

IZA DP No. 5037

Predictability of Asset Returns and the Efficient Market Hypothesis

M. Hashem Pesaran

June 2010

Predictability of Asset Returns and the Efficient Market Hypothesis

M. Hashem Pesaran

*University of Cambridge,
USC and IZA*

Discussion Paper No. 5037
June 2010

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Predictability of Asset Returns and the Efficient Market Hypothesis*

This paper is concerned with empirical and theoretical basis of the Efficient Market Hypothesis (EMH). The paper begins with an overview of the statistical properties of asset returns at different frequencies (daily, weekly and monthly), and considers the evidence on return predictability, risk aversion and market efficiency. The paper then focuses on the theoretical foundation of the EMH, and show that market efficiency could co-exist with heterogeneous beliefs and individual irrationality so long as individual errors are cross sectionally weakly dependent in the sense defined by Chudik, Pesaran, and Tosetti (2010). But at times of market euphoria or gloom these individual errors are likely to become cross sectionally strongly dependent and the collective outcome could display significant departures from market efficiency. Market efficiency could be the norm, but it is likely to be punctuated with episodes of bubbles and crashes. The paper also considers if market inefficiencies (assuming that they exist) can be exploited for profit.

JEL Classification: G12, G14

Keywords: market efficiency, predictability, heterogeneity of expectations, forecast averaging, equity premium puzzle

Corresponding author:

M. Hashem Pesaran
Faculty of Economics
University of Cambridge
Sidgwick Avenue
Cambridge, CB3 9DD
United Kingdom
E-mail: mhp1@econ.cam.ac.uk

* I am grateful to Elisa Tosetti for valuable help with the preparation of this paper, and to Alex Chudik, Bill Janeway, Ron Smith, Ansgar Walther, and Aman Ullah for helpful comments. Part of this paper is based on my presentation at the CFS symposium "Market Efficiency Today" held in Frankfurt/Main on October 6, 2005 in honor of Eugene F. Fama.

1 Introduction

Economists have long been fascinated by the nature and sources of variations in the stock market. By the early 1970's a consensus had emerged among financial economists suggesting that stock prices could be well approximated by a random walk model and that changes in stock returns were basically unpredictable. Fama (1970) provides an early, definitive statement of this position. Historically, the 'random walk' theory of stock prices was preceded by theories relating movements in the financial markets to the business cycle. A prominent example is the interest shown by Keynes in the variation in stock returns over the business cycle.

The efficient market hypothesis (EMH) evolved in the 1960's from the random walk theory of asset prices advanced by Samuelson (1965). Samuelson showed that in an informationally efficient market price changes must be unforecastable. Kendall (1953), Cowles (1960), Osborne (1959), Osborne (1962), and many others had already provided statistical evidence on the random nature of equity price changes. Samuelson's contribution was, however, instrumental in providing academic respectability for the hypothesis, despite the fact that the random walk model had been around for many years; having been originally discovered by Louis Bachelier, a French statistician, back in 1900.

Although a number of studies found some statistical evidence against the random walk hypothesis, these were dismissed as economically unimportant (could not generate profitable trading rules in the presence of transaction costs) and statistically suspect (could be due to data mining). For example, Fama (1965), concluded that "...there is no evidence of important dependence from either an investment or a statistical point of view.". Despite its apparent empirical success, the random walk model was still a statistical statement and not a coherent theory of asset prices. For example, it need not hold in markets populated by risk averse traders, even under market efficiency.

There now exist many different versions of the EMH, and one of the aims of this paper is to provide a simple framework where alternative versions of the EMH can be articulated and discussed. We begin with an overview of the statistical properties of asset returns at different frequencies (daily, weekly and monthly), and consider the evidence on return predictability, risk aversion and market efficiency. We then focus on the theoretical foundation of the EMH, and show that market efficiency could co-exist with heterogeneous beliefs and individual 'irrationality', so long as individual errors are cross sectionally weakly dependent in the sense defined by Chudik, Pesaran, and Tosetti (2010). But at times of market euphoria or gloom these individual errors are likely to become cross sectionally strongly dependent and the collective outcome

could display significant departures from market efficiency. Market efficiency could be the norm, but most likely it will be punctuated by episodes of bubbles and crashes. To test for such episodes we argue in favour of compiling survey data on individual expectations of price changes that are combined with information on whether such expectations are compatible with market equilibrium. A trader who believes that asset prices are too high (low) might still expect further price rises (falls). Periods of bubbles and crashes could result if there are sufficiently large numbers of such traders that are prepared to act on the basis of their beliefs. The paper also considers if periods of market inefficiency can be exploited for profit. We conclude with some general statements on new research directions.

We begin with some basic concepts and set out how returns are computed over different horizons and assets, and discuss some of the known stylized facts about returns by means of simple statistical models.

2 Prices and returns

2.1 Single period returns

Let P_t be the price of a security at date t . The absolute price change over the period $t - 1$ to t is given by $P_t - P_{t-1}$, the relative price change by

$$R_t = (P_t - P_{t-1})/P_{t-1}$$

the gross return (excluding dividends) on security by

$$1 + R_t = P_t/P_{t-1}$$

and the log price change by

$$r_t = \Delta \ln(P_t) = \ln(1 + R_t)$$

It is easily seen that for *small relative price changes* the log-price change and the relative price change are almost identical.

In the case of daily observations when dividends are negligible, $100 \cdot R_t$ measures the *per cent return* on the security, and $100 \cdot r_t$ is the continuously compounded return. R_t is also known as discretely compounded return. The continuously compounded return, r_t , is particularly convenient in the case of temporal aggregation (multi-period returns - see below), while the discretely compounded returns are convenient for use in cross-sectional aggregation, namely aggregation of returns across different

instruments in a portfolio. For example, for a portfolio composed of N instruments with weights $w_{i,t-1}$, ($\sum_{i=1}^N w_{i,t-1} = 1, w_{i,t-1} \geq 0$) we have

$$R_{pt} = \sum_{i=1}^N w_{i,t-1} R_{it}, \quad (\text{per cent return})$$

$$r_{pt} = \ln\left(\sum_{i=1}^N w_{i,t-1} e^{r_{it}}\right), \quad (\text{continuously compounded})$$

Often r_{pt} is approximated by $\sum_{i=1}^N w_{i,t-1} r_{it}$.

When dividends are paid out we have

$$R_t = (P_t - P_{t-1})/P_{t-1} + D_t/P_{t-1}$$

$$\approx \Delta \ln(P_t) + D_t/P_{t-1}$$

where D_t is the dividend paid out during the holding period.

2.2 Multi-period returns

Single-period price changes (returns) can be used to compute multi-period price changes or returns. Denote the return over the most recent h periods by $R_t(h)$ then (abstracting from dividends)

$$R_t(h) = \frac{P_t - P_{t-h}}{P_{t-h}}$$

or

$$1 + R_t(h) = P_t/P_{t-h}$$

and

$$r_t(h) = \ln(P_t/P_{t-h}) = r_t + r_{t-1} + \dots + r_{t-h+1}$$

where r_{t-i} , $i = 0, 1, 2, \dots, h-1$ are the single-period returns. For example, weekly returns is defined by $r_t(5) = r_t + r_{t-1} + \dots + r_{t-4}$. Similarly, since there are 25 business days in one month, then the 1-month return can be computed as the sum of the last 25 1-day returns, or $r_t(25)$.

2.3 Overlapping returns

Note that multi-period returns have overlapping daily observations. In the case of weekly returns, $r_t(5)$ and $r_{t-1}(5)$ have the four daily returns, $r_{t-1} + r_{t-2} + r_{t-3} + r_{t-4}$

in common. As a result the multi-period returns will be serially correlated even if the underlying daily returns are not serially correlated. One way of avoiding the overlap problem would be to sample the multi-period returns h periods apart. But this is likely to be inefficient as it does not make use of all available observations. A more appropriate strategy would be to use the overlapping returns but allow for the fact that this will induce serial correlations. For further details see Pesaran, Pick, and Timmerman (2010).

3 Statistical models of returns

A simple model of returns (or log price changes) is given by

$$\begin{aligned} r_{t+1} &= \Delta \ln(P_{t+1}) = p_{t+1} - p_t \\ &= \mu_t + \sigma_t \varepsilon_{t+1}, \quad t = 1, 2, \dots, T \end{aligned} \quad (1)$$

where μ_t and σ_t^2 are the conditional mean and the conditional variance of returns (with respect to the information set Ω_t available at time t) and ε_{t+1} represents the unpredictable component of return. Two popular distributions for ε_{t+1} are

$$\begin{aligned} \varepsilon_{t+1} | \Omega_t &\sim IID \mathcal{Z} \\ \varepsilon_{t+1} | \Omega_t &\sim \left(\sqrt{\frac{v-2}{v}} \right) IID \mathcal{T}_v \end{aligned}$$

where $\mathcal{Z} \sim N(0, 1)$ stands for a standard normal distribution, and \mathcal{T}_v stands for Student's t with v degrees of freedom. Unlike the normal distribution that has moments of all orders, \mathcal{T}_v only has moments of order $v - 1$ and smaller. For the Student's t to have a variance, for example, we need $v > 2$.

Since $r_{t+1} = \ln(1 + R_{t+1})$, where $R_{t+1} = (P_{t+1} - P_t)/P_t$, it then follows that under $\varepsilon_{t+1} | \Omega_t \sim IID \mathcal{Z}$, the price level, P_{t+1} conditional on Ω_t will be lognormally distributed. Note that $\Omega_t = (P_t, P_{t-1}, \dots)$ and $\Omega_t = (r_t, r_{t-1}, \dots)$ convey the same information and are equivalent. Hence, $P_{t+1} = P_t \exp(r_{t+1})$, and we have¹

$$\begin{aligned} E(P_{t+1} | \Omega_t) &= P_t E(\exp(r_{t+1}) | \Omega_t) \\ &= P_t \exp\left(\mu_t + \frac{1}{2}\sigma_t^2\right) \end{aligned}$$

¹Using properties of the moment generating function of normal variates, if $x \sim N(\mu_x, \sigma_x^2)$ then, $E[\exp(x)] = \exp(\mu_x + .5\sigma_x^2)$.

Similarly,

$$Var(P_{t+1} | \Omega_t) = P_t^2 \exp(2\mu_t + \sigma_t^2) [\exp(\sigma_t^2) - 1]$$

In practice, it is much more convenient to work with log returns, r_{t+1} , rather than asset prices.

The probability density functions of \mathcal{Z} and \mathcal{T}_v are given by

$$f(\mathcal{Z}) = (2\pi)^{-1/2} \exp\left[-\frac{\mathcal{Z}^2}{2}\right], \quad -\infty < \mathcal{Z} < \infty \quad (2)$$

and

$$f(\mathcal{T}_v) = \frac{1}{\sqrt{v}B(v/2, 1/2)} \left[1 + \frac{\mathcal{T}_v^2}{v}\right]^{-(v+1)/2} \quad (3)$$

where $-\infty < \mathcal{T}_v < \infty$, and $B(v/2, 1/2)$ is the beta function defined by

$$B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)}, \quad \Gamma(\alpha) = \int_0^\infty u^{\alpha-1} e^{-u} du$$

It is easily seen that

$$E(\mathcal{T}_v) = 0, \text{ and } Var(\mathcal{T}_v) = \frac{v}{v-2}$$

A large part of financial econometrics is concerned with alternative ways of modelling the conditional mean (mean returns), μ_t , the conditional variance (asset return volatility), σ_t , and the cumulative probability distribution of the errors, ε_{t+1} . A number of issues need to be addressed in order to choose an adequate model. In particular:

- Is the distribution of returns normal?
- Is the distribution of returns constant over time?
- Are returns statistically independent over time?
- Are squares or absolute values of returns independently distributed over time?
- What are the cross correlation of returns on different instruments?

The above modelling issues can be readily extended to the case where we are concerned with a vector of asset returns, $\mathbf{r}_t = (r_{1t}, r_{2t}, \dots, r_{mt})'$. In this case we also need to model the pair-wise conditional correlations of asset returns, namely

$$Corr(r_{it}, r_{jt} | \Omega_t) = \frac{Cov(r_{it}, r_{jt} | \Omega_t)}{\sqrt{Var(r_{it} | \Omega_t)Var(r_{jt} | \Omega_t)}}$$

Typically the conditional variances and correlations are modelled using exponential smoothing procedures or the multivariate generalized autoregressive conditional heteroskedastic models developed in the econometric literature.

3.1 Percentiles, critical values and Value at Risk

Suppose a random variable r (say daily returns on an instrument) has the probability density function $f(r)$. Then the p th percentile of the distribution of r , denoted by C_p , is defined as that value of return such that p percent of the returns fall below it. Mathematically we have

$$p = \Pr(r < C_p) = \int_{-\infty}^{C_p} f(r)dr$$

In the literature on risk management C_p is used to compute ‘Value at Risk’ or *VaR* for short. For $p = 1\%$, C_p associated with the one-sided critical value of the normal distribution is given by -2.33σ , where σ is the standard deviation of returns.

In hypothesis testing C_p is known as the critical value of the test associated with a (one-sided) test of size p . In the case of two-sided tests of size p , the associated critical value is computed as $C_{p/2}$.

3.2 Measures of departure from normality

The normal probability density function for r_{t+1} conditional on the information at time t , Ω_t , is given by

$$f(r_{t+1}) = (2\pi\sigma_t^2)^{-1/2} \exp\left[-\frac{1}{2\sigma_t^2}(r_{t+1} - \mu_t)^2\right]$$

with $\mu_t = E(r_{t+1} | \Omega_t)$ and $\sigma_t^2 = E[(r_{t+1} - \mu_t)^2 | \Omega_t]$ being the conditional mean and variance. If the return process is stationary, unconditionally we also have $\mu = E(r_{t+1})$, and $\sigma^2 = E[(r_{t+1} - \mu)^2]$.

Skewness and tail-fatness measures are defined by

$$\text{Skewness} = \sqrt{b_1} = m_3/m_2^{3/2}$$

$$\text{Kurtosis} = b_2 = m_4/m_2^2$$

where

$$m_j = \frac{\sum_{t=1}^T (r_t - \bar{r})^j}{T}, \quad j = 2, 3, 4$$

Table 1: Descriptive statistics for daily returns on SP 500, FTSE 100, German DAX, and Nikkei 225

Variables	SP	FTSE	DAX	NK
Maximum	14.11	10.05	12.83	20.70
Minimum	-9.88	-9.24	-8.89	-13.07
Mean (\bar{r})	-0.01	-0.01	-0.01	-0.01
S. D. ($\hat{\sigma}$)	1.39	1.33	1.65	1.68
Skewness ($\sqrt{b_1}$)	0.35	0.06	0.24	0.16
Kurtosis (b_2)	14.3	9.7	8.5	17.8
JB statistic	13453.6	4713.1	3199.2	23000.8

For a normal distribution $\sqrt{b_1} \approx 0$, and $b_2 \approx 3$. In particular

$$\hat{\mu} = \bar{r} = \sum_{t=1}^T r_t / T, \quad \hat{\sigma} = \sqrt{\frac{\sum_{t=1}^T (r_t - \bar{r})^2}{T - 1}}$$

The Jarque-Bera's test statistic for departure from normality is given by, see (Jarque and Bera 1980)

$$JB = T \left\{ \frac{1}{6} b_1 + \frac{1}{24} (b_2 - 3)^2 \right\}$$

Under the joint null hypothesis that $b_1 = 0$ and $b_2 = 3$, the JB statistic is asymptotically distributed (as $T \rightarrow \infty$) as a chi-squared with 2 degrees of freedom, χ_2^2 . Therefore, a value of JB in excess of 5.99 will be statistically significant at the 95 per cent confidence level, and the null hypothesis of normality will be rejected.

4 Empirical evidence: statistical properties of returns

Table 1 gives a number of statistics for daily returns ($\times 100$) on four main equity index futures, namely S&P 500 (SP), FTSE 100 (FTSE), German DAX (DAX), and Nikkei 225 (NK), over the period 3-Jan-00 to 31-Aug-09 (for a total of 2,519 observations).

The kurtosis coefficients are particularly large for all the four equity futures and exceed the benchmark value of 3 for the normal distribution. There is some evidence of positive skewness, but it is of second order importance as compared to the

magnitude of excess kurtosis coefficient given by, $b_2 - 3$. The large values of excess kurtosis is reflected in the huge values of the JB statistics reported in 1. Also under the assumption that returns are normally distributed, we would have expected the maximum and minimum of daily returns to fall (with 99% confidence) in the region of $\pm 2.33 \times$ S. D., which is ± 3.24 for SP500, as compared to the observed values of -9.88 and 14.11 . See also Figure 1.




Figure 1: Histogram and Normal curve for daily returns on SP500 (over the period 3-Jan-00 to 31-Aug-09)

The departure from normality is particularly pronounced over the past decade where markets have been subject to two important episodes of financial crises: the collapse of markets in 2000 after the dot-com bubble and the stock market crash of 2008 after the 2007 credit crunch. (see Figure 2).

However, the evidence of departure from normality can be seen in daily returns even before 2000. For example, over the period 3-Jan-94 to 31-Dec-99 (1565 daily observations) kurtosis coefficient of returns on SP500 was 9.5 which is still well above the benchmark value of 3. The recent financial crisis has accentuated the situation but can not be viewed as the cause of the observed excess kurtosis of equity returns.

Similar results are also obtained if we consider weekly returns. The kurtosis coefficients estimated using weekly returns over the period January 2000 to the end of August 2009 (504 weeks) were 12.4, 15.07, 8.9 and 15.2 for SP500, FTSE, DAX, and Nikkei, respectively. These are somewhat lower than the estimates obtained




Figure 2: Daily returns on SP500 (over the period 3-Jan-00 to 31-Aug-09)

using daily observations for SP500 and Nikkei, but are quite a bit higher for FTSE. For DAX daily and weekly observations yield a very similar estimate of the kurtosis coefficient.

For currencies the kurtosis coefficient of returns (measured in terms of US dollar) varies from 4.5 for euro to 13.8 for the Australian dollar. The estimates computed using daily observations over the period 03-Jan-00 to 31-Aug-09 are summarized in Table 2. The currencies considered are the British pound (GBP), euro (EU), Japanese yen (JPY), Swiss franc (CHF), Canadian dollar (CAD), and Australian dollar (AD), all measured in terms of US dollar.

The returns on government bonds are generally less fat-tailed than the returns on equities and currencies. But their distribution still show a significant degree of departure from normality.

Table 3 reports descriptive statistics on daily returns on the main four government bond futures: US T-Note 10Y (BU), Europe Euro Bund 10Y (BE), Japan Government Bond 10Y (BJ), and UK Long Gilts 8.75-13Y (BG) over the period 03-Jan-00 to 31-Aug-09.

It is clear that for all the three asset classes that there are significant departures from normality which needs to be taken into account when analysing financial time series.

Table 2: Descriptive statistics for daily returns on British pound, euro, Japanese yen, Swiss franc, Canadian dollar, and Australian dollar

Variables	<i>JPY</i>	<i>EU</i>	<i>GBP</i>	<i>CHF</i>	<i>CAD</i>	<i>AD</i>
Maximum	4.53	3.17	3.41	4.58	5.25	6.21
Minimum	-3.93	-3.01	-5.04	-3.03	-3.71	-9.50
Mean (\bar{r})	-.006	.016	.007	.012	.013	.022
S. D. ($\hat{\sigma}$)	.65	.65	.60	.70	.59	.90
Skewness ($\sqrt{b_1}$)	-.28	.01	-.35	.12	.09	-.76
Kurtosis (b_2)	5.99	4.5	7.2	4.9	9.1	13.8

Table 3: Descriptive statistics for daily returns on US T-Note 10Y, Europe Euro Bund 10Y, Japan Government Bond 10Y, and, UK Long Gilts 8.75-13Y

Variables	BU	BE	BG	BJ
Maximum	3.63	1.48	2.43	1.53
Minimum	-2.40	-1.54	-1.85	-1.41
Mean (\bar{r})	.0	.01	.01	.01
S. D. ($\hat{\sigma}$)	.43	.32	.35	.24
Skewness ($\sqrt{b_1}$)	-.004	-.18	.02	-.18
Kurtosis (b_2)	6.67	4.49	6.02	6.38

4.1 Other stylized facts about asset returns

Asset returns are typically uncorrelated over time, are difficult to predict and as we have seen tend to have distributions that are fat-tailed. In contrast the absolute or squares of asset returns (that measure risk), namely $|r_t|$ or r_t^2 , are serially correlated and tend to be predictable. It is interesting to note that r_t can be written as

$$r_t = \text{sign}(r_t) |r_t|$$

where $\text{sign}(r_t) = +1$ if $r_t > 0$ and $\text{sign}(r_t) = -1$ if $r_t \leq 0$. Since $|r_t|$ is predictable, it is, therefore, the non-predictability of $\text{sign}(r_t)$, or the direction of the market, which lies behind the difficulty of predicting returns.

The extent to which returns are predictable depends on the forecast horizon, the degree of market volatility, and the state of the business cycle. Predictability tends to rise during crisis periods. Similar considerations also apply to the degree of fat-tailedness of the underlying distribution and the cross correlations of asset returns. The return distributions become less fat-tailed as the horizon is increased, and cross correlations of asset returns become more predictable with the horizon. Cross correlation of returns also tend to increase with market volatility. The analysis of time variations in the cross correlation of asset returns is beyond the scope of this paper. But the interested reader might wish to consult (Pesaran and Pesaran 2010) where multivariate conditional volatility models are fitted to weekly returns on equities, bonds and currencies.

In the case of daily returns, equity returns tend to be negatively serially correlated. During normal times they are small and only marginally significant statistically, but become relatively large and attain a high level of statistical significance during crisis periods. These properties are illustrated in the following empirical application.

The first and second order serial correlation coefficients of daily returns on SP500 over the period 3-Jan-00 to 31-Aug-07 are -0.015 (0.0224) and -0.0458 (0.0224), respectively, but increase to -0.068 (0.0199) and -0.092 (0.0200) once the sample is extended to the end of August 2009 which covers the 2008 global financial crisis.² Similar patterns are also observed for other equity indices. For currencies the evidence is more mixed. In the case of major currencies such as euro and yen there is little evidence of serial correlation in returns and this outcome does not seem much affected by whether one considers normal or crisis periods. For other currencies there is some evidence of negative serial correlation, particularly at times of crisis. For example, over the period 3-Jan-00 to 31-Aug-09 the first order serial correlation of daily

²The figures in brackets are standard errors.

returns on Australian dollar amounts to -0.056 (0.0199), but becomes statistically insignificant if we exclude the crisis period. There is also very little evidence of serial correlation in daily returns on the four major government bonds that we have been considering. This outcome does not depend on whether the crisis period is included in the sample. Irrespective of whether the underlying returns are serially correlated, their absolute values (or their squares) are highly serially correlated, often over many periods. For example, over the 3-Jan-00 to 31-Aug-09 period the first and second order serial correlation coefficients of absolute return on SP500 are $0.2644(0.0199)$, $0.3644(0.0204)$; for euro they are $0.0483(0.0199)$ and $0.1125(0.0200)$, and for US 10Y bond they are $0.0991(0.0199)$ and $0.1317(0.0201)$. The serial correlation in absolute returns tend to decay very slowly and continue to be statistically significant event after 120 trading days. See Figure 3.

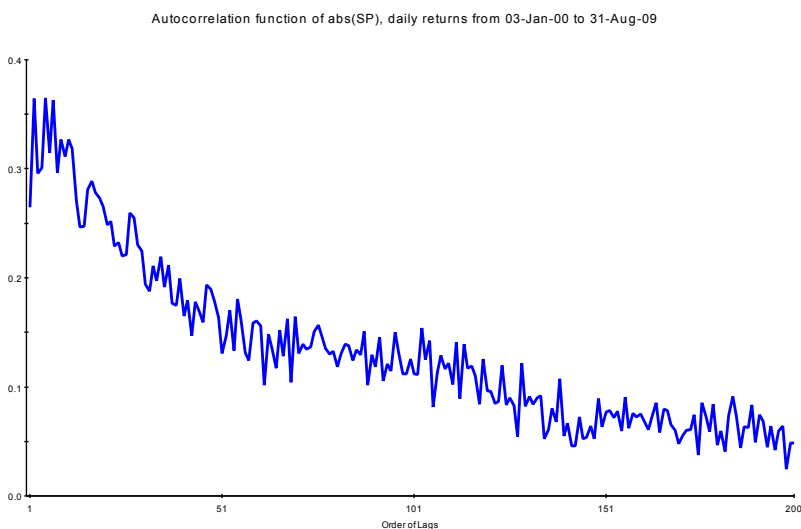


Figure 3: Autocorrelation function of the absolute values of SP500 (over the period 3-Jan-00 to 31-Aug-09)

It is also interesting to note that there is little correlation between r_t and $|r_t|$. Based on the full sample ending in August 2009, this correlation is $-.0003$ for SP500, 0.025 for euro, and 0.009 for the US 10Y bond.

4.2 Monthly stock market returns

Many of the regularities and patterns documented for returns using daily or weekly observations can also be seen in monthly observations, once a sufficiently long period is considered. For the US stock market long historical monthly data on prices and dividends are compiled by Shiller and can be downloaded from his homepage³. An earlier version of this data set has been analyzed in Shiller (2005). Monthly returns on SP500 (inclusive of dividends) is computed as

$$RSP_t = 100 \left(\frac{SP_t - SP_{t-1} + SPDIV_t}{SP_{t-1}} \right)$$

where SP_t is the monthly spot price index of SP500 and $SPDIV_t$ denotes the associated dividends on the SP500 index. Over the period 1871m1 to 2009m9 (a total of 1664 monthly observations) the coefficient of skewness and kurtosis of RSP amounted to 1.07 and 23.5 per cents, respectively. The excess kurtosis coefficient of 20.5 is much higher than the figure of 11.3 obtained for the daily observations on SP over the period 3-Jan-00 to 31-Aug-09. Also as before the skewness coefficient is relatively small. However, the monthly returns show a much higher degree of serial correlation and a lower degree of volatility as compared to daily or weekly returns. The correlation coefficients of RSP are 0.346 (0.0245) and 0.077 (0.027), and the serial correlation coefficients continue to be statistically significant up to the lag order of 12 months. Also, the pattern of serial correlations in absolute monthly returns, $|RSP_t|$, is not that different from that of the serial correlation in RSP_t , which suggests a lower degree of return volatility (as compared to the volatility of daily or weekly returns) once the effects of mean returns are taken into account.

Similar, but less pronounced, results are obtained if we exclude the 1929 stock market crash and focus on the post World War II period. The coefficients of skewness and kurtosis of monthly returns over the period 1948m1 to 2009m9 (741 observations) are -0.49 and 5.2, respectively. The first and second order serial correlation coefficients of returns are 0.361 (0.0367) and 0.165 (0.041), respectively. The main difference between these sub-sample estimates and those obtained for the full sample is the much lower estimate for the kurtosis coefficient. But even the lower post 1948 estimates suggest a significant degree of fat-tailedness in the monthly returns.

³See <http://www.econ.yale.edu/~shiller/data.htm>.

5 Stock return regressions

Consider the linear excess return regression

$$R_{t+1} - r_t^f = a + b_1x_{1t} + b_2x_{2t} + \dots + b_kx_{kt} + \varepsilon_{t+1} \quad (4)$$

where R_{t+1} is the one-period holding return on an stock index, such as FTSE or Dow Jones, defined by

$$R_{t+1} = (P_{t+1} + D_{t+1} - P_t)/P_t \quad (5)$$

P_t is the stock price at the end of the period and D_{t+1} is the dividend paid out over the period t to $t + 1$, and x_{it} , $i = 1, 2, \dots, k$ are the factors/variables thought to be important in predicting stock returns. Finally, r_t^f is the return on the government bond with one-period to maturity (the period to maturity of the bond should be exactly the same as the holding period of the stock). $R_{t+1} - r_t^f$ is known as the excess return (return on stocks in excess of the return on the safe asset). Note also that r_t^f would be known to the investor/trader at the end of period t , before the price of stocks, P_{t+1} , is revealed at the end of period $t + 1$.

Examples of possible stock market predictors are past changes in macroeconomic variables such as interest rates, inflation, dividend yield (D_t/P_{t-1}), price earnings ratio, output growth, and term premium (the difference in yield of a high grade and a low grade bond such as AAA rated minus BAA rated bonds).

For individual stocks the relevant stock market regression is the capital asset pricing model (CAPM), augmented with potential predictors:

$$R_{i,t+1} = a_i + b_{1i}x_{1t} + b_{2i}x_{2t} + \dots + b_{ki}x_{kt} + \beta_i R_{t+1} + \varepsilon_{i,t+1} \quad (6)$$

where $R_{i,t+1}$ is the holding period return on asset i (shares of firm i), defined similarly as R_{t+1} . The asset-specific regressions (6) could also include firm specific predictors, such as R_{it} or its higher order lags, book-to-market value or size of firm i . Under market efficiency, as characterized by CAPM,

$$a_i = 0, b_{1i} = b_{2i} = \dots = b_{ki} = 0$$

and only the ‘betas’, β_i , will be significantly different from zero. Under CAPM, the value of β_i captures the risk of holding the share i with respect to the market.

6 Market efficiency and stock market predictability

It is often argued that if stock markets are efficient then it should not be possible to predict stock returns, namely that none of the variables in the stock market regression

(4) should be statistically significant. Some writers have even gone so far as to equate stock market efficiency with the non-predictability property. But this line of argument is not satisfactory and does not help in furthering our understanding of how markets operate. The concept of market efficiency needs to be defined separately from predictability. In fact, it is easily seen that stock market returns will be non-predictable only if market efficiency is combined with risk neutrality.

6.1 Risk neutral investors

Suppose there exists a risk free asset such as a government bond with a known payout. In such a case an investor with an initial capital of $\$A_t$, is faced with two options:

- Option 1: hold the risk-free asset and receive

$$\$(1 + r_t^f)A_t$$

at the end of the next period.

- Option 2: switch to stocks by purchasing A_t/P_t shares, hold them for one period and expect to receive

$$\$(A_t/P_t)(P_{t+1} + D_{t+1})$$

at the end of period $t + 1$.

A risk-neutral investor will be indifferent between the certainty of $\$(1 + r_t^f)A_t$, and the his/her expectations of the uncertain payout of option 2. Namely, for such a risk neutral investor

$$(1 + r_t^f)A_t = E[(A_t/P_t)(P_{t+1} + D_{t+1}) | \Omega_t] \quad (7)$$

where Ω_t is the investor's information at the end of period t . This relationship is called the 'arbitrage condition'. Using (5) we now have

$$P_{t+1} + D_{t+1} = P_t(1 + R_{t+1})$$

and the above arbitrage condition can be simplified to

$$E[(1 + R_{t+1}) | \Omega_t] = (1 + r_t^f)$$

or

$$E\left(R_{t+1} - r_t^f \mid \Omega_t\right) = 0 \quad (8)$$

This result establishes that if the investor forms his/her expectations of future stock (index) returns taking account of all market information efficiently, then the excess return, $R_{t+1} - r_t^f$, should not be predictable using any of the market information that are available at the end of period t . Notice that r_t^f is known at time t and is therefore included in Ω_t . Hence, under the joint hypothesis of market efficiency and risk neutrality we must also have $E(R_{t+1} \mid \Omega_t) = r_t^f$.

The above set up can also be used to derive conditions under which asset prices can be characterised as a random walk model. Suppose, the risk free rate, r_t^f , in addition to being known at time t , is also constant over time and given by r^f . Then using (7) we can also write

$$P_t = \left(\frac{1}{1+r}\right) E[(P_{t+1} + D_{t+1}) \mid \Omega_t]$$

or

$$P_t = \left(\frac{1}{1+r^f}\right) [E(P_{t+1} \mid \Omega_t) + E(D_{t+1} \mid \Omega_t)]$$

Under the rational expectations hypothesis and assuming that the ‘transversality condition’

$$\lim_{j \rightarrow \infty} \left(\frac{1}{1+r^f}\right)^j E(P_{t+j} \mid \Omega_t) = 0$$

holds we have the familiar result

$$P_t = \sum_{j=1}^{\infty} \left(\frac{1}{1+r^f}\right)^j E(D_{t+j} \mid \Omega_t) \quad (9)$$

that equates the level of stock price to the present discounted stream of the dividends expected to occur to the asset over the infinite future. The transversality condition rules out rational speculative bubbles and is satisfied if the asset prices are not expected to rise faster than the exponential decay rate determined by the discount factor, $0 < 1/(1+r^f) < 1$. It is now easily seen that if D_t follows a random walk so will P_t . For example, suppose

$$D_t = D_{t-1} + \varepsilon_t \quad (10)$$

where ε_t is a white noise process. Then

$$E(D_{t+j} \mid \Omega_t) = D_t$$

and

$$P_t = \frac{D_t}{r^f} \quad (11)$$

Therefore, we also have

$$P_t = P_{t-1} + u_t \quad (12)$$

where $u_t = \varepsilon_t / r^f$.

The random walk property holds even if $r^f = 0$, since in such a case it would be reasonable to expect no dividends are also paid out, namely $D_t = 0$. In this case the arbitrage condition becomes

$$E(P_{t+1} | \Omega_t) = P_t \quad (13)$$

which is satisfied by the random walk model but is in fact more general than the random walk model. An asset price that satisfies (13) is a martingale process. Random walk processes with zero drift are martingale processes but not all martingale processes are random walks. For example, the price process

$$P_{t+1} = P_t + \lambda \{(\Delta P_{t+1})^2 - E[(\Delta P_{t+1})^2 | \Omega_t]\} + \varepsilon_t$$

where ε_t is a white noise process is a martingale process with respect to the information set Ω_t , but it is clearly not a random walk process, unless $\lambda = 0$.

Other modifications of the random walk theory is obtained if it is assumed that dividends follow a geometric random walk which is a more realistic than the linear dividend model assumed in (10). In this case

$$D_{t+1} = D_t \exp(\mu_d + \sigma_d \nu_{t+1}) \quad (14)$$

where μ_d and σ_d are mean and standard deviation of the growth rate of the dividends. If it is further assumed that $\nu_{t+1} | \Omega_t$ is $N(0, 1)$, we have

$$E(D_{t+j} | \Omega_t) = D_t \exp\left(j\mu_d + \frac{1}{2}j\sigma_d^2\right)$$

Using this result in (9) now yields [assuming that $(1 + r^f)^{-1} \exp(\mu_d + \frac{1}{2}\sigma_d^2) < 1$]

$$P_t = \frac{D_t}{\rho} \quad (15)$$

where

$$\rho = (1 + r^f) \exp\left(-\mu_d - \frac{1}{2}\sigma_d^2\right) - 1$$

The condition $(1 + r^f)^{-1} \exp(\mu_d + \frac{1}{2}\sigma_d^2) < 1$ ensures that the infinite sum in (9) is convergent and $\rho > 0$. Under this set up $\ln(P_t) = \ln(D_t) - \ln(\rho)$, and

$$\ln(P_t) = \ln(P_{t-1}) + \mu_d + \sigma_d \nu_t \tag{16}$$

which establishes that in this case it is log prices that follow the random walk model. This is a special case of the statistical model of return, (1), discussed in Section 3, where $\mu_t = \mu_d$, and $\sigma_t = \sigma_d$.

There are, however, three different types of empirical evidence that shed doubt on the empirical validity of the present value model under risk neutrality.

1. The model predicts a constant price-dividend ratio for a large class of the dividend processes. Two prominent examples being the linear and the geometric random walk models, (10) and (14), discussed above. For more general dividend processes the price-dividend ratio, $\rho_t = P_t/D_t$, could be time-varying, but it must be mean-reverting, in the sense that shocks to prices and dividends must eventually cancel out. In reality, the price-dividend ratio vary considerably over time, show a high degree of persistent, and in general it is not possible to reject the hypothesis that the processes for ρ_t or $\ln(\rho_t)$ contain a unit root. For the Shiller data discussed in 4.2 the autocorrelation coefficient of the log dividend to price ratio computed over the period 1871m1 to 2009m9 is 0.994 (0.024) and falls very gradually as its order is increased and amounts to 0.879 (0.111) at the lag order 12.
2. We have already established that under risk neutrality excess returns must not be predictable. See (8). Yet there is ample evidence of excess return predictability at least in periods of high market volatility. For example, it is possible to explain 15% of the variations in monthly excess returns on SP500 over the period 1872m2-2009m9 by running a linear regression of the excess return on a constant and its 12 lagged values - namely by a univariate $AR(12)$ process. This figure rises to 19% if we exclude the 1929 stock market crash and focus on the post 1948 period. See also the references cited in Section 7.1.
3. To derive the geometric random walk model of asset prices, (16), from the present value model under risk neutrality, we have assumed that innovations to the dividend process are normally distributed. This implies that innovations to asset returns must also be normally distributed. But the empirical evidence discussed in Section 4 above clearly shows that innovations to asset returns tend to be fat-tailed, and often significantly depart from normality. This anomaly between the theory and the evidence is also difficult to reconcile. Under the

present value model prices will have fat-tailed innovations only if the dividends that drive asset prices are also fat-tailed. But under the geometric random walk model for dividends (14), $E(D_{t+j}|\Omega_t)$ need not exist if the dividend innovations, ν_t , are fat-tailed. One important example arises when ν_t has the Student t distribution as defined by (3). For the derivation of the present value expression in this case we need $E(\exp(\sigma_d\nu_{t+j}))$, which is the moment generating function of ν_{t+j} evaluated at σ_d . But the Student t distribution does not have a moment generating function, and hence the present value formula can not be computed when innovations to the dividends are t distributed.

6.2 Risk averse investors

In addition to the above documented empirical shortcomings, it is also important to note that risk neutrality is a behavioral assumption and need not hold even if all market information is processed efficiently by all the market participants. A more reasonable way to proceed is to allow some or all of the investors to be risk averse. In this more general case the certain pay out, $(1 + r_t^f)A_t$, and the expectations of the uncertain pay out, $E[(A_t/P_t)(P_{t+1} + D_{t+1})|\Omega_t]$, will not be the same and differ by a (possibly) time-varying risk premium which could also vary with the level of the initial capital, A_t . More specifically, we have

$$E[(A_t/P_t)(P_{t+1} + D_{t+1})|\Omega_t] = (1 + r_t^f)A_t + \lambda_t A_t$$

where λ_t is the premium per \$ of invested capital required (expected) by the investor. It is now easily seen that

$$E\left(R_{t+1} - r_t^f \mid \Omega_t\right) = \lambda_t$$

and it is no longer necessary true that under market efficiency excess returns are non-predictable. The extent to which excess returns can be predicted will depend on the existence of a historically stable relationship between the risk premium, λ_t , and the macro and business cycle indicators such as changes in interest rates, dividends and various business cycle indicators.

In the context of the consumption capital asset pricing model λ_t is determined by the *ex ante* correlation of excess returns and changes in the marginal utility of consumption. In the case of a representative consumer with the single period utility function, $u(c_t)$, the first-order inter-temporal optimization condition (the Euler equation) is given by

$$E\left[\left(R_{t+1} - r_t^f\right) \frac{u'(c_{t+1})}{u'(c_t)} \mid \Omega_t\right] = 0 \tag{17}$$

where c_t denotes the consumer's real consumption in period t . Using the above condition it is now easily seen that⁴

$$\lambda_t = -\frac{Cov \left[R_{t+1}, \frac{u'(c_{t+1})}{u'(c_t)} \mid \Omega_t \right]}{E \left[\frac{u'(c_{t+1})}{u'(c_t)} \mid \Omega_t \right]} = -\frac{Cov [R_{t+1}, u'(c_{t+1}) \mid \Omega_t]}{E [u'(c_{t+1}) \mid \Omega_t]}$$

For a power utility function, $u(c_t) = (c_t^{1-\gamma} - 1)/(1 - \gamma)$, and $u'(c_{t+1})/u'(c_t) = \exp(-\gamma \Delta \ln(c_{t+1}))$, where $\gamma > 0$ is the coefficient of relative risk aversion. In this case λ_t is given by

$$\lambda_t = \frac{-Cov [R_{t+1}, \exp(-\gamma \Delta \ln(c_{t+1})) \mid \Omega_t]}{E [\exp(-\gamma \Delta \ln(c_{t+1})) \mid \Omega_t]} \quad (18)$$

This result shows that the risk premium depends on the covariance of asset returns with the marginal utility of consumption. The premium demanded by the investor to hold the stock is higher if the return on the asset co-vary positively with consumption. The extent of this co-variation depends on the magnitude of the risk aversion coefficient γ . For plausible values of γ (in the range 1 to 3) and historically observed values of the consumption growth, we would expect λ_t to be relatively small, below 1 per cent per annum. However, using annual observations over relatively long periods one obtains a much larger estimate for λ_t . This was first pointed out by Mehra and Prescott (1985) who found that in the 90 years from 1889 to 1978 the average estimate of λ_t in fact amounted to 6.18 per cent per annum, which could only be reconciled with the theory if one was prepared to consider an implausibly large value for the relative risk aversion coefficient (in the regions of 30 or 40). The large discrepancy between the historical estimate of λ_t based on $R_{t+1} - r_t^f$, and the theory-consistent estimate of λ_t based on (18) is known as the 'equity premium puzzle'. There have been many attempts in the literature to resolve the puzzle by modifications to the utility function, attitudes towards risk, allowing for the possibility of rare events, and the heterogeneity in asset holdings and preferences across consumers. For reviews see Kocherlakota (2003) and Mehra and Prescott (2003).

But even if the mean discrepancy between $E (R_{t+1} - r_t^f \mid \Omega_t)$ and λ_t as given by (18) is resolved, the differences in the higher moments of historically and theory-based risk premia, are likely to be important empirical issues of concern. It seems

⁴Let $X_{t+1} = R_{t+1} - r_t^f$ and $Y_{t+1} = u'(c_{t+1})/u'(c_t)$, and write the Euler equation (17) as

$$E [X_{t+1}Y_{t+1} \mid \Omega_t] = 0 = Cov [X_{t+1}Y_{t+1} \mid \Omega_t] + E [X_{t+1} \mid \Omega_t] E [Y_{t+1} \mid \Omega_t]$$

Then the required results follow immediately, also noting that r_t^f is known at time t and hence has a zero correlation with $u'(c_{t+1})/u'(c_t)$.

difficult to reconcile the high volatility of excess returns with the low volatility of consumption growth that are observed historically.

7 Return predictability and alternative versions of the efficient market hypothesis

In his 1970 review, Fama distinguishes between three different forms of the EMH:

1. The *weak* form asserts that all price information is fully reflected in asset prices, in the sense that current price changes can not be predicted from past prices. This weak form was also introduced in an unpublished paper by Roberts (1967).
2. The *semi-strong* form that requires asset price changes to fully reflect all publicly available information and not only past prices.
3. The *strong* form that postulates that prices fully reflect information even if some investor or group of investors have monopolistic access to some information.

Fama regarded the strong form version of the EMH as a benchmark against which the other forms of market efficiencies are to be judged. With respect to the weak form version he concludes that the test results strongly support the hypothesis, and considered the various departures documented as economically unimportant. He reached a similar conclusion with respect to the semi-strong version of the hypothesis; although as he noted, the empirical evidence available at the time was rather limited and far less comprehensive as compared to the evidence on the weak version.

The three forms of the EMH present different degrees whereby public and private information are revealed in transaction prices. It is difficult to reconcile all the three versions to the mainstream asset pricing theory, and as we shall see below a closer connection is needed between market efficiency and the specification of the model economy that underlies it.

7.1 Dynamic stochastic equilibrium formulations and the joint hypothesis problem

Evidence on the semi-strong form of the EMH was revisited by Fama in a second review of the Efficient Capital Markets published in 1991. By then it was clear that the distinction between the weak and the semi-strong forms of the EMH was

redundant. The random walk model could not be maintained either - in view of more recent studies, in particular that of Lo and MacKinlay (1988).

A large number of studies in the finance literature had confirmed that stock returns over different horizons (days, weeks, and months) can be predicted to some degree by means of interest rates, dividend yields and a variety of macroeconomic variables exhibiting clear business cycle variations. A number of studies also showed that returns tend to be more predictable the longer the forecast horizon. While the vast majority of these studies had looked at the US stock market, an emerging literature has also considered the UK stock market. US studies include Balvers, Cosimano, and MacDonald (1990), Breen, Glosten, and Jagannathan (1989), Campbell (1987), Fama and French (1989), and subsequently by Ferson and Harvey (1993), Kandel and Stambaugh (1996), Pesaran and Timmermann (1994), and Pesaran and Timmermann (1995). See Granger (1992) for a survey of the methods and results in the literature. UK studies after 1991 included Clare, Thomas, and Wickens (1994), Clare, Psaradakis, and Thomas (1995), Black and Fraser (1995), and Pesaran and Timmermann (2000).

Theoretical advances over Samuelson's seminal paper by Leroy (1973), Rubinstein (1976) and Lucas (1978) also made it clear that in the case of risk averse investors tests of predictability of excess returns could not on their own confirm or falsify the EMH. The neoclassical theory cast the EMH in the context of dynamic stochastic (general) equilibrium models and showed that excess returns weighted by marginal utility could be predictable. Only under risk neutrality, where marginal utility was constant, the equilibrium condition implied the non-predictability of excess returns.

As Fama (1991) noted in his second review, the test of the EMH involved a joint hypothesis - market efficiency and the underlying equilibrium asset pricing model. He concluded that "Thus, market efficiency *per se* is not testable." (see p. 1575). This, did not, however, mean that market efficiency was not a useful concept. Almost all areas of empirical economics are subject to the joint hypotheses problem.

7.2 Information and processing costs and the EMH

The EMH, in the sense of asset 'prices fully reflect all available information' was also criticised by Grossman and Stiglitz (1980) who pointed out that there must be "sufficient profit opportunities, i.e. inefficiencies, to compensate investors for the cost of trading and information-gathering."

Only in the extreme and unrealistic case where all information and trading costs are zero one would expect prices to fully reflect all available information. But if information is in fact cost-less it would be known even before market prices are

established.

As Fama recognized a weaker and economically more sensible version of the efficiency hypothesis would be needed, namely “prices reflect information to the point where the marginal benefits of acting on information (the profits to be made) do not exceed the marginal costs.” This in turn makes the task of testing the market efficiency even more complicated and would require equilibrium asset pricing models that allowed for information and trading costs in markets with many different traders and with non-convergent beliefs.

In view of these difficulties some advocates of the EMH have opted for a trade-based notion, and define markets as efficient if it would not be possible for the investors “... to earn above-average returns without accepting above-average risks.” Malkiel (2003) (see p.60). This notion can take account of information and transaction costs and does not involve testing joint hypotheses. But this is far removed from the basic idea of markets as efficient allocators of capital investment across countries, industries and firms.

Beating the market as a test of market efficiency also poses new challenges. Whilst it is certainly possible to construct trading strategies (inclusive of transaction costs) with Sharpe ratios that exceed those of the market portfolios *ex post*, such evidence are unlikely to be convincing to the advocates of the EMH. It could be argued that they are carried out with the benefit of hindsight, and are unlikely to be repeated in real time. In this connections the following considerations would need to be born in mind:

- (a) Data mining/data snooping (Pesaran and Timmermann (2005)).
- (b) Structural change and model instability (choice of observation window).
- (c) The positive relationship that seem to exist between transaction costs and predictability.
- (d) Market volatility and learning.
- (e) The ‘Beat the market’ test is not that helpful either in shedding light on the nature and the extent of market inefficiencies. A more structural approach would be desirable.

8 Theoretical foundations of the EMH

At the core of the EMH lies the following three basic premises:

1. *Investor rationality*: It is assumed that investors are rational, in the sense that they correctly update their beliefs when new information is available.
2. *Arbitrage*: Individual investment decisions satisfy the arbitrage condition, and trade decisions are made guided by the calculus of the subjective expected utility theory a la Savage.
3. *Collective rationality*: Differences in beliefs across investors cancel out in the market.

To illustrate how these premises interact, suppose that at the start of period (day, week, month) t there are N_t traders (investors) that are involved in act of arbitrage between a stock and a safe (risk-free) asset. Denote the one-period holding returns on these two assets by R_{t+1} and r_t^f , respectively. Following a similar line of argument as in section 6.2, the arbitrage condition for trader i is given by

$$\hat{E}_i \left(R_{t+1} - r_t^f \mid \Omega_{it} \right) = \lambda_{it} + \delta_{it}$$

where $\hat{E}_i \left(R_{t+1} - r_t^f \mid \Omega_{it} \right)$ is his/her subjective expectations of the excess return, $R_{t+1} - r_t^f$ taken with respect to the information set

$$\Omega_{it} = \Psi_{it} \cup \Phi_t$$

where Φ_t is the component of the information which is publicly available, $\lambda_{it} > 0$ represents trader's risk premium, and $\delta_{it} > 0$ is her/his information and trading costs per unit of funds invested. In the absence of information and trading costs, λ_{it} can be characterized in terms of the trader's utility function, $u_i(c_{it})$, where c_t is his/her real consumption expenditures during the period t to $t + 1$, and is given by

$$\lambda_{it} = \hat{E}_i \left(R_{t+1} - r_t^f \mid \Omega_{it} \right) = \frac{-\hat{C}ov_i(m_{i,t+1}, R_{t+1} \mid \Omega_{it})}{\hat{E}_i(m_{i,t+1} \mid \Omega_{it})}$$

where $\hat{C}ov_i(\cdot \mid \Omega_{it})$ is the subjective covariance operator condition on the trader's information set, Ω_{it} , $m_{i,t+1} = \beta_i u'_i(c_{i,t+1}) / u'_i(c_{it})$, which is known as the 'stochastic discount factor', $u'_i(\cdot)$ is the first derivative of the utility function, and β_i is his/her discount factor.

The expected returns could differ across traders due to the differences in the their perceived conditional probability distribution function of $R_{t+1} - r_t^f$, the differences in their information sets, Ω_{it} , the differences in their risk preferences, and/or endowments. Under the rational expectations hypothesis

$$\hat{E}_i \left(R_{t+1} - r_t^f \mid \Omega_{it} \right) = E \left(R_{t+1} - r_t^f \mid \Omega_{it} \right)$$

where $E\left(R_{t+1} - r_t^f \mid \Omega_{it}\right)$ is the ‘true’ or ‘objective’ conditional expectations. Furthermore, in this case

$$E\left[\hat{E}_i\left(R_{t+1} - r_t^f \mid \Omega_{it}\right) \mid \Phi_t\right] = E\left[E\left(R_{t+1} - r_t^f \mid \Omega_{it}\right) \mid \Phi_t\right]$$

and since $\Phi_t \subset \Omega_{it}$ we have

$$E\left[\hat{E}_i\left(R_{t+1} - r_t^f \mid \Omega_{it}\right) \mid \Phi_t\right] = E\left(R_{t+1} - r_t^f \mid \Phi_t\right)$$

Therefore, under the REH, taking expectations of the individual arbitrage conditions with respect to the public information set yields

$$E\left(R_{t+1} - r_t^f \mid \Phi_t\right) = E\left(\lambda_{it} + \delta_{it} \mid \Phi_t\right)$$

which also implies that $E\left(\lambda_{it} + \delta_{it} \mid \Phi_t\right)$ *must* be the same across all i , or

$$E\left(R_{t+1} - r_t^f \mid \Phi_t\right) = E\left(\lambda_{it} + \delta_{it} \mid \Phi_t\right) = \rho_t, \text{ for all } i$$

where ρ_t is an average market measure of the combined risk premia and transaction costs. The REH combined with perfect arbitrage ensure that different *traders* have the same expectations of $\lambda_{it} + \delta_{it}$. Rationality and market discipline override individual differences in tastes, information processing abilities and other transaction related costs and renders the familiar representative agent arbitrage condition:

$$E\left(R_{t+1} - r_t^f \mid \Phi_t\right) = \rho_t \tag{19}$$

This is clearly compatible with trader-specific λ_{it} and δ_{it} , so long as

$$\begin{aligned} \lambda_{it} &= \lambda_t + \varepsilon_{it}, \quad E(\varepsilon_{it} \mid \Phi_t) = 0 \\ \delta_{it} &= \delta_t + v_{it}, \quad E(v_{it} \mid \Phi_t) = 0 \end{aligned}$$

where ε_{it} and v_{it} are distributed with mean zero independently of Φ_t , and λ_t and δ_t are known functions of the publicly available information.

Under this setting the extent to which excess returns can be predicted will depend on the existence of a historically stable relationship between the risk premium, λ_t , and the macro and business cycle indicators such as changes in interest rates, dividends and various other indicators.

The rational expectations hypothesis is rather extreme which is unlikely to hold at all times in all markets. Even if one assumes that in financial markets learning takes

place reasonably fast, there will still be periods of turmoil where market participants will be searching in the dark, trying and experimenting with different models of $R_{t+1} - r_t^f$ often with marked departures from the common rational outcomes, given by $E\left(R_{t+1} - r_t^f \mid \Phi_t\right)$.

Herding and correlated behaviour across some of the traders could also lead to further departures from the equilibrium RE solution. In fact the objective probability distribution of $R_{t+1} - r_t^f$ might itself be affected by market transactions based on subjective estimates $\hat{E}_i\left(R_{t+1} - r_t^f \mid \Omega_{it}\right)$.

Market inefficiencies provide further sources of stock market predictability by introducing a wedge between a ‘correct’ *ex ante* measure $E\left(R_{t+1} - r_t^f \mid \Phi_t\right)$, and its average estimate by market participants, which we write as

$$\sum_{i=1}^{N_t} w_{it} \hat{E}_i\left(R_{t+1} - r_t^f \mid \Omega_{it}\right)$$

where w_{it} is the market share of the i^{th} trader.

Let

$$\bar{\xi}_{wt} = \sum_{i=1}^{N_t} w_{it} \hat{E}_i\left(R_{t+1} - r_t^f \mid \Omega_{it}\right) - E\left(R_{t+1} - r_t^f \mid \Phi_t\right)$$

and note that it can also be written as (since $\sum_{i=1}^{N_t} w_{it} = 1$)

$$\bar{\xi}_{wt} = \sum_{i=1}^{N_t} w_{it} \xi_{it} \tag{20}$$

where

$$\xi_{it} = \hat{E}_i\left(R_{t+1} - r_t^f \mid \Omega_{it}\right) - E\left(R_{t+1} - r_t^f \mid \Phi_t\right) \tag{21}$$

ξ_{it} measures the degree to which individual expectations differs from the correct (but unobservable) expectations, $E\left(R_{t+1} - r_t^f \mid \Phi_t\right)$. A non-zero ξ_{it} could arise from individual irrationality, but not necessarily so. Rational individuals faced with an uncertain environment, costly information and limitations on computing power could rationally arrive at their expectations of future price changes that with hindsight differ from the correct ones.⁵ A non-zero ξ_{it} could also arise due to disparity of

⁵This is in line with the premise of the recent paper by Angeletos, Lorenzoni, and Pavan (2010) who maintain the axiom of rationality, but allow for dispersed information and the possibility of information spillovers in the financial markets to explain market inefficiencies.

information across traders (including information asymmetries), and heterogeneous priors due to model uncertainty or irrationality. Nevertheless, despite such individual deviations, $\bar{\xi}_{wt}$ which measures the extent of market or collective inefficiency, could be quite negligible. When N_t is sufficiently large, individual ‘irrationality’ can cancel out at the level of the market, so long as ξ_{it} , $i = 1, 2, \dots, N_t$ are not cross sectionally strongly dependent, and no single trader dominates the market, in the sense that $w_{it} = O(N_t^{-1})$ at any time.⁶ Under these conditions at each point in time, t , the average expected excess returns across the individual traders converges in quadratic means to the expected excess return of a representative trader, namely we have

$$\sum_{i=1}^{N_t} w_{it} \hat{E}_i \left(R_{t+1} - r_t^f \mid \Omega_{it} \right) \xrightarrow{q.m.} E \left(R_{t+1} - r_t^f \mid \Phi_t \right), \text{ as } N_t \rightarrow \infty$$

In such periods the representative agent paradigm would be applicable and predictability of excess return will be governed solely by changes in business cycle conditions and other publicly available information.⁷

However, in periods where traders’ individual expectations become strongly correlated (say as the result of herding or common over-reactions to distressing news) $\bar{\xi}_{wt}$ need not be negligible even in thick markets with many traders; and market inefficiencies and profitable opportunities could prevail. Markets could also display inefficiencies without exploitable profitable opportunities if $\bar{\xi}_{wt}$ is non-zero but there is no stable predictable relationship between $\bar{\xi}_{wt}$ and business cycle or other variables that are observed publicly.

The evolution and composition of $\bar{\xi}_{wt}$ can also help in shedding light on possible bubbles or crashes developing in asset markets. Bubbles tend to develop in the aftermath of technological innovations that are commonly acknowledged to be important, but with uncertain outcomes. The emerging common beliefs about the potential advantages of the new technology and the difficulties individual agents face in learning how to respond to the new investment opportunities can further increase the gap between average market expectations of excess returns and the associated objective rational expectations outcome. Similar circumstances can also prevail during a crash phase of the bubble when traders tend to move in tandem trying to reduce their risk exposures all at the same time. Therefore, one would expect that during bubbles and

⁶Concepts of weak and strong cross section dependence are defined and discussed in Chudik, Pesaran, and Tosetti (2010).

⁷The heterogeneity of expectations across traders can also help in explaining large trading volume observed in the financial markets; a feature which has proved difficult to explain in representative agent asset pricing models. But see Scheinkman and Xiong (2003) who relate the occurrence of bubbles and crashes to changes in trading volume.

crashes the individual errors, ξ_{it} , to become more correlated, such that the average errors, $\bar{\xi}_{wt}$, are no longer negligible. In contrast, at times of market calm the individual errors are likely to be weakly correlated, with the representative agent rational expectations model being a reasonable approximation.

More formally note that since r_t^f and P_t are known at time t , then

$$\xi_{it} = \hat{E}_i \left(\frac{P_{t+1} + D_{t+1}}{P_t} \mid \Omega_{it} \right) - E \left(\frac{P_{t+1} + D_{t+1}}{P_t} \mid \Phi_t \right)$$

Also to simplify the exposition assume that the length of the period t is sufficiently small so that dividends are of secondary importance and

$$\xi_{it} \approx \hat{E}_i (\Delta \ln(P_{t+1}) \mid \Omega_{it}) - f_t,$$

where $f_t = E(\Delta \ln(P_{t+1}) \mid \Phi_t)$ is the unobserved price change expectations. Individual deviations, ξ_{it} , could then become strongly correlated if individual expectations $\hat{E}_i(\Delta \ln(P_{t+1}) \mid \Omega_{it})$ differ systematically from f_t . For example, suppose that

$$\hat{E}_i(\Delta \ln(P_{t+1}) \mid \Omega_{it}) = \theta_{it} \Delta \ln(P_t),$$

but $f_t = 0$, namely in the absence of heterogeneous expectations $\Delta \ln(P_{t+1})$ would have been unpredictable with a zero mean. Then it is easily seen that $\bar{\xi}_{wt} = \bar{\theta}_{wt} \Delta \ln(P_t)$, where $\bar{\theta}_{wt} = \sum_{i=1}^{N_t} w_{it} \theta_{it}$. It is clear that $\bar{\xi}_{wt}$ need not converge to zero if in period t the majority of market participants believe future price changes are positively related to past price changes, so that $\lim_{N_t \rightarrow \infty} \bar{\theta}_{wt} > 0$. In this simple example price bubbles or crashes occur when $\bar{\theta}_{wt}$ becomes positive over a relatively long period.

It should be clear from the above discussion that testing for price bubbles requires disaggregated time series information on individual beliefs and unobserved price change expectations, f_t . Analysis of aggregate time series observations can provide historical information about price reversals and some of their proximate causes. But such information is unlikely to provide conclusive evidence of bubble formation and its subsequent collapse. Survey data on traders' individual beliefs combined with suitable market proxies for f_t are likely to be more effective in empirical analysis of price bubbles.

An individual investor could be asked to respond to the following two questions regarding the current and future price of a given asset:

1. Do you believe the current price is (a) just right (in the sense that the price is in line with market fundamentals), (b) is above the fundamental price, or (c) is below the fundamental price?

2. Do you expect the market price next period to (a) stay about the level it is currently, (b) fall, or (c) rise?

In cases where the market is equilibrating we would expect a close association between the proportion of respondents who select 1a and 2a, 1b and 2b, and 1c and 2c. But in periods of bubbles (crashes) one would expect a large proportion of respondents who select 1b (1c) to also select 2c (2b).

In situations where the equilibrating process is well established and commonly understood, the second question is redundant. For example, if an individual states that the room temperature is too high, it will be understood that he/she would prefer less heating. The same is not applicable to financial markets and hence responses to both questions are needed for a better understanding of the operations of the markets and their evolution over time.

9 Exploiting profitable opportunities in practice

In financial markets the EMH is respected but not worshipped. It is recognized that markets are likely to be efficient most of the time but not all the time. Inefficiencies could arise particularly during periods of important institutional and technological changes. It is not possible to know when and where market inefficiencies arise in advance - but it is believed that they will arise from time to time. Market traders love volatility as it signals news and change with profit possibilities to exploit. Identification of exploitable predictability tend to be fully diversified across markets for bonds, equities and foreign exchange. Misalignments across markets for different assets and in different countries often present the most important opportunities. Examples include statistical arbitrage and global macro arbitrage trading rules.

Predictability and market liquidity are often closely correlated; less liquid markets are likely to be more predictable. Market predictability and liquidity need to be jointly considered in developing profitable trading strategies. Return forecasting models used in practice tend to be recursive and adaptive along the lines developed in Pesaran and Timmermann (1995) and recently reviewed in Pesaran and Timmermann (2005). Recursive modelling (RM) approach is also in line with the more recent developments in behavioural finance. The RM approach aims at minimizing the effect of hindsight and data snooping (a problem that afflicts all *ex post* return regressions), and explicitly designed to take account of the potential instability of the return regressions over time. For example, Pesaran and Timmermann (1995) find that the switching trading rule manages to beat the market only during periods of high volatility where learning might be incomplete and markets inefficient.

Pesaran and Timmermann (2005) provide a review of the recursive modelling approach, its use in derivation of trading rules and discuss a number of practical issues in their implementation such as the choice of the universe of factors over which to search, choice of the estimation window, how to take account of measurement and model uncertainty, how to cross validate the RM, and how and when to introduce model innovations.

The RM approach still faces many challenges ahead. As Pesaran and Timmermann (2005) conclude:

“Automated systems reduce, but do not eliminate the need for discretion in real time decision making. There are many ways that automated systems can be designed and implemented. The space of models over which to search is huge and is likely to expand over time. Different approximation techniques such as genetic algorithms, simulated annealing and MCMC algorithms can be used. There are also many theoretically valid model selection or model averaging procedures. The challenge facing real time econometrics is to provide insight into many of these choices that researchers face in the development of automated systems.”

Return forecasts need to be incorporated in sound risk management systems. For this purpose point forecasts are not sufficient and joint probability forecast densities of a large number of inter-related asset returns will be required. Transaction and slippage costs need to be allowed for in the derivation of trade rules. Slippage arises when long (short) orders, optimally derived based on currently observed prices, are placed in rising (falling) markets. Slippage can be substantial, and are in addition to the usual transactions costs.

Familiar risk measures such as the Sharpe ratio and the VaR are routinely used to monitor and value the potential of trading systems. But due to cash constraint (for margin calls etc.) it is large drawdowns that are most feared. Prominent recent examples are the downfall of Long Term Capital who experienced substantial drawdowns in 1998 following the Russian financial crisis, and the collapse of Lehman Brothers during the global financial crisis of 2008.

Successful traders might not be (and usually are not) better in forecasting returns than many others in the market. What they have is a sense of ‘big’ opportunities when they are confident of making a ‘kill’.

10 New research directions

We have identified two important sources of return predictability and possible profitable opportunities. One relates to the familiar business cycle effects and involves modelling ρ_t , defined by (19), in terms of the publicly available information, Φ_{t-1} .

The second relates to the average deviations of individual traders's expectations from the "correct" unknown expectations, as measured by $\bar{\xi}_{wt}$ and defined by (20)). As noted earlier this component could vary considerably over time and need not be related to business cycle factors. It tends to be large during periods of financial crisis when correlation of mis-pricing across traders rise, and tend to be negligible during periods of market calm when correlations are low. Over the past three decades much of the research in finance and macroeconomics has focussed on modelling of ρ_t , and by comparison little attention has been paid to $\bar{\xi}_{wt}$. This is clearly an important area for future research. Our discussions also point to a number of related areas for further research. There are clearly

- Limits to rational expectations (for an early treatment see Pesaran (1987)), also see the recent paper on Survey Expectations by Pesaran and Weale (2006).
- Limits to arbitrage due to liquidity requirements and institutional constraints.
- Herding and correlated behaviour with noise traders entering markets during bull periods and deserting during bear periods.

Behavioral finance, complexity theory and the Adaptive Markets Hypothesis recently advocated by Lo (2004) all try, in one way or another, to address the above sources of the departures from the EMH. Some of the recent developments in behavioural finance are reviewed in Baberis and Thaler (2003).

Farmer and Lo (1999) focus on the recent research that views the financial markets from a biological perspective and, specifically, within an evolutionary framework in which markets, instruments, institutions, and investors interact and evolve dynamically according to the 'law' of economic selection. Under this view, financial agents compete and adapt, but they do not necessarily do so in an optimal fashion.

Special care should also be exercised in evaluation of return predictability and trading rules. To minimize the effects of hindsight in such analysis recursive modelling techniques discussed in Pesaran and Timmermann (1995), Pesaran and Timmermann (2000) and Pesaran and Timmermann (2005) seem much more appropriate than the return regressions on a fixed set of regressors/factors that are estimated *ex post* on historical data.

References

- Angeletos, G., G. Lorenzoni, and A. Pavan (2010). Beauty contests and irrational exuberance: A neoclassical approach. Working Paper 15883 National Bureau Of Economic Research.

- Baberis, N. and R. Thaler (2003). A survey of behavioral finance. In G. M. Constantinides, M. Harris, and R. Stultz (Eds.), *Handbook of Behavioral Economics of Finance*, pp. 1052–1090. Elsevier Science.
- Balvers, R. J., T. F. Cosimano, and B. MacDonald (1990). Predicting stock returns in an efficient market. *The Journal of Finance* 45, 1109–1128.
- Black, A. and P. Fraser (1995). Uk stock returns: Predictability and business conditions. *The Manchester School Supplement* 63, 85–102.
- Breen, W., L. R. Glosten, and R. Jagannathan (1989). Economic Significance of Predictable Variations in Stock Index Returns. *Journal of Finance* 44, 1177–1189.
- Campbell, J. Y. (1987). Stock Returns and the Term Structure. *Journal of Financial Economics* 18, 373–399.
- Chudik, A., M. H. Pesaran, and E. Tosetti (2010). Weak and Strong Cross Section Dependence and Estimation of Large Panels. Mimeo, University of Cambridge.
- Clare, A. D., Z. Psaradakis, and S. H. Thomas (1995). An analysis of seasonality in the UK equity market. *Economic Journal* 105, 398–409.
- Clare, A. D., S. H. Thomas, and M. R. Wickens (1994). Is the gilt-equity yield ratio useful for predicting UK stock return? *Economic Journal* 104, 303–315.
- Cowles, A. (1960). A revision of previous conclusions regarding stock price behavior. *Econometrica* 28, 909–915.
- Fama, E. F. (1965). The behavior of stock market prices. *Journal of Business* 38, 34–105.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance* 25, 383–417.
- Fama, E. F. (1991). Efficient capital markets: II. *The Journal of Finance* 46, 1575–1617.
- Fama, E. F. and K. R. French (1989). Business Conditions and Expected Returns on Stocks and Bonds. *Journal of Financial Economics* 25, 23–49.
- Farmer, D. and A. Lo (1999). Frontiers of finance; evolution and efficient markets. In *Proceedings of the National Academy of Sciences*, Volume 96, pp. 9991–9992.
- Ferson, W. E. and C. R. Harvey (1993). The risk and predictability of international equity returns. *Review of Financial Studies* 6, 527–566.
- Granger, C. W. J. (1992). Forecasting stock market prices: Lessons for forecasters. *International Journal of Forecasting* 8, 3–13.

- Grossman, S. and J. Stiglitz (1980). On the impossibility of informationally efficient markets. *American Economic Review* 70, 393–408.
- Jarque, C. M. and A. K. Bera (1980). Efficient Tests for Normality, Homoscedasticity and Serial Independence of Regression Residuals. *Economics Letters* 6, 255–259.
- Kandel, S. and R. F. Stambaugh (1996). On the predictability of stock returns: An asset-allocation perspective. *Journal of Finance* 51, 385–424.
- Kendall, M. (1953). The analysis of economic time series - part 1: Prices. *Journal of the Royal Statistical Society* 96, 11–25.
- Kocherlakota, N. R. (2003). The equity premium: It’s still a puzzle. *Journal of Economic Literature* 34, 42–71.
- Leroy, S. (1973). Risk aversion and the martingale property of stock returns. *International Economic Review* 14, 436–446.
- Lo, A. (2004). The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management* 30, 15–29.
- Lo, A. and C. MacKinlay (1988). Stock market prices do not follow random walks: evidence from a simple specification test. *Review of Financial Studies* 1, 41–66.
- Lucas, R. E. (1978). Asset prices in an exchange economy. *Econometrica* 46, 1429–1446.
- Malkiel, B. G. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives* 17, 59–82.
- Mehra, R. and E. Prescott (1985). The equity premium: a puzzle. *Journal of Monetary Economics* 15, 146–161.
- Mehra, R. and E. C. Prescott (2003). The equity premium puzzle in retrospect. In M. H. G.M. Constantinides and R. Stulz (Eds.), *Handbook of the Economics of Finance*, pp. 889–938. Amsterdam: North Holland.
- Osborne, M. (1959). Brownian motion in the stock market. *Operations Research* 7, 145–173.
- Osborne, M. (1962). Periodic structures in the brownian motion of stock prices. *Operations Research* 10, 345–379.
- Pesaran, B. and M. H. Pesaran (2010). Conditional volatility and correlations of weekly returns and the VaR analysis of 2008 stock market crash. *Economic Modelling* Forthcoming.

- Pesaran, M. H. (1987). *The Limits to Rational Expectations*. Oxford: Basil Blackwell. Reprinted with corrections 1989.
- Pesaran, M. H., A. Pick, and A. Timmerman (2010). Variable Selection, Estimation and Inference for Multi-period Forecasting Problems. Mimeo, University of Cambridge.
- Pesaran, M. H. and A. Timmermann (1994). Forecasting Stock Returns: An Examination of Stock Market Trading in the Presence of Transaction Costs. *Journal of Forecasting* 13, 335–367.
- Pesaran, M. H. and A. Timmermann (1995). The robustness and economic significance of predictability of stock returns. *Journal of Finance* 50, 1201–1228.
- Pesaran, M. H. and A. Timmermann (2000). A recursive modelling approach to predicting uk stock returns. *The Economic Journal* 110, 159–191.
- Pesaran, M. H. and A. Timmermann (2005). Real time econometrics. *Econometric Theory* 21, 212–231.
- Pesaran, M. H. and M. Weale (2006). *Handbook of Economic Forecasting*, Chapter Survey Expectations. North-Holland.
- Roberts, H. (1967). Statistical versus clinical prediction in the stock market. Unpublished manuscript, Center for Research in Security Prices, University of Chicago.
- Rubinstein, M. (1976). The valuation of uncertain income streams and the pricing of options. *Bell Journal of Economics* 7, 407–425.
- Samuelson, P. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review Spring* 6, 41–49.
- Scheinkman, J. A. and W. Xiong (2003). Overconfidence and speculative bubbles. *Journal of Political Economy* 111, 1183–1219.
- Shiller, R. J. (2005). *Irrational Exuberance* (Second ed.). Princeton University Press.