Financial Economics*, 49:141–160.

Pontiff, J. and Woodgate, A. (2006). Share issuance and cross-sectional stock returns. Working Paper, Boston College.

Powell, J. G, Shi, J., Smith, T. and Whaley, R. (2006). Political regimes, business cycles, seasonalities and returns. Working paper, Vanderbilt University, Nashville.

Richardson, M. (1993). Temporary components of stock prices: a skeptic's view. *Journal of Business and Economic Statistics*, 11:199–207.

Roll, R. R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39:1127–1140.

Rouwehorst, G. (1998). International momentum portfolios. *Journal of Finance*, 53:267–284.

Rozeff, P. (1984). Dividend yields are equity premiums. *Journal of Portfolio Management*, 1(1):68–75.

Samuelson, P. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, 6(2):41–49.

Santa-Clara, P. and Valkanov, R. (2003). The presidential puzzle: political cycles and the stock market. *Journal of Finance*, 58:1841–1872.

Scholes, M. and Williams, J. (1977). Estimating beta from nonsynchronous data. *Journal of Financial Economics*, 5:309–327.

Schwert, G. W. (2003). Anomalies and market efficiency. In: G. M. Constantinides, M. Harris and R. M. Stulz (eds), *Handbook of the Economics of Finance*. Amsterdam: Elsevier, pp. 575–603.

Scruggs, J. (1998). Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: a 2-factor approach. *Journal of Finance*, 53:575–603.

Sharpe, W. F. (1964). Capital asset prices: a theory of market equilibrium under conditions of risk. *Journal of Finance*, 19:425–442.

Simin, T. (2006). The (poor) predictive performance of asset pricing models. *Journal of Financial and Quantitative Analysis*, forthcoming.

Solnik, B. (1993). The unconditional performance of international asset allocation strategies using conditioning information. *Journal of Empirical Finance*, 1:33–55.

Stambaugh, R. S. (1999). Predictive regressions. *Journal of Financial Economics*, 54:375–421.

White, H. (2000). A reality check for data snooping. *Econometrica*, 68:1097–1126.

Working, H. (1940). The investigation of economic expectations. *American Economic Review*, 39:150–166.

Yule, G. (1926). Why do we sometimes get nonsense correlations between time series? A study in sampling and the nature of time-series. *Journal of the Royal Statistical Society*, 89:1–64.

Zarowin, P. (1990). Size, seasonality and stock market overreaction. *Journal of Financial and Quantitative Analysis*, 25:113–125.