VALUE AVERAGING AND THE AUTOMATED BIAS OF PERFORMANCE MEASURES

Simon Hayley*

This version: 14 February 2012

Abstract

Value averaging (VA) is a popular investment strategy which is recommended to investors because it achieves a higher IRR than alternative strategies. However, this paper demonstrates that this is entirely due to a hindsight bias which raises IRRs for strategies which - like VA - link the scale of additional investment to the returns achieved to date. VA does not boost profits – in fact it suffers substantial dynamic inefficiency. VA can generate attractive behavioural finance effects, but investors who value these are likely to prefer the simpler Dollar Cost Averaging strategy, since VA imposes additional direct and indirect costs on investors as a result of its unpredictable cashflows.

JEL Classification: G10, G11

* Cass Business School , 106 Bunhill Row, London EC1Y 8TZ, UK. E-mail: simon.hayley.1@city.ac.uk. Tel +44 20 7040 0230. I am grateful to Stewart Hodges, Richard Payne, Giorgio Questa and Nick Ronalds for useful comments. The usual disclaimer applies.

VALUE AVERAGING AND THE AUTOMATED BIAS OF PERFORMANCE MEASURES

Value averaging (VA) is a very popular formula investment strategy which invests available funds gradually over time so as to keep the portfolio value 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 thousands of positive references to this strategy.

The use of the IRR to assess returns may seem intuitive, since it takes into account the varied cashflows that are inherent in the strategy. However, in this paper we find that the IRR is subject to a systematic hindsight bias for strategies which – like VA – link the size of any additional amounts invested to the performance achieved to date. This bias retrospectively increases the weight given in the IRR calculation to strong period returns and reduces the weight given to weaker returns. The modified internal rate of return (MIRR) is similarly biased

We demonstrate below that if returns follow a random walk then VA generates a higher IRR, but no increase in expected terminal wealth. The higher IRRs recorded for VA are instead entirely due to the hindsight bias.

Not only does VA fail to deliver the superior returns that its higher 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 offer imperfect time diversification. We demonstrate below

that VA is an inefficient strategy for any plausible investor utility function. We also quantify the resulting welfare losses. Certain types of weak form inefficiency in market returns could in principle justify the use of VA but we find that - after taking into account the biases which VA instills in most performance measures - the recorded historical returns suggest that the time structure of market returns has in fact tended to penalize VA.

The strategy has a number of other unattractive properties. It introduces a downward skew to cumulative returns which is likely to be welfare-reducing for many investors and it is likely to cause static inefficiency by requiring larger holdings of cash and liquid assets than would otherwise be optimal. VA's volatile and unpredictable cashflows may also increase management costs, transaction costs and tax liabilities compared to a buy-and-hold strategy. We conclude that not only does VA not generate the higher expected profits that are claimed, but using this strategy is likely to significantly reduce investor welfare.

This paper's contribution is: (i) Identifying the hindsight bias which affects the IRR for strategies such as VA where the scale of further investment is determined by returns to date; (ii) Identifying and quantifying the static and dynamic inefficiencies which are inherent in VA; (iii) Re-interpreting historical performance measures in the light of these biases to show that major markets have had time series properties which disadvantage VA.

The structure of this paper is as follows: the following section describes the VA strategy and related literature. Section II demonstrates that in contrast to its proponents' claims, VA cannot expect to generate excess profits when asset prices are unforecastable. Section III shows that IRRs are biased upwards for strategies such as VA. Section IV demonstrates that VA is

actually less efficient than alternative strategies, quantifies the dynamic inefficiency and identifies other sources of welfare loss. Section V finds that the time structure of historical returns in key markets has also tended to be detrimental to VA. Section VI finds that VA's popularity cannot be explained by wider behavioral finance effects. Conclusions are drawn in the final section.

I. The Value Averaging Strategy

VA is similar in some respects to dollar cost averaging (DCA), which is the strategy of building up investments gradually over time in equal dollar amounts. 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 I). 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 II). As long as there is any variation in prices DCA will achieve a lower average cost.

[Table I. here]

VA is a slightly more complex strategy which 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). The investor must make whatever additional investments are

necessary in each period to meet this target. Like DCA, VA purchases more shares after a fall in prices, but the response is more sensitive: In Table I 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 II. here]

VA could in principle be applied over any time horizon, but its originator suggests quarterly or monthly investments (Edleson, 2006). 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.

VA has so far been the subject of limited academic research. However, some results from the literature on DCA can be applied. Both VA and DCA commit 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 showed that strategies which pre-commit investors in this way will be dominated by strategies which instead allow investors to react to incoming news.

DCA and VA might seem to improve diversification by making many small purchases, but Rozeff (1994) shows that this is not the case for DCA. The strategy starts with a very low level of market exposure, so the terminal wealth will be most sensitive to returns later in the horizon, by which time the investor is more fully invested. Better time diversification is achieved by investing in one initial lump sum, and thus being fully exposed to the returns in each period.

Strategies which deliberately increase exposure over time (as DCA and VA both do) are likely to be sub-optimal in risk-return terms. An investor who has funds available should invest immediately rather than wait. Section IV confirms this and quantifies the resulting inefficiency of VA.

DCA has the benefit of stable cashflows, but 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].

Thorley (1994) compares DCA and VA with a static buy-and-hold strategy for the S&P500 index over the period 1926-1991 and finds that VA performs worse than other strategies 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 either DCA or VA. We consider these results further in section V, but they clearly do not support the claim that VA increases expected profits.

[1] Edleson and Marshall both calculate the IRR on the VA strategy without including returns on the side fund. We follow the same approach in this paper 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 the cash held in the side fund. This exclusion could be considered a misleading piece of accounting, since the availability of the side fund is a vital part of the strategy. However, including a side fund does not necessarily remove the bias: The modified IRR (MIRR) includes cash holdings, but section III shows that it is also a biased measure of VA's profitability.

However, VA's proponents continue to stress its demonstrable advantage: that it achieves a higher expected IRR than alternative strategies (Edleson (2006), Marshall (2000, 2006)). In this paper we focus explicitly on the reason for these higher IRRs, since this appears to be the key to VA's popularity. We find that the IRR is raised by a systematic bias inherent in the VA strategy which allows VA to generate attractive IRRs even though it does not increase expected profits.

II. Simulation Evidence

VA is presented by its proponents as a way of boosting returns in any market, even if the investor has no ability to forecast returns. Indeed, Edleson (2006) demonstrates that VA generates a higher expected IRR than alternative strategies even on simulated data which follow a random walk. This section uses simulation evidence to demonstrate that the IRR is a biased measure of the profitability of VA. Section III derives this result more formally and demonstrates how this bias arises.

We assume a random walk in the simulations below, although we subsequently relax this in section V to consider whether weak form inefficiencies in market returns could justify the use of VA. We also assume that this random walk has zero drift.[2] Investors presumably believe that

[2] The assumption of zero drift does not imply any loss of generality, since drift could be incorporated into this framework by defining prices not as absolute market prices, but as prices relative to a numeraire which appreciates at a rate which gives a fair return for the risks inherent in this asset ($p_i^*=p_i/(1+r)^i$, where r reflects the cost of capital and a risk premium appropriate to this asset). We could then assume that p_i^* has zero expected drift since investors who are willing to use VA or DCA will not believe that they can forecast short-term returns relative to other assets of equivalent risk level: those who do would again reject trading strategies which predetermine the timing of their investments. The results derived here would continue to hold for p_i^*, with profits then defined as excess returns compared to the risk-adjusted cost of capital. Indeed, this assumes that funds not yet needed for the VA strategy can be held in assets with the same expected return. This assumption is generous to VA – if instead cash is held on deposit at lower expected return, then VA's expected return is clearly reduced by delaying investment.

over the medium term their chosen securities will generate an attractive return, but they must also believe that the return over the short term (while they are building up their positions) is likely to be small. Investors who expect significant returns over the short term should clearly prefer to invest immediately in one lump sum rather than follow a strategy which invests gradually. Indeed, our assumption of zero drift clearly favors VA. A more realistic assumption of upward drift would see VA generating lower expected returns since funds are initially kept in cash and are only invested gradually.

We conducted 10,000 simulations in which, following Marshall (2000, 2006), the share price is initially $10 and evolves over 5 periods, with investments assumed to be liquidated in the final period. Returns are *niid* with zero mean and standard deviation of 10% per period. Table III shows the differential in the average costs, IRRs and profits achieved by VA and DCA. Both achieve very significant reductions in average purchase cost compared with a strategy which simply invests the available funds immediately in one lump sum. VA achieves a significantly larger reduction than DCA. VA and DCA also achieve significantly higher IRRs (by 0.28% and 0.08% respectively) than investing available funds in one immediately in one lump sum.

VA and DCA appear attractive when judged on their high IRRs and low average purchase costs, but this does not translate into higher expected profits. The simulations show that the profits made by these strategies are not significantly different from those of a lump sum investment strategy.

[Table III. here]

It is straightforward to confirm the results of these simulations by showing that under these conditions VA cannot generate a higher level of expected profit. The total dollar profit made by any investment strategy is the sum of the profits made on the amounts that are invested in each period i. The strategies we consider here give different weights to each period as they invest different amounts, but in a driftless random walk the expected profit is zero for investments made in *any* period, so altering the amount invested in each period cannot affect total expected profits. The weighted sum of a sequence of zeroes remains zero no matter how we change the weights attached to each.[3]

Thus in this situation VA generates seemingly attractive increases in the IRR and reductions in the purchase cost, but it generates exactly the same expected profit as a simple lump sum investment strategy. The average purchase cost and the IRR are biased indicators of expected profits.

The similar bias in DCA's average purchase cost has been covered elsewhere. DCA always achieves an average purchase cost which is lower than the average price, and this is the key advantage claimed for the strategy. Thorley (1994) notes that this lower average cost is "a seemingly plausible but irrelevant criterion", and Hayley (2010) shows that the comparison is

[3] Formally, we consider investing over a series of n discrete periods in an asset whose price in each period i is p_i. Alternative investment strategies differ in the quantity of securities q_i that are purchased in each period. We evaluate profits at a subsequent point T, after all investments have been made. If prices are then p_T, the expected profit made by any investment strategy is as shown below. Our assumption of a random walk implies that future price movements (p_T/p_i) are always independent of the past values of p_i and q_i. However, the random walk has zero drift, so $\mathrm{E}[p_T/p_i]=1$ for all i and expected profits are zero regardless of the amount $p_i q_i$ which is invested in each period.

$$\mathrm{E}[profit] = \sum_{i=1}^{n}\left(\mathrm{E}[p_T q_i]\right) - \sum_{i=1}^{n}\left(\mathrm{E}[p_i q_i]\right) = \sum_{i=1}^{n}\left(\mathrm{E}\left[\frac{p_T}{p_i}\right]\mathrm{E}[p_i q_i]\right) - \sum_{i=1}^{n}\left(\mathrm{E}[p_i q_i]\right) = 0$$

systematically misleading, since it effectively compares DCA with a counterfactual strategy which uses perfect foresight to invest more ahead of falling prices and less ahead of rising prices.

Table I can be used to illustrate this misleading comparison. DCA buys 100 shares in the first period. A strategy which bought the same numbers of shares in the next period would buy at an average cost equal to the average price. But 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 lower than the unweighted average price over this investment horizon. 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. As we saw, these strategies also achieve lower average costs when prices rise (Table II), 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. This is illustrated in Figure 1, which compares the total number of shares purchased by each strategy with the share price in the final period. DCA and VA purchase more shares after a fall in prices and fewer after a rise in prices. This alters the average purchase price, but these are retrospective responses to previous price movements - buying more shares after prices have fallen and fewer after prices have risen. These variations will not affect expected profits if share prices show no expected drift and returns are unforecastable.

[Figure 1 here]

Furthermore, if we introduce an upward drift in share prices it is clear that VA and DCA are inferior strategies. Investing the available funds immediately in one lump sum would be preferable, since the other strategies delay investing and so earn a lower risk premium. With a negative drift a superior strategy would be to avoid investing in this asset at all. We consider the impact of possible market inefficiencies in section V.

III. The Bias In The IRR

Edleson (2006) 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 inherent in VA. However, in this section we show why these IRRs are consistently misleading.

The 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 the investments equals the present value of the final liquidation value in period T:

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

Following Dichev and Yu (2009), we can substitute equation 1 into equation 2 to eliminate a_t. After some rearranging this 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:

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

The return in any period may be above or below the IRR, but we can further re-arrange equation 3 to show that the weighted average of these differentials is zero:

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

This 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:

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

The additional investment increases the weight given to future returns, compared to the weights based on the portfolio values K_t^* which would otherwise have been seen. Consistent with our assumption that returns cannot be forecast we can assume that periodic returns r_t are drawn from an underlying distribution with a fixed mean. If the r_t up to period m were on average below this mean, 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 and will tend to increase the IRR. Similarly, investing less (or even withdrawing funds) after a period of generally above-mean returns will tend to reduce the relative weight given to later $(r_t - IRR)$ terms, which would tend to be negative, again increasing the IRR.

This is precisely what VA does automatically, since the amount invested each period is determined by the degree to which the change in the value of the portfolio over the immediately preceding period ($r_m K_{m-1}$) fell short of the target. The first summation in Equation 5 includes r_m so the level of new investment will be negatively correlated with this summation. This correlation biases the overall IRR, since this additional 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) notes 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 time 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. Any strategy which targets a particular level of portfolio growth will tend to have this property – VA is just one example. 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. It is by doing this automatically that VA raises its expected IRR.

The IRR from investing in a given asset could be raised by either (i) positive correlation of the additional amounts invested with subsequent returns (good timing which will also increase expected profits), or (ii) a negative correlation of these additional investments with earlier returns (a hindsight bias which has no impact on expected profits). But when returns are assumed to follow a random walk the cashflows will be uncorrelated with future returns, and the good timing effect will average zero. By contrast, VA by construction ensures a negative correlation between cashflows and prior returns, so it must be this hindsight bias which accounts for the fact that even in random walk data VA achieves a higher IRR than investing in a single lump sum.

Including the 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 when back-testing. 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. At the end of period 1 an amount c is added to the risky allocation (or subtracted) from the side fund.

$$Terminal\ Wealth\ (TW) = a(1+r_1)(1+r_2) + c(1+r_2) + (b(1+r_f) - c)(1+r_f) \tag{6}$$

This measure is unbiased, 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:

$$IRR = \sqrt{\frac{TW}{a+b}} - 1 \tag{7}$$

This too is unbiased, since a and b are both fixed in advance. By contrast, the modified internal rate of return (MIRR) assumes the existence of a side fund which is just big enough to fund subsequent cash injections, and includes the cost of borrowing this amount (at rate r_b) in the denominator.

$$MIRR = \sqrt{\frac{a(1+r_1)(1+r_2) + c(1+r_2)}{a + c/(1+r_b)}} - 1 \tag{8}$$

The MIRR is biased because the weight ($a/(a+c/(1+r_b))$) given to r_1 is adjusted retrospectively. VA ensures that a low r_1 leads to a large subsequent cash injection c from the side fund, so the weight on r_1 is reduced after the event. If cash is added to the side fund following strong returns in period 1, then there is no bias since the MIRR adds the final value of this amount to the numerator, leaving the coefficient on r_1 unchanged. But in a multi-period setting the MIRR will only be unbiased if there are no additional cash injections in any period. To avoid this bias we would need to include a side fund which is big enough so that over no plausible paths is there ever a retrospective adjustment.

IV. Is Value Averaging Inefficient?

The analysis above showed that if asset returns follow a random walk VA does not generate higher expected profits than alternative strategies, despite its higher expected IRR. In this section we go one step further and consider whether VA is an inefficient strategy, with other strategies offering preferable expected return and risk characteristics. For this purpose we use the payoff distribution pricing model derived by Dybvig (1988a).

[Figure 2 here]

Figure 2 shows a simple model of the terminal wealth generated a VA strategy over four periods. The equity element of the portfolio is assumed to double in a good outturn and halve in a bad outturn. The investor has 100 initially invested in equities and has chosen a portfolio growth target of 40% each period. If the value of these equities rises in the first period to 200, then 60 is assumed to be transferred to the 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 equity component of the portfolio is adjusted back to the target level at the end of each period, so the potential profit/loss in any period is the same regardless of the path taken so far. This allows us to determine the cumulative profit/loss without considering the risky asset holdings and the side fund separately. All paths are assumed to be equally likely.

The key to this technique is comparing the terminal wealths with their corresponding state price densities (the state price divided by the probability – in this case $16(1/3)^u(2/3)^d$, where u

and *d* are the number of up and down states on the path concerned[4]). The better VA outturns generally correspond to the lower state price densities, but there are exceptions. 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 the VA strategy failing to make effective use of some relatively lucky paths (those with relatively low state price densities). This can be proved by deriving an alternative strategy which generates exactly the same 16 outturns at lower cost. This is done by changing our strategy so that the best outturns always occur in the paths with the lowest state price densities (and hence the largest number of up states), so we swap the 5th highest outturn in Figure 2 with the 6th and the 11th highest with the 12th. The state prices can then be used to determine the value of earlier nodes (in effect specifying the leverage at each point), and this in turn determines the initial capital required to generate these outturns. This alternative strategy is shown in Figure 3 and requires only 396.2 initial capital, compared to 400 above. This shows the degree to which the VA strategy is inefficient. A key advantage of this method is that it establishes this result without needing to specify the investor's utility function, since generating identical outturns at lower initial cost can be considered better under any plausible utility function (it assumes only that investors prefer more terminal wealth to less).

[4] More generally, the state price densities of one period up and down states are $(1/(1+r\Delta t))\left(1-\left((\mu - r)\Delta t/\sigma\sqrt{\Delta t}\right)\right)$ and $(1/(1+r\Delta t))\left(1+\left((\mu - r)\Delta t/\sigma\sqrt{\Delta t}\right)\right)$ 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 $\left(1+\mu\Delta t+\sigma\sqrt{\Delta t}\right)$ and $\left(1+\mu\Delta t-\sigma\sqrt{\Delta t}\right)$. See Dybvig (1988a).

[Figure 3 here]

Figure 3 also shows the component of total wealth which is held in equities at each point. Unlike VA (where it follows its pre-determined target), this varies depending on the path that has been followed. 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 being under-invested in early periods.

The doubling or halving of equity values at each step of this tree would - for volatility levels typical of developed equity markets - correspond to a gap of several years between successive investments. This extreme assumption allows us to illustrate dynamic inefficiencies in a very short tree, but it is unlikely to be realistic for most investors. For a more plausible 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 IV shows the degree of inefficiency of VA strategies estimated over this period using 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). The results are very similar for a range of different volatilities (not reproduced 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.

[Table IV here]

Two results are clear. 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 4 here]

Figure 4 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 simple lump sum strategy 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. By contrast, 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 an 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.

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 (2006) 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 IV 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 IV 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. The combination of

these two effects means that a longer time horizon sees much greater inefficiency. VA is intended for personal use, and will often be used for saving for retirement, so horizons of 10 to 20 years are likely to be more common than the 5 year horizon. Table IV shows that over such time horizons, and with r^* in the range 10-15%, the dynamic inefficiency can be very substantial.

However, these figures are likely to understate the true efficiency losses. This is because the constraint on the number of possible paths which can be computed means that the differences between each possible terminal wealth can be significant. Thus small potential inefficiencies will not be recorded if they are insufficient to lift the terminal wealth on one path by enough to exceed the terminal wealth achieved on a path with a lower 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, and results are shown in Panel B of Table IV. The continuous and discrete time estimates are compared in Figure 5. The discrete time estimates do indeed substantially underestimate the efficiency losses, especially for target returns close to r_f. The continuous time estimates ensure that even a very small range of variation in the terminal wealth achieved for each state price will – quite correctly – be recorded as an inefficiency. As a result, the continuous time estimates show an efficiency loss of 0.52% for an investment horizon of 10 years and a target growth rate only 5% above the risk free rate. This should be considered economically significant – an investment

manager who consistently underperformed by this margin would soon lose clients. But, as discussed above, this should be considered the lower end of the plausible range for target growth rates. Higher growth rates and longer time horizons would see massive efficiency losses.

[Figure 5 here]

Furthermore, even these continuous time figures are likely to be conservative estimates of the welfare loss to investors. They show how much more cheaply an investor could achieve the same potential 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 utility function (other than assuming that more wealth is preferred to less). However, there is no reason why of all the available strategies an investor who abandons VA should actually choose an alternative strategy with exactly the same potential payoffs. 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 here.

V. 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 assess 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 average 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.

This skew 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 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 increasing its Sharpe ratios.

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 biased up 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. Our simulations 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. Figure 6 shows that VA outperforms 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.

[Figure 6 here]

There has been some evidence of long-term reversals in asset returns (following 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 these returns tend to incorporate time structures which favor 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.

However, the static and dynamic biases outlined above mean that we should expect VA to achieve a higher expected Sharpe ratio even if returns follow a random walk. The fact that these historical studies tend to show VA underperforming corresponding lump sum investments despite these upward biases suggests that the time series properties of historical returns in these markets have been worse than random for VA, probably due to occasional positive autocorrelation of returns.

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 inefficiencies inherent in the strategy, as discussed above.

In sum, VA would be a poor choice for exploiting any identified market efficiencies and the poor performance of VA on historical data (despite the upward bias in performance measures such as the Sharpe ratio) suggests that the time structure of returns in key markets has actually tended to penalize VA. For these reasons, market inefficiency is not a convincing rationale for using VA.

VI. Behavioral Finance and Wider Welfare Effects

The sections above showed that VA does not generate the higher returns that its IRR appears to suggest. In this section we consider whether behavioral finance effects can 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. Section IV showed that VA must be a sub-optimal strategy regardless of the form of the utility function, 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 (compared to a buy-and-hold strategy). This would clearly be a very unwelcome property 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 from misguided attempts to time markets (investor timing has generally been shown to be poor), (ii) avoids the feeling of regret that might follow poorly-timed investments. VA could bring similar benefits. Our 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.

We should assess VA's performance on these wider criteria against those of DCA, its obvious alternative. If there is no time structure to asset returns then neither strategy will boost expected returns - in contrast to the claims made for them. Both commit the investor to add cash according to pre-specified targets. However, DCA's cashflows are by construction entirely stable, whereas VA's unpredictable cashflows are likely to require more active investor involvement. Large gains can force a VA investor to withdraw funds from the market, and large losses may leave an investor having to decide whether it is practical to achieve his chosen VA growth target. This suggests that VA is more likely to cause regret than an entirely stable and predictable DCA strategy.

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

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. 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.

We showed above that if investors just care about terminal wealth then VA must be a sub-optimal strategy. Behavioral finance factors may imply that this is not necessarily the case, but even if these are important to investors we find that VA is clearly inferior to DCA as a means of capturing these wider benefits.

Academics tend to be cautious about normative conclusions suggesting that investor behavior is misguided. But it is hard to avoid this conclusion in this case, since VA is clearly an inferior strategy in regard to both terminal wealth and wider welfare effects. Furthermore, VA's proponents recommend the strategy solely on the basis of its higher IRR, making no claim that it has any wider benefits. For both these reasons, the best explanation for VA's popularity is that investors are making a cognitive error in assuming that VA's higher IRR implies higher expected profits.

VII. Conclusion

VA is recommended to investors as a method for raising investment returns in any market - even when prices follow a random walk. We find that VA does indeed increase the expected IRR, but it does not increase expected profits. Instead the IRR is boosted by a hindsight bias which arises because VA invests more following poor returns and less following good returns. The same bias will be found for any other strategy which varies the level of new investment in response to the return achieved to date (strategies which take profits after hitting a specified target return are another example).

Not only does VA not achieve the outperformance that is claimed for it – it is also an inefficient strategy. We have identified six sources of inefficiency: (i) VA will be dynamically inefficient, except in the unlikely case that the target return is very close to the risk free rate; (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) Studies which back-test VA using historical data show comparatively poor Sharpe ratios, despite the fact that VA imparts an upward bias to such performance figures - this suggests that historic returns differ from a random walk in ways which have disadvantaged VA; (v) VA may increase management costs, transaction costs and tax liabilities compared to a buy-and-hold strategy; (vi) VA remains an inferior strategy even when we consider possible wider behavioral finance benefits.

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 hindsight bias. Furthermore, other properties of VA are likely to significantly reduce investor welfare. In short, VA has little to recommend it.

References

Constantinides, G.M. 1979. "A Note On The Suboptimality Of Dollar-Cost Averaging As An Investment Policy." *Journal of Financial and Quantitative Analysis*, vol.14, no. 2 (June): 443-450.

Chen, H. and Estes, J. 2010 "A Monte Carlo study of the strategies for 401(k) plans: dollar-cost-averaging, value-averaging, and proportional rebalancing." *Financial Services Review*, Vol. 19, 95-109

De Bondt, W.F.M. and Thaler, R, 1985. "Does the Stock Market Overreact?" *Journal of Finance*, vol. 40, no. 3, 793-805.

Dichev, I.D. and Yu, G. 2009 "Higher Risk, Lower Returns: What Hedge Fund Investors Really Earn." (July 2009). Available at SSRN: http://ssrn.com/abstract=1436759.

Dybvig, P.H. 1988a "Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market" *The Review of Financial Studies* 1988, vol.1, no.1, pp. 67-88.

Dybvig, P.H. 1988b "Distributional Analysis of Portfolio Choice" *Journal of Businesss* 1988, vol.61, no.3, pp. 369-393.

Edleson, M.E. 1988. "Value Averaging: A New Approach To Accumulation." *American Association of Individual Investors Journal* vol. X, no. 7 (August 1988)

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

Hayley, S. 2010 "Dollar Cost Averaging – The Role Of Cognitive Bias", Cass Business School working paper. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1473046

Ingersoll, J; Spiegel, M; Goetzmann, W and Welch, I. (2007) "Portfolio Performance Manipulation and Manipulation-proof Performance Measures," *Review of Financial Studies* 20-5, September 2007, 1503-1546.

Leggio, K. and Lien, D. "Comparing Alternative Investment Strategies Using Risk-Adjusted Performance Measures." *Journal of Financial Planning*, Vol. 16, No. 1 (January 2003), pp. 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*, vol. 13, no. 1 (Spring) 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*, Sept 2006.

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

Rozeff, M.S. 1994. “Lump-sum Investing Versus Dollar-Averaging.” *Journal of Portfolio Management*, vol. 20, issue 2: 45-50.

Statman, M. 1995. “A behavioral framework for dollar-cost averaging.” *Journal of Portfolio Management*, vol. 22, issue 2: 70-78.

Thorley, S.R. 1994. “The Fallacy of Dollar Cost Averaging.” *Financial Practice and Education*, vol. 4, no. 2 (Fall/Winter): 138-143.

APPENDIX

A Continuous Time Analysis Of The Inefficiency Of A Value Averaging Strategy

This appendix uses the payoff distribution pricing model of Dybvig (1988a) to derive the continuous time efficiency losses shown in Table 4[6]. We assume an equity index (with zero dividends) and which, relative to a constant interest rate bank account as numeraire, grows according to Geometric Brownian Motion as:

$$\frac{dS_t}{S_t} = \mu\, dt + \sigma\, dB_t \tag{1}$$

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 invested in the equity market in each period from its initial $V_t = V_0\, g(t)$, or specifically in this case a value averaging strategy with target portfolio growth of α per period: $V_t = V_0\, e^{\alpha t}$. 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:

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

This implies that the distribution of terminal wealth at any later time T is normal with mean and variance given by:

[6] I am very grateful to Stewart Hodges for this derivation.

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

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

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 just 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 will 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:

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

Thus the VA strategy 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 I. 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 II. Illustrative Comparison Of VA and DCA – Rising Prices

Strategies are as defined in Table I. The price of the asset is here assumed to rise over the three periods.

		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 III. 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. Prices start at $10 and then evolve in each of 10,000 simulations for five periods with *niid* price movements and a 10% standard deviation. DCA invests a fixed dollar amount ($400), VA invests whatever amount is required to increase the portfolio value by $400 each period, the lump sum strategy invests all $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.

	VA-Lump Sum	VA-DCA	DCA - Lump sum
Average cost (cents)	-20.16	-12.02	-8.15
	(1.08)	(0.12)	(1.10)
IRR (%)	0.278	0.201	0.076
	(0.023)	(0.003)	(0.023)
Profit ($)	0.086	-0.047	0.133
	(2.248)	(0.181)	(2.219)

Table IV. Assessing the Dynamic Efficiency of Value Averaging

This table shows the additional initial capital used 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 continuous time losses is 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: Total Number Of Shares Purchased By Each Strategy

The chart compares the total number of shares purchased across all periods by (a) an immediate lump sum investment, (b) DCA, (c) VA for the same simulated price paths. Prices start at $10 and evolve for five periods. As before, price movements are *niid* with zero mean and 10% standard deviation. The lump sum investment buys a fixed number of shares, but both DCA and VA buy more shares after a price fall and fewer after a price rise.

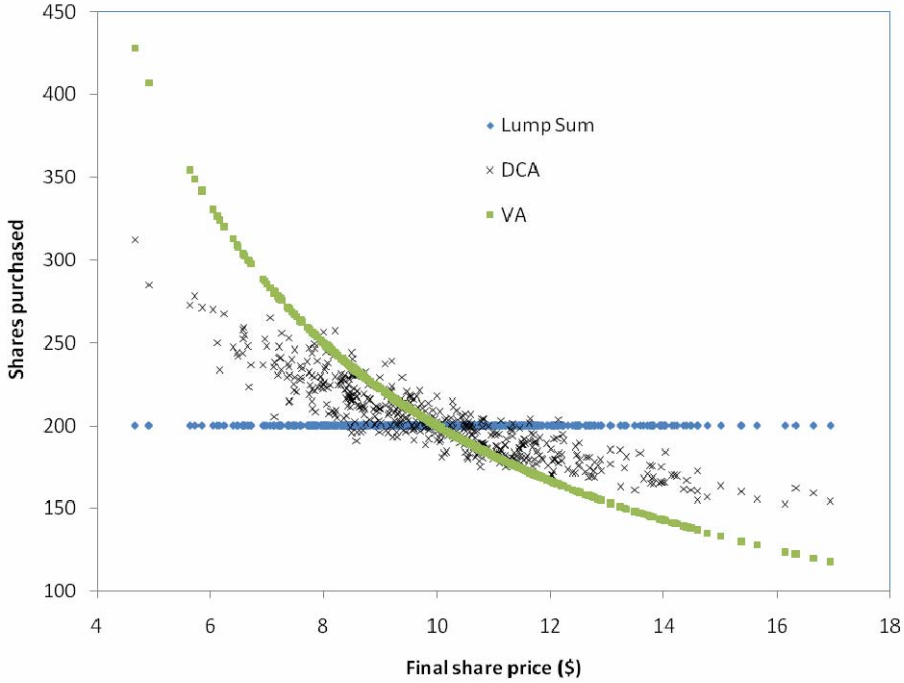

Figure 2. Simple Model of VA Strategy

This figure shows the total investor wealth at each point in a VA strategy with a portfolio growth target of 40% each period. 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 IV shows that inefficiencies persist with a higher risk free rate). The component of total wealth which is held in equities is identical for all paths and is noted at the bottom. All paths are assumed to be equally likely.

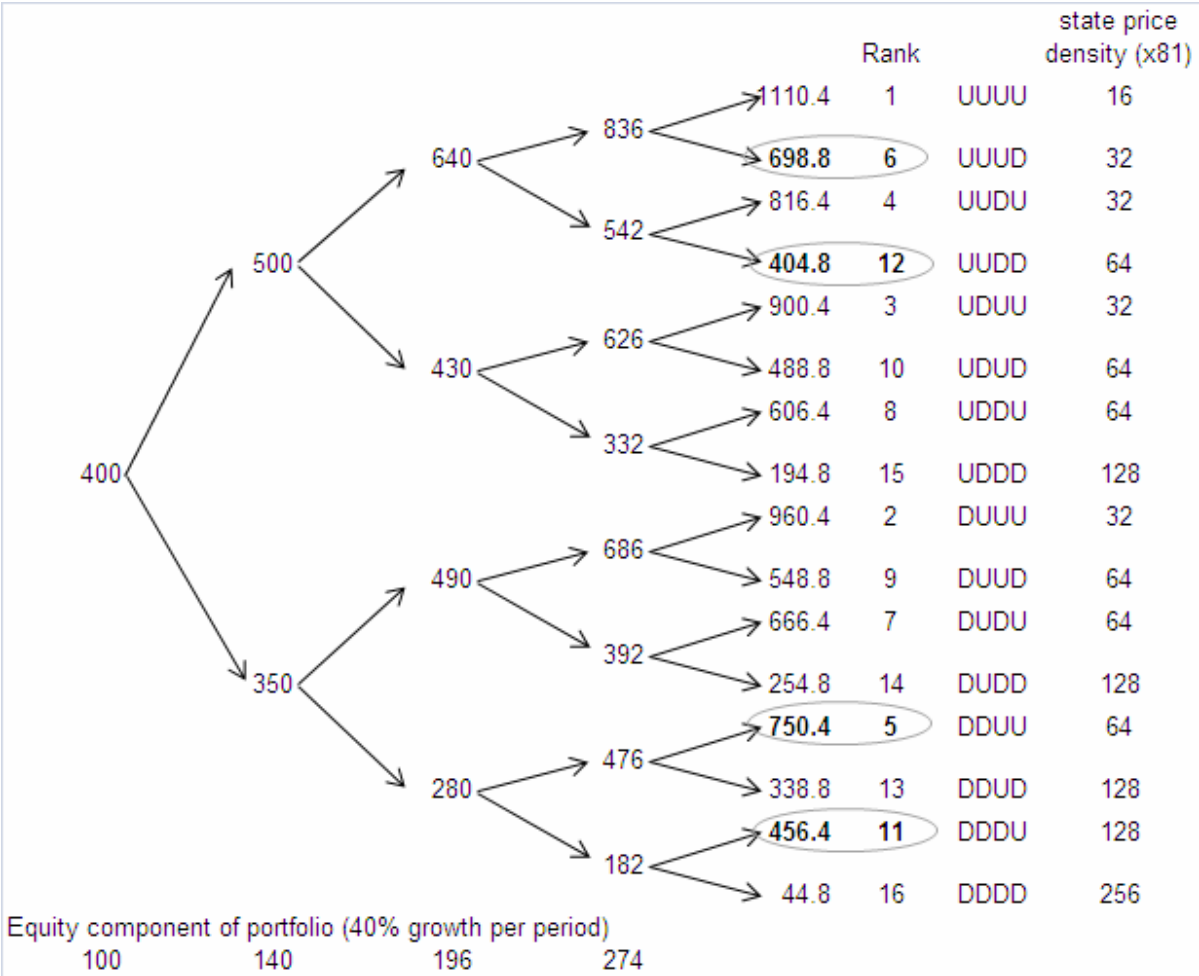

Figure 3. Optimized Strategy Giving Identical Outturns To VA Strategy

This figure shows the total investor wealth at each point in a strategy in which the equity exposure at each node has been set so as to replicate the outturns in Figure 2, but with these outturns optimized so that the largest outturns always come in the states with the lowest state price density. Compared with Figure 2, the outturns for UUUD and DDUU have been swapped, and the outturns for UUDD and DDDU. Equity returns are as assumed in Figure 2. The lower initial capital required for this optimized strategy shows the degree to which the VA strategy is inefficient. The amount of total investor wealth which is held as equity is shown immediately below the total wealth for each node.

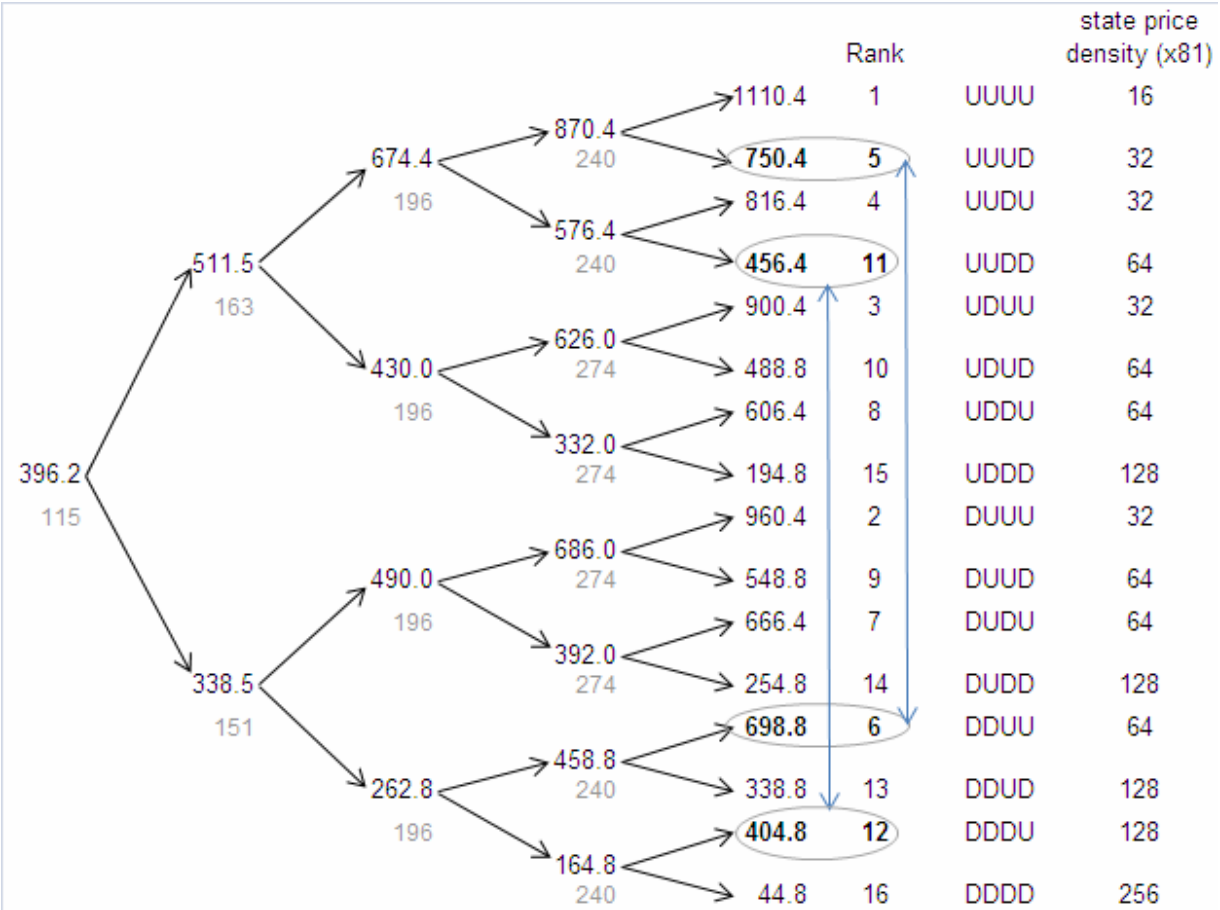

Figure 4: Terminal Wealth Achieved by VA vs. State Price

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.

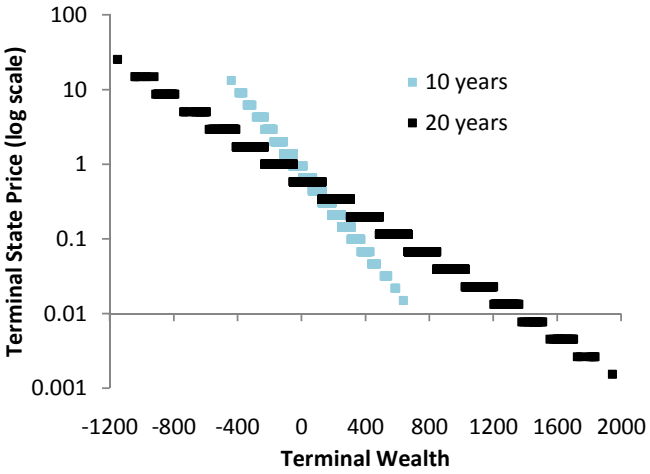

Figure 5: Dynamic Efficiency Losses Of VA Strategy

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 continuous time efficiency losses are derived in the Appendix.

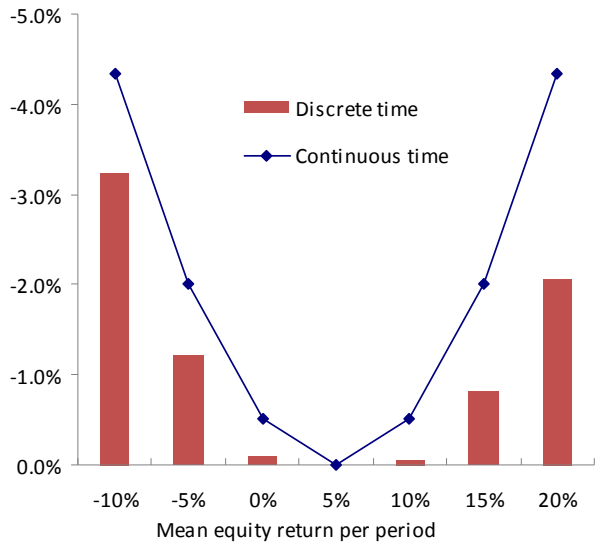

Figure 6: Profits of VA relative to DCA

This chart shows the differential between the simulated profits achieved by VA and DCA, as a function of the terminal share price of that simulated path. Investment is as set out in section 2, with asset price movements assumed *niid* with zero mean and 10% standard deviation per period.