Original Article

Optimal portfolio leverage

Received (in revised form): 21st September 2015

Paul van Rensburg

is the Frank Robb Professor of Finance at the University of Cape Town. He is also the founder and principal of Salient Quantitative Investment Management and the A-Dex Investment Toolkit.

Correspondence: Paul van Rensburg, University of Cape Town, Private Bag, Rondebosch, 7701, South Africa.
E-mail: paul.vanrensburg@uct.ac.za

ABSTRACT The Optimal Portfolio Leverage Ratio provides the level of leverage to use to attain the highest expected long-term terminal value of an investment and is calculated independently of investors' indifference curves. This article applies a discrete multi-period compounding framework to both discrete and continuous cross-sectional pay-off distributions. In both cases, an Optimal Portfolio Leverage Ratio is derived from first principles and in the case of the latter, a multi-asset solution is also presented. The primary implications for equilibrium asset pricing are considered and a multi-period analogue to the CAPM is derived. This version of the CAPM is to be tested as a joint hypothesis with a specified Optimal Growth Portfolio.

Journal of Asset Management advance online publication, 3 December 2015; doi:10.1057/jam.2015.36

Keywords: multi-period portfolio construction; Kelly rule; optimal portfolio leverage ratio; growth optimal asset pricing

INTRODUCTION

Mean–Variance optimised portfolio construction in a single-period framework is famously introduced in Markowitz (1952, 1959). However, despite common practice, for most applications a multi-period framework is needed instead. The single holding period assumption may appear a reasonable and 'neutral' simplifying assumption. However, its relaxation has substantial and surprising implications of practical relevance. This article considers optimal portfolio construction and its implications for asset pricing in a multi-period setting. Throughout, a framework where return compounding occurs after discrete intervals of time (rather than continuously) is adopted.

The practical motivation for the article is revealed through the paradox that inspired its origin: The author was considering an investment with approximately a 20-year time horizon for his young children. An all-equity portfolio seemed to have the highest long-term expected return. However, if this portfolio could be levered, portfolio theory would imply that a higher expected return could be achieved as leverage increases. This is, indeed, correct but consider the case where an equity portfolio was levered three times during the recent 2008 crisis, which had approximately a 45 per cent draw-down. The entire portfolio would be wiped out! How is it that very high levels of expected returns can simultaneously imply a near certainty of (eventual) ruin? This paradox, that is, as L (leverage)$\rightarrow\infty$, g (the compound growth rate and, hence, terminal portfolio value)$\rightarrow -\infty$ despite $E(R)\rightarrow\infty$ is to be more formally presented later in the article.[1] Clearly, with a leverage of zero in equities (that is, simply holding cash) zero excess returns are earned. With too much leverage the fund value will go to 0 (and below, if your account's losses are not capped at what you put in). Therefore, what is the optimal level of leverage to hold for a long-term investment in order to have the maximum expected terminal value?

The article is organised as follows. The next section brings the analysis at the most fundamental level by considering the difference between single- and multi-period returns and, thereby, motivating a revision of the multi-period objective function from 'expected return' to 'compound growth rate' of the investment over time. Thereafter, starting with the simplest case of repeatedly flipping a coin, optimal portfolio construction strategies for discrete distributions are derived step-by-step from first principles. Thereby, an Optimal Portfolio Leverage Ratio is derived in the subsequent section and a number of key applications illustrated. The section after that considers the case of continuous distributions of returns. A formula for optimal portfolio leverage in multi-period setting and multi-asset setting is derived. Implications of this revised framework for portfolio selection and asset pricing are then considered. The final section summarises and concludes.

CONTRASTING SINGLE AND MULTI-PERIOD RETURNS

The salient feature that differentiates a multi-period from a single-period framework is that in the multi-period framework *the proceeds from the previous period are what is available to invest at the start of the next period*. Thus, there is an inextricable link between periods that does not allow the dilution of a period's influence in an analogous manner to that achieved cross-sectionally across assets. This is in contrast to the example provided by Kelly (1956) where a gambler's wife gives him an allowance of $1 a week to bet. As this 'resets' his bet size each period, maximising long-run expected returns and letting time diversify away risk now becomes an optimal strategy – as if in a single-period framework. This insight is linked to the so called 'fallacy' of time diversification where the naive view is that good years will counterbalance poor years making equity returns less volatile over time. However, because of the absence of a 'reset' each year, a bad year (−100 per cent being an extreme example to clearly illustrate the point) stays in the system and does not get 'washed out' over time.[2]

To slightly more formally review some basics, returns, R, are additive within a particular time period:

$$R = \sum_{i=1}^{N} w_i R_i$$

where w_i = the weight in asset i

In contrast, over time, returns are multiplicative:

$$R_T = \left[\prod_{t=1}^{T}(1+R_t)\right] - 1$$

And for purposes of convenience and comparison, this can be restated in terms of a single-period compound growth rate(g) as follows:

$$g = \left[\prod_{t=1}^{T}(1+R_t)\right]^{1/T} - 1$$

Kelly (1956) seems to be first in realising that it is not an investment's single-period expected return but rather its growth rate over time (g) that needs to be maximised in order to obtain the highest expected long-term value of the portfolio. The next section solves for this optimal level in the discrete case from first principles. It does so in a manner outside Kelly's (1956) information theory context of symbol identification and the rate of signal transmission over a noisy communication channel.

THE OPTIMAL LEVERAGE RATIO (DISCRETE RETURN DISTRIBUTIONS)

In this section we will consider the simple case of flipping a coin in each period where there are only two discrete possible outcomes: we can win or lose. We have fixed odds of winning and fixed pay-off rates over time. It follows that if we assume a constant pay-off rate and constant odds in each time period, our betting fraction should also remain constant over time. Thus, a key characteristic of repeated betting is that we bet a constant fraction of our bankroll and not a constant currency amount in each period.

The value of a bankroll invested through time under the conditions of fixed-proportion continual betting can be expressed as follows:

$$V_T = \left[(1+\varphi_W \cdot L)^W (1-\varphi_L \cdot L)^{T-W}\right] V_0$$

Where:

V_t	the value of the bankroll in period $t=0, 1, 2, \ldots, T$
φ_W	the pay-off (expressed as a ratio of V) of a winning event, for example heads coin flip
φ_L	the pay-off (expressed as a ratio of V) of a losing event, for example tails coin flip
L	the leverage or fraction of bankroll bet, that is bet size
T	the total number of coin flips
W	the number of winning coin flip

Or expressed as a geometric return (g), that is, the long-term compound growth rate of the investment[3]:

$$(1+g)^T = (1+\varphi_W \cdot L)^W (1-\varphi_L \cdot L)^{T-W}$$

$$\text{where} \qquad (1+g)^T = \frac{V_T}{V_0} = (1+R_T)$$

Calculus can be used to find the L (leverage or bet size) with the highest value of g. Finding the derivative is eased by taking the log of both sides of the above equation thereby simplifying the exponential terms into multiplicative terms. If ln $(1+g)$ is maximised so too must be $(1+g)$ given the monotonic relationship between the two.

Taking the log of both sides and dividing through by T we obtain:

$$\ln(1+g) = \frac{1}{T} W \cdot \ln(1+\varphi_W \cdot L) + \frac{1}{T}\left[(T-W) \cdot \ln(1-\varphi_L \cdot L)\right]$$

Solving for the derivative of the above function requires use of the generalised log function rule[4]:

$$\frac{d}{dL} k \cdot \ln f(L) = k \cdot \frac{f'(L)}{f(L)}$$

where k = a constant term.

Applying the log function rule allows us to find the derivative below:

$$\frac{d\ln(1+g)}{dL} = \frac{W}{T} \cdot \frac{\varphi_W}{(1+\varphi_W \cdot L)} - \frac{1-\frac{W}{T} \cdot \varphi_L}{1-\varphi_L \cdot L}$$

Setting this value to zero so as to derive an optimal point, we find that:

$$\frac{W}{T} \cdot \frac{\varphi_W}{1+\varphi_W \cdot L} = \frac{1-\frac{W}{T} \cdot \varphi_L}{1-\varphi_L \cdot L}$$

Which, after a fair deal of simplification and rearrangement, can be expressed to solve for the optimal value of L:

$$L^\star = \left(\frac{W}{T}\right)\left(\frac{1}{\varphi_L}\right) - \frac{1-\frac{W}{T}}{\varphi_W}$$

Where $L^\star$ represents the optimal level of leverage. Finally, defining p, the probability of winning, such that

$$p = \lim_{T\to\infty}\left(\frac{W}{T}\right)$$

By simple substitution we obtain:

$$L^{\star} = \frac{p}{\varphi_L} - \frac{1-p}{\varphi_W} \qquad (1)$$

Where:

φ_W	the pay-off (return) of a winning event
φ_L	the pay-off of a losing event
p	the probability of a winning event

We refer to equation (1) as the '*generalised Kelly rule*' for fixed-proportion continually repeated betting.[5]

Equation (1) is a very handy tool that can be used in a wide variety of betting situations. The newly empowered punter would start off by allocating a certain sum of money as a bankroll for his speculative activities. Then he would apply equation (1) to decide what proportion of his bankroll to bet depending on the odds and pay-offs of the wager concerned. The examples below go through a few essential scenarios. In all examples it is assumed that the portion of one's bankroll that is not wagered is to have a return of zero.

Example 1: *The effect of pay-off rate on optimal leverage (even odds bets)* An inebriated casino owner offers to play a coin flipping game with you under varying pay-off rates. What proportion of your bankroll should you bet at each coin flip? In the examples below odds remain constant at even odds (where $p=0.5$ for a fair coin) and, for the purposes of comparison, the pay-off rate is allowed to vary:

1. If the pay-off rate (φ) is 2:1 then, substituting into equations (1), $L^{\star} = 0.5-[(1-0.5)/2]=0.25$ or 25 per cent of your bankroll should be bet in repeated bets.
2. If $\varphi = 10:1$ then $0.5-[(1-05)/10]=0.45$ or 45 per cent of your bankroll should be bet.

Note that as the pay-off increases the optimal bet size converges to *p,* i.e. your odds of winning (in this example 50 per cent). Thus, the interesting result is obtained that:

$$\lim_{\varphi\to\infty}(L^{\star}) = p$$

In other words, *no matter how high your pay-off rate, you should not bet a proportion higher than your odds of winning*. For example, even with a 1 000 000:1 pay-off, you should not bet more than 50 per cent of your bankroll on a fair coin in repeated bets.

Example 2: *The effect of odds on optimal leverage(even money bets)* Under an even money bet $\varphi_L = \varphi_W = 1$. Substituting into the generalised Kelly rule (equation (1)) the following formulation is derived:

$$L^{\star} = p-(1-p)$$

and defining $q=(1-p)$ as the probability of loss allows the following formulation:

$$L^{\star} = p-q$$

(1) If $p=50$ per cent then $L^{\star}=0$ per cent i.e. only bet if you have an odds advantage.
(2) If $p=60$ per cent then $L^{\star} = (60\%-40\%)=20\%$.
(3) If $p=90$ per cent then $L^{\star}=80$ per cent. Thus, in the case of even money bets, the general rule is to *bet the proportion that is your odds advantage*.

Example 3: *Balancing odds and pay-offs (for example, a horse race)* You are invited as one of Hugh Hefner's guests at the Kentucky Derby. You think that a particular horse has a 30 per cent chance of winning and the pay-off is 10:1. The optimal proportion of your bankroll to bet is $0.3-(1-0.3)/10=23\%$.

Example 4: *Combining two assets each with $g=0$* Consider repeated bets with a 50 per cent chance of doubling your stake and a 50 per cent chance of halving your stake. Note that the long-run growth of this investment, *g*, is zero. In this case: $p=50$ per cent and substituting these values into the generalised Kelly rule (equation (1)):

$$L^{\star} = \frac{p}{\varphi_L} - \frac{(1-p)}{\varphi_W}$$

gives: $L^{\star} = 0.5/0.5 - (1-0.5)/1 = 0.5$ or 50 per cent to optimally bet in the risky asset. The growth rate of a portfolio where there is a 'two possible state scenario' in each period can be expressed as follows:

$$(1+g) = \left(1 + \sum_{i=1}^{N} w_i R_{i,W}\right)^{p} \left(1 + \sum_{i=1}^{N} w_i R_{i,L}\right)^{1-p}$$

where: $R_{i,W}$ is the return on asset i in a 'winning' scenario and $R_{i,L}$ is the return in a 'losing' scenario.

In the case of above mentioned (continually rebalanced) 50 per cent bet:

$$(1+g) = [1 + 0.5(0\,\%) + 0.5(100\,\%)]^{0.5}$$
$$[1 + 0.5(0\,\%) + 0.5(-50\,\%)]^{0.5}$$

Simplifying:

$$(1+g) = \{[150\,\%] \cdot [75\,\%]\}^{0.5} = \sqrt{112.5\,\%} = 1.0607$$

Thus, $g = 6.07$ per cent per period. By adding a risky investment to a riskless investment (in each case where $g = 0$) the interesting result is obtained where $g > 0$ is obtained by the portfolio. This result illustrates point to be discussed in the section 'The compound growth rate (g)' that, unlike $E(R)$, the g of a portfolio is not a linear weighted average of its constituents.

CONTINUOUS RETURN DISTRIBUTIONS

Introduction

A fundamental conceptual question to address is why can't a long-term multi-period time horizon simply be converted to single-period horizon? For example, why can't 30-year expected returns just be inputted into the familiar Markowitz single-period framework, where 30 years is selected as the single-period time horizon? The section 'Compounding and non-normality'shows that because of the compounding of single-period normally distributed returns with a positive mean, an ever increasing right skew develops in the longer-term distribution of returns, which is a violation of the Markowitz assumptions. The alternative of using log returns is considered in the section 'Compounding and non-normality'but is not the approach applied in the remainder of this article. The section 'The compound growth rate (g)' derives an expression for the compound growth rate and proceeds to optimise its value as a function of leverage. Implications for portfolio selection (in section 'Optimal portfolio selection') and asset pricing (in section 'Multi-period asset pricing') are then considered.

Compounding and non-normality

Given the general formulation for cumulative portfolio value over time:

$$V_T = \prod_{t=1}^{T} (1 + R_t) V_0$$

a two period case ($T = 2$) can be used for purposes of illustration:

$$V_2 = (1 + R_1)(1 + R_2)$$

Expanding the above

$$V_2 = 1 + R_1 + R_2 + R_1 R_2$$

Where the final term ($R_1 R_2$) reflects the effect of compounding. If R_1 and R_2 are normally distributed ($R_1 R_2$) will not be and, hence, neither will V_2. Expanding a 3-period case becomes more dramatic:

$$V_3 = 1 + R_1 + R_2 + R_3 + R_1 R_2 + R_1 R_3 + R_2 R_3 + R_1 R_2 R_3$$

Four cross product terms are evident, all contributing to the non-normality of V_3. In general, as more periods are added, even if distributions are normal within each period, the resulting distribution of cumulative values (and cumulative returns) becomes less and less normal due to the effect of compounding. In particular, in the mean value of a single-period return is positive, the distribution will become increasingly skewed to the right as the number of period increases.[6] Thus, over longer periods, the mean and standard deviation do not adequately capture the nature of return distributions as required by the Markowitz framework.

The normality of log returns

If the cumulative portfolio value is logged then its distribution becomes normal, that is, multi-period returns are log normally distributed. This has often been stated in textbooks but without an extensive consideration of its implications. Taking the natural log of both sides of the formulation of cumulative portfolio value over time:

$$\ln(V_T) = \ln(V_0) + \sum_{t=1}^{T} \ln(1 + R_t)$$

Which can also be expressed as:

$$\ln(V_T) = \ln(V_0) + \sum_{t=1}^{T} \ln\left(\frac{V_{t+1}}{V_t}\right)$$

Given the central limit theorem that the mean of a set of *independent* random variables will have a distribution that tends to normality as the number of variables grows $\ln(V_T) = \ln(V_0(1+R_T))$, where $(1+R_T)$ is total period return, will tend to normality. As there are now no multiplicative terms in the logged version (for example, R_1R_2 in our earlier two period example), the distributions are independent across time. This independence allows multi-period logged return distributions to preserve normality. Note that 'gross returns' $(1+R_T)$ rather than R_T is logged so as to facilitate cumulation and avoid the impossible task of taking the log of a negative number.

However, it is important to note that logged returns remove the effect of compounding, which is the prime agent of return over expended periods and, thus, does not adequately represent the distribution of the actual return earned by a long-term investor. For this, the non-linear exponential function needs to be applied to the cumulative (natural) log returns to retrieve the (right skewed) distribution of future portfolio values. It is easy to forget that *logged returns are the log of returns and not actually returns.* In addition, the log returns of a portfolio are also not the weighted average of the log returns of its constituents and, thus, log returns cannot simply be substituted for arithmetic returns within the Markowitz framework as some practitioners erroneously do. Employing log returns together with a revised formula for portfolio returns (that be obtained though a Taylor series expansion) and an 'unlogging' at the end of the multi-period horizon (by the application of the exponential function) is an alternative approach to this problem (see Campbell and Viceira, 2002). To contain its focus, this article selects to restrict itself to adopt a discrete rather than continuous compounding approach.

The compound growth rate (*g*)

As demonstrated earlier, the compound growth rate, *g*, of an investment is directly related to its terminal value and is a more appropriate objective than expected mean return to optimise in order to maximise long-term expected investment value. Consider the framework of a binomial lattice where in each period (which may be made shorter and shorter to approach continuity) an asset has a 50 per cent chance of earning a return of $[E(R)+\sigma]$ and a 50 per cent chance of earning a return of $[E(R)-\sigma]$. In such a case the most likely (mode) and median outcome will have $T/2$ 'up' moves and $T/2$ 'down' moves:

$$\begin{aligned}(1+R_T) &= [1+E(R)+\sigma]^{\frac{T}{2}}[1+E(R)-\sigma]^{\frac{T}{2}} \\ &= \{[1+E(R)+\sigma][1+E(R)-\sigma]\}^{\frac{T}{2}} \\ &= \{[1+E(R)]^2-\sigma^2\}^{\frac{T}{2}}\end{aligned}$$

Thus, the compound growth rate, g, where $(1+g)^T = (1+R_T)$, can be expressed as follows:

$$(1+g)^T = \left\{[1+E(R)]^2 - \sigma^2\right\}^{\frac{T}{2}}$$

and as $T \to \infty$ can be rewritten as:

$$(1+g) = \left\{[1+E(R)]^2 - \sigma^2\right\}^{0.5}$$

Some simplification of terms can be obtained if a $T=2$ period (with one period being 'up' and the other 'down') case is considered as this would effectively remove the exponential term on the left hand side of the multi-period equation, that is, where $T=2$:

$$(1+g)^2 = \left\{[1+E(R)]^2 - \sigma^2\right\}$$

Multiplying out we obtain:

$$1+2g+g^2 = 1+2E(R)+E(R)^2 - \sigma^2$$

Rearranging results in the following expression:

$$E(R)-g = \left\{g^2 - \frac{\left[g^2 - E(R)^2\right]}{2}\right\} + \frac{\sigma^2}{2}$$

Considering the first term on the left-hand side of the equation, that is,

$$\left\{g^2 - \frac{\left[g^2 - E(R)^2\right]}{2}\right\}$$

It can be argued that this term is likely to be relatively small in magnitude. As short single period $E(R)$'s and g's are likely to be significantly less than one, squaring each is likely to reduce their size further. In addition, the difference between these squared terms is obtained (which is likely to be smaller still). Therefore, as a very convenient approximation, it can be stated that:[7]

$$g \cong E(R) - \frac{1}{2}\sigma^2 \qquad (2)$$

Note that compounding at g will provide the mode and median returns of the cumulative return distribution. In contrast, compounding at $E(R)$ will provide the mean outcome, which, in certain cases (for example, where extreme right skew is present), may be far less likely to occur. Equation 2 also shows why g cannot just be substituted instead of $E(R)$ in the single-period Markowitz framework. The second term on the right hand side of equation 2 shows that there will be 'diversification bonus' when two or more assets are combined together so that the g of a portfolio will be higher than (or equal to) a weighted average of its constituents.

The optimal portfolio leverage ratio (continuous return distributions)

If it is assumed that leverage is obtained by borrowing at the risk free rate R_f, the expected return of a levered portfolio can be expressed as follows:

$$E(R)_L = R_f + L\left(E(R) - R_f\right)$$

This can be substituted into equation 2 to obtain:

$$g = R_f + L\left[E(R) - R_f\right] - \frac{L^2\sigma^2}{2}$$

Using this equation to consider the implications of ever increasing leverage allows us to (at last!) formalise the 'Umngazi Paradox' mentioned in the introduction: as $L \to \infty$, $g \to -\infty$ despite $E(R) \to \infty$. A consideration of this situation brings it home that it is not an investment's single-period expected return but rather its growth rate over time that needs to be maximised.

To find an optimal value of L, the derivative of this expression with respect to L can be calculated as:

$$\frac{dg}{dL} = \left[E(R) - R_f\right] - \frac{2L\sigma^2}{2} = \left[E(R) - R_f\right] - L\sigma^2$$

Setting this value to 0 (given that the second derivative $= -\sigma^2$ is negative throughout, we will be solving for a global maximum), we can solve for the Optimal

Leverage Ratio,[8] $L^\star$:

$$L^\star = \frac{\lfloor E(R) - R_f \rfloor}{\sigma^2} \qquad (3)$$

Example 5: Assume that the annualised stock market premium is 6 per cent and the standard deviation of returns is 20 per cent. The optimal gearing is $6\%/4\% = 1.5$. This solves the type of problem mentioned in the introduction that originally motivated this paper.

Note that the Optimal Portfolio Leverage Ratio can be handily calculated from the (annualised) Sharpe Ratio by simply dividing it by the (annualised) standard deviation of returns. Indeed, the Optimal Leverage Ratio (equation 3) looks very much like the Sharpe Ratio except that its denominator is the variance rather than standard deviation of returns. However, unlike other risk adjusted performance measures such as the Sharpe and Treynor indices, the Optimal Portfolio Leverage Ratio has the following meaningful advantages:

(i) It is dimensionless in the sense of being comparable across varying measurement periods. (In contrast, for example, annualised Sharpe Ratios or Treynor Indices are not comparable to monthly versions).

(ii) It is sensitive to the degree of leverage applied. (In contrast, a portfolio with a high Sharpe Ratio may be over levered and destined to 'blow up'. The Sharpe Ratio has no way to distinguish this portfolio from an 'under-levered' version of itself).

(iii) It has a clear practical interpretation, that is, the proportion one should optimally invest in the portfolio concerned.

Optimal portfolio selection

Figure 1 below contrasts single and multi-period optimal portfolio selection in the absence of leverage constraints.

The straight line R_fM is the familiar Tobin (1958) single-period efficient frontier with unlimited borrowing an lending at R_f and M being the mean–variance efficient portfolio in the absence of leverage. As all of the points along this line are efficient, additional information is needed to select an optimal point. This is provided by investor's indifference curves (not drawn in Figure 1 for neatness). In contrast, the curved line labelled g is the geometric mean frontier of concern to the multi-period investor. It has an optimal point at $L^\star$, the Optimal Portfolio Leverage Ratio.

However, $L^\star$ is probably best viewed in practice not as a target level but rather as an extreme level of leverage not to go beyond. The first reason for this is the possibility of estimation errors for the inputs into $L^\star$, that is, $E(R)$ and volatility. For example, given that volatility varies over time, in a period of high volatility the investor may find himself over-levered. The second, is that the reward in terms of additional expected g per unit risk gets flatter as $L^\star$ is approached. Thus, it is likely to be prudent to adopt a more conservative level of leverage than $L^\star$ in practice and be alert to cases that are levered beyond this level. For example, giving a client 10 times leverage to trade single stock futures will most likely lead to a 'blow up' even where an odds and pay-off advantage may be present. Loans on houses that do require a deposit is another recent and topical example

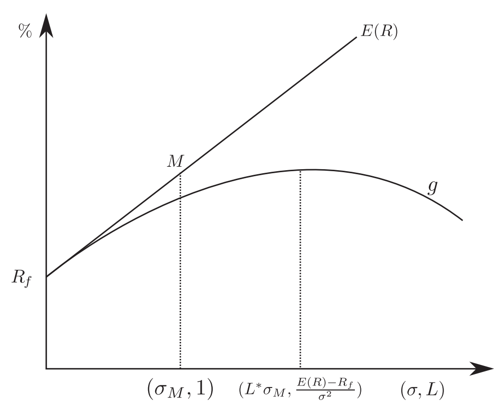

Figure 1: A comparison of single and multi-period portfolio selection (unconstrained leverage).

of imprudently violating the Optimal Leverage Ratio.

Perhaps, one of the most vital applications of the Optimal Portfolio Leverage Ratio is in the realm of low volatility strategies such as arbitraging the relative prices of very close substitutes. This ratio gives a vital indication of many such strategies that can safely be levered up in an attempt to 'turn crumbs into lunch'. The Long-Term Capital Management debacle shows what can happen if you, for whatever reason, find yourself beyond this limit.

The single asset solution can be generalised to a portfolio W_{opt} with weights $w_{opt,i}$ for $i = 1, \ldots, N$ assets with the use of linear algebra:

$$W_{opt} = C^{-1}E \tag{4}$$

Where:

W_{opt}	A (n×1) vector of growth optimal portfolio weights with typical element $w^{\star}_{opt,i}$
C^{-1}	The inverse of the (n×n) covariance matrix of returns
E	A (n×1) vector of expected excess returns with typical element $(E(R_i)-R_f)$

It may be worth considering the relationship between the Mean–Variance Efficient and Growth Optimal Portfolios under a situation without leverage constraints. It can be observed that, given a certain expected (arithmetic) mean return of a portfolio, there is no way to increase the geometric mean without reducing its standard deviation of returns. Similarly, for a given volatility level, a portfolio's geometric mean can only be increased by increasing its arithmetic mean. Thus, the Growth Optimal Portfolio will also be mean–variance efficient (although the reverse does not necessarily apply as the latter may be over-levered). Although the (expected) Sharpe Ratio does indicate what portfolio of risky assets to invest in, it does not tell us how much we should put into this portfolio.

However, in the situation where leverage is constrained, for example, different assets can be levered or shorted to different degrees, than a multi-asset version of equation 2 needs to be maximised subject to these particular constraints:

$$g_p = \sum_{i=1}^{N} w_i E(R_i) - \frac{1}{2}\sum_{j=1}^{N}\sum_{i=1}^{N} w_i w_j \sigma_{ij}$$

It is also worth noting that almost all practical experience and empirical research will be conducted in the unlevered left-hand side of Figure 1 where the differences between the single- and multi-period scenarios are not that extreme and, hence, have largely passed unnoticed. Nevertheless, an empirical researcher should be alert to the fact that if he inadvertently uses compound rather than mean returns in his analysis (there are many ways in which this error could creep in, for example, through annualising multi-year returns or Sharpe Ratios), he will erroneously observe a flatter relationship between risk and return than implied by single-period asset pricing models such as the CAPM.

Multi-period asset pricing

In an analogous manner to which the CAPM is derived as a positive theory of asset pricing based on the assumption that all investor's behave as the normative Markowitz portfolio theory suggests, it can be asked how assets would be priced if all investors follow the prescriptions of multi-period portfolio theory.

In a 'perfect world' of unconstrained leverage,[9] infinite time horizons and homogenous expectations, investors aiming to maximise the growth rate of their investments will purchase the Growth Optimal Portfolio until its Optimal Portfolio Leverage ratio approaches unity. Extra demand for this portfolio, if under-levered to maximise long-term growth, will increase its price thereby reducing its expected returns, E_{opt}, until its Optimal Portfolio Leverage Ratio is

1. This equilibrium pricing relationship implies that:

$$\left(E\left(R_{opt}\right)-R_f\right)=\sigma_{opt}^2 \quad (5)$$

To consider the implications of the multi-asset case, equation (4) can be rearranged as:

$$E=CW_{opt}$$

Or expressed in scalar form as:

$$\left(E(R_i)-R_f\right)=\sum_{j=1}^{N} w_{opt,j}\sigma_{ij}$$

For all assets $i=1, \ldots, N$. This can be simplified as:

$$\left(E(R_i)-R_f\right)=\sigma_{i,opt}$$

In other words, each share's expected excess returns are uniquely determined by their covariance with the Growth Optimal Portfolio. Using the fact that the Optimal Leverage Ratio for the Growth Optimal Portfolio is one, equation (5) can be substituted into the above equation to obtain:

$$E(R_i)=R_f+\beta_{i,opt}\left(E\left(R_{opt}\right)-R_f\right) \quad (6)$$

Where:

$$\beta_{i,opt}=\frac{\sigma_{i,opt}}{\sigma_{opt}^2}$$

Which is a multi-period analogue to the single-period CAPM. This version of the CAPM is to be tested as a joint hypothesis with a specified Growth Optimal Portfolio, which takes the place of the 'market' portfolio in the single-period CAPM. It is a stricter alternative to multi-factor testing as it requires that these factor exposures are optimised in the formation of the Growth Optimal Portfolio.

CONCLUSION

Although much time and analysis is conducted to ascertain the odds and associated pay-offs of various investment opportunities, this article considers how to to invest in these opportunities once the return distributions are estimated. The 'Umngazi Paradox' (that is, as $L\to\infty, g\to-\infty$ despite $E(R)\to\infty$) vividly highlights Kelly's (1956) argument that investors should aim to maximise the compound growth rate of their investment over time rather than expected return within a single-period framework. Although return distributions may be approximately normal instantaneously, over longer time horizons the influence of a positive mean leads to an ever increasing right skew because of the effect of compounding. This is recognised as a meaningful violation of the assumptions of the single-period Markowitz (1952) portfolio theory.

This article derives a generalised form of the Kelly Rule for discrete return distributions from first principles and illustrates its practical application in a number of examples. It is demonstrated that, unlike the single-period case, there is a long-term optimal level of leverage that maximises the investor's compound investment growth rate. In the case of continuous distributions, an Optimal Portfolio Leverage Ratio is similarly derived. Lack of familiarity aside, the latter form of the Optimal Portfolio Leverage Ratio is also shown to have certain clear advantages as follows over the Sharpe and Treynor Indices: (i) being comparable across different time periods, (ii) reflecting the effect of leverage and (iii) having a direct practical interpretation of its value. With the use of linear algebra, a multi-asset solution is also presented.

The implications for asset pricing under the assumption of optimal multi-period portfolio selection are then considered. It is argued that, in equilibrium, risk and expected return will be traded off such that the Growth Optimal Portfolio is priced in a way that its Optimal Portfolio Leverage Ratio will be unity. Using this condition, a multi-period version of the CAPM is derived that is to be tested as a joint hypothesis with a specified growth optimal blend of factor and asset exposures.

This article has defined its focus as optimising portfolio leverage (or, equivalently, individual asset weights in a multi-asset context) so as to maximise expected terminal portfolio wealth over a multi-period time horizon. An elaboration of this work that suggests itself is to introduce indifference curves in to the analysis and, hence, facilitate a trade-off between risk and expected terminal value. In this regard, it is worth noting that maximising the expected terminal value of the investment does not necessarily imply maximising the expected utility of the distribution of terminal values to the investor. Stated concisely, investors maximise $E(U(W_T))$ rather than $U(E(W_T))$. As illustrated in Figure 1, the lower reward per unit risk taken on as the growth optimal level of leverage is approached may motivate to more conservative level be adopted, and this can be formalised using indifference curves. However, the particular utility function to apply is a controversial topic and beyond the scope of this study.

With the caveats mentioned in the section 'The Normality of log returns', the analysis can alternatively be placed within a continuous compounding (log return) framework. Allowing investor contributions and later withdrawals to the portfolio may also be a useful step towards the aim of deriving a framework for optimum long-term portfolio construction that can be applied to the widespread practical need of investors' retirement planning.

ACKNOWLEDGEMENTS

Thanks to Bryce Gething, Jarryd Garisch, Kanshu Rajaratnum and Sammy Chun-Sung for assisting with the type setting of this article. Thanks to Hannes du Plessis for his comments on an earlier draft.

NOTES

1. This paradox went through the author's mind when driving his family to a holiday resort at Umngazi on the South African Wild Coast and, by association, is referred to as the 'Umngazi Paradox' in this article.
2. Time diversification is an interesting topic that is beyond the scope of this article. Suffice to say the author evolved from the naive view in about 1996 with the caveat that, given long-term (3–5 years) mean reversion in returns, there is still some merit in time diversification (which seems to be the current academic consensus). However, in the preparation of this article, he realised that the real benefit of time diversification is not so much the reduction in volatility (a second moment effect) that occurs but the powerful force of compounding a positive mean (a first moment effect) that results in an ever increasing right skew (third moment effect) in return distributions over time.
3. As an aside, it can be noted that this equation is likely to prove useful later once an optimal value for L is found as this may be substituted into the above to find its the associated rate of geometric return. It can also be used by setting g to -100 per cent and solving for the value of L (that is, 'excessive leverage') that results in long-run ruin.
4. As a refresher to non-mathematician readers: the generalised log function rule essentially states that the derivative of the log of a function equals the derivative of that function (not logged) divided by the (unlogged) function itself, which, to the extent that we know the rules of solving for derivatives for unlogged functions, allows us to calculate its values for the logged versions.
5. In the case where $L^{\star} = 0$ then

$$\frac{p}{\varphi_L} = \frac{1-p}{\varphi_W}$$

Which can be rearranged as:

$$\frac{\varphi_W}{\varphi_L} = \frac{1-p}{p}$$

It can be seen that, in this case where the Optimal Leverage Ratio is zero, the pay-offs are inversely related to their probabilities.
Equation (1) can also be contrasted to the special case where $\varphi_L = 1$, that is, exactly your entire bankroll is at stake in each successive bet. Here the formula reduces to:

$$L^{\star} = p - \frac{(1-p)}{\varphi}$$

Where $\varphi = (\varphi_W)/(\varphi_L)$ is the pay-off ratio of the investment. Note that as the pay-off ratio expressed here as a 'ratio to one' the formula makes no distinction between e.g. 100:10 versus 10:1 or 1000:100 pay-off sizes (all would be 10:1 as φ is currently defined). Clearly this 'leverage' of the pay-off sizes has a direct inverse effect on optimal 'exposure' leverage and, thus, the generalised version is necessary in those cases where φ_L does not equal one. As will be seen, the above equation is the discrete analogue to the Sharpe Ratio – while we would want it to be as high as possible, it does not inform us how much to bet.
6. Despite all the talk of 'Black Swans' the reality is that, over longer periods, the outliers are far more likely to be 'White Eagles'. See, for example, the positive skew in distribution of wealth in most societies and, indeed, the prevalence of power law distributions in almost all

measures of cumulative achievement as a manifestation of this effect.

7. This result is analogous to the drift obtained from applying Ito's Lemma on geometric Brownian motion in continuous time. See, for instance, Focardi and Fabozzi (2004, p. 274) or Racicot and Theoret (2006, p. 249) for more information on that subject. Thanks to an anonymous reviewer for pointing this out.
8. For an alternative derivation see Thorpe (1990).
9. Under a more realistic scenario of constrained leverage (due to most investment mandates not allowing it) it would be rational for investors to have greater demand for higher expected return portfolios as long as these increase the expected growth rate of their investment. In the absence of leverage, portfolios with the highest Sharpe Ratios may well not be growth optimal. The increased demand for these higher risk – higher expected return assets will increase their prices, lowering their expected returns. Thus, the underpriced, lower risk assets will have higher resulting Optimal Portfolio Leverage Ratios than their overpriced higher risk alternatives.

REFERENCES

Markowitz, H.M. (1952) Portfolio selection. *Journal of Finance* 45(1): 77–91.

Campbell, J. and Viceira, L. (2002) *Strategic Asset Allocation Portfolio Choice for Long Term Investors.* Oxford: Oxford University Press.

Focardi, S.M. and Fabozzi, F.J. (2004) *The Mathematics of Financial Modelling and Investment Management*. Haboken, NJ: John Wiley and Sons.

Markowitz, H.M. (1959) *Portfolio Selection: Efficient Diversification of Investments.* New York: John Wiley and Sons.

Kelly, J.L. (1956) A new interpretation of information rate. *Bell Systems Technical Journal* 35: 917–926.

Racicot, F.E. and Theoret, R. (2006) *Finance Computationnelle et Gestion des Risques*. Sainte Foy Quebec: Presses de l'Universite du Quebec.

Thorpe, E.O. (1990) The Kelly criterion in Black Jack, sports betting and the stock market. In: S.A. Zenios and W. Ziabed (eds.) *The Handbook of Asset and Liability Management III*, Oxford, UK: Elsevier.

Tobin, J. (1958) Liquidity preference as behaviour towards risk. *Review of Economic Studies* 26(1): 65–85.