



Geometric Return and Portfolio Analysis

Brian McCulloch

NEW ZEALAND TREASURY
WORKING PAPER 03/28

DECEMBER 2003



THE TREASURY
Kaitohutohu Kaupapa Rawa

**NZ TREASURY
WORKING PAPER
03/28**

Geometric Return and Portfolio Analysis

MONTH/YEAR

December 2003

AUTHOR

Dr Brian McCulloch
New Zealand Treasury
Email brian.mcculloch@treasury.govt.nz
Telephone +64 – 4 – 471 5077
Fax +64 – 4 – 499 0437

NZ TREASURY

New Zealand Treasury
PO Box 3724
Wellington 6015
NEW ZEALAND
Email information@treasury.govt.nz
Telephone 64-4-472 2733
Website <http://www.treasury.govt.nz>

This Treasury Working Paper is an expanded version of a practice note that was published on the Treasury internet at the time that the strategic investment policy for the New Zealand Superannuation Fund was announced (McCulloch 2003).

DISCLAIMER

The views expressed in this Working Paper are those of the author and do not necessarily reflect the views of the New Zealand Treasury. The paper is presented not as policy, but with a view to inform and stimulate wider debate.

Abstract

Expected geometric return is routinely reported as a summary measure of the prospective performance of asset classes and investment portfolios. It has intuitive appeal because its historical counterpart, the geometric average, provides a useful annualised measure of the proportional change in wealth that actually occurred over a past time series, as if there had been no volatility in return. However, as a prospective measure, expected geometric return has limited value and often the expected annual arithmetic return is a more relevant statistic for modelling and analysis. Despite this, the distinction between expected annual arithmetic return and expected geometric return is not well understood, both in respect of individual asset classes and in respect of portfolios. This confusion persists even though it is explained routinely in finance textbooks and other reference sources. Even the supposedly straightforward calculation of weighted average portfolio return becomes somewhat complicated, and can produce counterintuitive results, if the focus of future-orientated reporting is expected geometric return. This paper explains these issues and applies them in the context of the calculations underlying the projections for the New Zealand Superannuation Fund.

JEL CLASSIFICATION C53: Econometric Modelling – Forecasting and Other Model Applications
D84: Information and Uncertainty – Expectations
G10: General Financial Markets
H55: Social Security and Public Pensions

KEYWORDS Arithmetic; geometric; returns; portfolio; lognormal distribution.

Table of Contents

Abstract.....	iii
Table of Contents	iv
List of Tables	iv
List of Figures	iv
1 Introduction.....	1
2 Expected geometric return and expected arithmetic return.....	2
3 Expected portfolio return	5
4 Expected future size of an investment fund	7
5 Calculation of the required capital contribution rate for the New Zealand Superannuation Fund.....	8
6 Conclusion.....	9
Appendix One: Derivation of Median Returns	10
Appendix Two: Lognormal Distribution of Annual Returns	11
References	12

List of Tables

Table 1 – A Numerical Example of the Expected Value of Future Returns	3
---	---

List of Figures

Figure 1 – Arithmetic and Geometric Measures of Expected Return	4
Figure 2 – Alternative Calculations of Expected Portfolio Returns	6
Figure 3 – Expected Portfolio Geometric Return	7

Geometric Return and Portfolio Analysis

1 Introduction

Professional investment practitioners routinely report expected geometric return as a summary measure of the prospective performance of asset classes and investment portfolios. Expected geometric return has intuitive appeal because its historical counterpart, the geometric average, provides an annualised measure of the proportional change in wealth that actually occurred over a past time series, as if there had been no volatility in return. However, as a prospective measure, expected geometric return has limited value and often the expected annual (or arithmetic) return is a more relevant statistic for modelling and analysis. Despite this, the distinction between expected arithmetic return and expected geometric return is not well understood, both in respect of individual asset classes and in respect of portfolios.¹ This confusion persists even though it is explained routinely in finance textbooks and other reference sources.²

This paper addresses these issues in the context of the financial projections and capital contribution calculations for the New Zealand Superannuation Fund.³ Section Two provides an introduction to the issues by explaining the distinction between expected arithmetic return and expected geometric return. I show that, although average geometric return can be a useful measure of historical return, it is of only limited relevance in future-orientated analyses, where the expected arithmetic return is invariably more relevant. Section Three examines measures of portfolio return and illustrates that the portfolio geometric return is not (and is greater than) a weighted average of the geometric returns of the underlying asset classes. Indeed, it is possible for the portfolio geometric return to be greater than that of all of its constituent asset classes. Section Four then turns to measurement of the expected value of the stock of an investment portfolio, such as the New Zealand Superannuation Fund. Although the intuitive approach would be to compound the expected geometric return over time, this understates the expected portfolio size, whereas compounding the expected arithmetic return provides the correct result. Section Five shows that the calculation of the required contributions to the New Zealand Superannuation Fund also relies on the expected arithmetic return on the Fund's investment portfolio, not the expected geometric return. Finally, Section Six provides some concluding remarks.

¹ For example, see the discussion thread on the Casualty Actuarial Society website at <http://casact.org/discuss/discuss2.htm>.

² For example, see Brealey and Myers (2000 p 157) and Ibbotson (Ibbotson Associates 2002).

³ For a detailed description of the policy underlying the New Zealand Superannuation Fund, see McCulloch and Frances (2003).

In order to derive some of the mathematical results presented below, some assumptions are required about the time-series behaviour of returns. The usual assumptions are that returns are stationary ($E[r_i]$ is constant), homoscedastic ($Var[r_i]$ is constant and finite) and serially independent. The purpose here is not to defend these assumptions and not all of them are necessary to derive most of the results presented here. They are clearly not entirely descriptive of reality: expectations change, volatility can vary over time, and some serial dependence can be detected in some returns series. Nonetheless, these are standard assumptions that are adopted in financial analysis and they produce rigorous results. These assumptions can be relaxed without changing the general tenor of the results presented in this paper, but it would be at the cost of unnecessary added complexity in the calculations.⁴

2 Expected geometric return and expected arithmetic return

Measures of expected value provide essential information when preparing projections of the behaviour of financial investments into the future. However, there are two measures of return over time. The average of an observed set of returns can be measured either arithmetically, by summing the percentage return for each year then dividing by the number of years, or geometrically, by compounding the annual returns and putting this to the power of the inverse of the number of years.

The difference between arithmetic and geometric historical averages can be seen from a simple numerical example. Suppose returns in two years are +40%, followed by -40%. Starting with \$100, we have \$140 after the first year. In the second year, we lose 40% (\$56), giving an ending stock of \$84. The arithmetic average return is zero (being the simple average of +40% and -40%). The geometric average return is -8.3% (being $(1.40 \times 0.60)^{\frac{1}{2}} - 1$). In other words, the same ending wealth could have been achieved with a constant compounded annual return in both years of -8.3% (\$91.70 after year 1, then \$84 after year 2).

If returns are not constant and the time period of measurement is greater than one year, the geometric average will always be less than the arithmetic average.⁵ The difference will be greater the longer the time period and the greater the volatility. As a purely descriptive measure of historical return, the geometric average provides an annualised measure of the proportional change in wealth that actually occurred over the time horizon being examined, as if the wealth grew at a constant rate of return.

Like historical averages, the expected value of future returns can also be specified either in terms of an annual arithmetic mean or in terms of a geometric mean that is measured over a specified time horizon. This is illustrated in the following example. Suppose there is a 50% probability of a +40% return per year and a 50% probability of -40% per year. Starting with \$100, the probability tree of possible outcomes over three years is shown in Table 1. After one year, the wealth level is either \$140 or \$60, giving an expected wealth level of \$100 and expected return is 0%. After two years, there are three possible wealth

⁴ The analyses presented here also assume that the parameters of the distributions of returns (means, variances and so on) are known with certainty. This is a common assumption in financial analysis even though these parameters usually have to be estimated. The potential bias from this further layer of uncertainty has been ignored in the results presented here. It does not affect the general principles being discussed. This issue is discussed further below.

⁵ In the trivial situations where the time period is only one year ($N=1$) or where returns are constant, the geometric average and the arithmetic average will be the same.

levels and the expected wealth is still \$100. However, the expected geometric return is -4.2% per year over two years and, measuring over a three year time period, the expected geometric return declines further to -5.6% per year. The expected total wealth at any time in the future is calculated by compounding the initial wealth (\$100) by the expected annual arithmetic return. In this case, that is 0%, so the expected wealth stayed at \$100 over the three years. The expected annual return does not change with the time horizon. However, the expected geometric return does. It declines as the time horizon increases and it is not a particularly meaningful measure of the expected growth of wealth over time.

Table 1 – A Numerical Example of the Expected Value of Future Returns

Year 0	Year 1		Year 2				Year 3			
\$	\$	Prob	\$	Prob	Geometric Return	Arithmetic Return	\$	Prob	Geometric Return	Arithmetic Return
\$100	\$140	50%	\$196	25%	40.0%	40%	\$274	12.5%	40%	40%
	\$60	50%	\$84	50%	-8.3%	0%	\$118	37.5%	5.6%	13.3%
			\$36	25%	-40.0%	-40%	\$50	37.5%	-20.4%	-13.3%
							\$22	12.5%	-40%	-40%
Expected Value:	\$100		\$100		-4.2%	0.0%	\$100		-5.6%	0.0%

This example illustrates how the expected geometric return declines as the time horizon over which it is measured increases. Therefore, if a measure of long-term expected geometric return is reported, it will be relevant only to the specific time horizon to which it relates. It will understate expected geometric returns over shorter periods. In addition, the expected growth in wealth over time is obtained by compounding the expected annual arithmetic return. Compounding the expected geometric return will understate the expected growth in wealth. This is discussed further in Section Four. Conversely, when discounting back to present value, using the expected geometric return for the discount factor will overstate the present value, while using the expected annual arithmetic return will give the correct result.

The above example assumes that returns follow a discrete binomial distribution. It is more common in financial analysis to think of returns as following a continuous distribution. If returns are assumed to follow a lognormal distribution and are serially independent,⁶ then the exact relationship between the arithmetic mean ($E[r]$) and the locus of the geometric mean ($E[g_N]$) over time (N) is:

$$E[g_N] = (1 + E[r]) \left(1 + \frac{Var[r]}{(1 + E[r])^2} \right)^{\frac{1-N}{2N}} - 1 \quad (1)$$

⁶ These are standard assumptions derived from the central limit theorem and market efficiency, respectively (see Appendix Two). These assumptions are not necessary for the general conclusions presented in this paper. They are used here because they add some structure to the analysis, and they allow some exact results to be calculated for the illustration that follows.

For a given distribution of annual expected return and volatility, the geometric mean is smaller the longer the time period (N) that is being examined. This is illustrated in Figure 1 using the expected annual return on equities of 12.7% and standard deviation of 20.2% reported by Ibbotson Associates (Ibbotson Associates 2002). The expected annual return stays at 12.7%, regardless the time horizon. However, the expected geometric return starts at 12.7% if the time horizon is one year then declines so that for a twenty-year time horizon, the expected geometric return is about 11%.

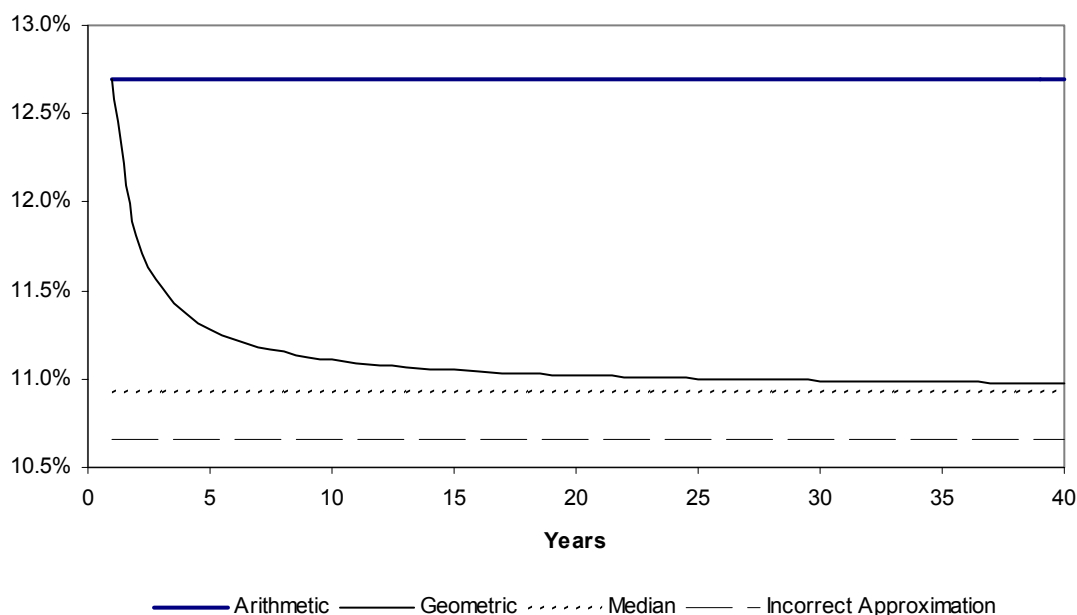


Figure 1 – Arithmetic and Geometric Measures of Expected Return

Stocks are more volatile than bonds, so most of this difference is also reflected in the risk premium. From the same source, the risk premium of stocks over bonds measured on an arithmetic basis was 7.0%, but the difference between the geometric averages was only 5.4%.⁷

There also is a frequently used approximation of the geometric mean. This is calculated as the expected annual arithmetic mean minus half its variance. This understates the true geometric mean for all time horizons, and it is especially wrong for shorter time horizons. This incorrect geometric approximation is also illustrated in the example in Figure 1.

Another central measure of returns is the median. Under lognormality, the median geometric return is a constant that does not decline over time, and it is equal to the median arithmetic return.⁸ It is also illustrated in Figure 1. It is the asymptote that the expected geometric return converges to as the time horizon is expanded out infinitely.

Understanding the distinction between geometric and arithmetic return is important because both metrics are used by commentators discussing issues of investment returns, such as the equity risk premium, and there is scope for confusion about which is relevant in any particular situation. Recognising this, some authors report their analyses on both bases (for example, Cornell 1999, Lally and Marsden 2002). Prospective (that is, future-

⁷ For further discussion regarding estimation of the expected market equity risk premium, see McCulloch (2002).

⁸ See Appendix One for derivation of the median return under standard assumptions.

orientated) applications, such as the capital asset pricing model⁹ and the assessment of the required capital contribution to the New Zealand Superannuation Fund, require an unbiased estimate of the expected annual return.¹⁰ As shown above, the expected geometric return over any period greater than one year will understate the expected annual return, while the expected arithmetic return provides the appropriate measure for this purpose.

A further complication with using any measure of prospective analysis is that the expected values (and other parameters of the return distribution) are not known with certainty and so must be estimated. However, Blume (1974) shows that an arithmetic average provides an unbiased and consistent estimate of the expected annual return, while the geometric average provides a downward biased estimate and it has a larger sample variance than the arithmetic average. In the related case of estimating discount factors for present value calculations, Cooper (1996) shows that both arithmetic and geometric averages provide downward biased estimates of the discount factor, and that the arithmetic average is least biased. This holds even if returns are serially correlated.

3 Expected portfolio return

Asset classes are often combined into portfolios and there is a need to calculate information about expected long-term portfolio returns. In order to illustrate the issues that this raises, suppose that a portfolio comprises two asset classes, equities and bonds. The annual return for the portfolio (r_t^p) is a weighted average of the annual returns on the two asset classes (r_t^e and r_t^b for equities and bonds, respectively):¹¹

$$r_t^p = \alpha r_t^e + (1 - \alpha) r_t^b \quad (2)$$

The portfolio can be thought of as a single asset with expected value and variance of annual returns being functions of the expected values, variances and covariance of the component asset classes:

$$E[r_t^p] = \alpha E[r_t^e] + (1 - \alpha) E[r_t^b] \quad (3)$$

$$Var[r_t^p] = \alpha^2 Var[r_t^e] + (1 - \alpha)^2 Var[r_t^b] + 2\alpha(1 - \alpha) Cov[r_t^e, r_t^b] \quad (4)$$

The same relationships between portfolio annual returns and portfolio geometric returns apply as described above for single assets. In particular, the expected geometric return over time is less than the expected annual return, the difference becomes greater as longer time periods are considered, and the expected portfolio annual return is a more meaningful measure of the expected growth in portfolio wealth than the expected geometric return.

Equation (2) illustrates that the annual portfolio return is a weighted average of the component asset returns. Similarly, Equation (3) illustrates that the expected annual

⁹ Sherris and Wong (2003) examine the merits of alternative measures of expected return in applications of the capital asset pricing model. They demonstrate that an arithmetic average of returns should be used.

¹⁰ McCulloch and Frances (2001) provides the derivation of the calculation of the required capital contribution rate for the New Zealand Superannuation Fund. The appropriate measure to use in that calculation is addressed below.

¹¹ Superscripts *e*, *b* and *p* are used to refer to equities, bonds and the whole portfolio, respectively, and the proportion of the portfolio held in equities is α . There is an implicit assumption that the portfolio is rebalanced each period so that α remains a constant over time.

arithmetic portfolio return is a weighted average of the component assets' expected annual arithmetic returns. However, it is a surprise to many practitioners that the geometric portfolio return is not equal to a weighted average of the component assets' geometric returns. It is greater than the weighted average. That is:

$$1 + g_n^p > \alpha(1 + g_n^e) + (1 - \alpha)(1 + g_n^b) \quad (5)$$

This is a common mistake made when computing portfolio returns. We can see this by decomposing each side of this equation. The result is:¹²

$$\prod_{t=1}^n (\alpha(1 + r_t^e) + (1 - \alpha)(1 + r_t^b))^{1/n} > \alpha \prod_{t=1}^n (1 + r_t^e)^{1/n} + (1 - \alpha) \prod_{t=1}^n (1 + r_t^b)^{1/n} \quad (6)$$

This also applies to the calculation of the portfolio *expected* geometric return. Therefore, if we want to calculate the expected geometric return for a portfolio from the expected geometric returns of the individual asset classes, it is necessary to start with the component asset classes' expected annual arithmetic returns, and take their weighted average to get the expected portfolio annual arithmetic return, then use that result (along with the portfolio volatility and time horizon) to derive the expected portfolio geometric return for the required time horizon.

The incorrect calculation of the expected geometric return (using the weighted average of the component asset classes' geometric returns) understates the true portfolio expected geometric return. If this incorrect result is *then* used when the expected annual return is more appropriate, the bias is even worse than if the correct geometric calculation had been used. Figure 2 illustrates how the correct and incorrect calculations of portfolio returns can vary over time.




Figure 2 – Alternative Calculations of Expected Portfolio Returns

¹² Jensen's Inequality can be used to prove this result. It is a strict inequality so long as $n > 1$ and α is not equal to 0 or 1. This two-asset example extends, with the necessary modifications, to the analysis of portfolios comprising multiple asset classes.

Another counterintuitive result of the nonlinear relationship between expected geometric asset class returns and expected portfolio geometric return is that it is possible for the expected portfolio geometric return to be greater than any of the individual asset class expected geometric returns. To illustrate, Figure 3 shows how the portfolio geometric return of a two-asset portfolio, comprising bonds and equities, changes as the portfolio allocation moves from 0% equities to 100%. In this example, there is a region of portfolio composition, from 55% equities to 100%, in which the portfolio geometric return becomes higher than that of either of the individual asset classes. Of course, this does not always happen – it depends on the structure of the return covariance matrix. Nonetheless, it is usual to find that the expected portfolio geometric return is at the upper end of the spread of the individual asset class expected geometric returns.

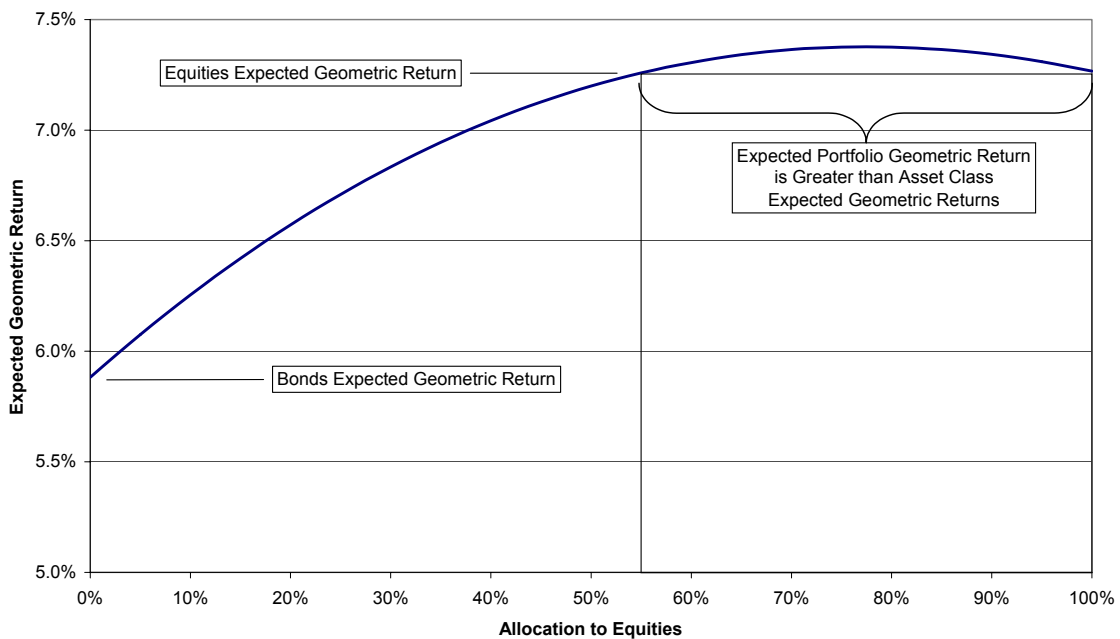


Figure 3 – Expected Portfolio Geometric Return

4 Expected future size of an investment fund

The issues discussed above surrounding the relative merits of geometric and arithmetic measures of expected return also extend to the calculation of projections of the expected size of an investment fund. Consider a stock that has compounding returns over n periods to a value of S_n (with $S_0=1$). The returns in each period are assumed to be random, serially uncorrelated with a constant annual expected value $E[r_t]$. Therefore:

$$S_n = \prod_{t=1}^n (1+r_t) = (1+E[r_t])^n \quad (7)$$

That is, the expected value of a \$1 stock that compounds for n periods at an expected annual arithmetic rate of $E[r]$ is $(1+E[r])^n$.¹³ This was illustrated in the numerical example of expected returns in Section Two.¹⁴

¹³ Another way to show this result would be with the Law of Iterated Expectations.

¹⁴ In that example, $E[r]=0$ and so $E[S_n]=100$ for all n .

In particular, the expected compounded value of the stock after n periods is not the expected geometric return over that time horizon to the power of n . That would be an understatement of the expected value of the stock (because the expected geometric return is less than the expected annual return). Similarly, the expected geometric return is not the n th root of the expected value of the stock at period n (because that would yield the expected annual arithmetic return).

This result that the expected size of a stock over time is calculated by compounding the expected annual arithmetic return over the time horizon has been applied in the projections of the expected growth of the New Zealand Superannuation Fund over time as illustrated in the Treasury's spreadsheet model of the New Zealand Superannuation Fund and in McCulloch and Frances (2001). Consistent with the above analysis, the expected Fund size is calculated by compounding the Fund balance (adjusted for capital contributions and withdrawals) by the expected annual arithmetic return.

5 Calculation of the required capital contribution rate for the New Zealand Superannuation Fund

The New Zealand Superannuation Act 2001 requires the Treasury annually to determine the capital contribution required to be made from the Crown to the Fund for the next financial year. This must be set so that, if that same proportion of forecast GDP were to be made to the Fund each year for the succeeding forty years, the Fund balance plus accumulated returns would be just sufficient to meet the expected net cost of entitlement payments over those forty years. This can be expressed as:¹⁵

$$E_0 \left[B_0 \prod_{t=1}^H (1+r_t) + \sum_{t=1}^H k_1 G_t \prod_{i=t+1}^H (1+r_i) - \sum_{t=1}^H P_t \prod_{i=t+1}^H (1+r_i) \right] = 0 \quad (8)$$

where:

- B_0 = Fund balance at the beginning of year 1.
- H = time horizon for the calculation. This is set at forty years.
- r_t = rate of return on the Fund in year t .
- k_1 = total contribution rate for year 1 as a proportion of GDP.
- G_t = GDP for year t .
- P_t = forecast entitlement payments in year t .

And the required capital contribution for the next period (in \$) is:

$$CapitalContribution_1 = k_1 G_1 - P_1 \quad (9)$$

¹⁵ This is a simplified version being used here to explain the principles at issue. A more detailed version is used for the actual calculation, taking into account such things as the fortnightly payment structure (McCulloch and Frances 2001).

Solving the above expectation equation for the total contribution rate (k_1) gives:¹⁶

$$k_1 = \frac{\sum_{t=1}^H \frac{P_t}{(1 + E[r])^{-t}} - B_0}{\sum_{t=1}^H \frac{G_t}{(1 + E[r])^{-t}}} \quad (10)$$

The summation terms in both the numerator and the denominator in this equation are analogous to present value calculations and the appropriate discount rate is $E[r]$, which is the expected annual arithmetic return on the investment portfolio of the Fund.¹⁷ If the expected geometric return were to be used in this calculation, the required contribution rate would be misstated.

6 Conclusion

Expected geometric return is routinely reported as a summary measure of the prospective performance of asset classes and investment portfolios. It has intuitive appeal because its historical counterpart, the geometric average, provides a useful descriptive measure of the annualised proportional change in wealth that actually occurred over a past time series, as if there had been no volatility in return. However, for applications that involve future projections or other prospective analyses, expected geometric return has limited value and often the expected annual arithmetic return is a more relevant statistic for modelling and analysis. Despite this, the distinction between expected annual arithmetic return and expected geometric return is not well understood, both in respect of individual asset classes and in respect of portfolios. This confusion persists even though it is explained routinely in finance textbooks and other reference sources.

Even the supposedly straightforward calculation of weighted average portfolio return becomes somewhat complicated, and can produce counterintuitive results, if the focus of reporting is expected geometric return. Simply calculating the portfolio expected geometric return for a particular time horizon as being the weighted average of the expected geometric returns of each asset class for that time horizon will understate the expected portfolio geometric return. The weighted average calculation should be carried out starting with the expected annual arithmetic returns of the individual asset classes. The true expected portfolio geometric return will be at the upper end of (and could possibly exceed) the spread of individual asset class expected geometric returns.

The issues are also interpreted in the context of the analysis underlying the New Zealand Superannuation Fund. Projections of the expected size of the Fund should be based on compounding the expected arithmetic return over time, not the geometric return. Similarly, calculation of the capital contributions the Crown is required to make to the Fund is a function of the expected arithmetic return on the Fund, not of the expected geometric return.

¹⁶ This solution is explained in detail in McCulloch and Frances (2001).

¹⁷ Note that the summation terms are not actually calculations of the present values of the cashflow streams, P_t and G_t . That would require the use of discount rates that reflected the risk inherent in those cashflow streams, and not the expected return on the investment portfolio.

Appendix One: Derivation of Median Returns

Let $X_i = 1 + r_i$, where r_i is the return in period i .

Assume that returns are distributed lognormally and are serially independent (which are standard assumptions supported by the Central Limit Theorem and market efficiency, respectively).

$$\text{Therefore, } \ln(X_i) \sim N(\mu, \sigma^2) \quad (\text{A1})$$

$$\text{Geometric return over } N \text{ years is defined as: } g_N = \prod_{i=1}^N X_i^{1/n} - 1 \quad (\text{A2})$$

$$\text{So its log is: } \ln(1 + g_N) = \frac{1}{N} \sum_{i=1}^N \ln(X_i) \quad (\text{A3})$$

The expected value of this is:

$$\begin{aligned} E[\ln(1 + g_N)] &= \frac{1}{N} \sum_{i=1}^N \mu \\ &= \mu \end{aligned} \quad (\text{A4})$$

$\ln(1 + g_N)$ is normally distributed (because it is the sum of $\ln(X_i)$, which are normally distributed). Therefore it has a symmetrical distribution and so its median equals its mean. So:

$$M[\ln(1 + g_N)] = \mu \quad (\text{A5})$$

That is, half of the distribution of $\ln(1 + g_N)$ is below μ . Therefore, half of the distribution of $(1 + g_N)$ is below e^μ . Hence:

$$M[1 + g_N] = e^\mu \quad \text{and} \quad M[g_N] = e^\mu - 1 \quad (\text{A6})$$

Therefore, although the expected geometric return declines as the time horizon increases, the median geometric return is a constant, invariant to the time horizon. It is the same as the median arithmetic return (because $g_i = r_i$), and it is less than both the expected geometric and expected arithmetic return.

It was noted above that the expected size of a stock over time ($E[S_n]$) is calculated by compounding the expected annual arithmetic return over the time horizon: $E[S_n] = (1 + E[r])^n$. A corresponding derivation to that used in this appendix can be used to show that the median stock also grows exponentially: $M[S_n] = e^{n\mu}$.

Appendix Two: Lognormal Distribution of Annual Returns

Stochastic analysis of investments requires an understanding of the statistical properties of returns. The assumption that annual returns follow a lognormal distribution is relatively robust. It is based on the Central Limit Theorem as follows.

Suppose that a year is made up of many (say $x=200$) trading days, with daily returns ($d_{year,day}$) that are serially independent and of finite variance, but the form of the distribution is unspecified. Daily returns compound into annual returns (r_{year}):

$$1 + r_t = \prod_{i=1}^x (1 + d_{t,i}) \quad (A7)$$

Taking the log of each side:

$$\log[1 + r_t] = \sum_{i=1}^x \log[1 + d_{t,i}] \quad (A8)$$

According to the Central Limit Theorem, the sum of n independent random variables with finite variance converges to a normal distribution when n is large. Since $d_{t,i}$ is an independent series with finite variance, so is $\log[1 + d_{t,i}]$. And 200 is large. Thus, $\log[1 + r_t]$ is approximately normally distributed and so $[1 + r_t]$ approximately follows the corresponding lognormal distribution.

$$\log[1 + r_t] \sim N(\mu, \sigma^2) \quad (A9)$$

This result requires no assumption about the shape of the distribution of daily returns. If daily returns, themselves, are lognormally distributed, then the annual returns will be exactly lognormally distributed (being the product of independent lognormally distributed variables).

The variable, $\log[1 + r_t]$, is also known as the continuously compounded rate of return.

The mean and variance of r_t can be expressed in terms of μ and σ^2 using the moment generating function of a normally distributed $(1 + r_t)$.¹⁸

$$E[r_t] = E[1 + r_t] - 1 = \exp\left[\mu + \frac{1}{2}\sigma^2\right] - 1 \quad (A10)$$

$$Var[r_t] = Var[1 + r_t] = E[1 + r_t]^2 \left(\exp[\sigma^2] - 1\right) \quad (A11)$$

¹⁸ See Aitchison and Brown (1957) for a detailed treatment on the lognormal distribution and its application in economics.

References

- Aitchison, John and J Alan C Brown (1957) *The lognormal distribution*. (Cambridge: Cambridge University Press).
- Blume, Marshall E (1974) "Unbiased estimators of long-run expected rates of return." *Journal of the American Statistical Association* 69(347): 634-638.
- Brealey, Richard A and Stewart C Myers (2000) *Principles of corporate finance*. (New York NY: Irwin.McGraw-Hill).
- Cooper, Ian (1996) "Arithmetic versus geometric mean estimators: Setting discount rates for capital budgeting." *European Financial Management* 2(2): 157-167.
- Cornell, Bradford (1999) *The equity risk premium: The long-run future of the stock market*. (New York NY: Wiley).
- Ibbotson Associates (2002) *Stocks, bonds, bills, and inflation 2002 yearbook*. (Chicago: Ibbotson). <<http://www.ibbotson.com>>
- Lally, Martin and Alastair Marsden (2002) "Historical market risk premiums in New Zealand: 1931-2000." Working Paper.
- McCulloch, Brian W (2002) "Estimating the market equity risk premium." New Zealand Treasury. <<http://www.treasury.govt.nz/release/super/market-risk.pdf>>
- McCulloch, Brian W (2003) "Long-term returns and the New Zealand superannuation fund projections." New Zealand Treasury. <<http://www.treasury.govt.nz/release/super/ltr-nzsf.pdf>>
- McCulloch, Brian W and Jane Frances (2001) "Financing New Zealand superannuation." New Zealand Treasury, Working Paper 01/20. <<http://www.treasury.govt.nz/workingpapers/2001/01-20.asp>>
- McCulloch, Brian W and Jane Frances (2003) "Governance of public pension funds: New Zealand superannuation fund." World Bank Public Pension Fund Management Conference. <<http://www.treasury.govt.nz/release/super/gppf-nzsf-wb.pdf>>
- Sherris, Michael and Bernard Wong (2003) "Continuous compounding, volatility and beta: Review of and response to fitzherbert." UNSW Actuarial Studies Working Paper Series, 2003/6. <<http://www.actuary.unsw.edu.au/cgi-bin/page.php?linkage=papers>>