

The risk and return of investment averaging: An option-theoretic approach

Robert Dobil*

Department of Finance, College of Business, San Jose State University, San Jose, CA 95192-0066, USA

Abstract

Automatic stock investment plans are touted to offer dollar cost averaging (DCA) advantages. Rather than focusing on potential lower costs and thus enhanced returns, we analyze the impact of the averaging on the *risk* of the investment. We use an option theory-based simulation model to compute the standard deviation of the realized return, the probability of shortfall and the conditional expected shortfall for a periodic DCA plan. The plan's terminal value depends on the total number of shares acquired over time. The number of shares is a stochastic variable that depends on the volatility of the underlying stock. We show that the risk reduction because of averaging is significant not only in terms of standard deviation, but also in terms of the expected shortfall of funds when the investment turns a loss. We show that the DCA benefits are greater, the longer the averaging, and the riskier the underlying investment.

JEL classification: D14; G12; G13

Keywords: Dollar cost averaging; Asian options pricing; Automatic investment plan; Shortfall

1. Introduction

Investment advisers and individual investors are familiar with the notion of dollar cost averaging (DCA). A practical tool to take advantage of the averaging is an automatic investment plan. By committing to investing a constant dollar amount, as opposed to a constant number of shares, a stock investor buys more shares when the price drops and fewer

* Corresponding author. Tel.: +1-408-924-3477; fax: +1-408-924-3463.
E-mail address: dobil_r@cob.sjsu.edu (R. Dobil).

shares when the price goes up. If the share price fluctuates a lot, this simple scheme can lead to a lower cost per share of the investment. But lower than what?

The return sequence matters crucially. If the stock always goes up, then spreading investment over time simply leads to missing the opportunity to buy the stock all upfront. If the stock always goes down, then spreading does lower the cost, but is still inferior to waiting and buying it all at the end. Of course, *ex ante* we do not know if the stock will go up or down. We may only have estimates of the *expected* return and its volatility and view the issue probabilistically. If the stock is expected to go up (positive mean), buying upfront makes sense; if the stock is expected to go down (negative mean), buying at the end makes sense. But *ex post*, neither may be the best strategy, if at some the random stock dipped well below the original or final level (volatility was high relative to the mean) and we missed the opportunity to buy lots of shares at a very low price.

The apparent advantages of DCA must then be relative to a strategy in which the investor spends the amount of money equal to the sum of all of his partial investments to buy the stock at an average price over the investment horizon.¹ While the strategy to buy at an average price may be a fair benchmark, it is important to realize that it is not executable. The stock never trades at an ‘average’ price, and even if it did, *ex ante* we would have no way of knowing when that is going to be.

The advantages of DCA can be more clearly seen through the prism of behavioral finance. In Kahneman and Tversky’s (1979) prospect theory, investors are not rational in the traditional sense of von Neumann and Morgenstern (1947). They suffer from *loss aversion*. The origin of their utility function is strictly zero. Gains have positive utility, while losses have negative utility, the latter being greater in absolute value for the same absolute level of return. Investors’ return benchmarks are governed by the past. An increase from previous period wealth is good, a decrease is bad. At least, the benchmark, equal to the past return, is known at the time of comparison. This is not so in regret theory, independently proposed by Bell (1982) and Loomes and Sugden (1982), according to which economic agents making decisions under uncertainty suffer *regret* if their decisions (e.g., invest in a U.S. stock) turn out to be inferior even relative to alternatives not considered *ex ante* (e.g., real estate in Bermuda). Investors observe not only the performance of their own portfolios, but also the performance of other stocks that they were able to invest upfront, but rationally decided not to. They feel disappointment if their picks underperformed relative to those alternatives. Investors’ performance benchmarks are not fixed *ex ante*, but rather depend on future states of the world.

We have a very similar situation with DCA. Investors want to minimize the per share cost of purchasing risky assets. In trivial cases—high positive (negative) mean combined with very low standard deviation—both rational and regret-driven investors choose to buy all upfront (at the end) to maximize their expected wealth. In non-trivial high volatility (risky) cases, regret-driven investors may deviate from rationally optimal strategies in favor of DCA to minimize their regret relative to a strategy that is definable *ex post* (“I should have bought on that day at that price. . .”). Note that in view of pure regret theory, the strategy to buy at an average price was not even in the choice set, that is, it was not executable.

The objective of this paper is focus on the risk of averaging. Because the underlying asset (stock or mutual fund) is subject to random fluctuations, the number of shares acquired

through DCA and thus the terminal value of our investment at a future date will be a complicated random function not only of the terminal price per share, but the entire price path of the asset. We will exploit option pricing theory as applied to path-dependent “Asian” options to develop simulation tools for the terminal values of DCA strategies. While we do not attempt to come up with the “best” return benchmark for DCA, we have to define the risk of an investment strategy. We look at traditional measures of total risk of an undiversified investment (standard deviation), as well as some non-traditional metrics (probability of shortfall and conditional expected shortfall) that are arguably of more concern to long-term individual investors. The latter metrics, defined *ex ante* in probability and dollar terms, reflect the risks that an investor may not be able to achieve his or her retirement goals. They are related to confidence interval statistics and thus take jointly into account the first and second moments of the distribution. Although they can be motivated by regret arguments, they simply measure the severity of the failure to achieve financial planning objectives (ending up with less than a risk-free investment). None of our results depend on utility theory assumptions.

What we find is that DCA not only reduces cost, subject to the above caveats, but it also offers significant risk reduction relative to non-averaged alternatives. The risk reduction is not only in terms of the standard deviation of the terminal value of the investment, but also in terms of the reduced probability of not achieving a terminal value target and the reduced expected shortfall conditional on falling short. In the most extreme cases of long averaging periods relative to the total investment horizon (five-year averaging with a five-year horizon), we show DCA reductions of 40% for standard deviations and 30% for expected conditional shortfall. Personal investors do care by how much (in dollars) they are likely to miss their retirement targets.

The paper is organized as follows. Section 2 provides definitions of DCA and upfront strategies, as well as risk metrics used to evaluate them. Section 3 exploits average (Asian) option pricing models to develop intuitions for DCA plans. Section 4 reviews the related Monte Carlo simulation techniques for path-dependent periodic investment plans. Section 5 summarizes our main numerical results for different mean stock return rates and their volatilities, and over different time horizons. Section 6 concludes.

2. Definitions

Let us assume that, at the beginning of each of the next N years, we spend a constant dollar amount PMT to buy shares of a stock (or mutual fund, or index basket). By the end of the buying program, we hold the number of shares equal to:

$$\left(\frac{PMT}{S_0} + \frac{PMT}{\tilde{S}_1} + \dots + \frac{PMT}{\tilde{S}_{N-1}} \right) \quad (1)$$

where \tilde{S}_t is the price of the stock t years from today. Once we have acquired the shares we hold them till some future time $T \geq N - 1$ years from today. The value of our DCA investment at that time is:

$$\tilde{V}_T = \left(\frac{PMT}{S_0} + \frac{PMT}{\tilde{S}_1} + \dots + \frac{PMT}{\tilde{S}_{N-1}} \right) \tilde{S}_T \quad (2)$$

Ex ante, all future stock prices, $\tilde{S}_1, \dots, \tilde{S}_{N-1}, \tilde{S}_T$, are unknown.

Let us consider a few simple strategies to which we may want to compare the DCA plan (Eq. 2). An upfront strategy would require us to spend the dollar amount of $N \times PMT$ once to buy stock at today's price S_0 . The terminal value of such a strategy would be equal to

$$\tilde{V}_T = \left(\frac{N \times PMT}{S_0} \right) \tilde{S}_T. \text{ Alternatively, we may choose to wait till time } N - 1 \text{ to buy all the}$$

shares, in which case the terminal value would turn out to be $\tilde{V}_T = \left(\frac{N \times PMT}{\tilde{S}_{N-1}} \right) \tilde{S}_T$. Both of these, the upfront strategy and the final time strategy, are executable: they have pre-determined times of purchase and amounts to spend. They do not require the investor to know in advance what will happen to the stock.

The strategy to buy at an average price over the investment plan period is one which asks the investor to know precisely that. The investor spends the dollar amount of $N \times PMT$ once (or in chunks) to buy stock at an average price of:

$$\frac{1}{N} \sum_{t=0}^{N-1} \tilde{S}_t \quad (3)$$

To execute this plan, the investor would have to know the average price of the stock over a future period and, ideally, the time the stock will trade at that price. Given that the investor does not know the ex post average (Eq. 3),² he cannot execute the plan.

All three benchmarks considered so far ignore the time value of money. Relative to DCA, an upfront buyer has to borrow money today to repay it in equal installments of PMT . Upfront he has available to invest an amount less than $N \times PMT$. Analogously, a final-time buyer is able to reinvest unused balances until time $N - 1$ giving him an amount to invest greater than $N \times PMT$. The strategy to buy at an average price (Eq. 3) has the same borrowing/reinvestment problem, but we do not know the purchase time.

To account for the time value of money, we amend our two executable benchmarks. In the upfront strategy, we invest the total present value of the amounts PMT :

$$PV_0 = \sum_{t=0}^{N-1} PV_t(PMT) \quad (4)$$

to buy the stock today at price S_0 . The terminal liquidation amount at time T is equal to:

$$\tilde{V}_T = \frac{PV_0}{S_0} \times \tilde{S}_T \quad (5)$$

In the final time strategy, we have the future value of amounts PMT at time $N - 1$ dividing the price \tilde{S}_{N-1} at time $N - 1$.

What discount rate should we use in the present value calculations? At first, we might be

tempted to choose the mean return on the stock following standard corporate finance theory, which dictates that the discount rate should be commensurate with the risk of the asset. This would require us to know $E[\bar{S}_t]$, rendering even the non-averaged strategies non-executable. The correct way is to resort to cash-and-carry arguments of futures pricing. The fair value of a financial futures contract is defined as the spot value of the underlying asset grossed up by the cost of financing the purchase of the asset spot until the maturity of the futures, net of any income earned (coupon or dividend). The financing rate used must be that charged on a loan collateralized by the asset (repo). Futures prices are not expected future spot prices, but today's spot prices adjusted for the cost of carry.³ For long-term individual investors, the financing cost is equal to the opportunity cost on investment alternatives that are (like the repo) close to risk-free, for example, Treasury bonds or mortgages on their own homes. This actually makes perfect sense if we think that investors engage in DCA strategies because they do not have extra savings (they plan to use future earnings to invest) or do not want to allocate them to the risky stocks (invest in near-cash balances or finance home purchases).

In the remainder of the paper, we abstract from the opportunity cost issues by focusing on the *excess return* of the stock over the risk-free rate. Without loss of generality, the latter is assumed to be zero.

Let us turn to the definition of the risk of an investment plan. The underlying asset is a risky stock with a price subject to random fluctuations. The number of shares acquired and the terminal value of a purchase plan are stochastic variables and functions of the path of the stock price. In continuous-time financial economics, it is common to define risk of an asset in terms the standard deviation of its (continuously compounded) returns, denoted by σ . It is also common to posit the log-normality of the asset price. We do so for convenience and clarity of the examples; we do not need it for the analysis.

The main focus of our analysis is the risk of our investment alternatives. First, we define as a measure of total risk the *standard deviation of the terminal dollar amount*, σ_{VT} . For the averaged strategy, this is the standard deviation of the total investment value (Eq. 2) at liquidation. For an up-front strategy, this is⁴ the standard deviation of the terminal amount (Eq. 5).

Next we look at two additional measures of risk of great relevance to individual investors. Bodie (1995) challenges the standard dogma that holding stocks over the long term is less risky than over a short term (because one may be able to ride out the ups and downs of the market). He notes that while the probability of shortfall, defined as the probability that the terminal value of the investment will be less than a known terminal value of a risk-free investment, declines with time, the size of the potential shortfall actually increases. This is reflected in the cost of insuring against the shortfall that increases with time. He defines that cost as the premium of the puts with strikes grossed up by the risk-free rate. This actually leads him to a conclusion that to avoid future retirement money shortfall, investors should only choose risk-free assets (preferably inflation-indexed) and shy away from stocks. Bodie (1996) argues similarly that defined-benefit pension plans should only engage in bond strategies, and, if they do not, any government-provided benefit guarantees be priced using (high) market prices of puts.

Where Bodie comes from is the often-missed observation that even a long-term investment in a risky asset can be viewed as a one-time gamble without the benefit of re-sampling.

The entire stock market history is a sample of *one* observation out of the random drawing of a stochastic process that drives the market. Given that an investor will have to live with the realization upon his retirement of his sample of one, he should focus on a different definition of risk (i.e., the probability that he may not be able to retire as planned because his sample turned to be a “bad” one). This argument is not new. Financial institutions and market regulators have long turned to confidence interval-based measures of risk as a way of insuring that capital held against market and credit risks does not fall below a prescribed level. Cumulative shortfall measures are used in mortgage servicing hedge arrangements. Published Value-at-Risk (VaR) statistics measure the amount of loss given the stated confidence level (95%). Unpublished conditional expected loss statistics estimate the expected loss once it has occurred. Where Bodie strays in our opinion is in his put value arguments and in putting the shortfall or pension guarantee issue at knife’s edge. As Wilkie (2001) elegantly demonstrates, it is only the 100% guarantee that requires risk-free investing; several non-perfect ones (absolute, proportional, or contingent) can be shown to be self-financing investments in stocks and put options. Also, the one-sample view of the universe is quite extreme. It is true that stock investors cannot invest in alternative universes of stocks following the same stochastic process to ensure an outcome.⁵ They can play with alternative stocks within the same universe to make it likely. The standard deviation of the returns can be used as a true indicator of risk if distributions are assumed stationary and time-independent. In fact, statistical shortfall measures rely on its estimation.

We adopt two metrics of risk that follow the spirit of Bodie’s arguments and that have been adopted to be used by financial institution; for best reference, see Longin (2001). The first is the probability of shortfall, which is the probability that the terminal value of the investment falls below a pre-specified level. This is set to the future value of the investment compounded by the risk-free rate. With a zero risk-free rate, it is equal to the upfront investment. The probability of shortfall p is defined as:

$$p = Pr\{\tilde{V}_T < PV_0\} \quad (6)$$

The conditional expected shortfall is the expected value of the shortfall conditional on the shortfall occurring. It measures the average amount by which an investor misses his retirement target, if he falls short. The conditional expected shortfall λ is defined as:

$$\lambda = E[PV_0 - \tilde{V}_T | \tilde{V}_T < PV_0] \quad (7)$$

In all of our computations, we use quantities that are continuously compounded. This includes excess returns, means, and standard deviations. Next we turn to the lessons of averaging of option payoffs.

3. Asian option pricing models—intuitions for automatic investment plans

The main observation of the standard option pricing theory is that a call or put option, even an out-of-the-money one with no intrinsic value, always has a positive value equal to the cost of the payoff-replicating strategy of a delta-hedger. There are three ways of valuing options in practice: (1) a closed-form solution based on the Black-Scholes hedge argument, limited

to European cases, (2) a binomial or trinomial tree that approximates the closed form through a recursive numerical induction, but can be extended to American cases and barrier conditions, or (3) a Monte Carlo simulation that approximates the risk-neutral expected value of the payoff based on a large number of simulated paths, and can handle path-dependent options, but cannot handle American exercise cases. For the latter two to be deemed numerically accurate, they must produce the same values as the first one for all options for which closed formulae are derivable. Theoretically, all three methods satisfy the same partial differential equation. Our analysis will use (3) Monte Carlo simulation, as we will be naturally dealing with a path-dependent quantity purchased.

To gain intuitions about the effect of averaging in an investment plan, we examine average options, that is, options whose payoff is a function of an average of the stock price over time. They are called Asian options and take on two forms: options whose underlying asset is an average, but the strike is not

$$\max[S_{Ave} - K, 0] \text{ or } \max[K - S_{Ave}, 0] \tag{8}$$

and options whose strike is an average, but the underlying asset is not

$$\max[S_T - S_{Ave}, 0] \text{ or } \max[S_{Ave} - S_T, 0] \tag{9}$$

It is the latter form that is of interest to us.

In most cases, see Hull (2000), Asian options are cheaper than standard European stock options, because the volatility of the average is always lower than the volatility of the underlying asset. Kemna and Vorst (1991) derive analytic formulae for geometric averages. Closed-form solutions cannot be obtained for options on arithmetic averages of log-normal variables, but, as Turnbull and Wakeman (1991) show, they can be approximated using a limiting argument. We rely on a less accurate, but more elegant and intuitive argument of Dubil and Dachille (1989). Any future stock price at time t can be written as the sum of today's price, S_0 , and non-overlapping contiguous daily increments $\Delta\tilde{S}_t$:

$$\tilde{S}_t = S_0 + \Delta\tilde{S}_1 + \dots + \Delta\tilde{S}_t \tag{10}$$

These increments are assumed to be identically independently normally distributed with a daily mean and standard deviation defined in dollars.⁶ An arithmetic average of N stock prices starting $T - N - 1$ days from today and ending T days from today can then be defined as:

$$S_{Ave} = \frac{1}{N} \sum_{t=T-N-1}^T \tilde{S}_t$$

$$= \frac{1}{N} \left\{ \begin{array}{cccc} S_0 & + \Delta\tilde{S}_1 & + \dots + \Delta\tilde{S}_{T-N-1} & \\ + S_0 & + \Delta\tilde{S}_1 & + \dots + \Delta\tilde{S}_{T-N-1} & + \Delta\tilde{S}_{T-N} \\ + \dots + & & & \\ + S_0 & + \Delta\tilde{S}_1 & + \dots + \Delta\tilde{S}_{T-N-1} & + \Delta\tilde{S}_{T-N} + \dots + \Delta\tilde{S}_T \end{array} \right\} \tag{11}$$

For example, if an option has 10 days left to maturity and the strike is equal to the average of the last five days then the average is equal to $S_{Ave} = 15\{S_6 + S_7 + S_8 + S_9 + S_{10}\}$. Each component of the average contains the daily price increments of the previous one plus one. Ex ante, the average is a random variable that is a weighted sum of the daily stock price increments between time 0 and time T . The variance of the average can be derived as the variance of the sum of independent normal variables, each having a variance of $1/365$ of the annual variance of the stock price.⁷ Dubil and Dachille (1989) show that the (annualized) volatility of the average, v , is related to the (annualized) volatility of the stock price, σ , through the following formula, which is the result of summing the number of increments squared:

$$v^2 = \left[\frac{T - N}{T} + \frac{(2N + 1)(N + 1)}{6NT} \right] \sigma^2 \quad (12)$$

The amount of money needed today to replicate the payoff of the average at maturity T is equal to:

$$Ave_0 = S_0 \times \frac{1}{N} \sum_{t=0}^{N-1} e^{-\frac{rt}{365}} \quad (13)$$

This formula reflects the strategy of splitting the money into N parts, depositing each part in an account bearing interest r until the time of each of the investment and then buying $1/N$ th of the stock. Once the two variables in Eqs. (12) and (13) are derived, valuing average options of type (Eq. 8) boils down to substituting (Eq. 13) for the stock price and the square root of (Eq. 12) for the volatility into the Black-Scholes (1973) formula. Valuing average options of type (Eq. 9) boils down to substituting (Eq. 13) and (Eq. 12) into a formula for an option to exchange one asset (stock) for another (average) modified from Black-Scholes by Margrabe (1978).

Having established the mathematics, let us turn to intuitions. An option on the average of stock prices is less valuable than an equivalent option on a stock price observed once. This is because any fluctuation in the price will be dampened by the averaging, thus lowering the volatility of the stochastic variable underlying the payoff. A 10-day option of type (Eq. 8) on the average of the last five days with a fixed strike will be less valuable than the same option on the price on Day 10. The longer the averaging relative to the option period, the greater is the reduction in the option value because of the volatility dampening. In the extreme example of an old option that has already entered the averaging period, the volatility will be close to zero as the averaging will include the already revealed prices of the last few days. On Day 9, our 10-day option on a five-day average will have four of the prices already known and only one to be revealed. Even if the stock is highly volatile, the final average will be unlikely to deviate much from the average so far. The same principle operates before the start of the averaging.

We can use Eq. (12) to construct a table of the ratio of the volatility of the average divided by the volatility of the stock for different averaging periods and different times to expirations, by taking the square root of the variance multiplier. The result is shown in Table 1.

Table 1

The ratio of the volatility of the average to the volatility of the underlying stock for different averaging periods and times to maturity

Averaging days	Time to maturity						
	10	30	90	180	360	720	1800
5	0.849	0.952	0.984	0.992	0.996	0.998	0.999
10	0.620	0.892	0.965	0.983	0.991	0.996	0.998
30		0.592	0.885	0.944	0.973	0.986	0.995
90			0.582	0.818	0.914	0.958	0.983
180				0.580	0.817	0.913	0.966
360					0.579	0.817	0.931
720						0.578	0.857
1800							0.578

The table shows that 360 days of averaging in a 720 day option results in a 28.3% reduction ($1 - 0.817$) in the volatility of the underlying, relative to a standard option. It is worth noting that the option value change will depend on the in-the-moneyness of the compared options. In most cases, the reduction in the value of the at-the-money calls will actually be greater than the volatility reduction.

Let us turn to the more instructive case of Asian call options of type (Eq. 9). Their payoff greatly resembles that of an automatic investment plan. We acquire the asset over time incurring the cost of S_{Ave} and liquidate it at some point (usually much) after the averaging has stopped. Our total gain is equal to

$$S_T - S_{Ave} \quad (14)$$

We can think of the two variables in Eq. (14) as two different assets. At liquidation, we exchange the average asset for the value of the stock at that time. We can use the same approach as before to come up with a formula for the volatility, v , of the *difference* of the stock at liquidation and the average as a function of the volatility, σ , of the stock itself. The two are related through the following equation:

$$v^2 = \left[\frac{(2N - 1)(N - 1)}{6NT} \right] \sigma^2 \quad (15)$$

Once again, we can construct a table of the volatility ratio, this time of the difference between the average and the stock to the stock. The result is shown in Table 2.

In this case, the longer the averaging period, the higher the ratio, reflecting the fact the average will more likely be different from the stock at expiry. The averaging over one day would make the two variables equal, making the ratio and the value of the option equal to zero. The highest volatility is obtained by averaging over the full period. It is always below 60%. Note that the case of the Asian options is different from the most common case of an automatic investment plan in that the averaging in options occurs at the end of the expiry period for N days. In the investment plan, it is more likely that an investor saves by buying portions of the stock first, then holds it for some time, and then sells it all at once (or also

Table 2

The ratio of the volatility of the difference between the average and the stock to the volatility of the underlying stock for different averaging periods and times to maturity

Averaging days	Time to maturity						
	10	30	90	180	360	720	1800
5	0.346	0.200	0.115	0.082	0.058	0.041	0.026
10	0.534	0.308	0.178	0.126	0.089	0.063	0.040
30		0.563	0.325	0.230	0.162	0.115	0.073
90			0.573	0.405	0.286	0.202	0.128
180				0.575	0.407	0.287	0.182
360					0.576	0.407	0.258
720						0.577	0.365
1800							0.577

gradually). This would imply the averaging at the beginning rather than at the end of the investment period.

Recently, the ideas of compound and average options have been applied to venture capital financing in a way related to how we analyze personal investments. Wang and Zhou (2004), Jones and Rhodes-Kropf (2003), and Chemmanur and Chen (2003) rely on moral hazard arguments to show that multistage financing is an optimal strategy for venture capital firms facing private information and incentives problems associated with original owners. Hsu (2002) uses Geske's (1979) compound option approach to value the options inherent in releasing capital in stages. The venture capitalist's strategy consists of an equity stake and a call on a call to acquire more equity. Thus it contains a real option to discontinue funding. But even without considering the multiple discontinuance option, Dubil (2004) shows that the venture capitalist reduces the risk of his investments simply by doling out capital in stages. He gets automatic DCA benefits, since for the same dollar investment he acquires more shares in the early financing stages, and fewer later when the business is more valuable. The individual DCA investor does not actively seek discontinuance options to eliminate any moral hazard, but obtains risk reduction passively because of the automatic nature of the investment. The big difference is that the venture capitalist tends to be much less risk averse than the personal investor and seeks cost reduction because of averaging and risk reduction due to information release, while the personal investor may seek both cost reduction and risk reduction. As we will show, often the risk reduction may be greater than cost savings.

4. Monte Carlo simulation of periodic stock purchases

Cox and Ross (1976) showed that a call option value is equal to its expected payoff discounted by the interest rate where the expectation is taken with respect to a special probability measure. Given this observation, an Asian call option of type (Eq. 8) or (Eq. 9) can be priced by evaluating the following expressions:

$$e^{-rT}E[\max(S_{Ave} - K, 0)] \text{ or } e^{-rT}E[\max(S_T - S_{Ave}, 0)] \quad (16)$$

If we have a numerical method of evaluating expected values, that is, probability-weighted averages of the payoff outcomes, then all we need to do is to discount them to today. In practice, the expectation evaluation is most easily performed with the use of a Monte Carlo simulation. The special risk-neutral probability measure requires that for a log-normally distributed traded asset, with the volatility σ , subsequent asset values are generated through the following recursive formula:

$$S_n = S_{n-1} e^{\left(r - \frac{1}{2}\sigma^2\right)\Delta t_n + \sigma \sqrt{\Delta t_n} \tilde{z}_n} \quad (17)$$

where Δt_n is the time (in years) between the consecutive stock price observations and \tilde{z}_n is a standard normal deviate generated using a random number generator. To generate the first unknown stock price tomorrow, \tilde{S}_1 , we take today's known price S_0 and one standard normal random number \tilde{z}_1 , and plug them into Eq. (17) to get $\tilde{S}_1 = S_0 e^{\left(r - \frac{1}{2}\sigma^2\right)\Delta t_{1day} + \sigma \sqrt{\Delta t_{1day}} \tilde{z}_1}$, where Δt_{1day} is equal to 1/365 or 1/250, whichever we feel appropriate. The Day 2 price is obtained using the Day 1 price and another generated standard normal random number, \tilde{z}_2 , plugged into Eq. (17) to get $\tilde{S}_2 = \tilde{S}_1 e^{\left(r - \frac{1}{2}\sigma^2\right)\Delta t_{1day} + \sigma \sqrt{\Delta t_{1day}} \tilde{z}_2}$. We generate prices for all days between now and the maturity of the option. We then evaluate the payoff of the option. If the option is a European call or put, then all we need is the last price \tilde{S}_T to compare it to the strike. In that case, we dispense with all the intermediate prices and use only one step to generate the final price $\tilde{S}_T = S_0 e^{\left(r - \frac{1}{2}\sigma^2\right)T + \sigma \sqrt{T} \tilde{z}}$, with T defined as a fraction of a year. If the option is Asian, then we do need the stock prices for all the days that are to be included in the average to determine the realized average $Ave_T = \frac{1}{N} \sum \tilde{S}_i$ and the final payoff. In both cases, we repeat the path and payoff generation many times and compute the expected payoff value as a simple arithmetic average of the payoff outcomes. The number of paths generated depends on the speed of the convergence of the expectation to a stable value. Many techniques ensure fast convergence, for example, antithetical variables (generate a set of random numbers and then reverse their signs to generate another path—the mean is assured to be zero) and conditioning (instead of computing payoffs directly, compute their deviations from a numerical outcome of a known closed-form value).

Let us turn to the evaluation of a simple constant-dollar investment plan. What we will try to determine is the risk of the plan relative to a one-time purchase. First, we want to know the standard deviation of the terminal value of our investment, assuming that we acquire stocks periodically over N purchases and then hold the stock till the final time T . We cannot derive a closed-form solution for the variance of the strategy, but we can compute it using the same Monte Carlo technique we described for average option valuation where, in place of the risk-neutral drift equal to the risk-free rate r , we substitute the mean excess return $\mu - r$. Similarly, we cannot derive the probability of shortfall or the conditional expected shortfall analytically, but we can obtain both numerically from the simulation. We compute the probability of shortfall, by counting the number of paths with the terminal value lower than the initial investment as defined in Eq. (6). We compute the conditional expected shortfall by

averaging the deviation of the terminal value from the upfront investment only over the paths with a shortfall.

We assume that we embark on an N -year DCA plan whereby we purchase stocks once a year *at the beginning of the year* by spending a constant dollar amount PMT . At the end of the purchase program, $N - 1$ years from now, we hold the number of shares given by Eq. (1). We use Eq. (17) to generate the realized stock prices based on today's starting stock value. For instance, the price in Year 6 is generated from the price in Year 5 through $\tilde{S}_6 = \tilde{S}_5 e^{(r - \frac{1}{2}\sigma^2) + \sigma\sqrt{1z_6}}$. Note that variance scaling issues are eliminated as time increments are annual, that is, $\Delta t_n = 1, - n$.

For example, in a 10-year plan with $PMT = \$10,000$, we buy \$10,000 worth of shares up-front in Year 0, \$10,000 worth of shares in Year 1, and so on, with the last purchase being in Year 9. We incur the randomly generated stock prices $S_0, \tilde{S}_1, \dots, \tilde{S}_9$. We reap the benefits of dollar cost averaging over 10 purchases, but, in advance, we do not know the average price of acquiring a share or the total number of shares we hold right after the last purchase in Year 9. In a five-year plan with a total $PV_0 = \$100,000$, we buy $PMT = \$20,000$ worth of shares in Years 0 through 4 incurring one known and four future random prices $\tilde{S}_1, \dots, \tilde{S}_4$.

We hold the acquired shares from the time of the last purchase in $N - 1$ years for $T - N + 1$ years until we liquidate the shares in Year T for the price \tilde{S}_T generated again using the correct time interval in Eq. (17). For example, for a 10-year DCA, if we plan to liquidate the shares in Year 20, then we hold the acquired shares for additional 11 years after the last purchase in Year 9. For a five-year DCA with a final horizon of 15 years, we hold the shares for 11 years after the last purchase in Year 4. We compute the terminal value of the investment plan in Year T using Eq. (2). We repeat the path generation thousands of times until convergence.

5. The reduced risk of automatic investment plans

In our tables, we consider two strategies. The first is an upfront investment of $PV_0 = \$100,000$. The second is a five-year DCA investment of $PMT = \$20,000$ per year. In each case, the opportunity cost of capital is $r = 0\%$ and the investor earns a positive excess continuously compounded return of $\mu - r$. We show two cases for the excess return: 4% and 8%. We run two different annual volatility scenarios $\sigma = 15\%$ and 30% (think of the first one as an index like the S&P or NASDAQ, and the second as an individual stock). We also consider two different final horizons $T = 5$ and 15 years. The first can be thought of as for people nearing retirement, the latter for younger investors who save early through automatic investment plans, hold through the middle age and then withdraw.

The tables show the expected value and the standard deviation of the investment value \tilde{V}_T at the time of withdrawal or liquidation for both the up-front and the periodic DCA investment plan. The tables also show the probability of shortfall and the conditional expected shortfall as defined in Eqs. (6) and (7) as well as the ratios of the standard deviations and conditional expected shortfalls for the two alternatives.

In Table 3, the final horizon is short relative to the averaging period. While the DCA

Table 3
Monte Carlo statistics for five-year final horizons

	$\mu-r = 4\%$		$\mu-r = 8\%$	
	Vol = 15%	Vol = 30%	Vol = 15%	Vol = 30%
DCA strategy				
Mean terminal value	113,594	114,840	127,791	129,217
SD of terminal value	26,827	57,830	30,687	66,264
Shortfall probability	34%	48%	18%	38%
Conditional expected shortfall	13,194	26,967	10,925	24,828
Up-front strategy				
Mean terminal value	122,816	124,332	148,323	150,154
SD of terminal value	42,947	96,824	51,866	116,932
Shortfall probability	33%	51%	16%	40%
Conditional expected shortfall	18,819	37,516	15,454	34,138
SD ratio	0.62	0.60	0.59	0.57
Conditional exp. shortfall ratio	0.70	0.72	0.71	0.73

An up-front investment of \$100,000 in Yr 0 vs. five-time \$20,000 DCA over Yrs 0–4.

Monte Carlo with 10,000 paths. Risk-free rate $r = 0$. Excess return $\mu-r$ equal to the mean stock return μ .

Shortfall threshold equal to the up-front investment.

Probability of shortfall defined as the percentage of paths with a terminal value below the threshold.

Conditional Expected Shortfall defined as the mean shortfall conditional on the occurrence of shortfall.

strategy yields a slightly lower average return (given positive mean risk premium), the risk reduction advantage is irrefutable. The standard deviation of the terminal value is only about 60% of that of the upfront strategy. The dollar values are more telling, especially for the riskier asset cases ($\sigma = 30\%$): the standard deviations for the upfront strategy are over three-fourths of the mean, while they are only one-half of the means for the DCA strategy. Bodie's point is even clearer. While the probability of shortfall is roughly the same, the size of the potential shortfall is not. The DCA strategy's expected shortfall is 30% lower. The intuition behind this result should be that of diversification. By engaging in DCA, a near-retirement investor buys five chunks of stock non-perfectly correlated to each other.⁸ A practical implication here may be that near-retirement investors consider stock ladders similar to bond ladders that match their withdrawal times. Comparing the numbers across the columns, given mean excess return, we can see that the risk ratios corresponding to different σ 's stay about the same. The implication is that the *absolute* reduction in risk is greater when investing in riskier assets (individual stocks rather than indexes).

In Table 4, the final horizon is three times as long as the averaging period. This accounts for the fact that although the results are qualitatively the same, the risk reduction of DCA is less significant. The standard deviation ratios are 69–83% and the shortfall ratios are in the 90%.

The most significant benefits of DCA are obtained in the cases of low excess returns combined with high risk. Given that the probability of shortfall is over 50%, any reduction in risk is important. What is striking in Table 4 is how large the standard deviations in dollars are relative to expected terminal values (often exceeding them) and how much larger the upfront standard deviations are than the DCA ones in high-risk cases. This reflects the fact that the upfront strategy has generally fatter tails on both sides, that is, it has a higher

Table 4
Monte Carlo statistics for 15-year final horizons

	$\mu-r = 4\%$		$\mu-r = 8\%$	
	Vol = 15%	Vol = 30%	Vol = 15%	Vol = 30%
DCA strategy				
Mean terminal value	170,151	178,618	279,181	293,169
SD of terminal value	101,575	292,468	167,275	483,793
Shortfall probability	24%	51%	5%	33%
Conditional expected shortfall	24,560	50,143	18,758	44,481
Up-front strategy				
Mean terminal value	183,930	193,951	323,972	341,624
SD of terminal value	122,976	398,038	216,608	701,101
Shortfall probability	24%	52%	5%	34%
Conditional expected shortfall	26,565	53,773	19,093	46,392
SD ratio	0.83	0.73	0.77	0.69
Conditional exp. shortfall ratio	0.92	0.93	0.98	0.96

An up-front investment of \$100,000 in Yr 0 vs. five-time \$20,000 DCA over Yrs 0–4.

Monte Carlo with 10,000 paths. Risk-free rate $r = 0$. Excess return $\mu-r$ equal to the mean stock return μ .

Shortfall threshold equal to the up-front investment.

Probability of shortfall defined as the percentage of paths with a terminal value below the threshold.

Conditional Expected Shortfall defined as the mean shortfall conditional on the occurrence of shortfall.

frequency of very good and very bad results. A risk-averse investor can view DCA as a dispersion- narrowing technique.

Tables 3 and 4 also illustrate to what extent greater excess returns make the leveraged upfront strategy more attractive, when we compare the expected terminal values across different excess returns. The increases in dollar values are greater for the upfront strategy than for the DCA strategy. By investing up-front rather than over time, the investor jump-starts years' worth of excess returns. At the same time, he increases his risk. Although the numbers are not clear-cut, they are telling that the extra return is not worth the extra risk, given that the volatility of the terminal value can be more than halved by using the periodic investment strategy. The buyers of stocks in the late 1990s who borrowed heavily against their home equities are surely keenly aware of these findings. Instead of waiting to invest through DCA as their disposable income allowed, they leveraged by borrowing to buy upfront. When negative excess return realized, their overall ex post return turned out lower than it would have been, and they had consciously taken on more risk (both in terms of standard deviation and expected shortfall). They lived through a single "bad" historical sample of a stochastic process.

In Table 5, we show additional results under the assumption that the opportunity cost of capital is $r = 6\%$. We consider 10-year DCA strategies with final holding horizons of 10 and 20 years. We show standard deviations only. The results confirm that the cost of capital assumption does not change our findings. DCA strategies result in lower expected terminal values if excess returns are positive, but offer significant risk benefits. The longer averaging here relative to Tables 3 and 4 results in a very significant risk reduction. This is especially so when dealing with highly volatile underlying assets.

The lifetime investment lessons here are not surprising. Investors should start with

Table 5. Expected \$ value, st. deviation and st.deviation ratio for up-front and 10-year DCA investment plans.

Borrow=Invest=6%				Borrow=Invest=6%			
Periodic		Up-front		Periodic		Up-front	
PMT	12,817.73			PMT	12,817.73		
FV(Vol=0)		Ratio		FV(Vol=0)		Ratio	
Withdraw Year 10		Withdraw Year 20		Withdraw Year 20		Withdraw Year 20	
Vol=10	Exp	179,005	178,980	Vol=10	Exp	320,515	320,399
	St. Dev	37,363	57,614		St. Dev	126,088	150,686
			0.649				0.837
Vol=20	Exp	178,728	178,622	Vol=20	Exp	320,333	320,024
	St. Dev	77,666	123,923		St. Dev	285,925	357,711
			0.627				0.799
Vol=30	Exp	178,198	177,964	Vol=30	Exp	320,787	320,672
	St. Dev	124,655	210,094		St. Dev	527,730	717,938
			0.593				0.735
Vol=50	Exp	176,050	175,144	Vol=50	Exp	321,762	324,076
	St. Dev	256,355	498,631		St. Dev	1,498,112	2,480,540
			0.514				0.604

Borrow=6%, Invest=8%				Borrow=6%, Invest=8%			
Periodic		Up-front		Periodic		Up-front	
PMT	12,817.73			PMT	12,817.73		
FV(Vol=0)		Ratio		FV(Vol=0)		Ratio	
Withdraw Year 10		Withdraw Year 20		Withdraw Year 20		Withdraw Year 20	
Vol=10	Exp	200,449	215,766	Vol=10	Exp	432,676	465,639
	St. Dev	42,645	69,455		St. Dev	171,237	218,993
			0.614				0.782
Vol=20	Exp	200,131	215,335	Vol=20	Exp	432,410	465,093
	St. Dev	88,742	149,393		St. Dev	388,835	519,864
			0.594				0.748
Vol=30	Exp	199,525	214,542	Vol=30	Exp	433,002	466,035
	St. Dev	142,696	253,275		St. Dev	719,427	1,043,386
			0.563				0.690
Vol=50	Exp	197,061	211,142	Vol=50	Exp	434,200	470,982
	St. Dev	294,953	601,116		St. Dev	2,054,273	3,604,990
			0.491				0.570

Borrow=6%, Invest=12%				Borrow=6%, Invest=12%			
Periodic		Up-front		Periodic		Up-front	
PMT	12,817.73			PMT	12,817.73		
FV(Vol=0)		Ratio		FV(Vol=0)		Ratio	
Withdraw Year 10		Withdraw Year 20		Withdraw Year 20		Withdraw Year 20	
Vol=10	Exp	251,809	310,403	Vol=10	Exp	781,929	963,685
	St. Dev	55,514	99,919		St. Dev	313,081	453,228
			0.556				0.691
Vol=20	Exp	251,395	309,783	Vol=20	Exp	781,390	962,555
	St. Dev	115,762	214,919		St. Dev	712,818	1,075,908
			0.539				0.663
Vol=30	Exp	250,604	308,641	Vol=30	Exp	782,399	964,503
	St. Dev	186,792	364,363		St. Dev	1,325,147	2,159,386
			0.513				0.614
Vol=50	Exp	247,372	303,751	Vol=50	Exp	784,218	974,742
	St. Dev	389,775	864,771		St. Dev	3,826,920	7,460,868
			0.451				0.513

well-endowed index portfolios (upfront). Over their working years, they may add riskier stocks through automatic investment plans (DCA). The choice of risky stocks for DCA will ensure the maximum risk benefit. Investors should extend their DCA period as much as possible, preferably all the way to the retirement time. This will make the DCA horizon long relative to the total holding horizon.

6. Conclusions

We have exploited the findings of Asian option pricing theory and developed a related simulation methodology to evaluate the risks of periodic DCA investment plans. We find that the averaging offers significant risk benefits over an upfront strategy in terms of traditional total risk statistic (standard deviation) and non-traditional shortfall-centered metrics. The risk benefit is the greatest for long DCA periods relative to the final holding horizon. It is also more significant for riskier underlying assets. We have demonstrated that while DCA may not lower the probability of shortfall, it *always* reduces the expected dollar amount of shortfall upon the liquidation of the investment; in the five-year DCA strategy by as much as 30%. The expected shortfall amount may be what matters most to an investor saving for retirement. Even at a cost of a slightly lower expected terminal value (as is the case with positive excess returns).

The results have several normative implications. Long-term investors should choose low risk assets for long buy-and-hold strategies and invest in them early. They should choose high-risk assets for DCA strategies and try to extend the averaging periods. Only investors with high risk tolerance should leverage by borrowing to invest upfront. The leveraged upfront strategy will not only expose them to the potential *ex post* negative mean (they will have “bought high” upfront instead of “low” later), but also to a higher *ex ante* standard deviation and conditional expected shortfall (even if the market does not trend down, if they end up with a deficit, it will be greater than that for a DCA investor).

We only touched on the analogy of DCA to staged venture capital investing. Yet there are other empirically exploitable applications of DCA. Well-established mutual funds that do not experience a lot investor turnover may use their investment inflows to attain DCA-like risk reduction benefits by adding to their well-diversified portfolio additional chunks of risky assets. In a sense, their averaging period is quite long relative to the holding horizon. Another DCA guise is the strategy of bond laddering that not only matches the investor’s ideal time of withdrawal, but also offers reduction of interest rate risk. We shall leave these topics for another paper.

Notes

1. Constantinides (1979) shows that blind averaging may not be the best approach as it is weakly dominated by sequential strategies where investors use the information about the past path of the stock to make additional purchases.
2. He may only be able to estimate his subjective *ex ante* mean $E\left[\frac{1}{N}\sum_{t=0}^{N-1}\tilde{S}_t\right]$.
3. Only for commodities or so-called convenience assets, that the cost-of-carry link becomes meaningless.
4. With a zero risk-free rate, (Eq. 4) reduces simply to the sum of all the partial amounts $PV_0 = N \times PMT$.
5. In a way, insurance companies do exactly that. They diversify their life policies among individuals who will follow different life paths, but *ex ante* can be viewed as generic.

6. The identically independently distributed assumption is standard for most stochastic processes. Normality (arithmetic Brownian motion) is assumed only to derive a proportionality relationship between the standard deviation of the stock and the average.
7. We use 365 days for clarity of equations. Variance annualization has been widely researched. For example, French and Roll (1986) show that non-trading (weekend) variance is significantly lower than trading variance, so the correct scaling should be closer to 250 trading days Gunthorpe and Levy (1993) show that monthly variance is close to 30% of annual variance. Tables 1 and 2 assume 360 days. Our Eqs. (12) and (15), and the analysis in subsequent sections do not depend on the assumed constant.
8. It can be shown mathematically that the correlation is equal to the square root of the ratios of times.

References

- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81, 637–659.
- Bodie, Z. (1995). On the risk of stocks in the long run. *Financial Analysts Journal*, May–June, 18–22.
- Bodie, Z. (1996). What the Pension Benefit Guarantee Corporation can learn from the Federal Savings and Loan Insurance Corporation. *Journal of Financial Services Research*, 10 (1), 83–100.
- Chemmanur, T. J., & Zhaohui, C. (2003). *Angels, Venture Capitalists and Entrepreneurs: A Dynamic Model of Private equity Investing*. Working Paper, AFA Washington, DC, meetings.
- 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.
- Dubil, R., & Dacheille, D. A. (1989). *The Valuation of Options on Averages*. J.P. Morgan Securities, Inc., Working Paper.
- Dubil, R. (2004). *The Optimality of Multistage Venture Capital Financing: An Option-Theoretic Approach*. Working Paper, Academy of Entrepreneurial Finance Washington, DC, meetings.
- Geske, R. (1979). The Valuation of Compound Options. *Journal of Financial Economics*, 7, 63–81.
- Hull, J. C. (2000). *Options, Futures, and Other Derivatives*, 4th Edition, Englewood Cliffs, NJ: Prentice Hall.
- Hsu, Y. (2002). *Staging of Venture Capital Investment: A Real Options Analysis*. Working Paper, EFMA London meetings.
- Jones, C., & Rhodes-Kropf, M. (2003). *The Price of Diversifiable Risk in Venture Capital and Private Equity*. Working Paper, AFA Washington, DC, meetings.
- Kemna, A., & Vorst, A. (1990). A Pricing Method for Options Based on Average Asset Values. *Journal of Banking and Finance*, 14, 113–29.
- Longin, Francois M. (2001). Beyond the VaR. *Journal of Derivatives*, 8, 36–48.
- Margrabe, W. (1978). The Value of Option to Exchange One Asset for Another. *Journal of Finance*, 33, 177–86.
- Rubinstein, M. (1991). One for Another. *RISK* magazine (July–August), 30–32.
- Turnbull, S. M., & Wakeman, L. M. (1991). A Quick Algorithm for Pricing European Average Options. *Journal of Financial and Quantitative Analysis*, 26, 377–89.
- Wang, S., & Zhou, H. F. (2004). Staged Financing in Venture Capital: Moral Hazard and Risks. *Journal of Corporate Finance*, 10, 131–155.
- Wilkie, D. (2001). On the Risk of Stocks in the Long Run: a response to Zvi Bodie. *Proceedings of the 11th International AFIR Colloquium*, Toronto, 6–7 September 2001, 2, 741–762.