Modelling of asset returns and the economy for the 2010 Reference Portfolio Review

Aaron Drew[1]
Guardians of New Zealand Superannuation

Abstract

This paper provides a technical description of the modelling framework used to generate portfolios, sensitivity analyses and performance metrics of interest for the NZSF's 2010 Reference Portfolio Review. Key features of the modelling that distinguish it from 'traditional' mean-variance analysis include the incorporation of mean reversion of risk premia, fat-tailed shocks, and linkages between macroeconomic variables and asset class returns. Mean reversion implies a time-varying pattern to risk, whilst the macroeconomic linkages enable the consideration of specific macroeconomic scenarios and the incorporation of extreme events into the distributions of potential portfolio returns. This provides a rich depiction of the risk and return trade-offs facing the Guardians of New Zealand Superannuation in its core risk profile choice.

[1] The author would like to acknowledge significant input from Tore Hayward (Manager Investment Strategy, ACC) in constructing the returns model for the 2010 Reference Portfolio Review and the expertise and advice of Sue Brake, Richard Frogley, and David Iverson (NZSF). Contact: adrew@nzsuperfund.co.nz.

1. Overview

The enabling Act of the New Zealand Superannuation Fund (NZSF) requires the Fund to periodically review its basic risk profile choice and associated asset allocation decisions. A full write up of the NZSF's 2010 Reference Portfolio Review, which includes motivation for the reference portfolio construct, the risk-profile choice of the Guardians, and how this choice, known as the Reference Portfolio, is used as benchmark for the Fund's actual portfolio and investment activities, is available on the NZSF's website. This paper more narrowly focuses on the modelling framework used for the review.

In Section 2 we outline key features of the model used to generate returns for the asset classes in the Reference Portfolio and then step through the model's equations. Section 3 discusses calibration of the model. In Section 4 three 'extreme' shocks we factored into the analysis for the Review are discussed. In Section 5 we describe how we construct portfolios, which are a function of the asset class returns and other relevant inputs (e.g. taxes, transaction costs, and the modelling of Crown net capital contributions to the Fund). Finally, in Section 6 we present simulation results for the Reference Portfolio and illustrate key model sensitivities.

2. The returns model

At a high level our returns modelling is in the spirit of the approach taken by the Dutch-based APG for their strategic asset allocation exercises (see Hoevenaars *et al.* 2008), which in turn is an extension of the VAR approach outlined in Campbell and Viceira (2005). As in these approaches, the model is a dynamic stochastic system of equations that has a well-defined equilibrium to which variables will converge following a shock.

Perhaps the key difference is that the model has been calibrated to reflect the long-run returns and risk premiums and single period (one year) variance and correlation outcomes, rather than have these outcomes being determined by estimation of the system. The calibration is presented in Section 3. The choice to calibrate in part reflects a desire to impose our 'priors' on these outcomes, and in part reflects the fact that estimation of the model would be a formidable exercise given its size, even using advanced Bayesian econometric techniques.

Key features that distinguish the returns model from the 'traditional' model used for mean-variance analysis, wherein returns are typically normally distributed and independent and identically distributed ('iid'), include mean reversion, fat-tailed shocks, and linkages between macroeconomic variables and asset class returns. These features are discussed in turn.

Mean reversion. Returns converge back to the long-run (steady state) returns and risk premium levels following a shock. An implication of this property is that there are what Campbell and Viceira (2005) term 'time diversification benefits'. That is, the standard deviation of returns over a single year is higher than the annualised standard deviation of returns over a longer horizon (see Section 3 for quantification of these effects). For a long-horizon investor such as the NZSF, this variance compression implies, for a given risk tolerance, a higher allocation to risky (growth) assets than investors with shorter investment horizons.

Fat-tails. We know that the historic returns data is fat-tailed and we think it is important our modelling incorporates fat-tailed returns in order to better describe the distribution of potential fund returns, particularly for metrics concerned with tail outcomes. To generate fat-tailed returns we subject the model to several 'extreme shocks' (described in Section 4) that have been calibrated to provide a stylised depiction of the type of large macroeconomic disturbances we see (fairly rarely) in the data.

Macroeconomic linkages. The model includes dynamic 'macro' equations for output, inflation, short-term interest rates, commodity prices and the real exchange rate. The equations have been calibrated such that their shock (impulse) responses are in line with the type of dynamics seen in the standard 'neo-Keynesian' models often employed in central banks (and elsewhere) for forecasting and policy analysis (e.g. see Clarida et al. 1999, Drew et al. 2008).

The resulting dynamics for short-term interest rates are the building-blocks for fixed interest returns. We also impose correlations between shocks to asset class returns and the macroeconomic variables (e.g. shocks to output and growth asset returns are positively correlated). However, it should be noted that most of the variance in the returns to growth assets is a function of shocks to these assets directly, rather than resulting from shocks to the macro equations. This distinction reflects the fact that it is notoriously difficult to find robust empirical linkages between macro variables and asset class returns. The exception is perhaps in times of extreme stress, which we explicitly capture to motivate fat-tailed returns (e.g. see Piplack and Straetmans 2009 and references therein).

The model has three 'blocks' of equations: one representing New Zealand, one for emerging markets and one for developed markets. We impose restrictions such that New Zealand specific shocks do not, in general, affect developed or emerging markets, whilst shocks to the latter do affect New Zealand. This restriction is based on the assumption that New Zealand is too small to meaningfully affect returns in the rest of the world; an assumption that is commonly applied in modelling of the New Zealand economy (e.g. see Buckle et al. 2007).

Model equations

The model has around 70 endogenous variables, of which around 30 are behavioural equations (i.e. not identities). Despite this relatively large size, the form of the equations is very similar for all growth assets and all fixed interest returns. Below we describe the general form for growth assets and fixed interest returns, and then step through the key macro equations.

Growth assets equations:

$$\overline{EY} = \overline{RR} + \overline{RP} \quad (1)$$

$$EY_t = \theta_1 \overline{EY} + (1 - \theta_1) EY_{t-1} + \varepsilon_{EY\,t} \quad (2)$$

$$RP_t = EY_t - RR_t \tag{3}$$

$$RE_t = RN_t + \theta_2 RP_t + (1 - \theta_2)\overline{RP} + \varepsilon_{RE\ t} \tag{4}$$

Where:

$\overline{EY}$ is the equilibrium earnings yield

$\overline{RR}$ is the real equilibrium short term interest rate

$\overline{RP}$ is the equilibrium risk premium

EY_t is the real earnings yield at time t

RP_t is the risk premium at time t

RN_t is the nominal short term interest rate at time t

RR_t is the real short term interest rate at time t

RE_t is the nominal return at time t

θ_1, θ_2 are fixed parameters

$\varepsilon_{EY\ t}, \varepsilon_{RE\ t}$ are exogenous shocks at time t

In equation (1) the equilibrium real earnings yield is defined as the equilibrium short-term real rate plus the equilibrium risk premium. Real returns for growth assets will converge to this earning yield following a shock. The foundation for this equation is the standard Gordon-growth model under the assumption that the return on equity (or newly invested capital) is equal to the earnings yield.

We specify the dynamic earnings yield (2) as an AR(1) process whose rate of convergence to the equilibrium level depends on the parameter θ_1. The actual return to equities is specified as an excess return to the short-term interest rate. The excess return is again modelled as an AR(1) that converges the risk premium onto the equilibrium level. Note that the equation for short-term nominal rates also converges onto an equilibrium level. This ensures that following a shock, total and excess returns eventually converge to their equilibrium levels.

We impose a high negative correlation between shocks to the earnings yield and returns - a shock that increases returns typically reduces the earnings yield (expands the price-earnings multiple) and hence reduces future returns, all else equal. However, the variance of the return shocks are an order of magnitude higher than the variance of the yield shocks. This implies that although the model encapsulates mean reversion in returns, the process is very noisy. Hence, over short-term horizons given returns generated from this process it would be very difficult to reject the hypothesis that returns were generated by an iid process.

Global sovereign fixed interest equations:

$$\overline{RN} = \overline{RR} + \overline{PI} \quad (5)$$

$$RN_t = \theta_3\left(\overline{RN} + \theta_4(PI_t - \overline{PI})\right) + (1 - \theta_3)RN_{t-1} + \varepsilon_{RN\,t} \quad (6)$$

$$E(RN_{t+N})_t = \sqrt[N]{(1 + RN_t)(1 + RN_{t+1}) \ldots (1 + RN_{t+N})} - 1 \quad (7)$$

$$B_t = E(RN_{t+N})_t + N\,\Delta E(RN_{t+N})_{t,t-1} + \overline{BRP} + \varepsilon_{B\,t} \quad (8)$$

Where

$\overline{PI}$ = the equilibrium inflation rate

PI_t = inflation rate at time t

$E(RN_{t+N})_t$ = expected nominal interest rate for period t + N formed at time t

RR_t = the real short term interest rate at time t

B_t = the nominal return on long − term fixed interest at time t

$\overline{BP}$ = equilibrium bond risk premium

θ_3, θ_4 are fixed parameters

$\varepsilon_{B\,t}, \varepsilon_{RN\,t}$ = exogenous shocks at time t

Equation (6) describes the short-term interest rate as a 'Taylor-type' rule: nominal interest rates are adjusted when inflation deviates from the target rate, subject to an interest rate smoothing constraint (1-θ_3). This ensures that short-term nominal rates and inflation are positively correlated, as in the historical data. In equation (7), an expectation of the path of the short-term interest rate is formed for N periods ahead. Finally, equation (8) describes the return from holding a zero-coupon bond. This bond is best thought of as a sovereign bond given the fixed risk-premium assumption.

Unanticipated changes in short-term interest rates (whether caused by a shock to inflation or interest rates themselves) will, in general, affect the future path of rates as well. This will cause a capital gain (or loss) via the duration term in equation (8), but lower (higher) subsequent returns, all else equal (i.e. a mean reversion effect). This is the source of most of the variance in the returns t sovereign bonds. In contrast, shocks to sovereign bond returns themselves, which may be thought of as representing liquidity shocks, have a relatively minor impact.

As per our assumptions, the overall correlation between sovereign bonds and growth assets is low. However, it should be noted that correlations are both time-varying and shock-specific. For example, under inflation shocks the correlation is higher given (positive) shocks to inflation negatively impact bond and growth asset returns.

The modelling of the returns to credit is slightly more complicated than the modelling of sovereign bond returns. Instead of assuming the risk premium is always at its equilibrium level, we model it as an AR(1) process:

$$CP_t = \theta_5 \overline{CP} + (1 - \theta_5) CP_t + \varepsilon_{CP\ t} \quad (9)$$

Where:

$\overline{CP}$ is the equilibrium credit risk premium

CP_t is credit premium at time t

θ_5 is a fixed parameters

$\varepsilon_{CP\ t}$ is an exogenous shocks at time t

Shocks to the credit risk premium are an important source of the variance in returns to credit. In line with our assumptions, we impose a high correlation between these shocks and shocks to other growth assets. As such, the correlation between the total return to credit and growth assets is considerably higher than the correlation between the return to sovereign bonds and growth assets.

Macro equations:

$$Y_t = \theta_6 \bar{Y} + (1 - \theta_6) Y_{t-1} + \varepsilon_{Y\ t} \quad (10)$$

$$PI_t = \theta_7 \overline{PI} + (1 - \theta_7) PI_{t-1} + \varepsilon_{PI\ t} \quad (11)$$

$$C_t = RN_t + \varepsilon_{C\ t} \quad (12)$$

$$Z_t = \theta_8 Z_{t-1} + \theta_9 (Z_{t+1} + RR_t - RR_t^*) + \varepsilon_{Z\ t} \quad (13)$$

Where:

$\bar{Y}$ is trend output growth and Y_t is output growth at time t

C_t is commodity prices at time t

Z_t is New Zealand's real exchange rate t

$RR - RR_t^*$ is the difference between NZ and foreign real short term interest rate at time t

$\theta_6, \theta_7, \theta_8$ are fixed parameters

$\varepsilon_{Y\ t}, \varepsilon_{PI\ t}, \varepsilon_{C\ t}$ and ε_{Zt} are exogenous shocks at time t

The equations for output and inflation are simple AR(1) processes while the equation for commodity prices can be thought of as very noisy cash. We impose a negative correlation between shocks to trend output and inflation, and a positive correlation between inflation and commodity prices to mimic the impact of a negative supply shock. We further impose a negative correlation between inflation and asset returns, and a positive correlation between trend growth and growth asset returns.

The equation for the real exchange rate is a modified form of the uncovered interest parity condition. The modifications include a lag of the exchange rate and some weight on its equilibrium value. These modifications prevent the exchange rate from acting as a 'jumper'. Instead, humped shaped responses are seen given an unanticipated shock to the short-term real interest rate differential. In other terms, the equations mimic some of the impact of carry on exchange rate behaviour.

Shocks to the real exchange rate determine most of the variance of this variable. We impose a positive correlation between these shocks and shocks to growth assets and commodity prices to reflect the fact that the New Zealand dollar is typically seen as a 'commodity currency' with some 'beta' characteristics.

We add foreign and domestic price levels to the real exchange rate equation to define the level of the nominal exchange rate. Given this, we can define the return to currency hedging as short-term nominal interest rate differentials combined with changes in the nominal exchange rate.

3. Model Calibration

To calibrate the model we require three sets of assumptions:

1) Equilibrium risk and return premia
2) The stand-alone volatility for all model variables
3) The correlations between model variables.

A full description of the methodology we use to pin down the assumptions and their rationale is provided in the Annex of the Reference Portfolio Review. Below we briefly summarise the approach and key assumptions.

Equilibrium risk premia and return assumptions

In the 2010 Review an internally consistent set of equilibrium risk and return assumptions for the various asset classes considered are developed: These assumptions are based on a blend of examining the historical data, empirical research, theory and judgement. In brief:

- First, equilibrium New Zealand and foreign real interest rate and inflation assumptions are established, giving us bill (cash) returns.
- Second, we determine the one-year standard deviation and correlation assumptions for cash and inflation.
- Third, we derive the equity risk premium ('ERP') as the expected return for developed market equities relative to the return for New Zealand Treasury Bills.
- Fourth, a set of expected excess returns (to Treasury Bills) are 'reverse optimised' on the assumption of a representative global investor holding an investable, market cap weighted combination of the various asset classes. Inputs into this problem include variances and correlations for all asset classes considered.

- Finally we determine the expected return for currency hedging from a New Zealand dollar perspective. We assume there is no trend change in the real exchange rate, implying the long run return to currency hedging is the differential between NZ and foreign cash rates.

The long-run expected returns and volatility assumptions for asset classes in the reference portfolio are summarised in Table 1, while Table 2 presents a correlation matrix. The model is calibrated to these assumptions. However, it is important to note that the variances and correlations represent average outcomes from Monte Carlo simulations over a one-year horizon only. As detailed later in this paper, there is a distinct time pattern to variances and correlations can materially change depending on the source of the shock. Finally, as outlined in Box 1 below, given the returns model has long-run mean reversion expected returns over long horizons will differ depending on the basis they are calculated. The model is calibrated such that expected returns over a 30-year horizon are consistent with the arithmetic return assumptions in Table 1.

Table 1: Asset class return and volatility assumptions

Asset classes	Arithmetic expected return (% p.a.)	Volatility (% p.a.)	Arithmetic excess returns (% p.a.)
Foreign bills in local currency	4.5	1.5	-1.50
New Zealand Treasury Bills	6	1.6	n.a.
Global sovereign bonds NZD hedged	6.4	4.5	0.40 0.20
Global credit spread	0.50	3.5	n.a
Global developed market equities NZD hedged	9.5	16	3.50
Emerging market equities NZD hedged	10.5	26	4.50
New Zealand equities[2]	8.5	18	2.50
Global listed property NZD hedged	8.8	16	2.80
Foreign currency return i.e. negative of currency hedging return	-1.5	11	n.a

[2] NZ equity risk premium includes an adjustment for the NZSF's specific tax status.

Table 2: Correlation matrix

	Global equities	Global Govt. bonds	Global credit spread return	Global listed property	Emerging market equities	NZ equities	Foreign currency return	Foreign bill yields	NZ Treasury Bill yields
Global developed market equities	1.0								
Global Govt. bonds	0.1	1.0							
Global credit spread return[3]	0.6	0.1	1.0						
Global listed property	0.8	0.1	0.6	1.0					
Emerging market equities	0.7	0.1	0.5	0.5	1.0				
NZ equities	0.7	0.1	0.6	0.5	0.5	1.0			
Foreign currency return[4]	-0.2	0.0	-0.2	-0.2	-0.2	-0.2	1.0		
Foreign bill yields	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	0.0	1.0	
NZ Treasury Bill yields	0.0	0.0	0.0	0.0	0.0	-0.1	-0.3	0.4	1.0

3 Correlations reported for credit spread returns shown in this table have the opposite sign to changes in the level of the credit spread. Credit bonds are a combination of duration (sovereign bond) exposure and credit spread exposure. The correlation for credit bonds reflects a combination of these two exposures.

4 The foreign currency return is the negative of the currency hedging return. The correlations with changes in the New Zealand currency have the same magnitude but the opposite sign

Box 1 Calibrating Expected Returns

The definition of expected returns we use, denoted below as *Expected Returns*, is over a horizon of 30 years, consistent with the horizon of our projections. To see how this is calibrated in the model, we start by defining the simple gross return for an asset class over 30 years as:

$$R_{30} = \prod_{t=1}^{30}(1 + r_t)$$

The expected value of the simple gross return is:

$$E\{ R_{30}\} = E\{ \prod_{t=1}^{30}(1 + r_t)\}$$

The Expected Return is then defined as the annualised expected return over 30 years:

$$\text{Expected Return} = [E\{ R_{30}\}]^{1/30} = [E\{ \prod_{t=1}^{30}(1 + r_t)\}]^{1/30}$$

If returns are IID, Expected Returns would correspond to the expected arithmetic average return (see Ansley 2002), defined as:

$$\text{Arithmetic average return} = 1/30 \sum_{t=1}^{30} r_t$$

But long horizon mean reversion in the return model implies Expected Returns lie below the expected arithmetic average return. Another common measure of return, the geometric average return, is defined as follows:

$$\text{Geometric average return} = [\prod_{t=1}^{30}(1 + r_t)]^{1/30}$$

It is well documented that the geometric average return lies below the arithmetic average return when annual returns are volatile. The geometric average return also lies below our measure of expected return.

To illustrate these differences, the table below compares these three measures of return, over horizons of one year and 30 years. The returns are for global large cap equities (NZD hedged), calculated from 50,000 Monte Carlo simulations of the returns model (before allowing for extreme shocks).

	One year	**30 years**
Arithmetic average return	10.2%	10.1%
Expected Return	10.2%	9.5%
Geometric average return	10.2%	8.75%

The equilibrium ERP of the returns model is calibrated to generate an expected return for global large cap equities of 9.5%, consistent with the Expected Return shown in Table 1. The same approach applies to the other asset classes.

4. Extreme shocks

The distributions for the shocks in the model's equations are all normal. Given the model is linear this implies that the returns will also be normal. In order to generate more realistic 'fat tailed' returns we add to the normally distributed returns the return outcomes from three 'extreme' shocks that have been calibrated to mimic the sort of historic fluctuations we see in times of macroeconomic stress. The addition of these extreme shocks to the normally distributed returns is randomly determined. In any single year, the probability of an addition of an extreme shock is very low (1-in-30 or less) and we further impose a restriction that only one extreme shock can hit in any 30-year period. As time increases, however, the probability that returns have been affected by an extreme shock, of course, rises.

The extreme shocks modelled are:

- a global credit/financial crisis;
- a large negative supply shock; and
- a large negative New Zealand specific shock.

The first two shocks represent, outside of wars, the largest easily identifiable negative shocks that the global economy has experienced over the past century. The rationale for the New Zealand specific shock is outlined below, along with more detail on these extreme events.

Shock 1: a global credit/financial crises

We simulate the impact of a credit crisis by imposing large negative shocks on output, inflation, interest rates and growth asset returns. Under this shock credit spreads blow-out, negatively impacting the returns to credit. The returns to sovereign bonds are, however, initially positive reflecting the reduction in interest rates and a presumed 'flight to safety'. We also assume that under this shock the New Zealand dollar falls sharply, given the country's reliance on foreign capital. This picture is qualitatively in line with what did occur in the recent financial crisis.

To get a better sense of this shock, Figure 1 below graphs its impact on growth asset returns, fixed interest returns and foreign currency returns and the real exchange rate. On impact of the shock developed market equities, listed property and the New Zealand dollar decline around 40%, while emerging markets equities decline 60%. The credit spread rises to around 500bps, causing the total return to credit to fall over 10%, despite sovereign bonds rallying over 15%. The magnitudes of these initial declines are in line with the sort of declines seen in the global financial crisis.

After the initial impact of the shock the equilibrating forces of the model take over. This results in an elevated period of returns to the growth assets as earnings yields take some time to normalise back to their equilibrium levels. In contrast, the returns to sovereign bonds and foreign currency are depressed for some time given the equilibration of interest rates and the exchange rate back to their equilibrium levels.

Figure 1: Stylised credit crisis

1A: Growth asset returns

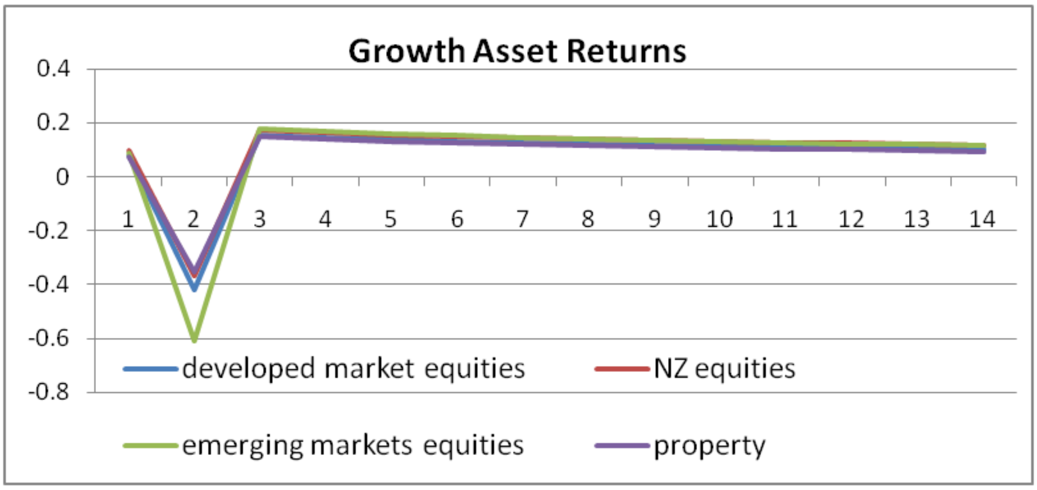

1B: Fixed interest returns

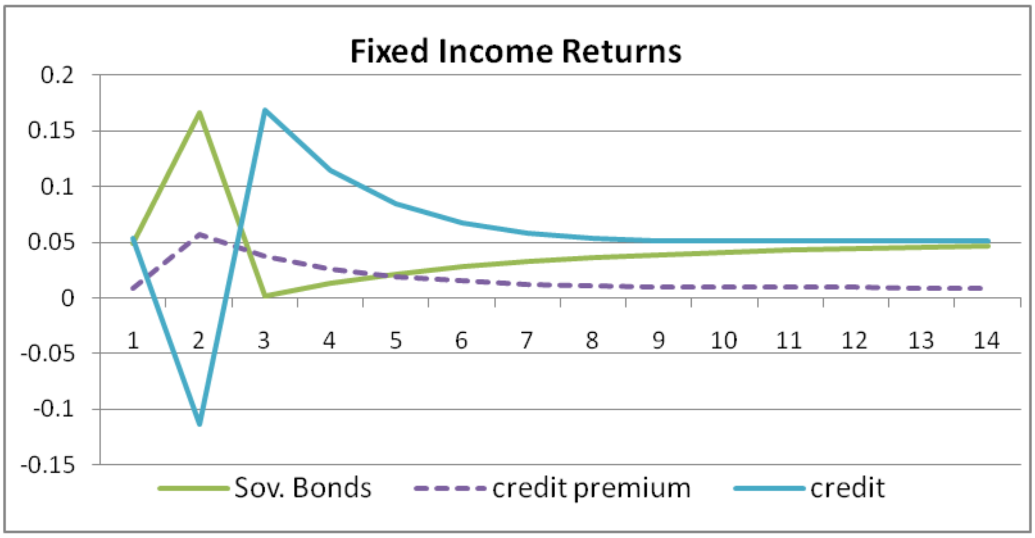

1C: Currency returns and the real exchange rate

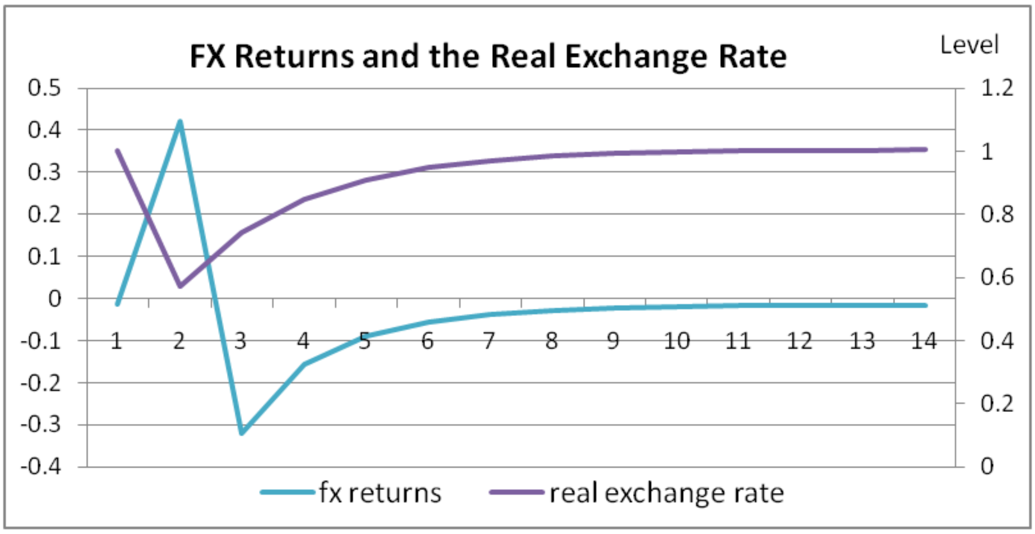

An important assumption of this and the other extreme shocks we consider is that they have no permanent impacts on risk premiums or returns. That is, the mean reversion forces in the model are not diluted by the imposition of these shocks. In one sense this may underplay the impact of such extreme shocks – it is certainly possible that corporate earnings and GDP get permanently 'knocked-off' their pre-shock trend growth paths even if growth rates in these variables themselves are not permanently affected (see Paniza *et al*. 2009). On the other hand, the elevated returns we see in the model following the initial impact of the shock are consistent with market behaviour typically observed following large declines.

Shock 2: A large negative supply shock (rise in global inflation)

The variance of the normally distributed shocks we apply to inflation and short-term interest rates are consistent with the sort of range we would see in these variables when inflation is low and stable and central banks are perceived as being committed to keeping it that way. However, history suggests we should not be sanguine – large negative supply shocks can occur, and inflation and inflation expectations can become de-stabilised.

We simulated the impact of such an event by applying large positive shocks to inflation and commodity prices, and large negative shocks to output, growth asset returns and the New Zealand dollar. We assume that interest rates initially decline as central banks worry more about the adverse output consequences of the shocks. However, further into the simulation they rise and stay above equilibrium levels for an extended period in order to generate positive real interest rates and reduce inflation. Empirical support for this type of policy response is seen, for example, in Cologni and Manera (2008).

Under this shock, all fixed interest and growth asset returns are negatively affected. The initial impact on growth asset returns and credit is much larger than the impact on sovereign bonds; however, these asset classes recover more quickly than sovereign bonds, which suffer a longer period of underperformance given the increasing interest rate environment. This pattern is qualitatively similar to what occurred after the first oil shock in 1973-4.

Shock 3: A New Zealand specific shock

New Zealand is a small open economy with a relatively undiversified production base, highly dependent on trade and international capital flows. Like any sovereign nation, laws may be changed at the local and national level, and such changes may affect the returns from investing in New Zealand assets. Much of the major population centres lie in geographically unstable regions, subject to risk of earthquakes, tsunamis and/or volcanic eruptions. Around 60% of New Zealand's export bundle is concentrated in a few key commodities, all of which have the potential to suffer catastrophic losses from bio-security risks and/or the imposition of international trade barriers. In short, there are a number of factors that give rise to idiosyncratic risks to the New Zealand economy and holding New Zealand assets.

We simulate the impact of a 'representative' New Zealand idiosyncratic shock by applying large negative shocks to New Zealand output, interest rates, the exchange rate and equities. These shocks, however, are presumed to have no impact on the rest of the world.

Impact of the extreme shocks on return distributions

At a conceptual level, we view our assumptions for returns and variances as applying to a world where extreme shocks can occur and where returns are fat-tailed. As such, we calibrate the normally distributed shock to the model such that when the extreme shocks are added we get back to our assumptions. In Table 3 below we present moments for the assets in the Reference Portfolio before and after the addition of the extreme shocks. We see that the average long-run returns are approximately the same, but the variances increases with the addition of these shocks, returns are negatively skewed and kurtosis is non-normal. Note that these latter outcomes are over a single year period.

Table 3: Impact of the extreme shocks on asset class moments

	Global Equities	New Zealand Equities	Property	Fixed Interest
Normal shocks				
Mean (long-run)	9.5%	8.5%	8.6%	6.6%
Std deviation (single year)	15%	16%	15%	4.0%
Skew (single year)	0.0	0.0	0.0	0.0
Kurtosis (single year)	3.0	3.0	3.0	3.0
Normal plus extreme shocks				
Mean	9.1%	8.4%	8.6%	6.7%
Std deviation (single year)	16%	18%	16%	5.0%
Skew (single year)	-0.5	-0.3	-0.2	-0.5
Kurtosis (single year)	4.3	3.6	3.4	5.9

Time diversification benefits

As discussed above, the returns model incorporates long-run mean reversion in returns and this property can be expected to deliver time diversification benefits in the sense that the annualised standard deviation of outcomes over a multi-year period will be lower than in a single year period. To quantify this impact Figure 2 presents time varying moments. We see that the annualised standard deviation of returns fall through time for all asset classes except Treasury Bills ('NZ Bills'). This fall does not, however, occur indefinitely. From around the year 15 annualised standard deviations stay fairly stable through time, reflecting that the mean reversion forces are medium-term in nature.

As in Table 3, asset class returns are initially negative skewed and fat-tailed. There is also a positive skew in the currency hedging returns and the credit spread given in the extreme shocks the New Zealand dollar and credit spreads blow out. As we move through time, however, returns become approximately log-normal. Kurtosis converges on a value of three as per a standard normal distribution while returns become positively skewed. As such, the impact of the extreme shocks on the shape of the return distributions is short-lived. Over a medium to long-term horizons the distributions appear approximately log-normal.

Figure 2: Moments of assets class returns over time

2A: Standard deviation of returns over time

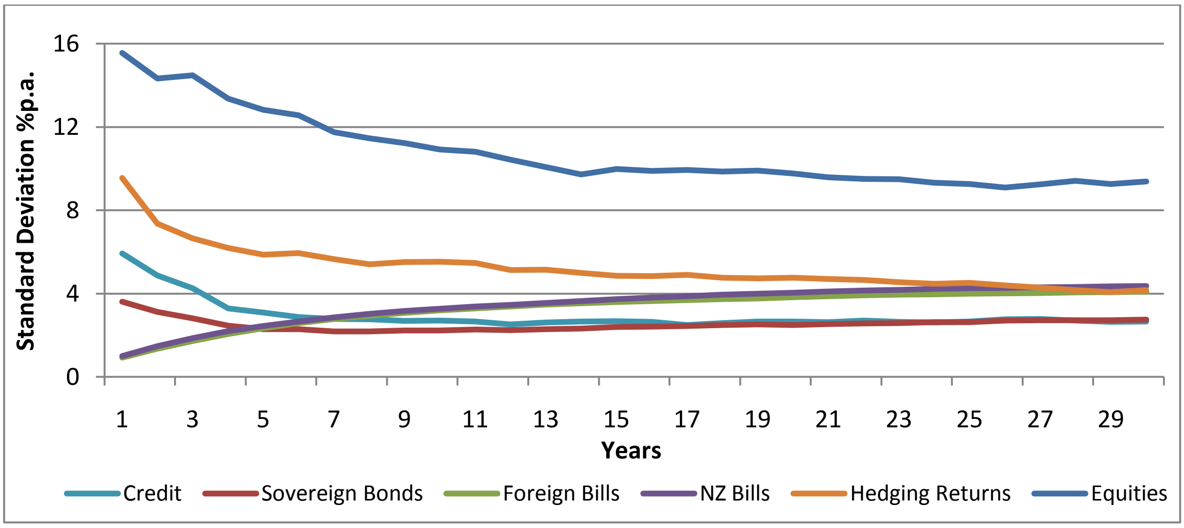

2B: Skewness in returns over time

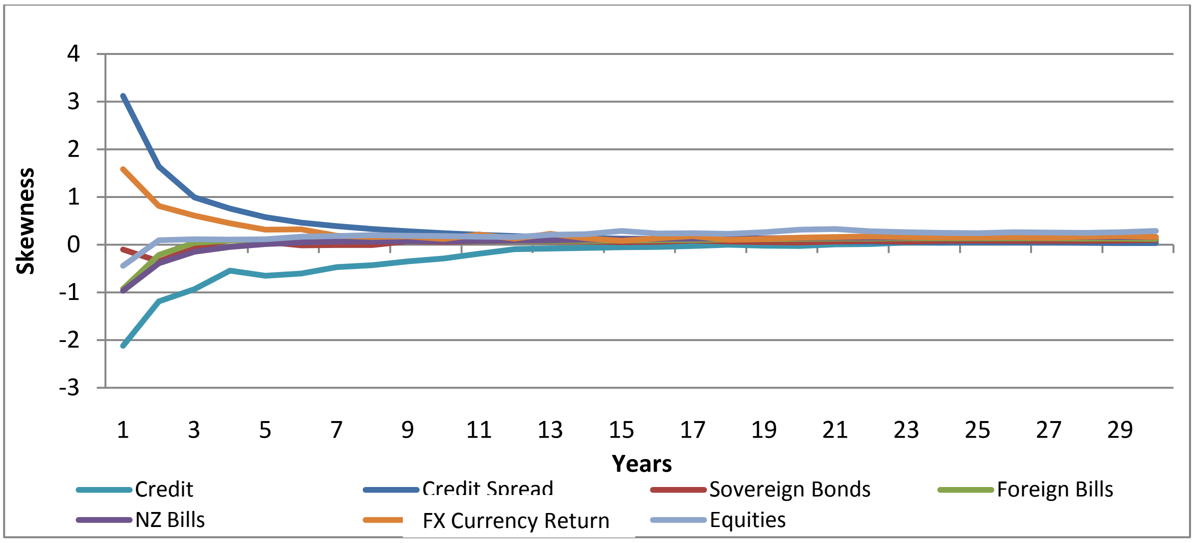

2C: Kurtosis in returns over time

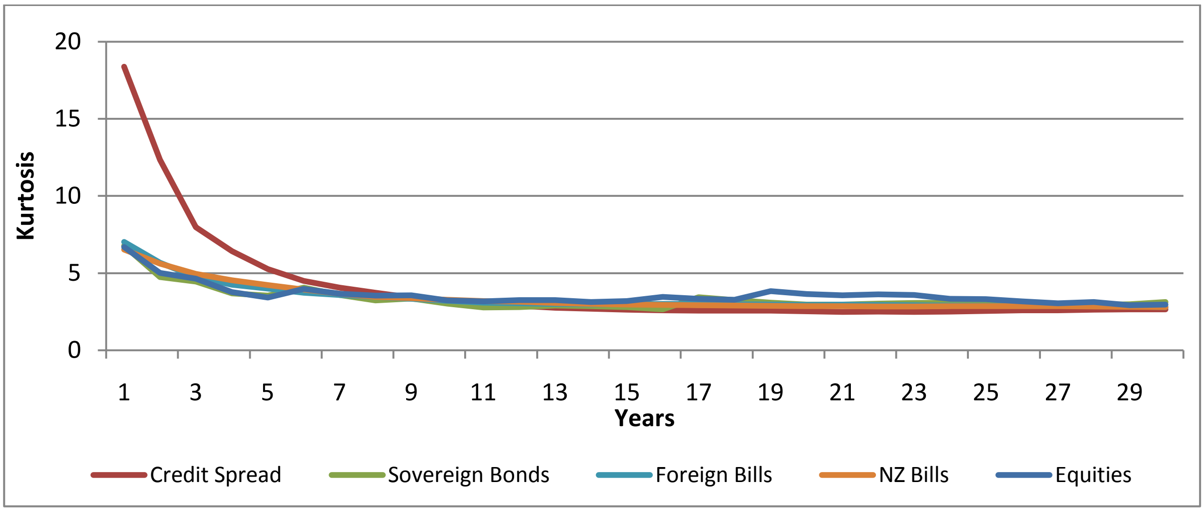

5. Generation of portfolio outcomes

In this section the methodology for generating expected portfolio outcomes or projections is described. We first discuss the approach to modelling the distribution of potential future outcomes in the presence of cash flows (including tax) between the New Zealand Government (Crown) and the Fund.[5] We then discuss our simulation assumptions.

Monte carlo modelling and incorporate Crown cash flows

There are significant cash flows between the Fund and the Crown. These flows arise from the Fund's status as a taxpayer and Crown capital contributions. The latter were suspended in the May 2009 Budget and it is the present government's policy to resume contributions when New Zealand's fiscal position improves from deficit to surplus. These flows present a challenge for the interpretation of projected future Fund outcomes. In the absence of cash flows, there would be a more direct relationship between the future value of the Fund and cumulative returns over the projection horizon. But large cash flows into and out of the Fund weaken this link.

There are two important relationships between the Fund returns and projected net Crown cash flows. One reflects the Fund's status as a taxpayer: higher returns generate higher tax payments to the Crown. The other reflects the linkage between Fund returns and the amount of gross funding from the Crown (from when it resumes). The funding formula used in the government's budget process has a link to Fund performance[6]. All else constant, if

---

[5] It is assumed that the Reference Portfolio takes the place of the Fund in determining flows to and from the Crown. We also measure Reference Portfolio outcomes as if the Reference Portfolio were an actual portfolio holding all of the Fund's assets.

[6] The funding model is designed to smooth over time the total Crown appropriations for NZ superannuation (i.e. the combination of meeting net NZ superannuation entitlement payments plus net flow into the Fund). The formula is described in McCulloch and Frances (2001).

Fund returns turn out stronger than currently projected, the Fund balance will turn out higher than currently projected. The funding formula compensates for a higher Fund value by generating lower gross contributions. Conversely, if Fund returns are weaker than projected, all else equal, the funding formula will generate higher gross contributions.

The funding formula reduces volatility in the projected future Fund value. But the funding formula also dampens the relationship between Fund returns and the future value of the Fund - in short, we need to do more than just look at the projected future Fund value.

The New Zealand Treasury produces a spreadsheet that projects out, over many years, the flows between the Crown and the Fund, and the growth in the value of the Fund over time[7]. This incorporates the funding formula, which relates the size of the contribution to the Fund in a given year to variables such as expected future NZ superannuation entitlement payments, growth in nominal GDP and Treasury's expectation for the future returns of the Fund. This model is 'deterministic', i.e. fund returns are held constant over the projection horizon.

We use the logic of this Treasury spreadsheet, and its key assumptions, as the basis for our monte carlo projections. The approach taken is to:

- incorporate the model of stochastic returns described in section 2;
- build up the stochastic asset class returns into total portfolio returns, taking into account the portfolio target weights, rebalancing back to these target weights, and transaction costs; and
- generate, for each monte carlo 'trial', a unique pattern of cash flows and the accumulated value of the Reference Portfolio. We run a large number of trials for each projection and collect up summary results that reflect the distribution of outcomes. Each projection lasts 30 years.

This modelling framework provides a dollar value for the portfolio at the end of the period for each monte carlo trial. We calculate a corresponding *Excess Dollars* by subtracting from the Fund value the net cash flows contributed by the Crown, with these net cash flows compounded forward at an 'opportunity cost return', which is taken to be the return on Treasury Bills. We express the Excess Dollars at the end of the period both in real terms (in 2009 dollars) and as a percentage of GDP (at the end of the horizon).

Simulation assumptions

Assumptions made in simulating portfolio outcomes include

i. *Transaction costs.* These include both brokerage and market impact costs. Transaction costs are set separately for each asset class. We assume one-sided (i.e. for a purchase or a sale, rather than a 'round trip') transaction costs of 0.25% for

[7] This spreadsheet and model guide can be downloaded from:
http://www.treasury.govt.nz/government/assets/nzsf/contributionratemodel/nzsf-model-v18.xls and
http://www.treasury.govt.nz/government/assets/nzsf/contributionratemodel/nzsf-model-guide-may09.pdf

global equities, 0.10% for fixed interest, 0.50% for New Zealand equities and 0.35% for property.

ii. *Taxable income method.* The taxable income method is the way taxable income is assessed for the purpose of determining tax payable by the Fund. Some assets are assumed to generate taxable income based on the change in their market value (the standard regime). Other assets are assumed to generate taxable income based on their average market value (the fair dividend regime). The method is set separately for each asset class. We assume global equities and property are assessed under the fair dividend regime. We assume fixed interest and New Zealand equities are assessed under the standard regime.

iii. *Tax rate.* Tax is payable at a fixed percentage of taxable income in each year. This tax rate is 30%.

iv. *Treasury funding model assumptions.* The Treasury's funding model assumes an 8.65% pre-tax return, and a 24% effective tax rate. We also take the Treasury's assumptions for forecasts of GDP and inflation.

v. *Start year of funding.* Zero funding is assumed until year 9 of the simulation for the Reference Portfolio, reflecting the funding freeze in place at the time of the Reference Portfolio Review.

vi. *FX hedging.* Foreign currency exposures are fully hedged for the Reference Portfolio.

Note that implications of changes in this the resumption of Crown funding and the degree of FX hedging are explored in the Reference Portfolio Review.

6. Simulation results and key modelling sensitivities

In this Section simulation results for the Reference Portfolio are shown along with how these results are affected by some of the key modelling assumptions, including the modelling of shocks and the degree of mean reversion in returns.

The Reference Portfolio

The Reference Portfolio embodies the basic risk profile choice of the Guardians of New Zealand Superannuation – a choice that has been taken to meet the Fund's investment purpose, taking into consideration key 'endowments' such as its long-term investment horizon and tax status. In Table 4 below we see the market exposures (asset classes) in the Reference Portfolio and the weights placed on these exposures.

Table 4: The Reference Portfolio

Exposure	Benchmark	Exposure Weight
Global equities NZD hedged	MSCI All Country World Investable Market Index NZD hedged	70%
New Zealand equities	NZX50 Capped index (custom NZSF Index)	5%
Global property NZD hedged	FTSE EPRA/NAREIT Developed Index NZD hedged	5%
Global fixed Interest	Customised index including: Barclays Capital Global Aggregate Index High Yield Debt and Inflation Indexed Bonds, NZD hedged.	20%
Foreign currency exposure	Not applicable	0%

Key summary measures from monte carlo projections of the Reference Portfolio are seen in Table 5. Over long horizons (30 years), it is expected to generate Excess Dollars of around $50b (in 2009 dollars) or around 12.5% of (2040) GDP. The distribution also suggests there is a fairly low probability (around 6.5%) of generating negative Excess Dollars over this long-term horizon. However, over short-term horizons the potential for significantly low short-term returns ('bumps along the way') appear are all but inevitable. For example, the probability of a negative return for at least one rolling 3-year period over the projection horizon is 100%. Over the first three years of the simulation this probability is around 12%.

Table 5: Key performance metrics for the Reference Portfolio

Long-term (30 year) outcomes: Excess Dollars (over T-Bills)		
	Excess Dollars	Excess Dollars (% 2040 GDP)
Expected Excess Dollars	49.4b	12.4%
probability Excess Dollars<0	6.5%	6.5%
5% of outcomes are below	-2.8b	-0.7%
Short-term risk: probability 3-year time weighted returns below thresholds		
Threshold	First 3-years	Any 3 year period
0%	12%	100%
-5%	4%	60%
Treasury Bill	35%	100%
New Zealand inflation	18%	98%

Part of the reason that the 'bumps along the way' are large is that returns (and by implication measures of value added) are negatively skewed and fat-tailed at short-term horizons (see Figure 2). However, as time progresses the distribution of returns significantly changes, as would be expected from the behaviour of individual asset class returns discussed. In particular, returns over longer horizons benefit from mean reversion. We can see from Figure 3 that the annualised standard deviation of returns for the Reference Portfolio compresses over time, from around 13% over one year to around 8% for horizons of 20 years or longer.

Figure 3: Standard deviation of returns for the Reference Portfolio over time

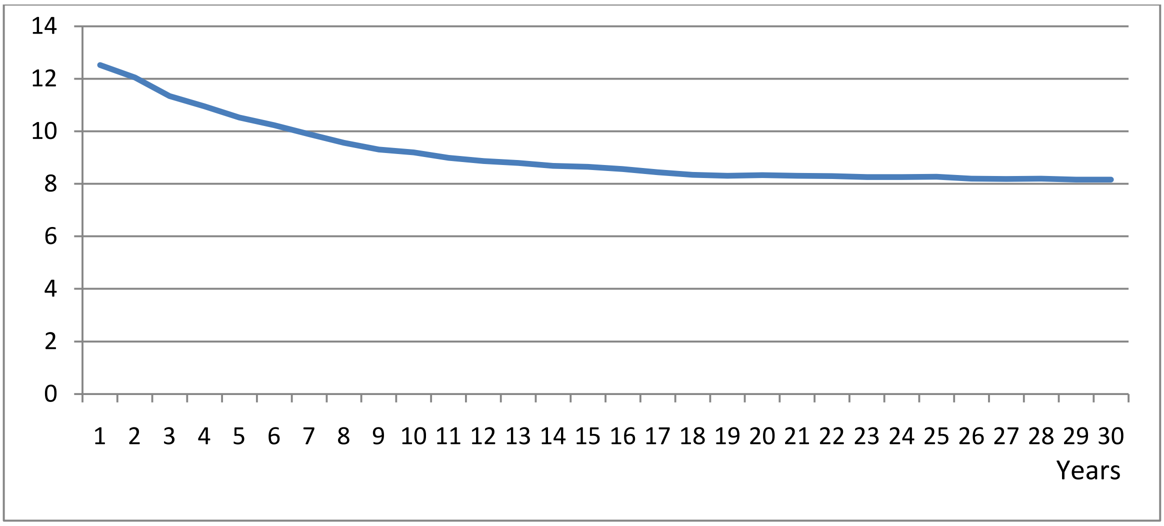

Sensitivity to the modelling of returns

In this section we provide sensitivity analysis around the Reference Portfolio to alternative modelling of the shocks and the degree of mean reversion in returns. Three alternative scenarios are considered:

Normal. Returns are normal, i.e. no extreme shocks are factored into the returns. In this scenario we have increased the variance of the normally-distributed shocks such that variances in a single year approximately match those of the base case.

Time independent. Returns are normal and there is no time dependence (mean reversion) in the pattern of returns. In this scenario we break the linkages between current yields and future returns in the returns model described in Annex E. This implies returns over time are independent and identically distributed (although there are still correlations between asset class returns in a single period).

Weak equilibrium. Returns are fat-tailed, as in the base case, but shocks have much more persistent impacts on yields and therefore returns. This scenario has been run to proximate what would occur if equilibrium returns and risk premiums were time varying rather than constant as in the base case and first two of these scenarios.

To provide a high level idea of the impact of these alternative modelling choices, Figure 7 shows moments (standard deviations, skewness and kurtosis) over time for the Reference Portfolio (the 'base case') and for the variants above.

The standard deviations of returns in the first year are similar, but as time progresses divergences occur. For the Reference Portfolio and normal returns variant the standard deviation of returns fall through time. This is due to long-run mean reversion in returns, as can be seen by comparing these results to the case where returns are normal and there is no mean reversion ('weak equilibrium'). When shocks to yields have persistent effects the standard deviation of returns initially compress, but further into the simulation they start increasing through time.

This is due to the mean reversion effect eventually being swamped by the impact of the persistent shock; in this scenario it is possible that the *average* level of returns drift from their equilibrium levels with time.

The initial skew and kurtosis in returns is non-normal for the two cases where the extreme shocks are incorporated into the analysis ('weak equilibrium' and Reference Portfolio 'base case'). However, with time returns become log-normal in all cases. The positive skew is largest in the cases where there is little or no mean reversion in yields ('weak equilibrium' and 'time independent').

In Figure 8 over the page we show the implications of these time-varying characteristics for the distribution of portfolio Excess Dollars relative to Treasury Bills, in both the first year of the simulations and at the end of year 30. In the first year outcomes are similar except for the left-tail of the distributions which show worse outcomes for the Reference Portfolio ('base case') and the 'weak equilibrium' case given the presence of extreme shocks. At year 30, however, the distributions differ markedly:

- The distribution of net valued added for the 'base case' versus 'normal' case is similar except for the upper percentiles. This is due to the fact that in the base case it is possible an adverse extreme shock hits the latter years of the simulation period, thereby reducing the simple time compounding benefit that the positive risk premium would otherwise confer.
- The distribution of outcomes for the scenarios where the standard deviation does not compress with time is much wider, particular on the right hand side of the distributions. This reflects the fact that in these simulations when yields are elevated relative to equilibrium there is little or no implications for future returns. Given returns are log-normal this implies that long-run outcomes are more positively skewed.

Overall, the results of these sensitivity analyses suggest our choice of modelling returns as fat-tailed with mean reversion is relatively conservative - short run outcomes show the potential for larger downside losses whilst long-run outcomes do not permit as large a positive skew as what can occur when there is no mean reversion and/or no potential for extreme shocks in the final years of the projection period.

Figure 7: Time varying moments under differing returns models

7A: Standard deviation of returns over time

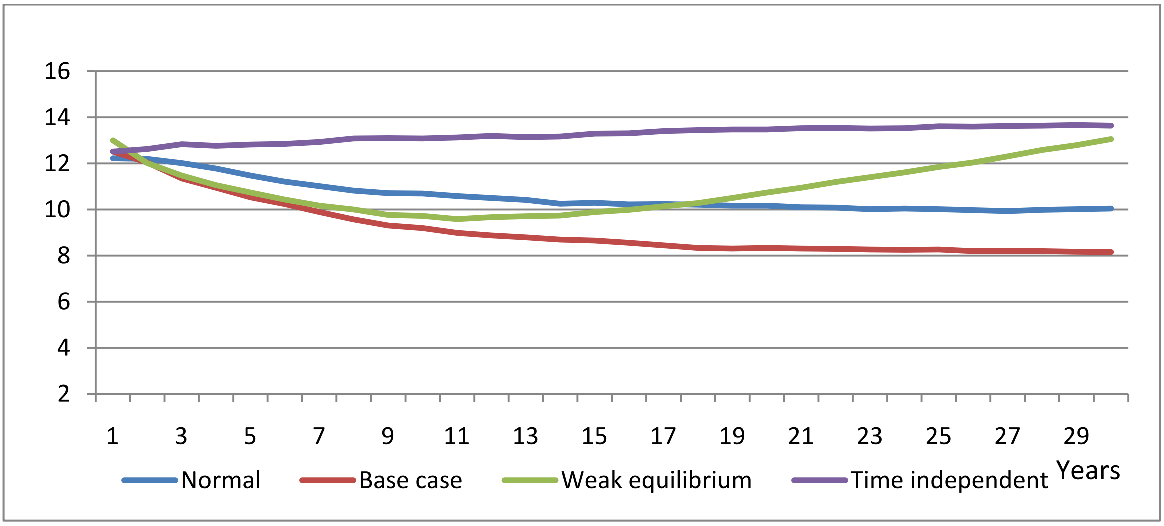

7B: Skewness in returns over time

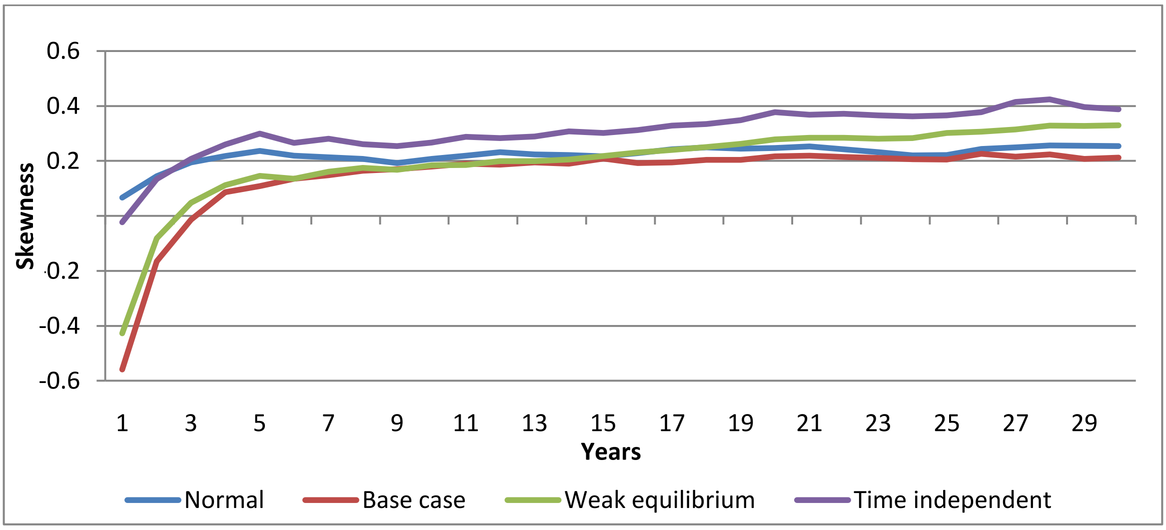

7C: Kurtosis in returns over time

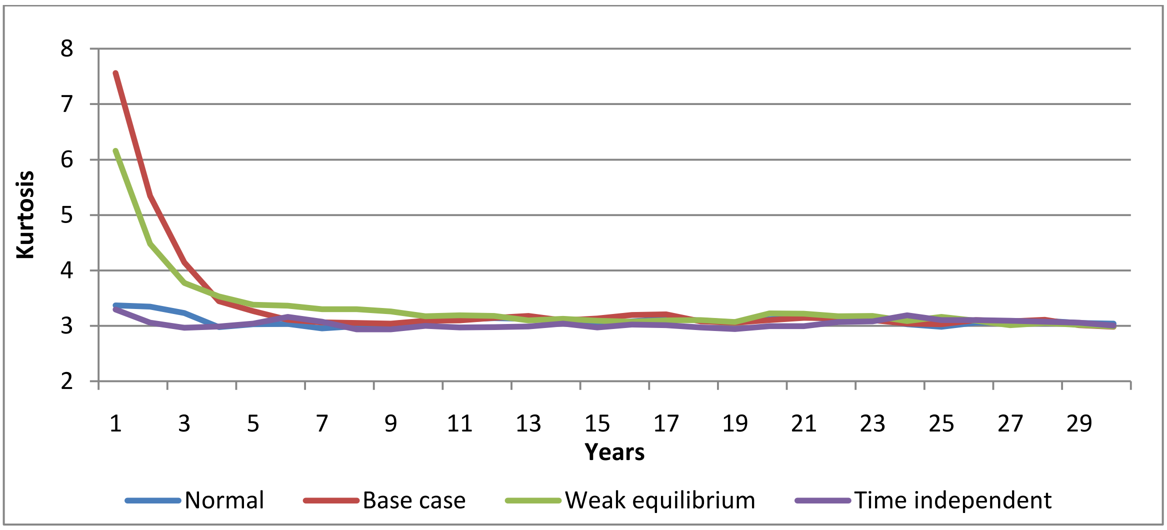

Figure 8: Excess Dollars for different return models

Relative to Treasury Bills in 2009 NZD billions

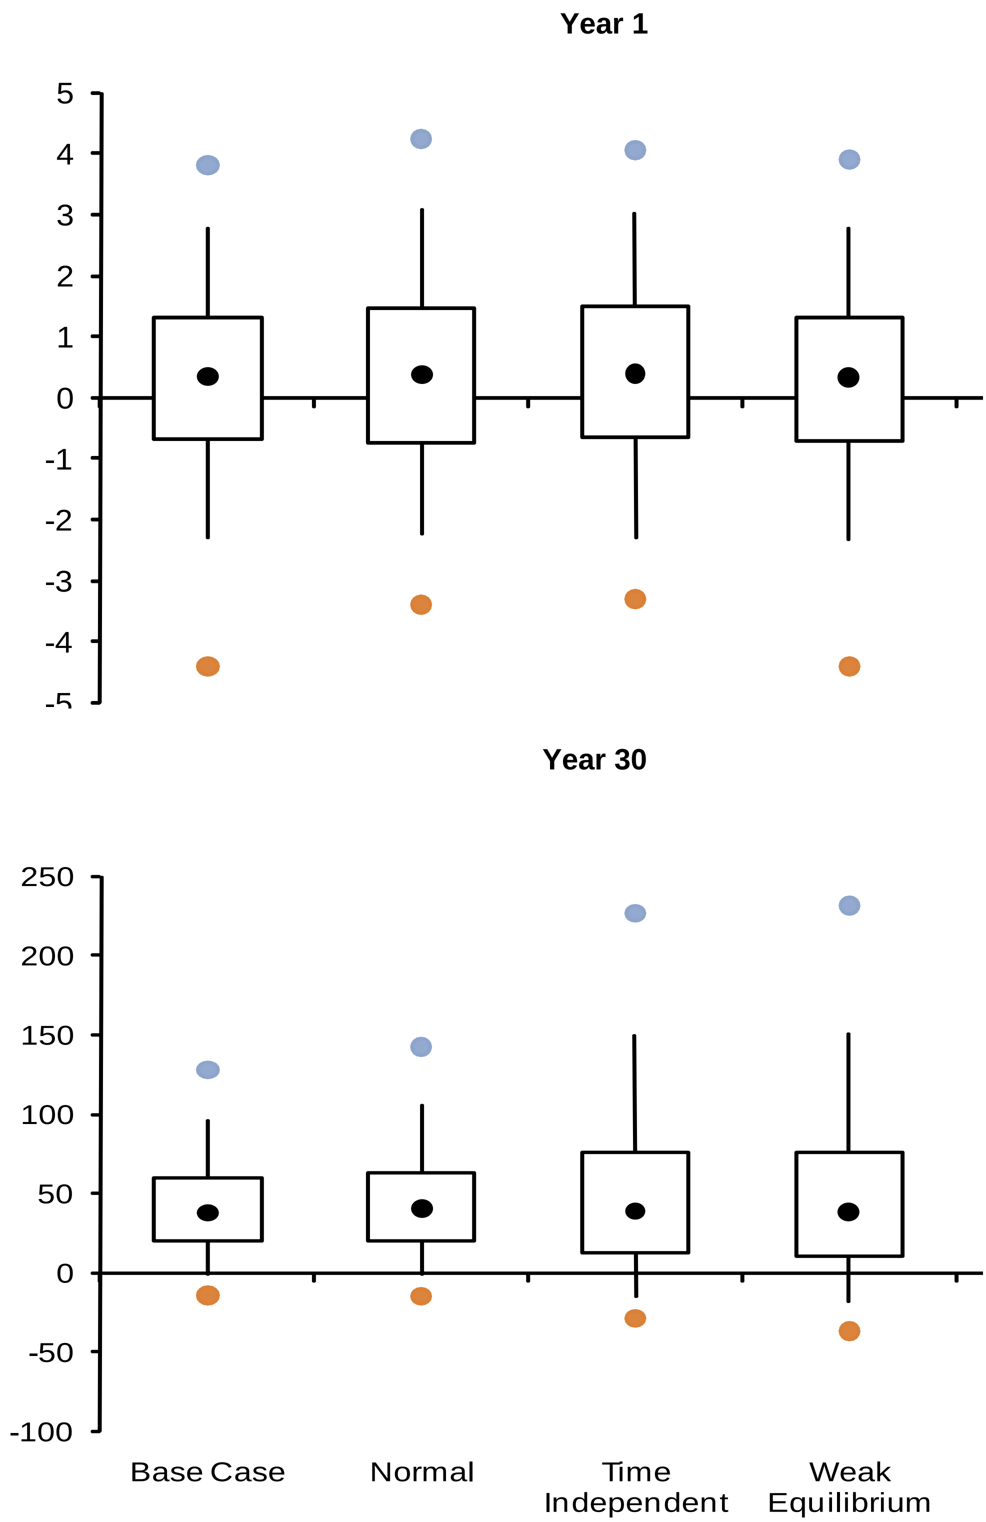

Note: The boxes represent the inter-quartile range, whiskers are to 5th and 95th percentiles, blue dots are 99th percentile, orange dots are 1st percentile outcomes, and black dots are median outcomes of the monte carlo simulations.

6 References

Ansley, C. (2002). "Expectations of Stochastic Interest Rate Functions", Frank Russell Company (NZ) Ltd, July 2002, see NZ Society of Actuaries website: http://www.actuaries.org.nz/Papers% 202002/P15%20Ansley.pdf

Buckle, R. A., K. Kim, H. Kirkham, N. McLellan and J. Sharma (2007) "A structural VAR business cycle model of a volatile small open economy," Economic Modelling, 24, 990–1017.

Campbell, J. Y. (2001) "Forecasting U.S. Equity Returns in the 21st Century". In papers by John Y. Campbell, Peter A. Diamond and John B. Shoven presented to the Social Security Advisory Board. 2001. http://www.ssab.gov/Publications/Financing/estimated%20rate%20of%20return .pdf)

Campbell, J. and L. Viceira (2005) "The term structure of the Risk-Return Trade-off", Financial Analysts Journal, 61, 34–44.

Clarida, R., J. Gali and M. Gertler (1999) "The Science of Monetary Policy: A New Keynesian Perspective," Journal of Economic Literature, American Economic Association, vol. 37(4), pages 1661-1707, December.

Cologni, M. and M. Manera (2008) "Oil prices, inflation and interest rates in a structural cointegrated VAR model for the G-7 economies", Energy Economics, Vol. 30, Issue 3.

Drew, A., O. Karagedikli, R. Sethi and C. Smith (2008) "Changes in the transmission mechanism of monetary policy in New Zealand", Reserve Bank of New Zealand Discussion Paper, DP 2008/03.

Hoevenaars, R., R. Molenaar, P. Schotman and T. Steenkamp (2008) "Strategic asset allocation with liabilities: Beyond stocks and bonds", Journal of Economic Dynamics and Control, Vol. 23, September 2008.

McCulloch, B. and J. Frances (2001) "Financing New Zealand Superannuation", New Zealand Treasury Working Paper 01/20

Paniza, U., V. Cerra and S. Saxena (2009) "International evidence on recoveries from recessions" IMF Working Paper 09/183.

Piplack, J. and S. Straetmans (2009) "Co-movement of different asset classes during market stress" Universiteit Utrecht Discussion paper 09/09.