Bad beta, Goodbye beta:
should governments alter the way they evaluate investment projects in light of modern macro-finance theory?

Andrew Coleman,
New Zealand Treasury.
August 2012

First draft. Please do not quote without the permission of the author.

This note is exploratory and does not reflect the views of the New Zealand Treasury.

Abstract

In the last thirty years, a combination of empirical evidence and theoretical developments has seen a revolution in finance theory. Where it was once understood that stock and bond returns were unpredictable, it is now recognized that they have a significant predictable component at long horizons. Where the CAPM model was once the benchmark for pricing risky assets, multifactor Fama-French models are now used. These developments may have significant implications for the way the Government calculates the weighted cost of capital to evaluate long term investment projects.

This seminar will outline these issues to explore whether the existing approach to capital charging needs modification. It will focus on the methodology used by John Campbell to distinguish between "Bad beta" and "Good beta" - that is, changes in asset values that are associated with long term changes in future earnings, and changes associated with temporary fluctuations in discount rates that affect the valuation of these earnings. This methodology challenges traditional private sector valuation approaches, raising questions about whether the government should distinguish between these two forms of risk when it evaluates investments, particularly if it has no intention of selling some classes of assets.

1. Introduction

The appropriate way for a government to evaluate and rank policies or investment projects that differ in the size, timing and riskiness of their prospective costs and benefits has long been a contentious subject. Since the costs and benefits that occur at different times or in different states of the world are not equally valuable, they must be discounted to a common basis to allow different projects to be sensibly compared. The choice of discount rate will therefore be crucial in determining which (if any) projects proceed. The choice of the appropriate discount rates is anything but straightforward, however. Every evaluation inevitably and quickly confronts a series of questions:

(i) Should the government use a discount rate that reflects average individual or social time preferences, and how are these time preferences calculated for people not yet alive?

(ii) Should the government use discount rates that reflect the returns available from alternative private investment opportunities?

(iii) If these discount rates differ, are there circumstances where one approach or other is preferable?

Despite a prodigious amount of debate, there is still no clear answer to the problem. Governments around the world use both approaches, depending on the project under consideration, or use a compromise methodology.

The New Zealand Government has traditionally chosen to evaluate investment projects by comparing a project's costs and benefits with alternative private investment opportunities. This approach involves discounting the flow of costs and benefits by the "weighted cost of capital," where the weighted cost of capital is meant to reflect the returns private sector agents demand from different types of investment. The intention is to ensure the same framework is used to assess risk irrespective of whether the project is in the private or public sectors. The approach adopted by the New Zealand Government is based on a standard variant of the Capital Pricing Model (or CAPM) developed by Sharpe (1964) and Lintner (1965) that takes into account taxation. The rationale for this approach has been investigated many times, and has found broad backing within the Treasury (Casey, Heerdegen, and Scobie (2007); Lally (1998); New Zealand Treasury, (2005), (2008); Young (2002)).

In contrast, in the last three decades, a large body of financial market and macroeconomic research has raised questions about the use of the CAPM as the benchmark approach for calculating discount rates to evaluate investments. For instance, Fama and French (2004) write

" The attraction of the CAPM is that it offers powerful and intuitively pleasing predictions about how to measure risk and the relation between expected return and risk. Unfortunately, the empirical record of the model is poor—poor enough to invalidate the way it is used in applications."
(Fama and French, 2004, p25)

More pointedly, in his 2010 Presidential address to the American Financial Association, Cochrane (2011, p1087) wrote
"*The first slide in a capital budgeting lecture looks something like this*

$$Value\ of\ investment = \frac{Expected\ \ payout}{R^f + \beta[ER^m - R^f]}$$

with a 6% market premium. All of which, we now know, is completely wrong. The market premium is not always 6%, but varies over time by as much as its mean. (And uncertainty about the market premium is also several percentage points.) Expected returns do not line up with CAPM betas, but rather with multifactor betas to the extent we understand them at all. And since expected returns change over time, the discount rate is different for cashflows at different horizons."

The critique has come from two different but linked directions. Following Campbell and Shiller (1989), one body of research has decomposed the volatility of asset prices into a component reflecting changes in the fundamental earnings of assets, and a component reflecting changes in the rates used to discount these earnings. This literature has not only demonstrated that a large fraction of asset price changes are caused by changes in discount rates, but that many of these discount rate changes are predictable. This empirical finding has many implications for macroeconomics, as it suggests that expected returns, risk premiums and discount rates are time-varying rather than constant, and vary systemically with economic conditions. As such, one reason why the government may not wish to use a *constant* risk adjusted discount rate to evaluate investment projects is that private agents do not. More fundamentally, however, it raises a deeper question. If private markets evaluate investments not only by how much they earn, but by their worth or resale value at different times and different stages of the economic cycle, is it appropriate to require governments to evaluate investments on a similar basis given that many of these investments (such as roads) will never be sold? This note explores this question. In terms of Campbell and Vuolteenaho's (2004) characterization of market beta as "Bad beta, Good beta" according to whether it reflects changes in fundamental earnings or changes in discount rates, the note can be interpreted as asking whether the Government should only evaluate "bad beta" riskiness, that is the beta associated with fundamental earnings.

The second body of research is associated with Fama and French's (1996) three factor model of asset pricing. They demonstrated that an asset's excess returns (compared to a risk free rate) could be explained not only by the covariance of their returns with the overall market (their market beta) but also by two other factors: the size (or market value) of the asset, and the extent to which an asset is a "growth" rather than "value" security. In itself, this demonstration is enough to show that market beta does not provide a sufficient description of asset returns. More significantly, however, these additional factors seem to be systematically related to macroeconomic factors correlated with times of economic distress. Growth stocks have low expected returns (conditional on their market beta) because their valuations fall relatively little in times of market distress, whereas the prices of value stocks are particularly prone to fall during these times. Again this suggests that a component of the value placed on assets by private markets is their ability to maintain their resale value during bad economic times, a factor not necessarily important to Governments.

This note is a preliminary exploration of the topic, and hopes to identify pertinent questions rather than to provide answers. It begins by providing a review of some of the main themes in the recent macro-finance literature that are relevant to how the government chooses discount rates to evaluate risky investments. As suggested above, the twin foci of this review are the literature based on the Campbell-Shiller decomposition and the literature based on the Fama-French three factor model. It then uses these literatures to evaluate the appropriateness of the Government continuing to use a CAPM based framework to value risky investments if it has different objectives and constraints than private agents.

2. The theoretical and empirical critique

2.1 Stochastic discount rates

The key idea of finance theory is that the price of assets is determined by expected discounted payoffs[1]

$$P_t^i = E_t[m_{t+1}(s)x_{t+1}^i(s)] \tag{1}$$

where

P^i_t is the price of an asset at time t

$E_t[\ ]$ is the expectations operator (based on information available at time t)

$m_{t+1}(s)$ is the rate used to discount resources available at time t+1 in state s of the world

$x_{t+1}^i(s)$ are the resources available at time t+1 in state s of the world

The stochastic discount rate $m_{t+1}(s)$ is low for states of the world in which resources are plentiful, and high for states of the world in which resources are scarce. The resources available will be equal to any dividend D_{t+1} paid at time t+1, plus the price the asset can be sold for:

$$x_{t+1}^i = D_{t+1}^i + P_{t+1}^i \tag{2}$$

A risk free asset pays the same amount R^f in all states of the world, so $R_t^f = 1/E[m_{t+1}]$ It follows from the definition of expectations,

$$P_t^i = \frac{E_t(x_{t+1}^i)}{R_t^f} + \text{cov}_t(m_{t+1}, x_{t+1}^i(s)) \tag{3}$$

[1] This section is based on Cochrane (2007)

or that the price of the asset is the expected future value discounted by the risk free rate (the risk neutral present value), plus a factor that adjusts for risk. The extent that the price is below the risk neutral expected future value depends on the covariance between the future value and the discount rate; if this is negative, meaning that the asset delivers few resources when the stochastic discount rate is high (because resources are scarce and particularly valuable), the asset will be valued less than its risk neutral present value.

Let the gross returns to an asset between *t* and *t+1* be

$$R_{t+1}^{i} = \frac{D_{t+1}^{i} + P_{t+1}^{i}}{P_{t}^{i}} \qquad (4)$$

Equation 3 can then be expressed in terms of the excess returns from an asset:

$$E_t[R_{t+1}^{ei}] = -Cov_t(R_{t+1}^{ei}, m_{t+1}) \qquad (5)$$

where $R_{t+1}^{ei} = R_{t+1}^{i} - R_{t}^{f}$ is the excess return to asset *i* relative to the risk free rate.

Using this framework, the fundamental questions in the asset pricing literature are readily apparent: (i) just what are these unobserved stochastic discount factors; and (ii) how can knowledge of them enable people to price new risky assets and investment opportunities?

The first question is being answered by a mixture of empirical and theoretical research. The empirical research has focused on uncovering empirical regularities about the pattern of returns in financial markets that can be used to make inferences about the values of these stochastic discount rates necessary to make equation 5 hold. The theoretical research has proceeded by suggesting possible theoretical relationships for stochastic discount rates, and then ascertaining if these relationships are consistent with the pattern of returns. This research shows it is difficult to reconcile the stochastic discount rates implied by either the CAPM model or the consumption CAPM model with the pattern of asset market returns, and that for this reason neither may be the appropriate way to value new risky assets. This research is discussed in section 2.3 below.

Much of the recent research is based around a rearrangement of equation 4 pioneered by Campbell and Shiller (1989). They decomposed asset price changes into two main components: those associated with changes in future earnings; and those associated with changes in the discount rates used to value these earnings. This rearrangement has shaped another stream of research investigating the cause of macroeconomic fluctuations in asset prices.

2.2 The Campbell-Shiller Decomposition

Equation 4 can be rearranged and then iterated forward:

$$P_t^i = \frac{D_{t+1}^i + P_{t+1}^i}{R_{t+1}^i}$$

$$= \frac{D_{t+1}^i}{R_{t+1}^i} + \frac{D_{t+2}^i + P_{t+2}^i}{R_{t+1}^i R_{t+2}^i}$$

$$= \sum_{j=1}^{\infty} \frac{D_{t+j}^i}{\Pi_{k=1}^j R_{t+k}^i} \qquad (6)$$

Campbell and Shiller note that if the dividend/price ratio is stationary, a Taylor expansion around the logarithm of equation 4 gives:

$$r_{t+1}^i \equiv \ln(R_{t+1}^i) = \ln(D_{t+1}^i + P_{t+1}^i) - \ln(P_t^i)$$
$$= k + \rho p_{t+1}^i + (1-\rho) d_{t+1}^i - p_t^i$$

where

$d_t = \ln(D_t)$

$p_t = \ln(P_t)$

$\rho = e^{\bar{p}-\bar{d}}$

$k = -\ln(\rho) - (1-\rho)\ln(1/\rho - 1)$

Using this approximation and applying to the infinite expansion (6):

$$p_t^i - d_t^i = \frac{k}{1-\rho} + \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}^i - \sum_{j=0}^{\infty} \rho^j r_{t+1+j}^i \qquad (7)$$

This equation says that the price dividend ratio depends on the future growth of dividends adjusted for changes in future asset returns. It follows that if asset prices are high relative to current dividends either dividends will subsequently increase or future returns will be low. In addition, a further rearrangement shows that the difference between actual returns and expected returns is equal to changes in expected future dividends plus changes in expected future returns.

$$r_{t+1}^i - E[r_{t+1}^i] = (E_{t+1} - E_t)\left[\sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}^i\right] - (E_{t+1} - E_t)\left[\sum_{j=1}^{\infty} \rho^j r_{t+1+j}^i\right] \tag{8}$$
$$= \quad N_{CF,t+1} \quad - \quad N_{DR,t+1}$$

Consequently, if returns between t and t+1 are higher than expected, it is either because future dividends are expected to increase, or expected future returns are expected to decrease. One or other must occur.

Equations 7 and 8 have used to motivate a variety of empirical research. First, by taking expectations, equation 7 can be used to explore the link between current price/dividend ratios and the predictability of dividend growth and future returns:

$$p_t^i - d_t^i = \frac{k}{1-\rho} + E_t\left[\sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}^i\right] - E_t\left[\sum_{j=0}^{\infty} \rho^j r_{t+1+j}^i\right] \tag{9}$$

Inverting this relationship, researchers have regressed measures of dividend growth and measures of future returns against the price dividend ratio. In theory, high asset prices could predict increasing dividends or low returns. In practice, they predict returns, but not dividend growth. To someone in the 1970s or 1980s, this would be surprising, as the prevailing view was that high asset prices predict rising dividends[2]. Yet, as Cochrane (1999) and Cochrane (2007) document, the relationship between price dividend ratios and future returns is robust and surprising large, in stark contrast to the relationship between price dividend ratios and dividend growth. A one percentage point increase in the dividend yield leads to a 4 percentage point increase in subsequent returns, rather than the one percentage point increase that would occur in a random walk model. It proves that the dividend yield typically increases because prices fall; they subsequently recover, increasing future returns in a predictable manner.

---

[2] Of course, individual stocks with high prices tend to pay high dividends, and changes in individual stock prices often predate changes in earnings and dividends. The relationship under investigation concerns aggregate returns through time.

More generally, a number of variables forecast aggregate stock, bond, and foreign exchange returns, meaning returns are time varying and to some extent predictable. Across asset classes, a measure of yield proves to be a good forecasting variable; moreover, it tends to be cyclical, meaning that prices are low and expected returns are predictably high when economic times are bad, when people appear less willing to hold risky assets. As Cochrane (1999 p54) notes,
"The price-based predictability patterns suggest a premium for holding risks related to recession and economy-wide financial distress.…Return forecasts are high at the bottom of a business cycle and low at the top of a boom.…to take advantage of the predictability strategies you have to buy stocks or long-term bonds at the bottom, when stock prices are low after a long and depressing bear market, in the bottom of a recession or the peak of a financial panic. This is a time when few people have the guts or the wallet to buy risky stocks or risky long-term bonds. "

These empirical results demonstrate two things relevant to how the government evaluates risky projects. First, discount rates vary through time, suggesting that a discount rate incorporating a constant risk premium is likely to be wrong. Secondly, discount rates vary systematically with the economic cycle, being lower in ordinary times than times of economic distress. The differences can be large, meaning that average risk premiums may be much higher than the risk premiums normally used by the private sector. Boyle (2005) estimated the annual market price of risk for New Zealand between 1970 and 2003. The average over this period was 6.4%; excluding the 4 years between 1988 and 1991, when New Zealand experienced a particularly severe recession and when numerous large companies collapsed, the average for the remaining 30 years was only 3.6%. It follows that if the government uses an average market risk premium to discount projects, it may well be discounting at a rate higher than that used by the private sector except at times of considerable economic stress.

The decomposition of equations 8 and 9 is related to Shiller's 1981 observation that stocks are excessively volatile compared to their fundamental earnings. If stock prices are not varying because earnings are varying, they must be varying because discount rates are varying; and if long run price-dividend ratios are bounded it must mean that price changes stemming from changes in discount rates are subsequently reversed. Lettau and Ludvigson (2001, 2004) use this observation to argue that the effect of changes in wealth on consumption depends on the cause of the change in wealth. Changes in wealth stemming from changes in fundamental earnings should have a much larger effect on consumption than those stemming from changes in discount rates as they are long lasting whereas the latter tend to be temporary, albeit still quite persistent. Empirically, in U.S. data most aggregate changes in wealth are transitory and unrelated to consumption. As discussed in section 2.3 below, this observation is linked to the equity premium puzzle, which notes that the ratio of average excess returns to the standard deviation of excess returns from equities is much greater than readily be explained in terms of consumption volatility. If investments are valued by the market for more than just their ability to deliver consumption benefits, it seems important to understand why they are valued and whether these factors are important to the government.

Campbell and Vuolteenaho (2004) pushed the dividend growth/ discount rate decomposition further by estimating a CAPM model in which the deviation of asset returns from their expected values comprise the two terms in equation 8. Before examining this issue, it is useful to summarise the strand of the literature that has directly documented the empirical difficulties of the CAPM model.

2.3 The CAPM and the Consumption CAPM

The basic CAPM model derives the relationship between the expected returns of a particular asset *i* and other assets when investors only care about the mean and variance of returns, when they have common knowledge about the distribution of returns, and when they can freely borrow or lend at the risk free rate. In these circumstances:

$$E(R^i_{t+1}) = R^i_{t+1} + \beta_{im}\left[E(R^i_{t+1}) - R^f_{t+1}\right]$$
$$\beta_{im} = \frac{\text{cov}(R_i, R_M)}{\sigma^2(R_M)} \tag{10}$$

where *E[ ]* is the unconditional expectations operator

- R^i is the one period return on asset i
- R^M is the one period return on the market portfolio, the value weighted portfolio of all risky assets
- R^f is the one period return on a riskless asset.

β_{iM} is the beta coefficient , measuring the extent that the asset covaries with the market portfolio.

This is an equilibrium relationship: it says that in order to hold an additional amount of risky asset *i* at the margin, investors must be compensated with a return that exceeds the risk free rate by an amount proportional to the extent the asset returns covary with the market. The coefficient of proportionality is the risk premium, the amount that average market returns exceed the risk free rate.

A comparison of equations 5 and 10 shows that the CAPM model is the stochastic discount model where the discount rates $m_t(s)$ are proportional to the return from the market portfolio. If agents are only concerned about the mean and variance of returns, and there are states of the world where expected market returns are very high, these must be states of the world where discount rates are high. In addition, equation 10 states that an asset's beta is the only information needed to identify expected excess returns: all other information is redundant.

This last implication has formed the basis of the tests of the CAPM. If variables other than an asset's beta coefficient (its covariance with the market portfolio) can explain its expected excess return, then the CAPM model is not a complete description of the way assets are valued on financial markets. Two other implications of the model have also been tested: that assets with returns uncorrelated with the overall market should have no excess return; and that these assets should return the risk free rate.

The tests of the CAPM are typically performed on portfolios of similar assets, as this significantly reduces the measurement error involved with estimating an asset's beta coefficient. Once the beta coefficients for different portfolios are estimated, the relationship between the estimated beta and returns is examined. The tests also depend on the definition of the market portfolio, which typically has been average U.S. equity returns, although often wider classes of assets (including bonds, real estate, and non-U.S. assets) are used. In general, early tests using cross sectional data showed that while returns were related to beta, the relationship was too "flat": that is the coefficient on beta was less than average market returns, meaning low beta assets returned more than could be expected from the CAPM theory and high beta assets returned less than could be expected (Fama and French 2004). They also showed that assets with low earnings-price ratios and assets of large stocks typically produced lower returns than can be expected, given their beta coefficients.

Additional research in the 1990s confirmed this analysis. Culminating in the three factor model summarized by Fama and French (1996), researchers demonstrated that two factors in addition to an asset's market beta could explain returns: the size or market value of the company; and the ratio of the book value of the asset to the market value of the asset. Small firms, and firms with high book to market value have returns that are not only higher than large firms and firms with low book-to-market-value, but these excess returns are unrelated to the variance of returns with the overall market (market beta). The differences are large. When 25 portfolios are formed by sorting assets by size and book-to-market-value, the portfolio of small and high book-to-market-value firms outperforms the portfolio of large and low book-to-market-value firms by approximately 0.7 percent per month (see Table 1)

There have been at least three responses to this is evidence. The first is to argue that the evidence does not so much undermine the CAPM as shows that researchers have not been able to capture the appropriate market portfolio that investors actually use. This is possible, but a large range of portfolios have now been used, and even if true it still raises questions as to the appropriate risk adjustment government officials should use when thinking about risk premiums. Few researchers consider this the most likely explanation.

Secondly, there has been a search for more complex, theoretically derived, asset pricing models. It is well known that the CAPM is based on many unrealistic assumptions, such as the assumption that investors only care about the mean and variance of returns. Moreover, it assumes investors only care about their wealth at the end of the period, rather than their consumption possibilities or their future investment opportunities. As early as 1973, the intertemporal CAPM was developed to examine these concerns (Merton 1973); it suggests that investors will be concerned with the covariance of returns with future state variables, such as consumption or the future price of assets. This approach has an obvious affinity with the stochastic discount

Table 1: Mean monthly returns for portfolios sorted by size and Book/Market value (Fama and French 1996: Table 1)

Book/Market value→ Size↓	Low	2	3	4	High
Small	0.31	0.70	0.82	0.95	1.08
2	0.48	0.71	0.91	0.93	1.09
3	0.44	0.68	0.75	0.86	1.05
4	0.51	0.39	0.64	0.80	1.04
Large	0.37	0.39	0.36	0.58	0.71

---

model, and indeed the search for state variables or factors that explain the ICAPM was part of the motivation for the empirical research showing that factors other than market beta explain returns.

Thirdly, there has been an attempt to understand the fundamental reasons why size and book-to-market-value are systematically related to returns. The leading contention is based on the observation that small stocks and stocks with high book-to-market-value do badly at times when the market is particularly distressed. If investors are only prepared to buy stocks that do badly at these times at a price that is sufficiently low that they have high average excess returns, it suggests that investors are particularly concerned about returns in these conditions. This evidence is consistent with the time series evidence that returns are significantly correlated with the state of the economic cycle.

Liquidity factors: Acharya and Pedersen (2005)

A different but related perspective on this issue is from Acharya and Pedersen (2005). They examined how liquidity risk affects asset prices in equilibrium by deriving a liquidity adjusted CAPM where there are price risks stemming from the risk of changes in the liquidity of individual assets and changes in liquidity of the market as a whole. Their theoretical model suggests that investors require a return premium for securities that are illiquid, and an additional premium for securities whose illiquidity increases when the market as a whole is illiquid. Moreover, the model indicates investors should be prepared to accept lower returns for securities that are liquid when the market as a whole is illiquid, or that are liquid when asset prices are low.

Acharya and Pedersen found reasonable empirical support for their liquidity-adjusted capital asset pricing model. They found that stocks that on average had low liquidity had high liquidity risk: that is their liquidity was highly correlated with market liquidity overall. The returns to these stocks were sensitive to market liquidity (prices fell when the market became illiquid, reducing current returns but raising future returns), while their illiquidity was sensitive to market returns (liquidity fell when the asset prices fell overall.) When asset portfolios were sorted by size and liquidity, the lowest liquidity portfolio had a large annual return premium of 4.6 percent. This comprised an average premium of 3.5 percent, reflecting the low average liquidity of the stocks, and an additional 1.1 percent that reflected the covariance of these illiquid stocks with overall market liquidity. The most illiquid stocks were small in the sense of having low market value.

This paper, and the subsequent literature, provides evidence that investors are concerned not just about the earnings of securities, but also about the resale value of these assets in market downturns and on occasions when liquidity dries up. As such the model provides further insights into why the additional factors in the Fama-French three factor model matter. Small stocks, and value stocks (stocks with low book-to-market value) do particularly badly at times of market distress, when illiquid assets are difficult to sell, but when there may be a premium for raising funds.

The Consumption CAPM

The consumption CAPM is a theory that discount factors $m_t(s)$ reflect the marginal value of consumption in different states of the world. In this case, equation 5 becomes

$$\begin{aligned} E_t[R^{ei}_{t+1}] &= -Cov_t\left(R^{ei}_{t+1}, \frac{u'(c_{t+1})}{u'(c_t)}\right) \\ &= \gamma Cov_t\left(R^{ei}_{t+1}, \frac{c_{t+1}}{c_t}\right) \quad \text{if} \quad u(c) = \frac{c^{-\gamma}}{\gamma} \end{aligned} \tag{11}$$

If this theory is correct and consumption is very low (so marginal utility is very high), then the discount rate will be high: i.e. assets that pay off in these circumstances will be particularly valuable. Conversely, assets that pay off when consumption levels are relatively high and marginal utility is low will not be so valuable and will have low prices.

This reasoning is appealing. However, as an empirical matter, the theory appears deeply flawed as risky asset have much higher returns than can be explained by variation in consumption levels for standard utility (preference) models. This, of course is the primary observation of the equity premium puzzle associated with Mehra and Prescott (1985)[3]. The literature this article spawned is vast. Despite some interesting research exploring a wide range of utility functions and financial market frictions, no entirely satisfactory resolution of the puzzle has emerged other than that the consumption CAPM is not an adequate description of the way people value assets.

2.4 Bad Beta, Good Beta (Campbell and Vuolteenaho 2004)

As described in section 2.2, the CAPM model is an application of equation 5

$$E_t[R^{ei}_{t+1}] = -Cov_t(R^{ei}_{t+1}, m_{t+1})$$

in which the stochastic discount rate *m* is the return on the market portfolio.

Using equation 9 to decompose the return on the market portfolio into the components $r^M_{t+1} - E[r^M_{t+1}] = N_{CF,t+1} - N_{DR,t+1}$, a variant of the CAPM can be constructed that has two beta terms and two risk premiums corresponding to the dividend growth (or cashflow) and discount rate terms respectively. The cashflow and discount rate betas are

$$\beta_{i,CF} = \frac{Cov(r^i_t, N_{CF,t})}{Var(r^{eM}_t - E_{t-1}[r^{eM}_t])}$$

$$\beta_{i,DR} = \frac{Cov(r^i_t, N_{DR,t})}{Var(r^{eM}_t - E_{t-1}[r^{eM}_t])} \qquad (12)$$

with $\beta_{iM} = \beta_{i,CF} - \beta_{i,DR}$ (13)

The two beta coefficients are then estimated across 25 Fama French portfolios sorted by size (market equity) and the book-to-market ratio.

[3] Again, see the discussion in Cochrane (2007).

The authors also solve a multi-period intertemporal CAPM that captures the way changes in asset prices stemming from discount rate shocks but not cashflow shocks are subsequently reversed, and thus which captures the persistence as well as the size of different shocks. This model shows the cashflow risk premium should exceed the discount rate risk premium by a factor equal to the agent's relative risk aversion[4]. Discount rate shocks should not affect investors as much as cashflow shocks, because they tend to unwind over time.

To interpret the results, Campbell and Vuolteenaho first estimate a vector autoregression of monthly returns using U.S. data from 1928 to 2001. The cashflow component of returns has a standard deviation of 2.5% per month, half the standard deviation of the discount rate component. Moreover the two components were almost uncorrelated. They then estimated the beta coefficients. While these differed systematically with the portfolios, with growth stocks having lower beta coefficients than value stocks, the average estimated value of the cashflow beta was less than a quarter of the discount rate beta: $\hat{\beta}_{i,CF} = 0.26$ and $\hat{\beta}_{i,DR} = 1.09$. The estimates are presented in Tables 2a and 2b; few of the cashflow betas were statistically different from zero, although all of the discount rate betas were. It therefore appears that not only is the variance of asset prices primarily driven by changes in discount rates rather than earnings factors, but asset prices vary in systematically different ways to cashflow shocks and discount rate shocks.

It is reasonable to ask how much of this difference was due to the economic shocks associated with the Depression and World War II. To ascertain this, Campbell and Vuolteenaho re-estimated over the 1929-1963 and 1964-2001 sub periods. There were significant differences in the two periods; in particular the cashflow betas were much lower in the second period than the first, averaging $\hat{\beta}_{i,CF} = 0.09$ not $\hat{\beta}_{i,CF} = 0.35$, whereas the average discount rate betas were nearly the same. In the second period, the cashflow betas were not significantly different from zero.

[4] They use the Epstein-Zin utility function which distinguishes time preferences from risk aversion. The risk premium associated with cashflow beta is γ times the risk premium associated with discount rate beta, where γ is the relative risk aversion parameter in the utility function.

Table 2a: Estimated cashflow betas for portfolios sorted by size and Book/ Market value (Campbell and Vuolteenaho (2004): Table 4)

Estimated cashflow beta: 1929 - 2001					
Book/Market value→ Size↓	**Low**	**2**	**3**	**4**	**High**
Small	0.36	0.31	0.29	0.3	0.36
2	0.2	0.25	0.26	0.29	0.33
3	0.2	0.22	0.24	0.27	0.35
4	0.14	0.2	0.24	0.26	0.36
Large	0.14	0.15	0.21	0.25	0.3

Table 2b: Estimated discount rate betas for portfolios sorted by size and Book/ Market value (Campbell and Vuolteenaho (2004): Table 4)

Estimated discount rate beta: 1929 - 2001					
Book/Market value→ Size↓	**Low**	**2**	**3**	**4**	**High**
Small	1.45	1.43	1.27	1.22	1.22
2	1.22	1.17	1.08	1.08	1.17
3	1.23	1.04	1.03	0.97	1.15
4	1.01	1	0.94	0.96	1.19
Large	0.92	0.84	0.83	0.91	1

In the final step, the authors estimate the risk premiums associated with each type of risk and beta by regressing the returns on the 25 portfolios against the estimated betas. Unfortunately, the risk premiums were not precisely estimated, although they are qualitatively as predicted. The estimated combined market risk premium (the usual risk premium) was estimated to be 0.68% per month (8.2% annual) with a standard deviation of 0.47%. This is similar to other estimates for the period. The estimated discount rate premium was 0.3% per month (3.6% annual), with a standard deviation of 0.03%; and the estimated cashflow risk premium was 2.6% per month (31% annual) with a standard deviation of 3.1%.

In combination, these results suggest that:

1. most variation in asset prices is driven by shocks to discount rates, not fundamental earnings;
2. the covariance between individual asset returns and market returns (the market betas) is dominated by the common movement to discount rate shocks, not cashflow shocks, as the discount rate betas are much higher than the cashflow betas;
3. the price of discount rate risk is much lower than the price of cashflow risk.

Although the cashflow risk premium was not precisely estimated in the paper, the difference in the risk premiums for cashflow and discount rate beta is considerable. This difference led the authors to describe the two as “bad beta” and “good beta.” The bad beta is the cashflow risk, for while this only comprises a small part of the overall fluctuation in asset prices, it tends to have permanent effects on asset prices and is thus feared. In contrast, the discount rate beta is good beta, for while discount rate fluctuations comprise the largest component of asset price fluctuations, and while the discount rate beta is much larger than the cashflow beta, these fluctuations tend to be temporary, and so have a smaller effect in the long run. It is worth quoting Campbell and Vuolteenaho in full (2004, p1249-1250)

The value of the market portfolio may fall because investors receive bad news about future cash flows; but it may also fall because investors increase the discount rate or cost of capital that they apply to these cash flows. In the first case, wealth decreases *and investment opportunities are unchanged, while in the second case, wealth decreases but future investment opportunities improve. These two components should have different significance for a risk-averse, long-term investor who holds the market portfolio. Such an investor may demand a higher premium to hold assets that covary with the market's cashflow news than to hold assets that covary with news about the market's discount rates, for poor returns driven by increases in discount rates are partially compensated by improved prospects for future returns. To measure risk for this investor properly, the single beta of the Sharpe- Lintner CAPM should be broken into two different betas: a cash-flow beta and a discount-rate beta. We expect a rational investor who is holding the market portfolio to demand a greater reward for bearing the former type of risk than the latter. In fact, an intertemporal capital asset pricing model (ICAPM) of the sort proposed by Robert Merton (1973) suggests that the price of risk for the discount-rate beta should equal the variance of the market return, while the price of risk for the cash-flow beta should be γ times greater, where γ is the investor's coefficient of relative risk aversion."*

Investment horizons (Lettau and Watcher 2007)

Lettau and Watcher (2007) developed a model that explained the different betas and returns of growth and value stocks in terms of the timing of their cashflows. They observed that growth stocks are characterised by having a greater proportion of their (discounted) cashflows in the distant future rather than the near future; or, in other words, that their cashflows have greater duration, similar to difference between short horizon and long horizon bonds. In turn, they have greater discount rate betas than value stocks, for as their cashflows are further in the future they have greater sensitivity to the way these cashflows are discounted. In contrast, value stocks are more sensitive to near term cashflow shocks. Since the price of discount rate risk and the price of cashflow risk are different, a value stock and a growth stock with the same overall beta will have different returns, because their overall beta comprise different mixtures of cashflow and discount rate beta. In particular, since the risk premium associated with discount rate beta is relatively low, as investors fear it less as changes in discount rates tend to be subsequently reversed, the risk premium associated with growth stocks is lower than the premium associated with value stocks. Long horizon stocks may be riskier than short horizon stocks, but investors are more prepared to bear these risks and so don't demand such a high premium.

Lettau and Watcher simulate a model that endogenously derives risk premiums from the duration of earnings associated with different securities. Portfolios created from sorting these stocks by maturity have very similar characteristics to the value and growth portfolios analysed by Fama and French. Their simulated discount rates declined sharply with horizon, from 18 percent per annum for dividends two years into the future to 4 percent per annum for dividends 40 years distant. They note that the sharpness of this decline depends on the correlation between dividends and the unobserved price of risk. They assume this correlation is zero, based on its correlation with the three macroeconomic series that modern theory suggests is linked to aggregate risk aversion: the consumption/wealth ratio, the size of non-housing consumption in total consumption, and the difference between measures of human wealth and outstanding home mortgages[5]. If this correlation is positive, the premium associated with value stocks decreases. Unfortunately, this parameter has not been well estimated in the literature, and more research is needed to get more accurate information on the size of this correlation.

"The results of this paper provide a suggestive and intuitive link between the Campbell-Shiller and Fama-French literatures. As the authors write (p88) *"Value stocks, as short-horizon equity, vary more with fluctuations in cashflows, the fluctuations that investors fear the most. Growth stocks, as long horizon equity, vary more with fluctuations in discount rates, which are independent of cashflows and which investors do not fear. As we show, such a resolution accounts for the time series behaviour of the aggregate market, the relative returns of growth and value stocks, and the failure of the capital asset pricing model to explain these returns."*

3. Implications for Government discount rates

"*The first slide in a capital budgeting lecture looks something like this*

Value of investment $= \frac{Expected\ payout}{R^f + \beta[ER^m - R^f]}$, *with a 6% market premium. All of which, we now know, is completely wrong. The market premium is not always 6%, but varies over time by as much as its mean. ... Expected returns do not line up with CAPM*

[5] For an extended discussion of the link between these macroeconomic variables and risk aversion, see Cochrane (2007).

betas, but rather with multifactor betas to the extent we understand them at all. And since expected returns change over time, the discount rate is different for cashflows at different horizons (Cochrane 2011, p1087).

In this section some of the implications of the macro-finance literature for the choice of the government discount rate(s) are considered. The fundamental contention is that there are two separate reasons why the current method of assessing investment risk is flawed. First, there is no need for the government to copy the private sector approach to risk as it has different objectives and concerns – in particular, it has much less concern about the risk of fluctuations to the resale value of assets. Secondly, even if it were appropriate for the government to have the same framework as the private sector, if it uses discount rates based on historic average private sector risk premiums, it will obtain systematically different results than the private sector as this is not how the private sector evaluates projects. For both reasons, if the government adopts a discount rate based on average market risk premiums, the discount rate will be too high most of the time, possibly by several percentage points.

The macro-finance literature suggests there are three related reasons why the traditional approach by the New Zealand Government based on the capital asset pricing model may not be suitable.

(1) There is no need for the Government to use the same risk premiums as the private sector, as private sector discount rates reflect concern for fundamental earnings risk *and* the risk that changes in discount rates will affect the resale value of the assets. In most cases discount rate risk is of little relevance to the government, as it has no intention, or no need, or no ability to sell the assets. If the government is only concerned about fundamental earnings risk, it should evaluate risky investments using a discount rate that reflects the market valuation of earnings risk (the cashflow beta and the cashflow risk premium), not a discount rate that reflects both cashflow and discount rate risk.

(2) There is no need the Government to use a discount rate based on the *mean* risk premium averaged across different market securities to evaluate the earnings of risky investments. Returns systematically differ by the size and the book-to-market ratios of securities, reflecting the risk that the underlying firms will suffer from significant economic distress. The Government is large and

unlikely to suffer from distress because of the assets it owns; if private sector risk premiums are to be used to evaluate risky investment projects, the premiums should be those pertaining to large and low book-to-market value firms, not the market average.

(3) There is no need the Government to use a discount rate based on the *mean* risk premium averaged across time to evaluate the earnings of risky investments. The risk premium used by the private sector varies substantially over time and may have sharp upward spikes in times of economic distress. If the median risk premium of the private sector is significantly lower than the mean risk premium, and the government uses the mean premium, it will be using a discount rate that is too high most of the time.

These arguments are expanded below.

(1) Cashflow risk and discount rate (valuation) risk.

The literature in section 2 makes a compelling case that the private sector is concerned not just about the earnings of securities or investments, but also about fluctuations in the values of these securities because of changes in discount rates. Illiquid stocks and "value" stocks (those with high book-to market value) have higher returns that can be explained by their covariance with market returns (their beta), and the estimates by Fama and French (1996) and Acherya and Pedersen (2005) of these additional returns are large. The systematic investigation by Campbell and Vuolteenaho (2004) shows that the covariance of individual security and market returns primarily reflects highly correlated changes in discount rates rather than highly correlated changes in fundamental earnings; or in their terminology, that discount rate beta is a large fraction of total beta. But their investigation also suggests that the risk premium on discount rate beta is much lower than the risk premium of cashflow beta.

When contemplating investments, it is not at all clear that the government needs to be as concerned about discount rate fluctuations as private agents. Finite lived individuals are clearly likely to be concerned that the market value of their investments might plummet at exactly the time they need to sell these investments, and are therefore likely to demand a premium for holding them. In many if not most circumstances, the government does not have the same concerns about the resale value of its assets. Many investments such as roads or schools are unlikely ever to be sold. If the government is evaluating or ranking potential investment projects that are not to be sold, it is inappropriate to value them using the same criteria as private markets. If they do so, they will use a higher risk-inclusive discount rate than that which reflects the fundamental riskiness of the projects. If they want to use a CAPM approach, they should use a modified CAPM that only encompasses cashflow risk, the risk that the fundamental earnings of the project will be different than expected. This will be a smaller discount rate than the one used by a private sector that is concerned about valuation or discount rate risk as well as fundamental earning risk.

There may be some circumstances where the government is concerned about discount rate risk. If it were nearly bankrupt, for instance, the resale value or liquidity of an asset may be of concern. If it were managing a large sovereign fund that was expected to be realized at some point, discount rate risk would be important[6]. But these examples seem to be an exception to the rule. If the government uses a high market discount rate rather than a lower one pertaining to fundamental earnings risk alone, they are likely to invest too little or rank projects in an inappropriate order.

It is unfortunate that Campbell and Vuolteenahos' (2004) estimate of the price of earnings risk by is so imprecise. If the government wants to adopt a discount rate based on fundamental earnings risk, it would appear to be a priority for it to estimate the price of cashflow risk, and the size of cashflow beta, so that a more realistic cost of capital can be calculated.

(2) Investor size, and investment duration

Even if the Government decides that it wishes to discount investment projects using a risk premium based on discount rate risk as well as fundamental earnings risk, there is no need for it to calculate using the mean market beta and the mean market risk premium. The private sector does not use a uniform discount rate; instead, mean returns and therefore mean discount rates differ systematically by the size and the

[6] For example, if the government created a fund with the intention of selling it at a later date to temporarily smooth the effects of a demographic bubble, it would appear reasonable to value the fund taking discount rate risk into account. In contrast, a fund being accumulated to permanently fund a save-as-you-go pension scheme is not for sale and thus discount rate risk is unimportant.

book-to-market ratios of securities. The Fama-French three factor model clearly demonstrates that market beta is not a sufficient statistic that explains returns. Rather, other factors reflecting the risk that the underlying firms will suffer from significant economic distress are important. The differences between the returns associated with large “growth” securities rather than small “value” securities are large, perhaps eight percent per annum. The difference between large growth securities and the market average is smaller, but still could be four percent per annum.

The Government is large. The government’s assets have very long duration and therefore have features of growth stocks. The government is unlikely to suffer from economic distress because of its investments, although economic distress is possible if it adopts poor policies. Consequently, there is a prima facie case that if the government wishes to use private sector risk premiums to evaluate risky investment projects, it should use the premiums associated with large and low book-to-market value firms, not the premiums associated with the market average. In turn, this means using a risk premium to discount projects that is likely to be significantly lower than that currently in use.

(3) Time varying returns.

If the Government uses a constant risk premium to evaluate investment options, it is also diverging from private sector behaviour. Consider the estimates of private sector risk premiums for New Zealand from Boyle (2005). He estimated the average risk premium was 3.6 percent for 30 of the 34 years between 1970 and 2003, just over 25 percent between 1988 and 1991, and 6.4 percent overall. Consider how the private sector would approach an investment paying a 5 percent return. It would undertake the investment on 30 out of the 34 years, but between 1988 and 1999 it would have avoided these investments, not just because many companies were experiencing considerable economic distress, but those not experiencing distress would have had better opportunities. In contrast, if the government has adopted the average discount rate, it would have never undertaken the investment. As this example shows, the use of high *average* discount rates results in the blanket avoidance of moderate yielding investments that the private sector would be prepared to undertake on many if not most occasions, based on the reluctance of the private sector to undertake these investments on a few occasions. This seems to be a fundamental misinterpretation of the way the private sector actually behaves.

This issue is potentially quite serious. When risk premiums and discount rates vary through time, as they clearly do, using the average discount rate systematically biases against long duration projects with high returns, for these projects are most sensitive to the choice of discount rate. Since many of the investment projects considered by governments are of this type, the economy will be inefficient relative to its potential. This suboptimal outcome could be avoided by choosing a discount rate based on the lower median private sector risk premium rather than mean premiums.

Governance issues.

One of the reasons for the government to adopt a CAPM framework to assess risk is to ensure that the same evaluation technique is used irrespective of whether the project is in the private or public sectors. Each of the above arguments suggests that the government is justified using lower rates to evaluate risky projects than the private sector. If the government were to do this, there is a risk that it would crowd out the private sector, as it would be able to justify undertaking projects that the private sector could not. Given a general view that economic performance is likely to be compromised if the government sector grows too large, if the government were to adopt lower discount rates, it would need to adopt a different governance structure to ensure the government sector did not become too large.

From this perspective, it appears that a high discount rate is a substitute for sophisticated capital investment governance mechanisms. By ensuring only very high returning projects get approved, a high discount rate automatically deters the government from expansion. Unfortunately this mechanism comes at a high distortionary cost, as it means the economy systematically avoids potentially high returning, long duration projects. To avoid these costs, and avoid the problems of private sector crowd out, the government needs to simultaneously reduce its discount rates in line with actual rather than hypothesized private sector practice, and adopt governance rules to prevent the government from undertaking projects for which it is not well suited.

Better governance need not be a problem. Often the government only uses its discount rate to rank projects, having already established an overall budget. In this case, the choice of a lower discount rate would not lead to an expansion in the size of government, but would avoid the current incentive to invest in relatively low yielding short duration projects rather than higher yielding long duration ones[7]. In other cases, where the government may be tempted to increase total investment activity as a result of using a lower discount rate, the government needs to adopt other procedures to avoid excessive expansion.

4.Conclusion

Many governments around the world, including the New Zealand government, use a variant of the private sector capital asset pricing model as a basis for evaluating the value of risky investments. Their intention, to ensure the same framework is used to assess risk irrespective of whether the project is in the private or public sectors, seems reasonable. Unfortunately the practice is based on an outmoded understanding of the way the private sector actually prices assets. As Fama and French (2004) argued, the empirical record of the CAPM is sufficiently poor to invalidate the way it is used by the government. The government needs to take notice of the empirical and theoretical research of the last three decades and adopt a better approach that takes into account the way the private sector actually behaves.

This note argues that this research suggests the Government should be adopting a lower risk-premium and discount rate than it currently uses. There is clearly some merit in adopting a risk premium that varies with an investment's covariance with the relevant stochastic discount factors. Nonetheless, current practice seems badly misguided. There is no obvious reason why the government should always be concerned with the resale value of its assets, a factor that is very important to the private sector. There is no obvious reason why the government should ignore the fact that it is a large player that invests in long duration assets, factors that reduce risk

[7] There is a second bias towards short duration projects. When the government simply calculates the expected value of costs and returns at different dates, and then discounts these values by a risk adjusted discount rate, it ignores the different covariance between costs and the stochastic discount rate and benefits and the stochastic discount rate. Risk means the discount rates for risky benefits should be higher, but the discount rate for risky costs should be lower, and possibly less than one. See Bazelon and Smetters (2001).

premiums in the private sector. And there is no reason why the government should use time-averaged risk premium when risk premium substantially through time.

Together, these problems with current practice means the government could potentially adopt a real discount rate three or four percent lower than the current recommended eight percent. The exact amount would need proper research, though a reduction in the real discount rate to four of five percent would bring New Zealand back in line with standard international practice. If it were to do so, it would need to adopt stronger governance conventions to ensure the government sector did not unnecessarily crowd out the private sector. If it does not change, it will perpetuate a bias against high return, long duration assets, a practice that has significant potential to undermine the performance of the New Zealand economy.

References

Acharya, Viral V. and Lasse Heje Pedersen (2005) "Asset pricing with liquidity risk," *Journal of Financial Economics* 77(2) 375-410

Bazelon, Coleman and Kent Smetters (1999). " Discounting inside the Washington D.C. Beltway," *Journal of Economic Perspectives* 13(4) 213-228

Boyle, Glenn (2005) "Risk, expected returns, and the cost of equity capital," *New Zealand Economic Papers* 39(2) 181-194

Campbell, John Y. and Robert J. Shiller (1989) " The dividend-price ratio and expectations of future dividends and discount factors," *Review of Financial Studies* 1(3) 195-228

Campbell, John Y. and Tuomo Vuolteenaho (2004) "Bad Beta, Good beta," *American Economic Review* 94(5) 1249-1275

Casey, Anthony, Wayne Heerdegen and Grant M Scobie (2007) "The Public Sector Discount Rate: Review and Recommendations," mimeo, *New Zealand Treasury*.

Cochrane, John H. (1999) "New Facts in Finance," *Economic Perspectives* 23 36-58

Cochrane, John H.(2007) "Financial markets and the real economy," in Rajnish Mehra, ed. *Handbook of the Equity Risk Premium* (Elsevier, Amsterdam).

Cochrane, John (2011) "Presidential Address: Discount Rates," *Journal of Finance* 46(4) 1047-1108

Fama, Eugene F. and Kenneth R. French (1996) "Multifactor explanations of asset pricing anomalies," *Journal of Finance* 51:1 55-84

Fama, Eugene F. and Kenneth R. French (2004) "The capital pricing model: theory and evidence," *Journal of Economic Perspectives* 18(30) 25 – 46.

Lettau, Martin and Sydney C. Ludvigson (2001) "Consumption, aggregate wealth and expected stock returns," *Journal of Finance* 56(3) 815-849

Lettau, Martin and Sydney C. Ludvigson (2004) "Understanding trend and cycle in asset values: reevaluating the wealth effect on consumption." *American Economic Review* 94(1) 276-299

Lettau, Martin and Jessica Wachter (2007) "Why is long-horizon equity less risky? A duration based explanation of the value premium," *Journal of Finance* 62(1) 55-92

Lintner, John (1965) "The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets," *Review of Economics and Statistics* 47(1) 13-37

Merton, Robert C. (1973) "An intertemporal capital pricing model," *Econometrica* 41(5) 867-887.

Shiller, R.J. (1981) "Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71 (3) 421-436

Treasury, New Zealand (2005) *Cost benefit analysis primer (*Wellington: New Zealand Government*)*

Treasury, New Zealand (2008) *Public sector discount rates for cost benefit analysis (*Wellington: New Zealand Government*)*

Seigel, Jeremy J. (1999) "The shrinking equity premium: historical facts and future forecasts," *The Journal of Portfolio Management Fall* 1999 10 -17

Sharpe, William F. (1964) " Capital asset prices: a theory of market equilibrium under conditions of risk," *Journal of Finance* 19(3) 425-442.

Young, Louise (2002) "Determining the discount rate for government projects," *New Zealand Treasury Working Paper* 02/21