VALUE AVERAGING AND HOW DYNAMIC STRATEGIES BIAS THE IRR AND MODIFIED IRR

Simon Hayley
Cass Business School,
106 Bunhill Row,
London EC1Y 8TZ, UK

E-mail: simon.hayley.1@city.ac.uk
Tel +44 20 7040 0230
Fax +44 20 7040 8881

This draft: 30 August 2013

I am grateful to Ilya Dichev, Maela Giofré, Aneel Keswani, Ian Marsh, Richard Payne, Giorgio Questa, Nick Ronalds and participants at the 2012 Behavioural Finance Working Group and the May 2012 CFA UK seminar for useful comments, and in particular to Stewart Hodges for his insight in the derivation of the appendix. The usual disclaimer applies.

VALUE AVERAGING AND HOW DYNAMIC STRATEGIES BIAS THE IRR AND MODIFIED IRR

Abstract

This paper demonstrates that the IRR is a biased indicator of expected profits for Value Averaging (VA) and for any other dynamic strategy which is based on a target return or profit level, or which takes profits or "doubles down" following losses. The modified IRR is similarly biased. VA is popular, but this paper demonstrates that it is an inefficient investment strategy (for any plausible investor risk preferences) and quantifies the resulting welfare losses. VA's popularity appears to be due to investors making a cognitive error in assuming that the strategy's attractive IRR implies greater expected terminal wealth.

EFMA Classification: 320, 370, 380, 720

Keywords: value averaging; internal rate of return; behavioral finance; cognitive error; portfolio performance evaluation

VALUE AVERAGING AND HOW DYNAMIC STRATEGIES BIAS THE IRR AND MODIFIED IRR

1. Introduction

Value averaging (VA) is a popular formula investment strategy which invests available funds gradually over time so as to keep the portfolio growing at a pre-determined target rate. It is recommended to investors because it demonstrably achieves a higher internal rate of return (IRR) than plausible alternative strategies. An online search on "value averaging" and "investment" shows many thousands of references to this strategy. These references, and those in other media, are overwhelmingly positive, recommending the strategy to investors as a means of boosting expected returns.

The use of the IRR to evaluate investor returns may seem intuitive, since it takes into account the varied cashflows that are inherent in VA. However, this paper demonstrates that the IRR recorded for any VA strategy is systematically biased up. This bias retrospectively increases the weight given in the IRR calculation to periods with strong returns and reduces the weight given to weaker returns. This paper also demonstrates that the higher IRRs recorded for VA are entirely due to this retrospective bias: VA does not increase expected terminal wealth.

Nor is this problem specific to VA − the bias will affect the IRR of any dynamic strategy which links the scale of future investment to the returns achieved to date. This includes any strategy which is based on a target return or profit level, or which includes any systematic element of taking profits, or "doubling down" after taking losses. Section 4 below shows that the modified internal rate of return (MIRR) is similarly biased.

Not only does VA fail to deliver the superior returns that the IRR suggests, it is also systematically inefficient. As Dybvig (1988a) notes, a common misconception is that if markets are efficient then strategies which alter portfolio exposure over time will do no harm. But such strategies will be inefficient if they needlessly concentrate risk in particular periods. Section 5 demonstrates that VA is an inefficient strategy for any plausible investor risk preferences and quantifies the resulting welfare losses. Certain types of weak form inefficiency in market returns could in principle justify the use of VA but it would be an inefficient means of profiting from such inefficiencies. Historical data shows that VA has underperformed alternative strategies.

In addition, VA's volatile and unpredictable cashflows are likely to force investors to hold more cash and liquid assets than would otherwise be optimal, and may increase management costs, transaction costs and tax liabilities compared to a buy-and-hold strategy. VA may bring behavioral finance benefits, but investors who value these are likely to prefer the simpler Dollar Cost Averaging strategy (DCA).

VA's proponents recommend the strategy on the grounds of its higher IRR. The contribution of the present paper is (i) to show that this higher IRR is entirely due to a retrospective bias; (ii) to demonstrate that VA is an inefficient strategy for any plausible investor risk preferences; (iii) to quantify the associated efficiency losses; (iv) to show that this bias affects both the IRR and MIRR for a wide range of dynamic strategies. The analysis in this paper thus contradicts the claims made by VA's proponents, and finds that not only does the strategy not generate the higher expected profits that are claimed, it is also likely to signifiantly reduce investor welfare.

The structure of this paper is as follows: the next section describes the VA strategy and discusses the academic research which relates to it. Section 3 shows that VA does not outperform alternative strategies despite its superior IRR. Section 4 identifies the bias in the IRR which explains this apparent contradiction. Section 5 demonstrates that far from outperforming, VA is in fact a systematically inefficient strategy. Section 6 examines whether VA might outperform in the presence of specific types of market inefficiency. Section 7 explores whether non-profit motives such as behavioral finance effects might justify the use of VA. Conclusions are drawn in the final section.

2. The Value Averaging Strategy

VA is similar in some respects to dollar cost averaging, which is the strategy of building up exposure gradually by investing an equal dollar amount each period. DCA automatically buys an increased number of shares after prices have fallen and so buys at an average cost which is lower than the average price over these periods (Table 1 shows an example). Conversely, if prices rose DCA would purchase fewer shares in later periods, again achieving an average cost which is lower than the average price over this period (Table 2). As long as there is any variation in prices DCA will always achieve a lower average cost.

[Table 1 here]

VA is a more complex strategy because the additional investment made each period is not constant. The investor sets a target increase in portfolio value each period (assumed here to be a rise of $100 per period, although the target can equally well be defined as a percentage increase) and at the end of each period must make whatever additional investments are necessary in order to meet this target. Like DCA, VA purchases a larger number of shares after a fall in prices, but the response is more sensitive since the dollar amount invested by VA also rises after a fall in prices. Thus in Table 1 VA buys 122 shares in period 2, compared to 111 for DCA. The greater sensitivity of VA to shifts in the share price results in an even lower average purchase cost. Again, this is true whether prices rise, fall or merely fluctuate.

[Table 2 here]

VA could in principle be applied over any time horizon, but its originator suggests quarterly or monthly investments (Edleson, 1991). It appears to be aimed largely at private investors, although a mutual fund has recently been established which is explicitly based on VA, shifting investor funds from money markets to riskier assets according to a VA formula.

Despite its popularity, VA has so far been the subject of limited academic research. VA commits the investor to follow a fixed rule, allowing no discretion over subsequent levels of investment. As a result, both are subject to the criticism of Constantinides (1979), who shows that strategies which pre-commit investors in this way will be dominated by strategies which instead allow investors to react to incoming news. VA might seem to improve diversification by making many small purchases, but Rozeff (1994) shows that this is not the case for DCA. Both strategies start with a very low level of market exposure, so the terminal wealth will be much more sensitive to returns later in the horizon, by which time the investor is more fully invested. Better diversification is achieved by investing in one initial lump sum, and thus being fully exposed to the returns in each period. An investor who has funds available should invest immediately rather than wait.

VA's cashflows are volatile and unpredictable. Each period investors must add whatever amount of new capital is required to bring the portfolio up to its pre-defined target level, so these cashflows are determined by returns over the most recent period. Edleson envisages investors holding a 'side fund' containing liquid assets sufficient to meet these needs[1].

Although VA generates impressive IRRs, empirical studies show no corresponding outperformance on other performance measures. Thorley (1994) compares VA with a static buy-and-hold strategy for the S&P500 index over the period 1926-1991 and finds that it performs worse in terms of mean annual return, Sharpe ratio and Treynor ratio. Leggio and Lien (2003) find that the rankings of these three strategies depend on the asset class and the performance measure used, but the overall results do not support the benefits claimed for VA.

However, VA's proponents continue to stress its demonstrable advantage: achieving a higher expected IRR than alternative strategies (Edleson (1991), Marshall (2000, 2006)). This

---

[1] Edleson (1991) and Marshall (2000, 2006) both calculate the IRR on the VA strategy without including returns on the side fund. We follow the same approach here in order to demonstrate that even in the form used by its proponents VA does not generate the higher returns that are claimed. Thorley (1994) rightly criticises the exclusion of the returns on cash in the side fund. However, including a side fund does not remove the bias: The modified IRR includes cash holdings, but I demonstrate below that this too is a biased measure of VA's profitability.

appears to be the key to VA's popularity. The following sections demonstrate that the IRR is raised by a systematic bias which allows VA to generate attractive IRRs even without increasing expected profits.

3. Simulation Evidence

VA is recommended by its proponents as a strategy which boosts expected returns in any market, even if the investor has no ability to forecast returns. Thorley (1994) presents a binomial tree which shows VA generating lower average dollar profits than investing in one initial lump sum (IRRs are not calculated). By contrast, Edleson (1991, 2006), Marshall (2000, 2006) and other proponents demonstrate that VA generates a higher IRR than alternative strategies even on simulated random walk data (corresponding dollar profits are not calcultated). This section uses a consistent set of simulations to demonstrate that the IRR is a biased measure of the profitability of VA. The following section derives this result more formally and demonstrates how this bias arises.

We assume here that returns follow a random walk. This is consistent with the fact that investors who use VA are unlikely to believe that they are able to forecast short-term returns. Those who (rightly or wrongly) believe that they have such forecasting ability should prefer alternative strategies which – unlike VA – allow them some discretion over the timing of their investments. Section 6 below considers whether weak form inefficiencies in market returns could justify the use of VA.

The simulations also assume that this random walk has zero drift. This is the simplest assumption, and it is in fact generous to VA. A more realistic assumption of upward drift would penalize VA since its relatively large initial holdings of cash would then earn a lower expected return than those invested in risky assets. For simplicity we also assume that the security that is purchased pays no dividend or other income. This assumption is similarly generous to VA.

Table 3 shows how the average costs, IRRs and profits achieved by VA and DCA differ from those obtained by a simple strategy of investing in one initial lump sum. Both VA and DCA achieve significantly lower average purchase costs and higher IRRs. VA clearly appears to be the most attractive strategy when judged on these criteria (in particular, it raises the IRR by 0.28% compared to a buy-and-hold strategy, whilst DCA achieves an increase of only 0.08%), but there is no significant difference between the dollar profits generated by these three strategies.

[Table 3 here]

DCA responds to the fall in prices by buying more shares (111) in the second period. By buying more shares when they are relatively cheap, DCA always achieves an average purchase cost which is below the average price. VA responds more aggressively than DCA, since in order to achieve its target portfolio value it must also make up for the $10 loss suffered on its earlier investment by investing an additional $10 in period 2. VA thus achieves an even larger reduction in its average purchase cost than DCA. These strategies also achieve lower average costs when prices rise (Table 2), but none of this makes any difference to expected profits. All else equal, lower average costs would lead to higher profits, but all else is not equal here since the different strategies invest different total amounts.

4. The Bias In The IRR

Edleson (1991) and Marshall (2000, 2006) focus exclusively on the IRRs achieved by VA. This might seem a reasonable approach, since the IRR takes account of the fluctuating cashflows that are an inherent part of the strategy. However, these IRRs are systematically misleading. Hayley (2013) shows that in aggregate IRRs tend to be biased down as equity investors "chase returns" by increasing their exposures following strong returns. This section uses the same approach to demonstrate that, by contrast, VA automatically biases the IRR up.

An investor's portfolio value at the end of period t (K_t) is determined by the return in the previous period plus any additional top-up investment a_t made at the end of this period:

$$K_t = K_{t-1}(1 + r_t) + a_t \tag{1}$$

By definition, when discounted at the IRR, the present value of these periodic investments equals the present value of the final liquidation value in period T:

$$K_0 + \sum_{t=1}^{T} \frac{a_t}{(1 + IRR)^t} = \frac{K_t}{(1 + IRR)^T} \tag{2}$$

Substituting equation 1 into equation 2 allows us to eliminate a_t (following Dichev and Yu, 2009). Rearranging shows that the IRR is a weighted average of the individual period returns, where the weights reflect the present value of the portfolio at the start of each period:

$$IRR\sum_{t=1}^{T}\frac{K_{t-1}}{(1+IRR)^{t}}=\sum_{t=1}^{T}\left(\frac{K_{t-1}}{(1+IRR)^{t}}\times r_{t}\right) \tag{3}$$

We can re-arrange equation (3) further to show the deviation of periodic returns from the DW return. Periodic returns r_t will be either above or below the IRR, but the weighted sum of these differences must be zero:

$$\sum_{t=1}^{T}\left(\frac{K_{t-1}}{(1+IRR)^{t}}\left(r_{t}-IRR\right)\right)=0 \tag{4}$$

Dividing the horizon in two gives a convenient form in which to show the effect on the IRR of a single additional investment at the end of period *m* which has a value equal to *b*% of the portfolio at that time:

$$\sum_{t=1}^{m}\left(\frac{K_{t-1}}{(1+IRR)^{t}}\left(r_{t}-IRR\right)\right)+(1+b)\sum_{t=m+1}^{T}\left(\frac{K_{t-1}^{*}}{(1+IRR)^{t}}\left(r_{t}-IRR\right)\right)=0 \tag{5}$$

Additional investment after period *m* increases the weight given to later returns, compared to the weights based on the portfolio values K_t^* which would otherwise have been seen. If, for example, the periodic returns r_t up to period *m* were low, then these early $(r_t - IRR)$ terms will tend to be negative, and subsequent terms will tend to be positive. A large new investment at this point would increase the weight given to subsequent $(r_t - IRR)$ terms relative to the earlier terms so the IRR must increase in order to keep the weighed sum at zero. Similarly, investing less (or even withdrawing funds) after a period of strong returns will tend to reduce the relative weight given to later *(r_t - IRR)* terms, which would tend to be negative. This too would increase the IRR.

However, the impact on the IRR could reflect two very different effects. The IRR could be raised by relatively large additional investments taking place ahead of periods with relatively high returns. This would represent good investment timing and would clearly increase expected profits, but this effect cannot explain the high IRRs in the simulations since our assumption of a random walk means that future returns are unforecastable and investments will on average be badly timed as frequently as they are well timed.

However, a large new investment will not only increase the weight which the IRR calculation gives to future returns, it will also reduce the weight given to earlier returns (equation (3) shows that these weights sum to unity). This would be a retrospective adjustment which has no impact on expected profits as long as these intermediate cashflows are not correlated with future returns (and by construction there is no such correlation in our simulations). We know that this effect is inherent in VA, since by construction disappointing returns are followed by larger net investments.

Specifically, the net investment demanded by VA each period is determined by the degree to which organic growth in the value of the portfolio over the immediately preceding period ($r_m K_{m-1}$) fell short of the investor's target. The first summation in Equation 5 includes r_m so the level of new investment *b* will tend to be large (small) when the first summation is negative (positive). The second summation will be correspondingly positive (negative) and will be given more (less) weight as a result of this additional investment. All else equal, the weighted sum over all periods would become positive, but the IRR then rises to return the sum to zero. Thus VA biases the IRR up by automatically ensuring that the size of each additional investment is negatively correlated with the preceding return.

Phalippou (2008) shows that the IRRs recorded by private equity managers can be deliberately manipulated by returning cash to investors immediately for successful projects and extending poorly-performing projects. VA cannot change the end of the investment horizon in this way, but it achieves its bias by reducing the weight given to returns later in the horizon following good outturns, and increasing it following poor returns.

More generally, any performance measure which is in effect a weighted average of individual period returns can be biased by following a strategy which retrospectively reduces the weight given to bad outturns and increases the weight given to good outturns. This will be a property shared by any strategy which targets a particular level of portfolio growth, systematically takes profits after strong returns or "doubles down" after weak returns. Ingersoll et al. (2007) show that a fund manager could give an upward bias to the Sharpe ratio, the Sortino ratio and Jensen's alpha by reducing exposure following a good outturn and increasing exposure following a bad outturn (although they do not cover the IRR in their analysis). It is by doing this automatically that VA raises its expected IRR.

Including Edleson's "side fund" in the calculation is not sufficient to avoid this bias. We must also ensure that the size of this side fund is fixed in advance and not adjusted retrospectively. This can be seen from the bias in the modified internal rate of return (MIRR), and can be illustrated with a simple two period example. Suppose an investor initially allocates a to risky assets and b to the side fund, where it earns a risk-free return r_f. At the end of period 1 an amount c from the side fund is used to buy additional risky assets.

$$\textit{Terminal Wealth (TW)} = a(1+r_1)(1+r_2) + c(1+r_2) + (b(1+r_f) - c)(1+r_f) \tag{5}$$

This measure is not affected by any retrospective adjusment, since the weight attached to r_1 is fixed in advance. Including the side fund in the calculation of the IRR means that intermediate cashflows just become a shift from one part of the portfolio to the other, leaving just the initial and terminal cashflows. Thus the IRR simply becomes the geometric mean return:

$$IRR = \sqrt{\frac{TW}{a+b}} - 1 \tag{6}$$

This too is unbiased if a and b are both fixed in advance. The bias comes about because the side funds must be sufficiently large to meet the VA strategy's future cash needs, but this is a function of future returns and so is unknown. This tends to lead to the size of the side fund being set retrospectively to ensure that it is sufficient. This can be illustrated by considering the modified internal rate of return (MIRR), which assumes the existence of a side fund which is just big enough to fund subsequent cash injections (implying that $b(1+r_f)=c$ in the expression above for terminal wealth). Hence:

$$MIRR = \sqrt{\frac{a(1+r_1)(1+r_2) + c(1+r_2)}{a + c/(1+r_f)}} - 1 \tag{7}$$

The MIRR is biased because the relative weight $a/(a+c/(1+r_f))$ given to r_1 is adjusted retrospectively. VA automatically ensures that a low r_1 will be followed by a large cash injection c, so the relative weight given to r_1 is automatically reduced, increasing the MIRR. The weight on r_1 is changed after the event, so although this alters the MIRR it has no effect on expected terminal wealth. Thus VA increases the expected MIRR because of a retrospective bias.[2]

5. The Inefficiency of Value Averaging

The analysis above showed that VA does not generate higher expected profits than alternative strategies, despite its higher expected IRR. In this section we go one step further and demonstrate that VA is an inefficient strategy, with other strategies offering preferable risk-return characteristics. For now we maintain our assumption that asset returns follow a random walk. In Section 6 we relax this assumption to consider the implications of a predictable component in asset returns.

[Figure 1 here]

We use the payoff distribution pricing model derived by Dybvig (1988b) to demonstrate that VA is inefficient. Figure 1 shows the simplest possible illustration of this technique, using a binomial model of the terminal wealth generated over four periods by a VA strategy. In a good outturn equity prices are assumed to double, whilst they halve in a bad outturn. The investor has chosen a portfolio growth target of 40% each period and initially invests 100 in equities. If the value of these equities rises in the first period to 200, then 60 is assumed to be transferred to the

[2] There is no bias if strong returns in period 1 lead to assets being sold and the proceeds added to the side fund. This is because only cash injections (new investments) are added in the denominator: the $c/(1+r_f)$ term is omitted if $c<0$. But in a multi-period setting the MIRR will only be unbiased if there are no additional cash injections in any period.

side fund, which for simplicity we assume offers zero return. Conversely, a loss in the first period sees the equity portfolio topped up from the side account to the target 140.

The inefficiency of this strategy can be demonstrated by comparing the ranking of the terminal wealth outturns and the state price densities (the state prices divided by the probability – for this tree they are $16(1/3)^{u}(2/3)^{d}$, where u is the number of up states and d the number of down states on the path concerned[3]). Higher terminal wealth outturns generally come in the paths with lower state price densities, but not always. The best outturn is in the UUUU path, which has the lowest state price density. The second, third and fourth best outturns see three ups and one down. But the fifth best is DDUU, which beats UUUD into sixth place. Similarly, DDDU in eleventh place beats UUDD.

These results show that the VA strategy fails to make effective use of some relatively lucky paths (those with relatively low state price densities). This can be proved by generating a strategy which produces the same 16 outturns with a smaller initial investment. We do this by altering our strategy so that the paths with the lowest state price densities (the largest number of up states) always generate the greatest terminal wealth, so we swap the 5th highest outturn in Figure 1 with the 6th and the 11th highest with the 12th. We then work backwards through the tree using the state prices to calculate the equity and cash which must be held at each prior point. Ultimately this determines the initial capital which is needed. This new strategy is shown in Figure 2 and needs

---

[3] More generally, the state price densities of one period up and down states are $(1/(1+\mathrm{r}\Delta t))(1-((\mu-\mathrm{r})\Delta t/\sigma\sqrt{\Delta t}))$ and $(1/(1+\mathrm{r}\Delta t))(1+((\mu-\mathrm{r})\Delta t/\sigma\sqrt{\Delta t}))$ respectively, where r is the continuously compounded annual risk-free interest rate and the risky asset has annual expected return μ and standard deviation σ. The corresponding one period risky asset returns are $(1+\mu\Delta t+\sigma\sqrt{\Delta t})$ and $(1+\mu\Delta t-\sigma\sqrt{\Delta t})$. See Dybvig (1988a).

only 496.2 initial capital, rather than the 500 above, thus demonstrating the extent to which VA is inefficient. By generating the same set of possible outturns this alternative strategy must be taking the same level of risk as VA, no matter which measure of risk we use.

A strategy which generates the same outturns with less initial capital is clearly more efficient. The reduction in the initial capital required is a measure of VA's inefficiency compared to our alternative strategy. This is a powerful result because it demonstrates VA's inefficiency without needing to specify the investor's risk preferences, since producing the same set of outturns with less initial capital must be preferable regardless of the investors' risk preferences (this assumes only that terminal wealth is what investors care about, and that they prefer more terminal wealth to less).[4]

This analysis assumes that returns follow a binomial distribution, but the inefficiency of VA extends to other distributions. Rieger (2011) generalizes Dybvig's results to show that path-dependent strategies which generate outturns which have a non-monotonic relationship with market returns will be sub-optimal no matter what distribution these market returns follow. VA is an example of such a path-dependent strategy.

[Figure 2 here]

Figure 2 also shows the component of total wealth which is held in equities at each point. Unlike VA, where it follows its pre-determined growth target, the equity exposure of this

[4] Hayley (2012) uses the same technique to demonstrate that DCA is also inefficient regardless of investor risk preferences.

alternative strategy depends on the path taken. The equity holding of the optimized strategy is higher than for VA in the first and second periods, equal to VA's in the third, and equal or lower in the fourth. This confirms our intuition that the inefficiency of VA stems from investing gradually, and thus being insufficiently exposed to equity returns in early periods.

The doubling or halving of equity values in each period is an extreme assumption – for typical levels of equity volatility this would imply several years between successive investments. This allows us to illustrate dynamic inefficiencies in a short tree, but it is unrealistic for most investors. For a more realistic strategy we consider an eighteen period tree. This has 2^{18} paths, and is the largest that was computationally practical.[5]

Panel A in Table 4 shows the degree of inefficiency in VA strategies over a range of different time horizons and target returns (r^*). These were derived using a risk free rate of 5%, and risky asset returns with mean 10% and standard deviation 20% (all per annum). These efficiency losses remain very similar for a range of different volatilities (to save space these are not reproduced here). Correspondingly, the continuous time efficiency losses calculated in the annex are entirely unaffected by the level of volatility.

[Table 4 here]

[5] Dybvig (1988a) used this technique to demonstrate the inefficiency of stop-loss and target return strategies which are invested either fully in the risky asset, or fully in the risk-free asset. The number of paths involved is thus limited since the tree is generally recombinant, and collapses to a single path on hitting the target portfolio value. By contrast, VA varies the exposure in successive periods so DU and UD paths will not result in the same portfolio value. Thus an n period tree has 2^n paths and computation rapidly becomes impractical as n rises.

Two results are clear in Table 4. First, VA becomes increasingly inefficient if the target growth rate is set at a level which is significantly above or below the risk-free rate. Second, inefficiency increases dramatically as the time horizon is increased.

[Figure 3 here]

Figure 3 helps illustrate the reasons for these effects by showing the range of different terminal wealths which may be achieved for each of the possible final state prices. An inefficiency arises when the lines for different state prices overlap, showing that for some paths VA achieves lower terminal wealth than other paths which were less fortunate (those with a higher state price). By comparison, a strategy which simply invests in one initial lump sum would show only a single terminal wealth for each of the 19 possible state prices, since these outturns are then determined solely by the number of up and down states in each path, regardless of the order they occur in.

The order matters for VA. For example, if we have a high target growth rate then terminal wealth will be greater if the highest returns come late in the horizon (similar returns early in the horizon would have much less impact since a large proportion of the investor's wealth is at that time held as cash). VA is dynamically efficient when $r^*=r_f$, since a good outturn then has the same effect on expected terminal wealth whichever period it takes place in. Such an outturn in an early period will increase the amount of cash held by the investor, which will earn interest at rate r_f. A similar outturn in a later period will boost the value of the equity portfolio which will have grown at rate r^* in the meantime. If $r^*=r_f$, then these two effects offset each other, and the terminal wealth is not affected by the order in which U and D states occur. Only a single possible terminal wealth is then associated with each state price density and there is never an inefficient underutilization of a relatively lucky path. This result is confirmed by the continuous time efficiency losses derived in the appendix.

However, in practice r^* is likely to be substantially in excess of r_f, for three reasons. First, investors will naturally expect to earn a risk premium on their exposure to risky assets. Second, they are likely to overestimate their expected returns in the mistaken belief that VA will boost returns above what could normally be expected on these assets. Third, VA is generally used as a means of investing new savings as well as generating organic portfolio growth, so r^* is likely to be set above the expected rate of organic growth. Consistent with this, Edleson (1991) explicitly envisages that periodic cashflows will generally be additional purchases of risky assets rather than withdrawals of funds. Taking the risk premium to be 5% (as a very broad approximation), when we add investor overestimation of this risk premium and the desire to make further net investments, target growth rates are likely to be at least 5% higher than r_f, and quite plausibly 10% higher. Table 4 is calculated with r_f =5%, so the outturns shown for target growth rates in the range 10-15% are likely to be most representative.

Table 4 also shows that VA is much more inefficient over longer time horizons. Even with a fixed number of steps in the tree a longer time horizon allows greater variation in exposure over time, increasing the range of overlap in the terminal wealth levels that are possible for each state price. The differences between terminal state prices will also be larger. These two effects mean that a longer time horizon sees much greater inefficiency. VA is generally recommended as a long-term investment strategy (in particular for saving for retirement), so horizons of 10 to 20 years are likely to be more common than a 5 year horizon. Table 4 shows that over such time horizons, and with r^* in the range 10-15%, the dynamic inefficiency can be very substantial.

Furthermore, these figures are likely to understate the true efficiency losses for two reasons. First, the limited number of paths which can be computed results in comparatively large differences between ranked terminal wealth outturns. Thus small potential inefficiencies will not be recorded if they do not reduce the terminal wealth on one path sufficiently for it to fall below the terminal wealth achieved on at least one path with a higher state price density. This problem can be avoided by shifting to continuous time. This represents a simplification, since VA is intended to make any required additional investments at discrete (eg. monthly) intervals. But it has the advantage that all inefficiencies will be recorded since there will be an indefinite number of different paths with terminal wealths which differ only minutely.

An expression for the efficiency losses resulting from VA is derived in the appendix. The results are shown in Panel B of Table 4, and the continuous and discrete time estimates are compared in Figure 4. The discrete time estimates do indeed substantially underestimate the efficiency losses, especially for target returns close to r_f.

[Figure 4 here]

Furthermore, all the figures in Table 4 should be regarded as conservative estimates of the welfare loss to investors. They show how much more cheaply an investor could achieve the same distribution of outturns as a VA strategy. This method allows us to derive these welfare losses without needing to make any assumption about the form of the investor's risk preferences, but there is no reason why an investor who abandons VA should actually choose an alternative strategy with exactly the same payoff distribution. The investor is likely instead to find other strategies even more attractive, implying that the actual welfare benefits of abandoning VA are higher than shown in Table 4.

Thus whilst we can calculate plausible lower limits for the efficiency losses associated with VA, more realistic estimates would be more judgmental. However, the fact that there are efficiency losses for any distribution of returns and for any form of investor risk preferences comes in stark contrast to proponents' claims that VA outperforms alternative strategies.

6. Value Averaging In Inefficient Markets

In this section we consider whether VA could outperform in markets where asset returns contain a predictable time structure. However, it is worth stressing at the outset that this would be a much weaker argument in favor of VA than the outperformance in all markets (including random walks) which is claimed by VA's proponents. We also consider VA's performance against historical data.

This analysis is complicated by the fact that many popular performance measures will be inappropriate for assessing whether VA outperforms. The level of risk taken by VA depends on the growth target used, so differences in the expected return achieved by comparison strategies might simply reflect a different risk premium. This could normally be corrected for by comparing Sharpe ratios, but VA introduces a negative skew into the distribution of cumulative returns (compared to a lump sum investment) since larger additional investments are made following losses. For example, a series of negative returns could result in a VA strategy losing more than its initial capital as additional investments are made to keep the risk exposure at its target level. This would of course be impossible for a lump sum investment. Conversely, VA reduces exposure following strong returns, restricting the upside tail. This negative skew will be welfare-reducing under many plausible utility functions, and it also means that the Sharpe ratio will be misleading, since the comparatively small upside risk reduces the standard deviation of a VA strategy, even though investors are likely to prefer a larger upside tail. In addition, Ingersoll et al. (2007) show that performance measures such as the Sharpe ratio will be biased upwards when investment managers reduce exposure following good results and increase it following bad results. VA automatically adjusts exposures in this way, so there is also a dynamic bias which increases its Sharpe ratio.

Chen and Estes (2010) derive simulation results which explicitly include the cost of VA's side fund. These show that VA does indeed generate higher Sharpe ratios, but with greater downside risk. Given the negative skew, the Sortino ratio might be considered a more appropriate performance measure, but Chen and Estes show that VA generates a lower Sortino ratio than a lump sum investment. This is particularly discouraging since Ingersoll et al (2007) show that this ratio is also increased by the same dynamic bias as the Sharpe ratio.

Relaxing our previous assumption of weak-form efficiency, mean reversion in prices will tend to favor VA. Additional simulations (not reproduced here) suggest that single period autocorrelation has little impact on profits, but multi-period autocorrelation has a larger effect. Successive periods of low (high) returns result in large (small) cumulative additional investments which leave the portfolio well positioned for subsequent periods of high (low) returns. Consistent with this, VA outperforms DCA in our earlier simulations when the terminal asset price ends up close to its starting value, and it underperforms DCA when prices follow sustained trends in either direction.

There is evidence of long-term reversals in some asset returns (eg. de Bondt and Thaler, 1985) but, conversely, there is also a large literature documenting positive autocorrelation in other markets (momentum or 'excess trending'). The most relevant test for our purposes is whether VA outperforms when back-tested using historical returns – this will show whether any time structure in these market returns is sufficient to offset the innate inefficiency of VA.

Studies using historical data have not found that VA outperforms. Thorley (1994) calculates the returns to a VA strategy which invests repeatedly in the S&P500 index over a 12 month horizon for the period 1926-1991. He finds that the average Sharpe ratio of this strategy is below that of corresponding lump sum investments. Similarly, Leggio and Lien (2003) find that VA generates a Sharpe ratio which is lower than for lump sum investment in large capitalization US equities, corporate bonds or government bonds, with VA generating a larger Sharpe ratio only for small firm US equities. These results hold for both 1926-1999 and the more recent 1970-1999 period. The lower Sharpe ratios achieved by VA are particularly striking given the static and dynamic biases outlined above, which tend to bias the Sharpe ratio up.

This does not rule out the possibility that there are some markets which show time structures in their returns that VA could exploit but, as Thorley (1994) points out, even where suitable market inefficiencies can be detected, VA would be a very blunt instrument with which to try to profit from them. Other strategies are likely to be much more effective at extracting profits from such market inefficiencies, such as long/short strategies with buy/sell signals calibrated to the particular inefficiency found in historic returns in each market. Furthermore, any advantage gained by VA in such markets would have to outweigh the inherent inefficiency of the strategy, as demonstrated above. For all these reasons, market inefficiency is unlikely to be a convincing reason for using VA.

7. Behavioral Finance and Wider Welfare Effects

We saw above that far from generating the higher returns that its IRR appears to suggest, VA is actually an inefficient strategy. VA's proponents recommend the strategy solely on the basis of its higher IRR, making no claim that it has any wider benefits, but in this section we nevertheless consider whether wider welfare effects, such as behavioral finance effects, might explain why VA nevertheless remains very popular.

Statman (1994) proposed several behavioral finance effects which might explain DCA's popularity. First, prospect theory suggests that investors' utility functions over terminal wealth may be more complex than in traditional economic theory. However, this cannot explain VA's popularity. We saw above that VA must be a sub-optimal strategy regardless of the form taken by investor risk preferences, since alternative strategies can duplicate VA's outturns at lower initial cost. Indeed, VA produces a distribution of terminal wealth which has a downward skew, which would clearly be very unwelcome for investors who are loss averse or highly sensitive to extreme outliers.

Statman also suggested that by committing investors to continue investing at a pre-determined rate DCA prevents investors from exercising any discretion over the timing of their investments, and so: (i) stops investors misguidedly attempting to time markets (investor timing has generally been shown to be poor); (ii) by giving investors no discretion over timing it avoids the feelings of regret that might follow poorly-timed investments. VA could plausibly bring similar benefits[6], but even in the light of such wider possible benefits, it is likely to remain a less attractive strategy than DCA. Both strategies commit the investor to adding cash according to a pre-specified rule, but VA's cashflows are unpredictable so this is likely to require more active investor involvement (compared to DCA's entirely stable and predictable cashflows), implying more potential for regret.

Furthermore, the need for a side fund of cash or other liquid assets to fund VA's uncertain cashflows is likely to lead investors to hold a higher proportion of their wealth in such assets than would otherwise be optimal, with correspondingly less invested in risky assets. Investors' holdings of liquid assets are driven by the needs of the VA strategy and so cannot be set to

[6] The results above assumed that investors always prefer greater terminal wealth to less, but this might not be true if regret is important, since investor utility would then depend on the path taken, rather than just the terminal wealth ultimately achieved.

maximize investor welfare. This would imply a static inefficiency in addition to the dynamic inefficiency seen above.

The required size of the side fund will depend on the volatility of risky assets, but is likely to be substantial. With aggregate equity market volatility of around 15-20% per annum, a side fund of at least this fraction of the risky assets might be considered a bare minimum since we should anticipate occasional annual market returns substantially in excess of 20% below their mean. An alternative perspective is that another decade like 2000-2009 would see many markets stay flat or fall. For plausible levels of r^* this would leave investors trying to find additional cash worth more than the original value of their investments.

Furthermore, VA requires investors to sell assets after any period in which organic growth in the portfolio exceeds r^*. This may result in increased transaction costs compared to a buy-only strategy and, worse, could trigger unplanned capital gains tax liability. Edleson (2006) suggests that investors could reduce these additional costs by delaying or ignoring entirely any sell signals generated by the VA strategy, and that investors should in any case limit their additional investments to a level they are comfortable with. However, this re-introduces an element of investor discretion, implying possible bad timing and regret. By avoiding this DCA again appears to be the preferable strategy.

8. Conclusion

VA is recommended to investors as a method for raising investment returns in any market – even when prices follow a random walk. This paper shows that VA does indeed increase the expected

IRR, but it does not increase expected profits. Instead the IRR is boosted by a retrospective bias which arises because VA invests more following poor returns and less following good returns. The same bias will be found for any strategy which varies its exposure in response to the return achieved to date. This includes all strategies based on a target return or profit level, and also those which systematically take profits following strong returns or "double down" following weak returns.

VA does not achieve the outperformance that is claimed for it – in fact it is an inefficient strategy. This paper identifies four sources of inefficiency: (i) VA is dynamically inefficient, except in the unlikely case that the target return is very close to the risk free rate (this is a powerful result since it applies regardless of the form taken by investor risk preferences); (ii) VA also introduces a downside skew to cumulative returns which is likely to be welfare-reducing for many investors; (iii) VA is likely to cause static inefficiency by requiring larger holdings of cash and liquid assets than would otherwise be optimal; (iv) VA may increase management costs, transaction costs and tax liabilities compared to a buy-and-hold strategy. Behavioral finance effects may be important enough to some investors that they outweigh all these inefficiencies, but for such investors VA is likely to be an inferior strategy to DCA, which has stable cashflows.

Thus the central claim that is put forward for VA is illusory – it does not increase expected profits. Instead, the high IRRs that it generates are due to a retrospective bias. Furthermore, other properties of VA are likely to significantly reduce investor welfare. In short, VA has very little to recommend it. VA's popularity appears to be due to investors making a cognitive error in assuming that its higher IRR implies higher expected profits.

References

Constantinides, G.M., 1979. “A Note On The Suboptimality Of Dollar-Cost Averaging As An Investment Policy.” *Journal of Financial and Quantitative Analysis* 14, 443-450.

Chen, H. and J. Estes, 2010. “A Monte Carlo Study Of The Strategies For 401(K) Plans: Dollar-Cost-Averaging, Value-Averaging, And Proportional Rebalancing.” *Financial Services Review* 19, 95-109.

De Bondt, W.F.M. and R. Thaler, 1985. “Does the Stock Market Overreact?” *Journal of Finance* 40, 793-805.

Dichev, I.D. and G. Yu, 2011. “Higher Risk, Lower Returns: What Hedge Fund Investors Really Earn.” *Journal of Financial Economics* 100, 248–263.

Dybvig, P.H., 1988a. “Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market.” *Review of Financial Studies* 1, 67-88.

Dybvig, P.H., 1988b. “Distributional Analysis of Portfolio Choice.” *Journal of Business* 61, 369-393.

Edleson, M.E., 1991. “Value Averaging: The Safe and Easy Strategy for Higher Investment Returns.” Wiley Investment Classics, revised edition (second edition 2006).

Hayley, S., 2012. “Dollar Cost Averaging – The Role Of Cognitive Bias.” Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1473046

Hayley, S., 2013. “Hindsight effects in dollar-weighted returns.” Journal of Financial and Quantitative Analysis, forthcoming.

Ingersoll, J; Spiegel, M; Goetzmann, W and I. Welch, 2007. “Portfolio Performance Manipulation and Manipulation-proof Performance Measures.” *Review of Financial Studies* 20, 1503-1546.

Leggio, K. and D. Lien, 2003. “Comparing Alternative Investment Strategies Using Risk-Adjusted Performance Measures.” *Journal of Financial Planning* 16, 82-86.

Marshall, P.S., 2000. “A Statistical Comparison Of Value Averaging vs. Dollar Cost Averaging And Random Investment Techniques.” *Journal of Financial and Strategic Decisions* 13, 87-99.

Marshall, P.S., 2006. “A multi-market, historical comparison of the investment returns of value averaging, dollar cost averaging and random investment techniques.” *Academy of Accounting and Financial Studies Journal*, September.

Phalippou, L., 2008. "The Hazards of Using IRR to Measure Performance: The Case of Private Equity." *Journal of Performance Measurement*, Fall.

Rieger, M. O., 2011. Co-monotonicity of optimal investments and the design of structured financial products, *Finance and Stochastics* 15, 27-55.

Rozeff, M.S., 1994. "Lump-sum Investing Versus Dollar-Averaging." *Journal of Portfolio Management* 20, 45-50.

Statman, M., 1995. "A Behavioral Framework For Dollar-Cost Averaging." *Journal of Portfolio Management* 22, 70-78.

Thorley, S.R., 1994. "The Fallacy of Dollar Cost Averaging." *Financial Practice and Education* 4, 138-143.

APPENDIX: A Continuous Time Analysis Of VA's Inefficiency

This appendix uses the payoff distribution pricing model of Dybvig (1988b) to derive the continuous time efficiency losses shown in Table 4 and Figure 4. We assume an equity index (with zero dividends) which, relative to a constant interest rate bank account as numeraire, grows according to Geometric Brownian Motion as:

$$\frac{dS_t}{S_t} = \mu\, dt + \sigma\, dB_t \qquad (1)$$

This market offers a risk premium of μ and a Sharpe Ratio of μ/σ. We consider the degree of inefficiency by an investor who invests according to a fixed rule which determines the growth in the value $V_t = V_0\, g(t)$ invested in the equity market in each period from its initial V_0. Specifically in this case a value averaging strategy with target portfolio growth of α per period implies that V_t=$V_0 e^{\alpha t}$. These amounts are also relative to the bank account as numeraire, so a constant g (or $\alpha = 0$) corresponds to a value which grows at the interest rate. The investor's total wealth W_t grows according to:

$$dW_t = V_0 g(t)\left[\mu\, dt + \sigma\, dB_t\right]. \qquad (2)$$

We assume that the investor's initial wealth W_0 is sufficient to keep Vt on its target path, or that the investor can borrow enough for this purpose (indeed, we could set W_0=0 and assume that the strategy is entirely debt financed). These assumptions favour VA, since in practice no finite W_0 or credit line will ever be able to guarantee that adverse market outturns will not result in the VA strategy demanding more funds than the investor has available. This assumption implies that the distribution of terminal wealth at any later time T is normal with mean and variance given by:

$$E[W_t] = W_0 + V_0 \int_0^T \mu g(t) dt \tag{3}$$

$$Var[W_t] = V_0^2 \int_0^T \sigma^2 g^2(t) dt \tag{4}$$

The normal distribution of these outturns is due to the fact that the equity market exposure follows a pre-determined target path, and does not depend on the returns made to date. This opens up the possibility of total losses exceeding the initial wealth W_0, as following earlier losses the strategy demands that the investor borrows to top the portfolio up to its required level (this is in contrast to the lognormal distribution of a buy-and-hold strategy). We now need to work out the cost of the cheapest way to buy a claim with this normal distribution. For fixed horizon T the future index value is:

$$S_T(u) = S_0 \exp\left\{(\mu - ½\sigma^2)T + \sigma\sqrt{T}u\right\} \tag{5}$$

where u is a standard normal variate. The pricing function for this economy is:

$$m(u) = \exp\left\{-½\left(\frac{\mu}{\sigma}\right)^2 T - \left(\frac{\mu}{\sigma}\right)\sqrt{T}u\right\} \tag{6}$$

This has expectation of one, and integrates with S_T to give $E\left[m(u)S_T(u)\right] = S_0$. or, scaling to a payoff equal to the normal variate u: $E[u\ m(u)] = \mu\sqrt{T}/\sigma$

The exponential case

We now explicitly evaluate the minimum cost where $g(t) = e^{\alpha t}$. In this case:

$$E[W_t] = W_0 + M \text{ where}$$

$$M = V_0 \int_0^T \mu e^{\alpha t} \, dt$$

$$= \begin{cases} V_0 \, \mu \left[e^{\alpha T} - 1 \right] / \alpha; & \alpha \neq 0 \\ V_0 \, \mu T; & \alpha = 0 \end{cases}$$

$$Var[W_t] = S^2 = V_0^2 \int_0^T \sigma^2 e^{2\alpha t} \, dt$$

$$= \begin{cases} V_0^2 \sigma^2 \left[e^{2\alpha T} - 1 \right] / (2\alpha); & \alpha \neq 0 \\ V_0^2 \sigma^2 T; & \alpha = 0 \end{cases}$$

Dybvig (1988b) shows that the minimum cost of obtaining a specified set of terminal payoffs is given by the expected product of these payoffs with the corresponding state prices, where the payoffs and state prices are inversely ordered, so that the highest payoffs come in the lowest state price paths. Thus the minimum cost of obtaining the normally-distributed payoff $W_0 + M + S\,u$ is:

$$\begin{aligned} & W_0 + M + S\,E[u\,m(u)] \\ & = W_0 + M - S\,\mu\sqrt{T} / \sigma. \end{aligned}$$

This compares to the W_0 cost assumed for the VA strategy, so VA is inefficient by the magnitude of $S\,\mu\sqrt{T} / \sigma - M$ which simplifies to:

$$V_0 \mu \left\{ \sqrt{\frac{T}{2\alpha}\left[e^{2\alpha T} - 1 \right]} - \left[e^{\alpha T} - 1 \right] / \alpha \right\}.$$

Note that there is no inefficiency if μ or α are zero, and the inefficiency is small if T is small. Furthermore, σ cancels out, so volatility plays no role in determining the size of the inefficiency. Intuitively, the inefficiency is also proportional to V_0 and the initial wealth W_0 plays no role at all.

Table 1. Illustrative Comparison Of VA and DCA – Declining Prices

DCA and VA strategies are used to buy an asset whose price varies over time (the price could also be interpreted as a price index, such as an equity market index). DCA invests a fixed dollar amount ($100). VA invests whatever amount is required to increase the portfolio value by $100 each period.

		Dollar Cost Averaging (DCA)			Value Averaging (VA)		
Period	Price	Shares bought	Investment ($)	Portfolio ($)	Shares bought	Investment ($)	Portfolio ($)
1	1.00	100	100	100	100	100	100
2	0.90	111	100	190	122	110	200
3	0.80	125	100	269	153	122	300
Total		336	300		375	332	
Avg.price 0.90		Avg.cost:	0.893		Avg.cost:	0.886	

Table 2. Illustrative Comparison Of VA and DCA – Rising Prices

Strategies are as defined in Table 1. The price of the asset is here assumed to rise.

		Dollar Cost Averaging (DCA)			Value Averaging (VA)		
Period	Price	Shares bought	Investment ($)	Portfolio ($)	Shares bought	Investment ($)	Portfolio ($)
1	1.00	100	100	100	100	100	100
2	1.10	91	100	210	82	90	200
3	1.20	83	100	329	68	82	300
Total		274	300		250	272	
Avg.price 1.10		Avg.cost:	1.094		Avg.cost:	1.087	

Table 3. Simulation Results: Performance Differentials

This table compares strategies which invest in an asset whose returns are assumed to follow a random walk with no drift. Following Marshall (2000, 2006), equity prices are assumed to start at $10 and then evolve in each of 100,000 simulations for five periods. In each period returns are *niid* with a 10% standard deviation. DCA invests a fixed $400 each period; VA invests whatever amount is required to increase the portfolio value by $400 each period; the lump sum strategy invests $2000 in the first period. These parameters were chosen so that the VA and DCA strategies will be identical if prices remain unchanged. Standard errors are shown in brackets.

	Average Cost (cents)	IRR (%)	Profit ($)
DCA - Lump Sum	-8.15	0.076	0.133
	(1.10)	(0.023)	(2.219)
VA-Lump Sum	-20.16	0.278	0.086
	(1.08)	(0.023)	(2.248)

Table 4. Measuring the Dynamic Efficiency of Value Averaging

This table shows the additional initial capital required by a VA strategy compared with an optimized strategy which generates an identical set of final portfolio values. These figures are derived using the Dybvig PDPM model applied to a VA strategy over an 18 period tree with risk free rate 5%, expected market return 10% and volatility 20% (all per annum). The inefficiency is shown as a percentage of the average monthly portfolio exposure of the VA strategy. For the discrete time calculation an 18 period tree is used throughout, with the length of each period varied to achieve the total time horizon shown. The derivation of the corresponding continuous time losses is shown in the appendix.

Panel A: Discrete time estimates of efficiency losses

Target growth (per annum)	5 years	10 years	15 years	20 years
-10%	0.33%	3.25%	9.80%	20.01%
-5%	0.06%	1.21%	4.12%	8.90%
0%	0.00%	0.09%	0.66%	1.79%
5%	0.00%	0.00%	0.00%	0.00%
10%	0.00%	0.05%	0.43%	1.18%
15%	0.03%	0.81%	2.68%	5.25%
20%	0.18%	2.06%	5.69%	10.12%

Panel B: Continuous time estimates of efficiency losses

Target growth (per annum)	5 years	10 years	15 years	20 years
-10%	0.57%	4.33%	13.43%	28.73%
-5%	0.26%	2.01%	6.50%	14.59%
0%	0.06%	0.52%	1.72%	4.02%
5%	0.00%	0.00%	0.00%	0.00%
10%	0.06%	0.52%	1.72%	4.02%
15%	0.26%	2.01%	6.50%	14.59%
20%	0.57%	4.33%	13.43%	28.73%

Figure 1. Simple Model of VA Strategy

This figure shows the total investor wealth (the upper figure at each point) for a VA strategy with a portfolio growth target of 40% each period. The lower figures show the amount of this total wealth which is held in equities. Equity values are assumed to double in a good outturn and halve in a bad outturn. Equity investment is adjusted back to the target value after each period using transfers into and out of the side account. For illustrative purposes funds in the side account are assumed to earn zero interest (Table 4 shows that inefficiencies persist with a higher risk free rate). All paths are assumed to be equally likely.

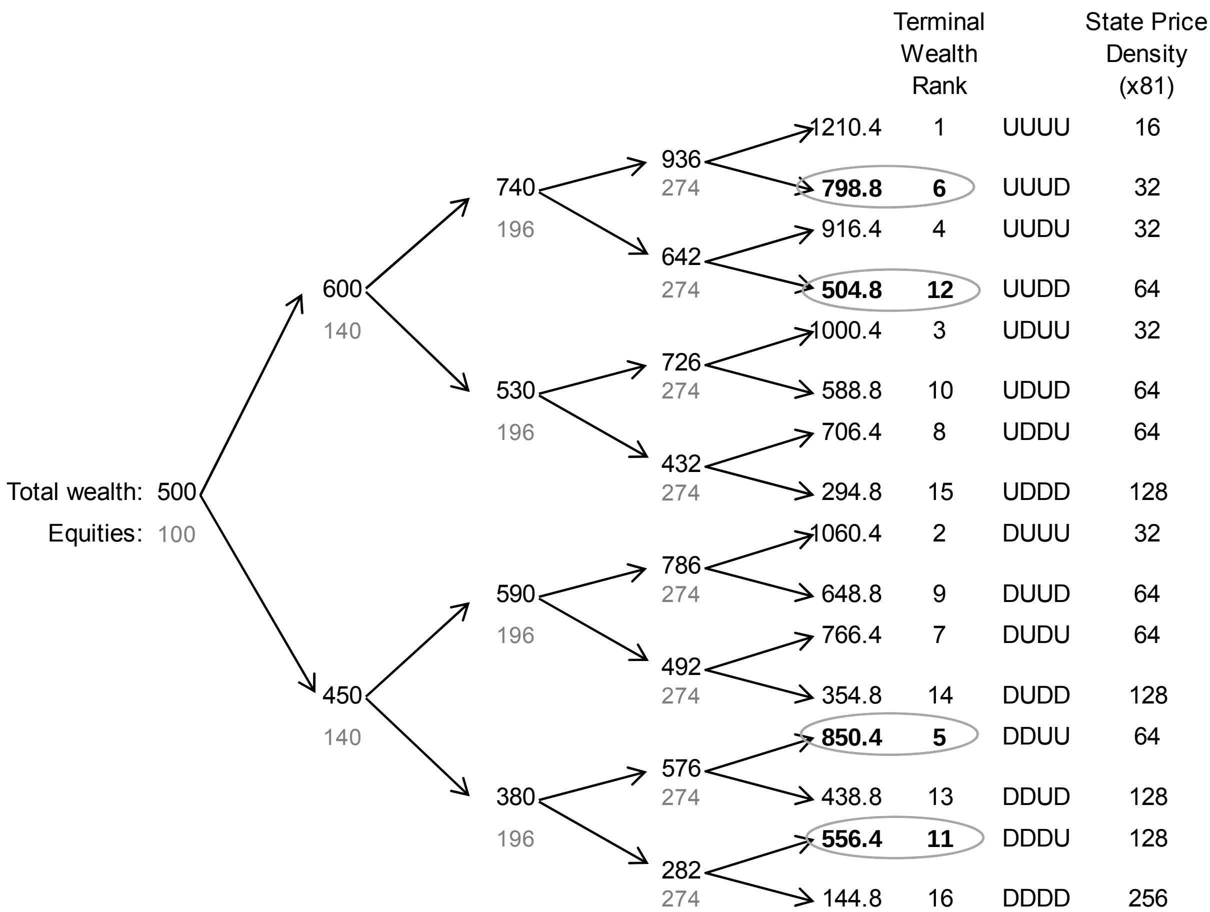

Figure 2. Optimized Strategy Giving Identical Outturns To VA Strategy

The upper figure shows the total investor wealth at each point in a strategy in which the equity exposure (the lower figure at each node) has been set so as to replicate the total wealth outturns in Figure 1, but with these outturns optimized so that the largest outturns always come in the states with the lowest state price density. Compared with Figure 1, the outturns for UUUD and DDUU have been swapped, and the outturns for UUDD and DDDU. Equity returns are as assumed in Figure 1. The lower initial capital required for this optimized strategy to generate an identical set of outturns shows the degree to which the VA strategy is inefficient.

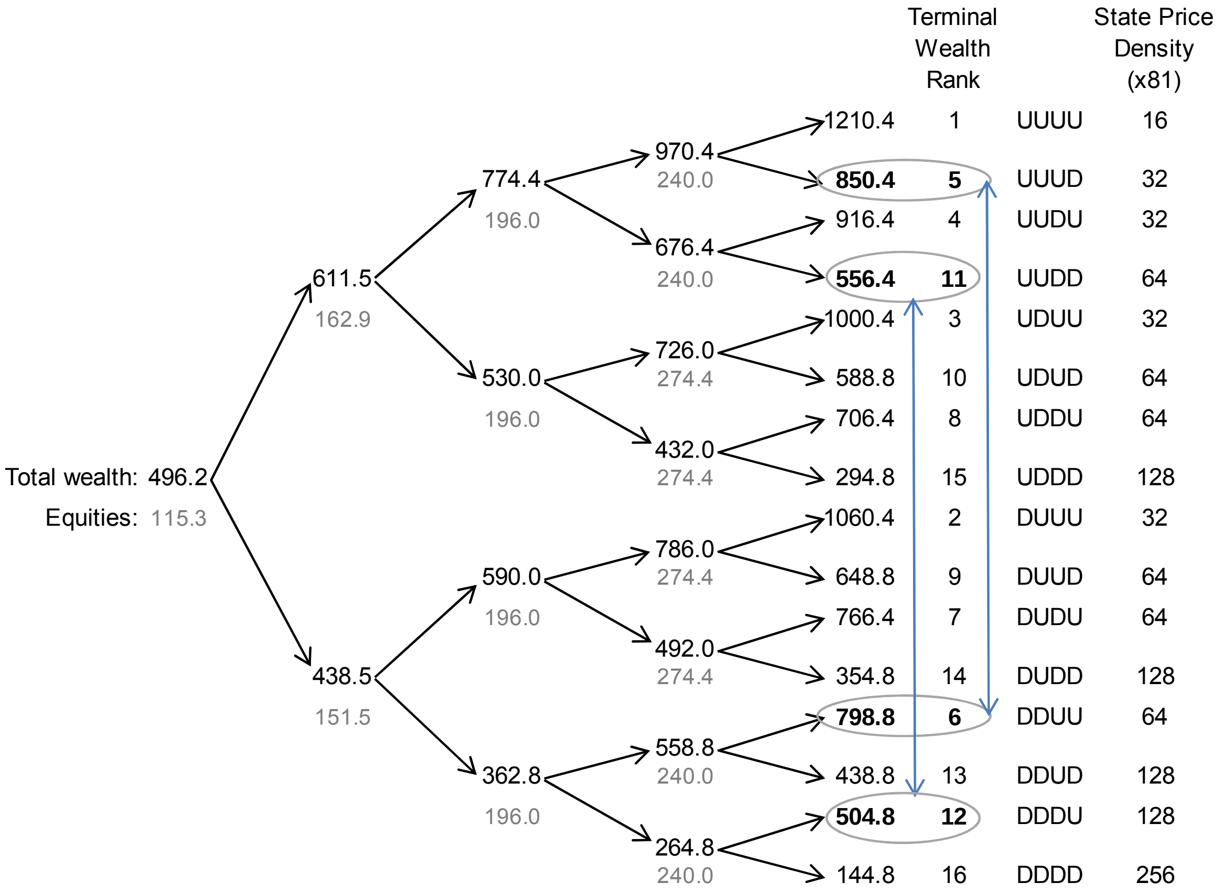

Figure 3. Terminal Wealth Achieved by VA vs. State Price

This shows the range of terminal wealth levels achieved by VA for each of 19 possible terminal state prices. The strategy is run over 18 periods with target return 10%, risk free rate 5%, expected market return 10% and volatility 20% (all per annum). An overlap, where any path achieves a greater terminal wealth than a path with a higher terminal state price, represents an inefficiency: The strategy could then be changed to achieve identical outturns at lower cost. Results are shown for investment horizons of 10 and 20 years.

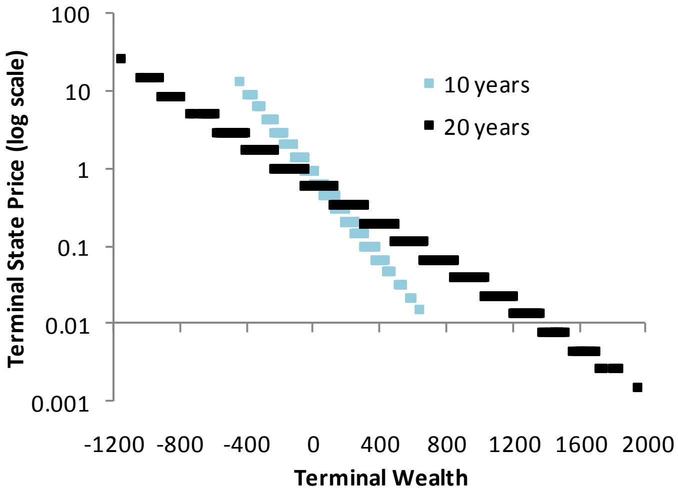

Figure 4. Dynamic Efficiency Losses Of VA Strategy

This shows the efficiency losses of a VA strategy (calculated in both discrete and continuous time) as a percentage of the average equity exposure of the strategy. The investment horizon is 10 years, risk free rate 5% and volatility 20% per annum. The derivation of the continuous time efficiency losses is shown in the Appendix.