A REVIEW OF THE MARKET RISK PREMIUM

and

COMMENTARY ON TWO RECENT PAPERS

A REPORT PREPARED FOR

THE ENERGY NETWORKS ASSOCIATION

STEPHEN GRAY AND R.R. OFFICER

15 August, 2005

EXECUTIVE SUMMARY

In this paper, we review two recent papers on the estimation of the Australian market risk premium (MRP). We begin by reviewing the relevant aspects of the theoretical framework of the Capital Asset Pricing Model in so far as it is relevant to the estimation of MRP. This helps to provide a framework for thinking about how one should best estimate the MRP, and the advantages and disadvantages of various approaches.

Then we turn to a specific review of Hathaway (2005) and Hancock (2005). Both of these papers primarily examine historical data – the excess return of a broad stock market index relative to the yield on government bonds.

Both introduce slight modifications to the standard statistical methods that are usually applied to analyse this type of data. We conduct a number of statistical exercises and simulation experiments to examine the strengths and weaknesses of the proposed techniques. We make three conclusions about this:

1. The mean historical excess return on the market portfolio is substantially above 6% whether data from the last 30, 50, 75, 100, or 120 years is used.
2. The alternative statistical techniques employed in these two papers confirms this – the mean historical excess return on the market portfolio over recent periods is 6% or more.
3. Some of the statistical techniques that are proposed in these two papers are inferior to the simple mean estimate in important respects.

Even though we raise some statistical issues in relation to the techniques that are proposed in these two papers, this is largely a moot point in that they both confirm that the historical data supports a MRP estimate of 6% or more.

Both authors then proceed to make ad hoc “adjustments” to the historical estimates. Hathaway (2005) makes an adjustment for the increase in the price-earnings ratio that has occurred over the last 30 years. Hancock (2005) makes adjustments based on arguments that discount rates have fallen over the last 30 years and that the introduction of dividend imputation caused a massive appreciation in stock prices in 1987. In our view, such ad hoc adjustments should not be made to the historical data.

Both authors argue that events that are unanticipated and unlikely to repeat should be removed from the data set or the subject of adjustments to the historical data. Our response is that there are many events that are both unexpected and unlikely to repeat, and yet are not the subject of adjustment in either paper. The terrorist attacks of 2001 and the Asian crisis of 1997 are some examples.

There are many economic events that affect stock returns. To eliminate those that are claimed to be unexpected and non-recurring would be to leave a scant and practically useless data set. Indeed it is precisely because there are unexpected events that affect markets in different ways that there exists a MRP in the first place! Rather than selectively eliminate from the data events that are considered to be unexpected, the preferred approach is to analyse a longer data set that contains both positive and negative shocks. Moreover, in a regulatory setting, this would invite an avalanche of submissions on which events were expected and which were not.

Our conclusion is that there is nothing in the recent data nor in these papers that justifies a change in the regulatory precedent of using 6% as an estimate of the market risk premium. Indeed the mean excess market return is substantially above 6% over relatively short or long historical periods. Estimates below 6% can only be achieved by making selective adjustments to the historical data.

We also note that the effect of franking credits on the estimate of MRP is small relative to both estimation error and the way in which other evidence is reflected in the final MRP estimate. We conclude that (i) it is appropriate to combine data from before and after the introduction of imputation and to express an estimate of the MRP that ignores any adjustment for the value of franking credits, and (ii) that the estimate of 6% that has been adopted by regulatory and market practice is such an estimate. We believe an adjustment to the MRP for franking credits is likely to be less than 50 basis points and to take the MRP to a decimal point, in view of general measurement errors, in our opinion would give a spurious impression of precision in the estimate.

CONCEPTUAL FRAMEWORK

Introduction

One of the most critical and yet elusive measures of modern approaches to valuation is the ex-ante market risk premium (MRP). Ex-post it is impossible for this variable to be a constant because if it was constant there would be no risk and no risk premium. However, this does not mean the ex-ante estimation of this variable cannot be represented by a stable distribution with constant parameters; it is the ex-post measure which is stochastic. Moreover, the inherent stochastic nature of the ex-post MRP and its importance to estimating the ex-ante MRP, inevitably, will make its estimation a subject of controversy and debate.

Our view is that the ex-ante MRP is probably not constant and cannot be adequately represented by a stable distribution. Unfortunately, however, the theory as to what might cause the parameters of the distribution (and thus the mean ex-ante MRP) to change is not well developed. This makes forecasting changes difficult if not impossible. Moreover, given the volatility of ex post market excess returns, even detecting such a change after the event is almost impossible. Given this state of knowledge about the MRP we recommend caution before changing the MRP estimates without strong evidence. Otherwise changes will just increase the variability of estimates and the risks of valuation – a message that is particularly important for regulators of infrastructure assets[1].

The CAPM and MRP

The importance of the MRP in valuations has arisen largely because of the prominence of the use of the capital asset pricing model (CAPM). The popularity of the CAPM is due to the fact that it provides a relatively simple model for valuing equity where there is no contractual rate of return. The assumptions underlying the CAPM are not consistent with market reality and the notion that risk should be confined to one parameter of covariance ignores other important risk parameters such as liquidity. Nonetheless, the simplicity of the CAPM has meant that it is widely used to estimate the required return to equity which, unlike debt, does not have a contractual rate.

The CAPM is defined as:

$$E(R_{jt}) = R_{ft} + \beta_{jt}[E(R_{mt} - R_{ft}] \ldots (1)$$

Where:

$E(R_j)$ is the expected return on asset (stock) *j* at time *t*;

R_f is the "risk free rate" of return, usually estimated as the yield on a long term government bond at time *t*;

β_{jt} is the estimate of this asset's covariance risk at time *t*. Usually estimated as the slope coefficient of an ordinary least squares regression of the asset's return against the returns of a market index; and

$E(R_{mt})$ is the expected return on the market factor at time *t*, usually estimated as the return on a broad index of the stock market.

The ex-ante market risk premium is defined as:

$$MRP = E(R_{mt}) - R_{ft} \qquad \ldots (2)$$

It is important to recognise that the CAPM and MRP are ex-ante, as distinct from ex-post[2], models or variables. The role of the CAPM is to forecast the return that is expected (or required) from the asset, E(R_{jt}). The MRP is a forecast of the expected or required premium or spread relative to the risk free return that is required to induce investors to hold stocks rather than riskless government bonds. As with most ex-ante models in economics and finance, the empirical validation of the CAPM and the MRP are notoriously difficult and subject to considerable debate. There have been a number of other models suggested, usually justified through the arbitrage pricing theory, and there have been a number of alternative suggestions and additions to β (the covariance risk of

[1] If the MRP is allowed to change on a whim, for example, the political pressure that will be exerted on regulators to meet short term political expediencies in their pricing decisions will become enormous and to the ultimate detriment of consumers and regulated entities alike.

[2] The ex-post version of the *MRP* (= $R_m - R_f$) is the observed difference between the return on the market (R_m) and the yield on a government bond (R_f) over a particular time period, usually one year.

an asset) and by implication the MRP. Nonetheless, in spite of the uncertainties surrounding the models there has been no real successor in the sense that it provides a simple and superior means of forecasting expected or required returns relative to the CAPM.

The equations for CAPM and MRP have been defined in the context of a particular point in time, but this does not imply there is any defined time period for the models. The theory does not specify any time period for the models – the CAPM applies to a single time period of unspecified length. As a consequence, there has been quite a deal of controversy and variation in the implied time period that has formed the basis of parameter estimates for these models. Insofar as the surrogate for the risk free rate has been a 10-year government bond yield, this would imply a 10-year planning horizon. The reason that a 10-year yield has been adopted is because most of the projects for which CAPM has been used as a means of estimating the required return on equity have been long term projects and it would be a mistake, in these circumstances, to use short dated government securities as the risk free rate surrogate. On the other hand, those who do use shorter-term government bill rates often point to the fact that traders in equity markets are basically setting prices and these people have short term planning horizons. The response is that the investment planning by corporations should not be affected by short term market movements and, insofar as their planning is long term, an extended period for the risk free surrogate such as the 10 year bond yield is the most appropriate.

The most critical and least understood parameter of the models is the expectations operator (E). Most statisticians and economists, when they see such a representation, immediately assume that it refers to the mean of a known distribution function. This usually leads them to assume a normal distribution and to use historical returns to estimate the parameters of the distribution. The problem with this approach is that there is not any single known distributional form that adequately captures the stochastic process generating stock market returns or variables encompassing such returns. There is no natural law that says returns have to be characterised or represented by any mathematical function. This does not mean that one should not use distributions but simply they, like many models in finance, should be used with a degree of discretion because the distributions of stock market returns are rarely so well behaved that parameters can be estimated from historical returns and then used with any confidence to forecast future returns.

The expectations operator (E) should be correctly thought of as the market's forecast of future or required returns before they will invest in the equity of this risk class. Ideally, what we need is some method of forecasting investor's expectations or equivalently their required returns for the different risk class of assets. Unfortunately, while such models exist, they rarely have very much to offer in the way of forecast ability. In an investment environment, this is perhaps not surprising insofar as if there were forecast abilities in these models then this would remove elements of risk and make the models redundant insofar as they are based on risk or stochastic returns.

In such circumstances, it is perhaps inevitable that forecasts, in order to be objective, rely heavily on historical data. Indeed it is common for many economic forecasts to be based on projections of historical data. The argument for relying on such data is that the expectations of investors will be framed on the basis of their experiences, which are of course historical. Therefore, the argument runs, the mean of historical distributions of returns or models framing returns could be expected to have had the greatest influence on investors' expectations about the future. Hence the reliance on some average of historical excess market returns in order to settle on an estimate of the investor's expected or required MRP.

However, having a justification for choosing an average of an historical series does not overcome the problem of which average from the distribution of historical excess returns is appropriate to reflect investors' expectations. Theory is of little help and the conventional practice has been to choose a mean of annual market excess returns basically because these observed excess returns are usually publicised as annual rates and one would expect them to have a more profound effect on investor's expectations than shorter periods such as monthly rates return or indeed longer periods such as 5 or 10 year rates of return. Of course, if the distribution of excess returns was a stationary over time then it would not matter whether monthly, annual or 10-yearly periods were chosen as a stationary or stable distribution would mean that the parameters were invariant over various time periods and that one year rates would be simply a product of monthly rates and in turn 10-year rates a product of annual rates.

Unfortunately the evidence is not consistent with such stability, even when the returns are expressed as natural logs and the mean of log normal distributions is adopted as the estimate. As a characterisation of ex-post investment returns logs are superior to normal returns because the geometric mean, which is the simple average of logarithmic returns, is equal to the compound rates return that are earned in the market place. Unfortunately, the log normal distribution is not an accurate approximation of the actual return distribution nor is it likely that investors think in terms of geometric means even though they approximate the compounded return earned on an investment. It is likely that investors think and are influenced by arithmetic average returns and hence it is this reason that arithmetic averages tend to be used to forecast future MRP's.

Moreover, the arithmetic average is usually used on the basis that we are seeking an estimate of the expected return on the broad market over and above the yield on government bonds over the next year. If all annual observations are independent draws from the same distribution, the appropriate estimate of the expected value is the arithmetic mean. In some circumstances, a geometric mean is computed. This is appropriate when estimating the aggregated return from a buy and hold strategy over a long period, but that is not the purpose here. The MRP is to be used in the CAPM to compute the cost of equity expressed in annual terms. Therefore, we require an estimate of the expected return, over the next year, on the market portfolio over and above the risk-free rate. What return do we expect on the market portfolio over the next year, relative to the risk-free rate? The historical data provides us with many observations on what the market returned relative to the risk-free rate over a one-year period. To the extent that each of these should be given equal weight, a simple arithmetic average is appropriate.

There are further problems in the lack of a stable distribution that can readily approximate returns and MRPs. The lack of stability means that the standard statistical tests of significance of a mean from observed values cannot be relied upon and while measures of dispersion such as standard deviation may give a good approximation of the variability of outcomes it is unlikely that the probabilities of errors in assessing the hypothesis under test can be applied with any accuracy.

What does all this amount to? Clearly, it is difficult to frame estimates of MRPs in a regulatory setting with any degree of certainty. The problems of measurement, distributional assumptions, the significance of hypothesis testing, the general vagaries of investment markets and the expectations of investors make the task of framing a MRP as part of a regulatory process for determining returns and then prices of products or services, extremely difficult and the results tenuous. Nonetheless, estimates have to be made! However, the message coming from the tenuous nature of the estimates is that regulators should be careful about adjusting estimates of the MRP without strong evidence of a shift in the MRP. Regulators will induce another element of uncertainty and costs into the regulatory process if they allow MRP estimates to vary widely on the basis of short-term movements in ex-post observations. This is especially the case if such changes are at variance with what turns out to be, with the benefit of hindsight, market expectations.

We recognise that it is likely that the MRP is not stationary and likely to vary under different economic conditions. However, the fact that there is no adequate theory underlying the variability of MRPs makes it dangerous to adjust an MRP estimate simply because another year or two or three of data alter the estimated mean. For example, a year ago the 30-year mean excess return was less than 6%, leading some to call for a reduction in the MRP used by Australian regulators. Now, the most recent 30-year mean excess return is 7.7%. We do not advocate increasing the MRP now for the same reason we did not advocate reducing the MRP estimate last year. The problems of the theory and measurement of MRPs suggest a conservative approach – a regulator should be very careful about making any changes without compelling evidence.

The MRP and the Value of Dividend Imputation Franking Credits

Under a dividend imputation tax system, there are potentially three components to the return received by equity holders – dividends, capital gains, and franking credits. In this setting, the appropriate measure of MRP is one that includes all three components of the equity return. This point is clearly demonstrated in Officer (1994). However, standard stock market accumulation indexes reflect dividends and capital gains only. Consequently, the value of franking credits should, in theory, be added to the historical estimates of stock index returns.

However, this is problematic for two reasons. First, the value of franking credits differs across investor types.[3] Second, even if we could precisely aggregate the value of franking credits over all investors (and potential investors) in the Australian market, the value of franking credits cannot be observed. In fact, the value of franking credits depends on an unobserved parameter, gamma, which must be estimated from market data. A number of approaches have been proposed to estimate gamma, each of which tends to produce a different estimate. Because gamma cannot be observed, but must be estimated from noisy data using indirect means, there is no single precise and robust estimate that is universally viewed as being correct.

For these reasons, it is common to ignore the value of franking credits when constructing stock return indexes. This is also consistent with market practice. In a

[3] Here, we refer to the value after company tax but before personal tax. In the case of franking credits, the identity of the recipients of franking credits has an impact on the amount of corporate tax that is rebated. Thus, different investors place a different after-company-tax value on franking credits. Although different investors have different personal tax obligations in relation to dividends and franking credits, the value after company tax but before personal tax is the same for all investors. It is this basis of valuation that we use in this paper, since that is the valuation basis on which assets trade in the market.

recent paper, Truong, Partington and Peat (2005)[4] survey 356 listed Australian firms about various corporate finance practices. All firms were included in the All Ordinaries Index in August 2004, Australian and not in the finance sector. On the question of how franking credits were treated, 85% of respondents indicated that they made no adjustment to their estimate of MRP to reflect the value of franking credits.

Moreover, the size of the adjustment is likely to be small relative to the uncertainty in the estimate of MRP in any event. This is because the size of the adjustment is limited by the assumed value of franking credits (gamma) and the rate at which they are distributed (which in turn depends on dividend yields). For example, Hathaway (2005)[5] values the effect of franking credits at around 50 basis points and Hancock (2005)[6] values the effect of franking credits at zero. This must be contrasted with the standard deviation of the mean excess market return which over the last 30 years is 415 basis points. That is, the likely adjustment that would be made for the value of franking credits is about 12% of a single standard error of the estimate of MRP using the last 30 years of data. Even when using 120 years of data, the standard error of the mean excess market return estimate is over 150 basis points, which is more than three times the size of the likely adjustment for franking credits.

In addition, regulatory and market practice[7] is to compute an estimate of MRP based on historical data, but to adopt a final estimate that reflects appropriate judgment about other information such as recent trends, changes in the market, survey evidence, evidence from various economic models and so on. These judgments explain why regulatory and market practice has been to use an estimate of 6% even though historical data from the last 30, 50, 75, or 100 produce estimates that are considerably higher. Any likely

[4] Truong, G., Partington, G. and Peat, M. (2005) "Cost of Capital Estimation and Capital Budgeting Practice in Australia," Working Paper, University of Sydney, and Conference Proceedings, AFAANZ 2005.

[5] Hathaway, Neville. "Australian Market Risk Premium", Capital Research Pty Ltd, January 2005.

[6] The South Australian Centre for Economic Studies, "The Market Risk Premium for Australian Regulatory Decisions," April, 2005.

[7] See again Truong, G., Partington, G. and Peat, M. (2005).

adjustment to reflect the value of franking credits is economically small and likely to be subsumed in the estimation error from historical data and the economic judgment that is applied in adopting a final value for MRP. In short, taking the MRP to a decimal point would give an impression of accuracy in the estimate that is misleading.

For this reason, we consider that (i) it is appropriate to combine data from before and after the introduction of imputation and to express an estimate of the MRP that ignores any adjustment for the value of franking credits, and (ii) the estimate of 6% that has been adopted by regulatory and market practice is such an estimate. Indeed, the average excess market return over relatively short or long historical periods, ignoring franking credits, is above 6%. To the extent that an adjustment for franking credits is required, this adjustment must be economically small and is likely to be subsumed in the estimation error.

STATISTICAL TECHNIQUES

Standard techniques

The MRP is most commonly estimated with reference to historical data. This data takes the form of excess market returns – the return on a broad stock market index less the yield on government bonds. In Australian regulatory determinations, it has been common to observe data at an annual frequency and to use 10-year government bonds.

The standard means of analysing this data is to take the simple mean over some historical period. The arithmetic average is usually used on the basis that we are seeking an estimate of the expected return on the broad market over and above the yield on government bonds over the next year. If all annual observations are independent draws from the same distribution, the appropriate estimate of the expected value is the arithmetic mean. In some circumstances, a geometric mean is computed. This is appropriate when estimating the aggregated return from a buy and hold strategy over a long period, but that is not the purpose here. The MRP is to be used in the CAPM to compute the cost of equity expressed in annual terms. Therefore, we require an estimate of the expected return, over the next year, on the market portfolio over and above the risk-free rate. What return do we expect on the market portfolio over the next year, relative to the risk-free rate? The historical data provides us with many observations on what the market returned relative to the risk-free rate over a one-year period. To the extent that each of these should be given equal weight, a simple arithmetic average is appropriate.

Applying this standard approach to the most recent data that is available produces the results that are documented in Table 1. In each case, we have computed the average excess return of a broad stock market index[8] over the yield on 10-year government bonds reported at the beginning of the year.

[8] We have used a series of stock market indexes from 1882 to the present, as described in Hathaway (2005), p.6 and www.capitalresearch.com.au. For each year, the major stock market index at that time is used.

Table 1: Average observed excess market returns

Length of period (years)	Period	Mean excess return
30	1975-2004	7.70%
50	1955-2004	6.43%
75	1930-2004	6.58%
100	1905-2004	7.15%
120	1885-2004	7.17%

Clearly, whether the most recent 30-year period or a long historical period of 120 years is taken, the data supports a mean excess return estimate of at least 6%.

Hathaway (2005) [9]

Hathaway presents two types of empirical estimate. First, he reports standard mean market excess returns over various historical periods. These results correspond to those presented in Table 1 above. All of the empirical work being done in this area uses essentially the same data set, so there is little disagreement about the fact that the mean market excess return over the range of historical periods presented in Table 1 is above 6%.

Hathaway also proposes the analysis of 10-year returns. This involves comparing the yield on 10-year government bonds with the observed market return over the subsequent 10-year period. When estimating the MRP, it is standard to observe data at an annual frequency because an estimate of the MRP at an annual frequency is required to be used in the CAPM to compute the annual cost of equity capital. However this is done more for convenience than for any conceptual reason. Indeed the CAPM is silent on the appropriate time horizon to be examined. It is quite valid to set this time horizon to 10 years and to observe excess market returns over this horizon.

[9] Hathaway, Neville. "Australian Market Risk Premium", Capital Research Pty Ltd, January 2005.

However, there are two potential problems with this approach. First, there are very few independent observations of 10-year market excess returns available. Only 12 observations are available from the full data set and this produces an estimate with such large standard errors that it is essentially unusable. Moreover, Hathaway also examines the post-1960 data for which only four observations are available. The precise estimation of the mean of such a volatile series using only four observations is essentially impossible. Of course, it is possible to use overlapping data sets and to include a new observation every year. But these observations are no longer independent as successive observations have nine of the ten years of market data in common. Including these additional observations, therefore, does nothing to reduce the true standard errors of the estimate. Hathaway's focus is on the point estimates of the mean of the series, so he does not report any standard errors. However, if these standard errors were reported on the basis of the number of non-overlapping observations or after allowing for the inevitable serial correlation that results from the use of over-lapping observations, they would indicate that the estimates are so imprecise as to require considerable caution in their interpretation.

The second problem is that the interpretation of these 10-year estimates differs from the standard interpretation of the MRP in an important way. The usual interpretation of the MRP is that this is the expected return on the market portfolio, over the next year, in excess of the risk-free return. Suppose that under this interpretation the MRP is 6%. Then we can observe the risk-free rate at the beginning of the year and then expect that the market return will be, on average, 6% higher than this over the course of the year. The following year we can do the same thing – observe the risk-free rate at that time and expect that the market return over the following year will be 6% higher than this, on average. This interpretation allows us to compute the mean of the observed excess market returns as these excess returns are all drawn from the same distribution – the expected excess return is 6% in every year.

The use of 10-year returns involves a different interpretation. Under this interpretation we observe the risk-free rate at the start of a 10-year period. We then expect the market return over the next year to be 6% higher than this, on average. The following year, we expect the same market return regardless of the risk-free rate at the time. That is, every year for 10 years we expect the market to return 6% more than the risk-free rate observed at the beginning of the 10-year period. This occurs independently of any rise or fall in interest rates over the 10 years.[10] Again, the CAPM is silent on the appropriate time horizon to be examined. The point here is simply that the interpretation under Hathaway's 10-year return approach is quite different from the standard interpretation.

Consider, as an example, the most recent 10 years of data. Annual observations of the yield on 10-year government bonds (the usual proxy for the risk-free rate in this setting) are reported in Table 2.

Table 2: Observed risk-free rates

Year	**Risk-free rate (% p.a.)**	**Expected market return**
1995	10.04	16.04
1996	8.18	14.18
1997	7.37	13.37
1998	6.05	12.05
1999	5.01	11.01
2000	6.96	12.96
2001	5.46	11.46
2002	6.01	12.01
2003	5.16	11.16
2004	5.60	11.60
Mean	**6.58**	**12.58**

Yields on 10-year Australian Government bonds at the beginning of the year.

Under the standard interpretation, the expected market return is reported in the final column based on a 6% MRP. Under this interpretation, the average expected return over the 10-year period is 12.58%. Now suppose that the actual market return over this period matched this expectation from the usual interpretation of MRP, as would be the case, on average, if this interpretation were correct. The standard approach of subtracting the risk-

[10] Of course, some allowance must be made for the compounding of returns over the 10 years, so the expected market return each year is not exactly 6% more than the initial risk-free rate. But the point here is

free rate at the beginning of the year from the market return observed over the year and finding the mean, would produce an estimate of 6%. However, the Hathaway 10-year approach would produce an estimate of 12.58 – 10.04 = 2.54% using this same data. Which statistical technique is appropriate depends on whether one views data in which the market return is, on average, 6% p.a. above the risk-free rate at the beginning of the year as consistent with a MRP of 6% or 2.5%.

In any event, this is somewhat of a moot point, as even applying the 10-year return approach, Hathaway reports mean market excess return estimates of more than 6%. It is only the subsequent adjustment to these estimates from market data that produces a lower MRP estimate. We examine the adjustments in the subsequent section.

Hancock (2005) [11]

Hancock introduces a number of methodological innovations into the estimation of the MRP. Each of these is examined in turn, with a view to establishing the strengths and weaknesses of each approach and the results of that approach when applied to the most recent data.

Hodrick-Prescott Filter

The Hodrick-Prescott (HP) filter is designed to filter high-frequency noise out of macroeconomic time series. The HP filter is based on the data and a smoothing parameter (lambda) that is selected by the analyst. In Footnote 21 on Page 32, Hancock states that setting the smoothing parameter to 1600 provides superior results (in terms of mean square forecast error) than when it is set to 6400. However, in all of the figures, that parameter is set to 6400. When set to 1600, the most recent value is 5.7%, as shown in Figure 1 below.

---

simply that expected market returns are linked to the risk-free rate at the beginning of the period and not the observed risk-free rate each year.

[11] The South Australian Centre for Economic Studies, "The Market Risk Premium for Australian Regulatory Decisions," April, 2005.

Figure 1

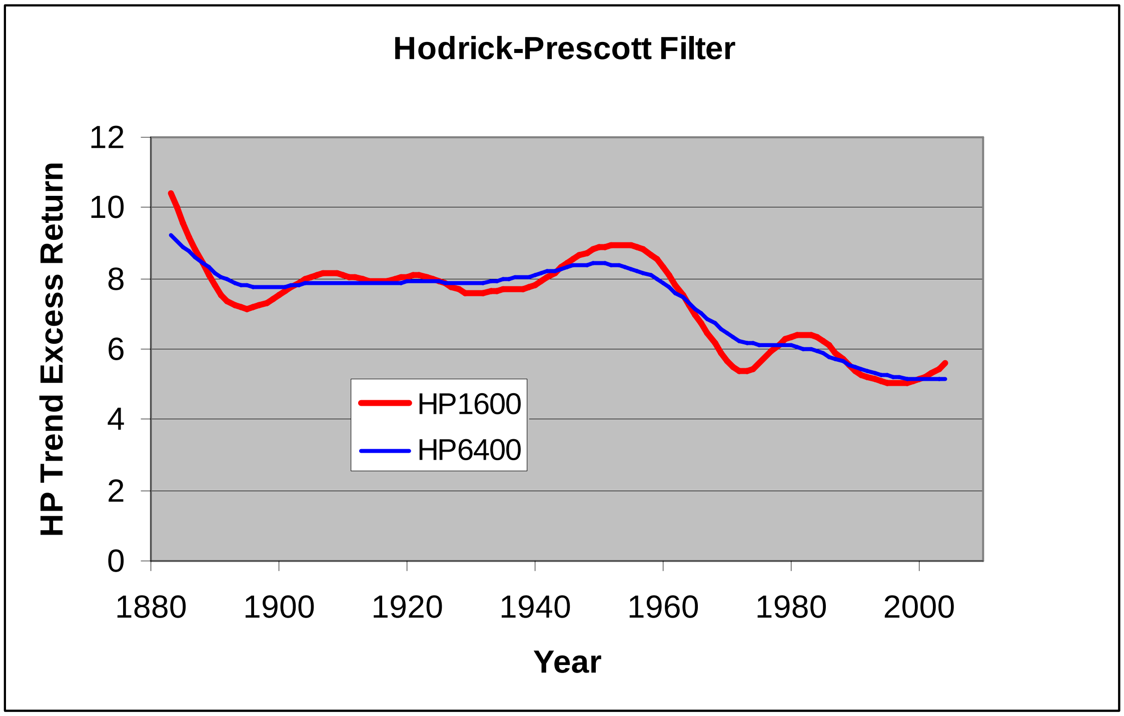

This figure seems to suggest that the average MRP was around 8% for the first part of the sample and has been oscillating around 6% since about 1970.

To determine how much weight should be placed on the estimate from the HP filter, we performed the following simulation experiment. First, we randomly generated a series of 120 data points from a normal distribution with mean 6% (consistent with Australian regulatory precedent) and standard deviation of 17% (consistent with historical data on excess market returns). We then applied the HP filter to the resulting series, setting the smoothing parameter lambda to 1600 and to 6400. We also computed the mean from the last 10 and the last 30 observations. This whole procedure was repeated 10,000 times. The results are summarised in Table 3.

Table 3: Proportion of last MRP estimates in the time series in various ranges when data is N(6,17)

Range	HP 1600	HP 6400	10-year Mean	30-year Mean
< 0	18.55	14.27	12.8	2.46
0 – 1	4.28	4.52	4.38	2.70
1 – 2	4.55	5.40	5.27	4.53
2 – 3	5.15	5.80	6.20	7.04
3 – 4	5.74	6.24	6.86	9.56
4 – 5	6.10	6.68	6.98	11.14
5 – 6	5.80	7.06	7.25	12.69
6 – 7	6.13	6.96	7.34	12.53
7 – 8	5.92	7.21	7.16	11.34
8 – 9	5.30	6.39	6.66	9.21
9 – 10	5.40	5.67	6.22	7.38
10 – 11	4.81	5.12	5.36	4.28
11 – 12	4.21	4.06	4.16	2.24
> 12	18.06	14.62	13.36	2.90

Each simulation contains 120 data points. We apply the HP filter to the series and record the filter estimate at the end of the series. We also compute the mean excess return over the last 10 and last 30 observations in the series. Table 3 reports the proportion of the 10,000 simulations falling into each range. For example, for 18.55% of the simulations, the HP 1600 filter produced a filter value at the end of the series of less than zero.

These results suggest that the results from the HP filter are relatively unreliable and that this technique might not be appropriate for the purpose of estimating MRP from historical data. Even when applied to data that are known to be normally distributed with a mean of 6%, the HP 1600 filter produces estimates, at the end of the sample period, less than zero on nearly one in five occasions. Symmetrically, there is nearly a 20% chance of reporting an estimate above 12% at the end of the sample period – even when the data is known to have a mean of 6%. The present estimate of 5.7%, therefore, certainly cannot be interpreted as providing any support whatsoever for the notion that the MRP is below 6%.

We can contrast these results with the simple mean over the last 10 and 30 years of the sample. There is a degree of academic debate about what length of period should be used to provide the best estimate of the MRP. A long period of data provides better statistical precision (the mean estimate has a lower standard error), but data from long ago may be less representative of current circumstances. It is generally agreed, however, that the minimum period required to provide sensible estimates is 30 years. For example, if annual excess returns are normally distributed with a mean of 6% (consistent with Australian regulatory practice) and standard deviation of 17% (consistent with historical data), there is a 23% chance that 10-years of randomly generated data would have a mean below 2%. Even when we know that the true MRP is 6%, the volatility in the data series is such that there is nearly one chance in four of observing a 10-year average less than 2%. Symmetrically, there is a 23% chance of observing a 10-year average greater than 10%. That is, even when the true MRP is known to be 6%, there is almost a 50% probability of observing a 10-year average lower than 2% or higher than 10%. This probability is less than 20% for a 30-year mean (and less than 10% for a 50-year mean).

Use of monthly data

Hancock develops a technique to use monthly data to estimate mean one-year excess market returns (pp. 7-10). Under this technique, monthly excess returns are assumed to be log-normally distributed. While this assumption can be statistically rejected, it does provide a mathematically convenient way of aggregating monthly returns into an annual estimate, which is what is required for use in the CAPM.

Hancock applies this technique to monthly data from the 30-year period 1974 – 2003 and reports an estimate of 5.6%. He adopts this estimate of "about 5.5%" and subsequently applies two "adjustments" that are described below. Thus, this technique, and the resulting estimate, form the basis of Hancock's final conclusion on the appropriate estimate of the MRP.

Of course, if any technique is to be used, it should be applied to the most recently available data. Hancock's analysis of annual data uses a data set that includes 2004, yet his analysis of monthly data considers the 30-year period from 1974 – 2003. This is probably because the monthly data is released by the AGSM with some time delay and may have been unavailable when the Hancock paper was first drafted. The 2004 data has now been released by the AGSM. Consequently, it is possible to apply Hancock's lognormal estimator to the most recently available 30-year period. We have done this, and we have also replicated Hancock's estimates from the earlier period, and the results are summarised in Table 4.

Table 4: Estimates of MRP using lognormal estimator

Period	**Data Frequency**	**Lognormal Estimator**
1974 – 2003	Monthly	5.6%
1975 – 2004	Monthly	7.6%
1883 - 2004	Annual	7.4%

Table 4 demonstrates quite clearly that Hancock's estimate of 5.6%, which is low by historical standards, is not driven by the statistical technique that is used but by the particular data period that is analysed. Applying exactly the same technique to the most recent 30-years of monthly data produces an estimate of 7.6%. This is economically indistinguishable from Hancock's estimate of 7.4% using the longest annual data set available (see his p. 10). Thus, this statistical technique reinforces the conclusion from Table 1 above that there is no statistical evidence of a decline in the observed excess market return in recent times.

Statistical evidence of a change in the MRP

Hancock addresses the question of whether there is evidence of a change in the distribution of market excess returns in his Section 3. He begins by reviewing a short paper by Gray (2001) which demonstrates that the high volatility in observed excess returns makes it difficult to document statistically significant changes in the mean excess return over time. In particular, Gray (2001) applied a standard test for the difference between two means to various sub-periods of annual excess return data and documented a lack of statistical significance even when the mean excess return in the two sub-periods were economically quite different.

The whole point of that paper is one that we reinforce here – the market excess return series is highly volatile and one must be very careful before concluding that the MRP has changed.

Hancock then goes on to demonstrate the lack of power of the standard statistical tests for the difference between two means. He presents a simulation exercise in which the market excess return is shocked downward in the second sub-period. He demonstrates that the standard statistical test is often unable to detect even quite large changes in excess returns. This is all true, but that is precisely the point – the data are such that one must be very careful before concluding that the MRP has changed.

The same point can be made by another simulation exercise. Suppose that the distribution of market excess returns has not changed over time – what is the likelihood of estimating a mean excess return that is economically different from the current regulatory precedent of 6%. This is precisely what is reported in the last two columns of Table 3 above. There, we randomly generated a series of 120 data points from a normal distribution with mean 6% (consistent with Australian regulatory precedent) and standard deviation of 17% (consistent with historical data on excess market returns). We then applied the HP filter to the resulting series and we also computed the mean from the last 10 and the last 30 observations. This whole procedure was repeated 10,000 times. Table 3 shows that even when the data are known to come from a stationary distribution with a mean of 6%, there is a high chance of estimating a mean excess return less than 2% or above 10%. One must be very careful before concluding, on the basis of a short historical data set, that the MRP has changed.

Conclusions

Standard statistical procedures and both of these two recent papers confirm that any thorough analysis of historical data produces an estimate of the mean market excess return of at least 6%. The relevant results are summarised in Table 5 below.

Table 5: Summary of MRP estimates from historical data

Author	Data Period	Data Frequency	Estimation Technique	Estimate
Historical Data	1975-2004	Annual	Mean	7.70%
Historical Data	1955-2004	Annual	Mean	6.43%
Historical Data	1930-2004	Annual	Mean	6.58%
Historical Data	1905-2004	Annual	Mean	7.15%
Historical Data	1885-2004	Annual	Mean	7.17%
Hancock	1974 – 2003	Monthly	Lognormal	5.6%
Hancock	1975 – 2004	Monthly	Lognormal	7.6%
Hancock	1883 - 2004	Annual	Lognormal	7.4%
Hancock	1883 - 2004	Annual	Exponential Smoothing α=0.2	6.6%
Hancock	1883 - 2004	Annual	Exponential Smoothing α=0.02	7.0%
Hancock	1883 - 2004	Annual	HP Filter	5.7%
Hathaway	1875-2005 (March)	Annual	Mean	7.0%
Hathaway	1875-2005 (March)	10-years	10-year returns vs. 10-year bond yield	7.2%
Hathaway	1960-2005 (March)	Annual	Mean	5.6%

To produce an estimate of the MRP below 6%, it is necessary to apply some sort of downward adjustment to the observed market data. We examine some suggested adjustments, and the justification that has been proposed, in the subsequent section.

AD HOC ADJUSTMENTS TO EMPIRICAL ESTIMATES

Introduction

We have established in the previous section that both Hathaway (2005) and Hancock (2005) confirm that the mean market excess return over various historical periods is 6% or more depending upon the particular empirical technique and data period that is employed.

Both studies then go on to reach a conclusion, apparently independently, that the MRP is 4.5%, rejecting the empirical evidence and the most widely adopted number of 6% in finance industry practice, regulatory reviews and elsewhere. Both studies, again independently, make different adjustments to the data series by leaving out what they consider unrepresentative data in order to reach their conclusion of the 4.5% MRP. It is important to note that without the adjustments to the data their conclusions would readily support a 6% MRP.

Both studies comprehensively review aspects of the estimation issues or problems that arise whenever MRP's or similar rates of return in capital markets are being measured and estimated. The studies offer some interesting alternative means of estimation in the examination of what might be an appropriate MRP and these are reviewed in the previous section. However, the alternative estimation procedures, while interesting and certainly worth examining, do not fundamentally change the estimation of a mean MRP of at least 6% thrown up by the historical data. It is the adjustments the authors make to the data series, or more particularly the information that they believe they are justified in leaving out of the series, that leads them to the conclusion of a 4.5% MRP.

Therefore, it is critical to the point estimate of the MRP, whether it is 4.5% or 6%, to examine in some detail the reasons the authors use to adjust the data and whether these reasons are justified in order to reach their conclusion.

Price inflation adjustment

The Hathaway study adjusts the data for what he refers to as "price inflation for equities" arguing "… that it has added an annual 52bp to the 1882-2005 compound average growth rate and an annual average 145bp to the 1961-2005 compound average growth rate (the price inflation has only happened in the post 1960 data)".[12] The Hathaway Report, reduces the 6% whole period (1882-2005) geometric average to 4.55% on the basis of the 145bp effect of what he refers to as the PER (price/earnings ratio) inflation effect that has occurred in the post-1960 period.[13]

It is difficult to find the reasoning for the adjustment that Hathaway makes other than where he notes that "…we are reasonably confident that the PER inflation is a one-off historical event as it has also happened in the USA market of approximately the same period as in Australia and it only happened once there."[14] This is hardly convincing evidence. To remove or adjust for such an effect, we must be confident that we know exactly how the event or period affected investors' expectations. There are many one-off events in the economic life of the data series and to reject them would leave the series full of gaps. It may be legitimate to omit data which is not likely to appear as part of ex-ante MRP estimates but one would want to see stronger and more objective evidence that the share price inflation of the 1960's was a once off historical event which will not occur again (a profound forecast).

Dividend imputation adjustment

Similarly, Hancock (2005) warns the reader that before inferring anything about the market risk premium from the historical data one should first question whether there are any obvious biases in the data. He refers to three "biases" which he believes may be of concern. The bias which he places greatest weight on in adjusting the data is the price appreciation that he argues resulted from the introduction of dividend imputation in 1987:

---

[12] Hathaway (2005), p. 8.
[13] Ibid, Section 4.4 p. 28.
[14] Ibid, Section 2.2 p.9.

> *The introduction of dividend imputation in 1987 produced a large unanticipated excess return. Over the period July to September 1987 there was an excess return of 21%, far above the 1½% that might have been expected based on the average. Spread over a 30 year period covering the introduction of imputation, it is thus possible that the average excess returns have an upward bias from their true mean of the order of 2/3 %..*[15]

In effect, Hancock has adjusted the results for an economic event which he believes creates bias in a not dissimilar manner to the Hathaway study, although they would presumably disagree on whether to include or exclude this particular event.

But, of course, there are many unique economic events that affect stock returns. To eliminate them all on the basis that they would be non-recurring would be to leave practically a useless data set.

Moreover, the notion that a change in the tax system should be omitted from the data, even though we could expect that the "large unanticipated excess return" would be the market's capitalisation of the benefits of the tax system over a number of years, would lead to rejecting changes in tax rate generally and there have been many of those. There is no valid reason other than the author's conjecture that this particular event creates a bias in the data unless adjusted.

In summary, there are two problems with this "adjustment". First, it is not at all clear that dividend imputation was the cause of these returns. Global markets were all sharply up over this period and the Australian market was pushed higher by a resources boom and substantial takeover activity. Second, there are many other events that could not have been anticipated by investors. Presumably the Asian crisis of 1997 and the terrorist attacks of 2001 were unanticipated. Rather than selectively eliminate from the data events that are considered to be unexpected, the preferred approach is to analyse a longer data set that contains both positive and negative shocks.

---

[15] Hancock (2005), p. 11.

Reduction in discount rates

The second bias that Hancock believes is important is a long term downward move in discount rates. Hancock apparently believes this second bias has led to an unanticipated capital gain on stocks and therefore a bias into an ex-ante MRP. The argument for the long term decline in discount rates (real interest rates) is not dissimilar to the Hathaway's argument about PER inflation. The problem with both these arguments is that they are not backed by any strong theoretical or any empirical justification linking real discount rates to the magnitude of the ex-ante MRP. While it might be tempting to do so, as the Hathaway paper points out[16], to treat the real interest rate and the MRP as linearly related would be to imply on occasions a negative MRP which is a ridiculous number.

In this case, rather than eliminate the data in question (as was done for the high returns in 1987), an ad hoc adjustment of 0.33% is proposed. But again, there are many economic events that affect stock returns. To eliminate those that are claimed to be unexpected and non-recurring would be to leave a scant and practically useless data set. Indeed it is precisely because there are unexpected events that affect markets in different ways that there exists a MRP in the first place!

Effect of adjustments

These two adjustments cause Hancock to adjust the mean historical excess return downward by 1% (two thirds of this in relation to dividend imputation and one third in relation to the change in discount rates). Applying this adjustment to his estimate of "about 5.5%" based on his analysis of 30-years of monthly data yields his final recommendation of 4.5%.[17]

But if these adjustments are to be made to the 30-year average observed monthly market excess return, we should use the most recent data available. This same technique applied

[16] Hathaway (2005), p. 27.
[17] Hancock (2005), p. 13.

to the most recent 30-year period produces an estimate of 7.6%, as documented in the previous section. Therefore, even after applying these adjustments, the estimated MRP would only be reduced to 6.6%, still above the regulatory precedent of 6%.

Moreover, since the majority of this adjustment is argued to be due to the introduction of dividend imputation, it would be inconsistent to ignore the contribution of franking credits to the MRP. Under the CAPM, the MRP includes all forms of returns to equity holders. To the extent that franking credits are assumed to have value, this must also be included in the estimate of MRP. It is quite inconsistent to assume that franking credits have such value that their anticipated introduction drove stock prices up by more than 20%, but then to assume that those same franking credits are irrelevant when they are actually paid.

Summary

In short, without the *ad hoc* adjustments made to the data set by the authors there is little to suggest that their data would not have accepted a mean market risk premium of 6% or more.

We believe that such *ad hoc* adjustments detract from what are otherwise two quite good papers[18] which explore a number of important issues in relation to estimation and forecast problems of using capital market data. Similarly, it would be a mistake to believe (because of the relative depth of the discourse on issues relating to measurement) that the authors' single point estimate of an MRP has much in the way of any credibility given the *ad hoc* nature of the adjustments that were made.

[18] This does not imply we agree with the authors of the papers on all the issues raised. In fact, there are a number of estimation procedures of MRPs concerning the distributional assumptions employed and the time period used to estimate ex-ante MRPs that we disagree with or would qualify but this does not affect the point estimate of 6% we believe is the reasonable conclusion drawn from the data without the adjustments made by the authors.