B.A.J. **15**, III, 573-677 (2009)

DYNAMIC ASSET ALLOCATION TECHNIQUES

By S. Jarvis, A. Lawrence and S. Miao

[Presented to the Institute of Actuaries, 27 April 2009]

ABSTRACT

Investment strategy is often static, punctuated by infrequent reviews. For most long-term investors, this practice results in large risks being taken that could otherwise be managed with a more dynamic investment policy. The bulk of this paper is aimed at analysing and describing two multi-period investment strategy problems — in order to derive potential dynamic strategies. Along the way, we show how static investment strategies can fail to deliver an investor's long-term objectives and describe the relationship of our work to other areas of the finance literature. This paper does not cover trading strategies such as Tactical Asset Allocation.

This paper sets out two main approaches to the multi-period problem. The first approach optimises a utility function. The second approach uses partial differential equation (PDE) technology to optimise a target statistic (in this case, TailVaR) subject to return and long-only constraints.

KEYWORDS

Investment Strategy; Stochastic Modelling; Dynamic Programming; Multi-Period Optimisation; Tail VaR; Utility

CONTACT ADDRESS

S. J. Jarvis, Blackrock, 33 King William Street, London EC4R 9AS. Tel: +44(0) 20 7743 3000. E-mail: stuart.jarvis@blackrock.com

1. Introduction

1.1 *Beyond Single-Period Investment Theory*

1.1.1 Modern portfolio theory, as introduced by Markowitz and developed by many others, provides practical techniques for analysing — and optimising — the trade off between risk and reward. This body of work is applied in a number of contexts and in a number of ways (for example, Sharpe (1964), Ross (1976), Rom & Ferguson (1994)). A common strand is that these are single-period theories: there is a single time horizon, risk and return estimates are made over that horizon and these enable a choice of investment strategy to be made.

1.1.2 An equally important but less widely understood stream of work within the financial academic literature examines the investment problem

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

Table 1. Approaches compared

Single-period approach	Multi-period approach
Short-term focus predominates	Long-term objectives can play a role alongside short-term risk management
Strategy will only evolve when another investment review takes place	Methodology for strategy evolution is put in place on day one
Symmetric view of risk likely to be chosen for simplicity or computational reasons	Asymmetric, goal-focussed, approach naturally leads to asymmetric risk measures

over multiple periods. A challenge — and a key attraction — of these papers is that the objectives of the investor come much closer to the surface. For a real-world investor — a long-term institution such as a pension plan or an insurance company, or an individual seeking to manage their savings so as to smooth their lifetime income — the multi-period experience is much more relevant, and the investment strategy can and should evolve through time in sympathy with their objectives.

1.1.3 To see the importance of this, consider performance measurement as a simple example. Modern portfolio theory provides clear and useful guidance on the construction of risk-adjusted performance: Sharpe ratios, Sortino ratios and the like are heavily used across the capital markets industry. In contrast, for a specific investor in the real world, it is the money-weighted returns that are important: not these ratios that depend upon time-weighted returns generated by a notional investment.

1.1.4 This paper seeks to describe, and apply, some techniques inspired by this multi-period approach: we are concerned with investment strategies that are dynamic in some way. These are strategies where the investment policy changes through time in some defined manner, in order to better meet an overall investment goal. We spend a fair amount of time expounding approaches to the multi-period problem that deserve to be better known. The problems faced by long-term investors are often multi-period in nature; breaking such a problem down into a series of single-period problems often means that some opportunities for a better solution are missed (see Table 1).

1.2 *Policy Evolution Over Time*

1.2.1 The usual output from Modern Portfolio Theory is an efficient frontier of investment strategies: there is a hedge portfolio, a growth portfolio, and the choice of a particular point on this frontier (the amount invested in the growth portfolio) is a matter of investor judgement. This choice might typically be expressed in terms of either a return target, risk target or risk aversion. When it comes to updating this view at the next review (which might be between a week and three years later), should this judgement change? Describing a framework in which it is sensible for the

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

point on the efficient frontier to move, and the manner of its movement over time, is one of our objectives here.

1.2.2 Dynamic strategies offer the opportunity to manage risk across time, but this also leads to a sudden explosion in the number of possible strategies. Fortunately, this is not a new problem in financial theory: quite the reverse. The dynamic investment problem for individuals has been the subject of a huge volume of financial theory, starting, in its modern form, with Merton (1969) and Merton (1971). While the initial model economies had constant interest rates, more recent work has moved this into more complex (stochastic interest rate) models, where the risks of assets and liabilities can be more fairly treated together.

1.2.3 The actuarial community has also contributed to this body of work. For example, Appendix A in Smith (1996) contains an exposition of various techniques. Kemp (1996) discusses how these techniques could be applied in practical situations, and shows how important it is to work within a model that does not allow anomalously attractive investment strategies — dynamic optimisation is often much more able to find and exploit these opportunities than a static approach. This leads to results that say more about the model than about sensible investment strategy, and may have acted as a disincentive to dynamic optimisation. Cairns (2000) solves a dynamic optimisation problem in a pension fund context, by solving a Bellman equation, and discusses the impact of different preference structures on the appropriate strategy. Neilsen (2006) presents a good overview of the academic literature. The recent EDHEC paper, Martellini *et al.* (2009) takes a similar perspective to ours here. On the numerical side, the Dempster *et al.* paper (2003) describes numerical approaches that form the inspiration for the work described in Section 6 below.

1.3 *Market Evolution over Time*

1.3.1 We should emphasise that the evolutions in portfolio allocation that we discuss in this paper are triggered by the investor's objectives. Although the investor's preferences are assumed to stay constant, they say different things at different values of the portfolio. In contrast, it is very common for investors to change their portfolio allocation in response to changes in markets. The market opportunity set does change over time, and an informed investor should of course be wise to this and adjust accordingly.

1.3.2 This adjustment could take several forms. At its simplest, it would be awareness that market risk is not constant — and therefore risk forecasts should be updated regularly to reflect new information. A much further move would be to actively allocate assets, based on an investor's views of the returns from the changing opportunity set. This is hard, and usually expensive. This form of tactical asset allocation is outside the scope of this paper.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

1.4 *Motivations and Scope*

1.4.1 As discussed above and more fully in Section 4, one motivation behind our work was to assess the potential benefit to long-term investors of dynamic strategies. We develop strategies in relatively simple model settings in order better to analyse these benefits. These candidate strategies can then be tested in more complicated models and in more practical market contexts.

1.4.2 A related motivation behind our work was to provide a framework within which option strategies can be considered. There has been an increased interest among the advisers to pension schemes in the use of options to tailor the distribution of future funding levels. Although these can be fairly easily presented with plausible arguments in their favour, it is usually very difficult to see why a particular strategy has been chosen above another, unless perhaps the profit to the writer of the options is taken into account. We wanted to support the general thrust behind these option strategies, while providing a means whereby less expensive and potentially more liquid options can be used to meet the same objectives.

1.4.3 This paper therefore seeks to meet several goals:

(i) to provide a brief description of the utility-based approach often used to analyse multi-period problems, and to give an overview of some of the results in this significant vein of the finance literature;
(ii) to provide a parallel description of a non-utility based solution;
(iii) to provide examples from the finance literature, from life insurance and from pensions where a dynamic policy is often considered;
(iv) to show how some of the insights from the utility approach can be applied to these problems; and
(v) to analyse the improvement that could result from the use of more dynamic investment strategies, and in particular the use of options within pension plans.

2. Risk and Return Measures

2.1 *Quantifying the Trade-off between Risk and Return*

2.1.1 When choosing an investment strategy, investors select assets that are likely to meet their particular return goals while managing the likelihood and/or severity of returns falling short of these goals. Assessing the right trade-off between return and risk therefore requires careful thinking about:

— *Target outcome*. The investor's goals should be clear. This may be expressed in terms of a benchmark or a return target, for example. Ideally these goals will be embedded in the wider situation facing the investor — the investor's asset allocation is likely to be only one of several levers that can be adjusted to meet these goals. This goal then helps determine the average return required of the investor's assets.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

— *Risk tolerance*. A judgement of the extent to which the investor is prepared to tolerate the return falling short of this target return. In the simplest situation, this could be expressed as an aversion to volatility, although more explicitly focusing on the probability or severity of poor returns may be more appropriate.

2.1.2 With these in place, an investor will then seek to carry out a form of optimisation — finding a strategy which gives as good a trade off as possible between the average return and the variability relative to this average. In carrying out this optimisation, further choices have to be made: what time horizon is relevant for the investor? What instruments are available in the investment universe? Is the strategy static or is it able to evolve dynamically in response to events?

2.2 *Choosing Appropriate Measures*

The particular choices of return and risk measure will vary according to the investor, or user, of the statistics. Some possible considerations include:

— Regulators are likely to focus on tail risk; the risks may be measured via deterministic stress tests in addition to more probabilistic measures such as value at risk.
— An investor with a longer-term perspective may be happier to focus on the dispersion of returns around the general level, rather than on the tail. This may lead to worse short-term performance, e.g. if market liquidity dries up, but the investor will expect to get paid a sufficient premium for bearing this risk.
— Particular portfolio characteristics may be more intuitive in specific circumstances. For example, active duration or PV01 would often be a more natural risk metric in a fixed income setting than a statistical measure such as tracking error.
— A risk metric agreed between an investor and an investment manager will need to be verifiable. Ex-ante risk, used in portfolio construction, and ex-post risk, derived from actual returns achieved, may both need to be tracked to ensure that they are compatible. The latter may be especially hard for some measures of tail risk (we don't want to wait 200 years to see if a tail event only happened 1 year in 200) or for strategies where the risk budget varies dynamically, so some visibility into the ex-ante risk methodology — and perhaps the use of a third party model — may be required.
— The timescale(s) for risk measurement should be clear, particularly for instruments or strategies whose exposures are not constant over time. For example, the characteristics of an option in a portfolio held for a short period would often be materially different if the option is held until maturity, and risk is measured over this longer timeframe.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

2.3 *Mean-variance Approaches*

2.3.1 The classic Markowitz approach involves trading the expected return outcome off against the volatility of returns around this outcome. This can be framed in multiple equivalent ways:

— maximise expected return for a given level of volatility;
— minimise volatility for a given level of average return; and
— maximise risk-adjusted return = return − aversion × variance.

These are equivalent: varying the level of volatility, return or aversion respectively in the three versions above sweeps out the same 'efficient frontier' of portfolios that are mean-variance efficient.

2.3.2 In a single-period situation, this can work very well. If we use one of the simplest models of asset price uncertainty, namely the assumption asset returns are distributed according to a multivariate normal distribution, we can readily create risk/return trade-offs for a variety of risk measures. For many risk measures, the joint-normality assumption means that these generally give identical frontier portfolios, so the simplicity of mean-variance optimisation makes it an attractive approach.

2.3.3 The situation in a multi-period situation is rather different. Joint-lognormality is the natural simple choice in a multi-period framework. If we work in a continuous time dimension, then this results in multivariate Brownian motion with constant drift and covariance.

2.3.4 Efficient frontiers can be established for different rebalancing frequencies: from static portfolios, through portfolios that rebalance n times over a fixed period, through to a theoretical limit of continuous rebalancing.

2.3.5 Figure 1 (from Goetzmann *et al.*, 2002 — see Appendix B for a derivation) shows the distribution of outcomes from a portfolio that is continuously-rebalanced in order to optimise the realised Sharpe ratio.

2.3.6 It is readily seen that:

(a) despite the underlying process being normally distributed, the resulting optimal distribution is far from normal;
(b) the left-hand tail of 'bad' outcomes is significantly fatter than the static portfolio comparator.

2.3.7 The Sharpe ratio measure, based on standard deviation, may fail to adequately capture an investor's risk preferences once we move away from normally distributed outcomes. A metric which appears to adequately select between 'good' and 'bad' strategies in a simple situation may no longer be an adequate description of the investor's goals in situations that allow more complex outcomes.

2.4 *Utility-Based Approaches*

2.4.1 In principle, therefore, we would like to find a way to rank

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

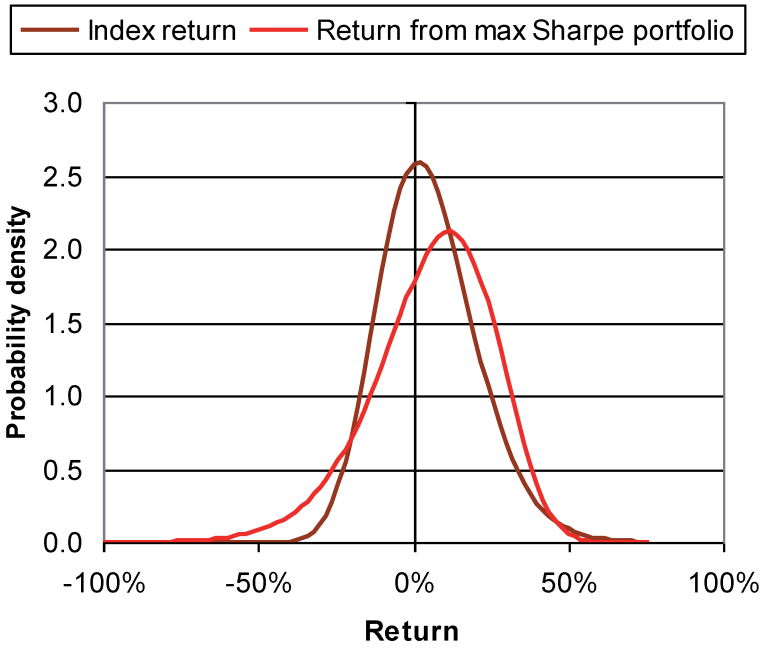


Figure 1. Probability density of a mean-variance optimal dynamic strategy

possible sets of outcome — we might express these outcomes as cash flows or values that arise from our strategy over time. Formally, consider the set $C = \{c : T \times \Omega \to \mathbb{R}\}$ of real-valued consumption processes — that is, $c_t(\omega) = c(t, \omega)$ is the cash flow at time t in scenario ω. Here T might be a subset of the positive integers or positive real numbers according to whether problems are being considered in discrete or continuous time. A scoring function is then just a map from this set to the real numbers:

$$u : C \to \mathbb{R}.$$

2.4.2 Although this is not crucial, a common approach is to consider expected-utility scoring functions, i.e. of the form

$$u\big(\{c_t(\omega)\}\big) = \sum\nolimits_t \mathbb{E}(U_t(c_t)),$$

where $U_t(w)$ is the utility score from wealth w. The utility function applies a score to each consumption and wealth outcome; two strategies can then be directly compared by comparing the expected utility, using the distributions of outcomes from following the two strategies.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

2.4.3 In this basic set-up, utility is acting as both return and risk measure. Alternatively, utility could be combined with a risk measure such as volatility: maximising utility subject to a cap on the variance of outcomes would then provide a way to focus separately on the desired outcomes and express tolerance for risk.

2.4.4 In principle, this is in fact little different to the expected return/standard deviation trade off, which can be expressed in terms of wealth w at terminal time T as:

$$\text{Maximise } \mathbb{E}(w_T) - \lambda \mathbb{E}\left(w_T^2\right). \tag{1}$$

Maximising expected utility, U, is simply

$$\text{Maximise } \mathbb{E}(U(w_T)). \tag{2}$$

2.4.5 Clearly a quadratic utility function will suggest the same investment strategies as those for an investor using the Sharpe ratio as their risk measure. It is easy to mock the quadratic utility function (e.g. it makes no sense for utility to fall beyond a certain level of wealth — the so-called bliss point) however if the time period is short and the impact of rebalancing choices small, a quadratic approximation to any utility function will be likely to do a good job.

2.4.6 The utility function formulation is explicitly more general, so the question arises: what choices of utility function make sense/give 'sensible' portfolio outcomes? Time-homogeneity is a natural requirement to seek. This means that maximising expected utility over two adjacent time periods separately should give the same answer as maximising over the whole time period, at least if the two sub-periods are independent and distributionally identical (i.i.d.). Such considerations lead to a family of 'constant relative risk aversion' (CRRA) utility functions (see Figure 2). These are of the form (for an arbitrary positive parameter γ):

$$U(w) = \frac{w^{1-\gamma} - 1}{1 - \gamma}. \tag{3}$$

2.4.7 Surprisingly, although (1) is amazingly successful in the single period case, this does not generalise neatly to multi-period investment policies, and maximising a function of the form (3) becomes the more natural choice. The results of using (3) in simple models are also much more intuitive, of which more below.

2.4.8 Investment strategy problems using the CRRA utility function are well-studied. Although this leads to mathematically-precise, aesthetically-pleasing results, the challenge in practical situations that interest us is to

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

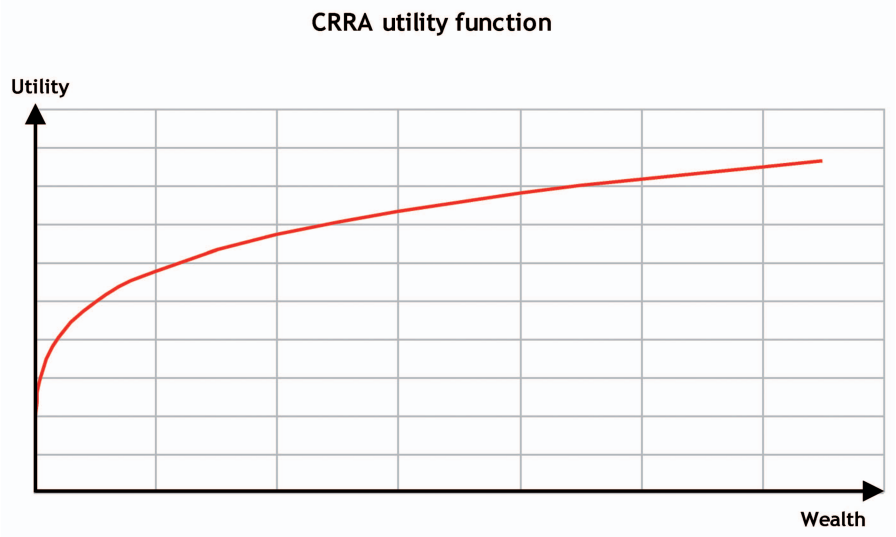


Figure 2. A constant relative risk aversion utility function

identify how a CRRA function fails to properly express investor preferences. Utility functions are in practice impossible to calibrate to an investor (indeed a significant strand of the behavioural finance literature argues that this is not possible even in theory) — however we can hope to use utility functions to express key features and then use optimisation techniques to learn what investment strategies make sense as a result. This is just a case of George Box's oft-quoted statement (Box *et al.* (1987), p. 424) "All models are wrong, but some models are useful".

2.4.9 Appendix A provides some additional background on utility functions.

2.5 *Asymmetric Outcomes and Tail Risk Measures*

2.5.1 Variance has the notorious property that it penalises above-average outcomes as much as below-average outcomes. As we have seen, volatility, being derived from variance, is most appropriate if the underlying distributions to which it is applied are symmetrical. If the distributions have material skewness — for example if options form part of the allowable investment universe or dynamic strategies could give rise to option-like payoffs — then using a symmetric risk measure may distort results. The average outcome plus dispersion around this may give a misleading sense of the real outcome distribution.

2.5.2 This simply re-emphasises the point that the focus should be on

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

the investor's goals — if these are properly captured by an average return, well and good. If not, an appropriately structured utility function may provide more control. Alternatively, measures of risk have been developed that focus instead on the region of interest, i.e. losses — the 'left hand tail' of returns well below the average outcome.

2.5.3 Value at risk (VaR) is such a measure that has seen wide use for many years. More recently, academics have tended to use other risk measures, in a growing literature that was kick started by the famous article of Artzner *et al.* on coherent risk measures, which pointed out the lack of coherence of VaR. Whereas VaR essentially reports quantiles such as the '1 in 200 event' TailVaR, 'tail value at risk' or 'conditional value at risk' looks at the average of say the worst 0.5% of possible outcomes. Due to its coherence, it has attractive theoretical and analytical properties that make it more suitable for modern risk management.

2.5.4 Therefore, trading off the mean outcome against a tail risk measure such as TailVaR or the mean utility against this measure can be an attractive means to compare tail-risk management procedures.

2.6 *Value at Risk (VaR) and Tail Value at Risk (TailVaR)*

2.6.1 Value at risk (VaR) measures the potential loss over a particular time frame. The loss is stated at a particular probability level, e.g. 99.5% of the time the annual loss will be no more than 20%. Although authors generally agree upon the meaning in practice, there is considerable scope for confusion when writing down a definition: do you think of the probability level as 0.5% or 99.5%? Is the focus on the portfolio return or the portfolio loss? For clarity, the definition we use here is as follows. Suppose X is a random variable representing the portfolio value at the time horizon we are interested in, and ε (e.g. $\varepsilon = 5\%$) is the tail probability we are interested in:

$$VaR_{\varepsilon}(X) = -\inf\{x : P(X \leq x) > \varepsilon\}.$$

2.6.2 This definition is illustrated in Figure 3, which is based on a Student (t) distribution with 4 degrees of freedom. In this continuous-density example, the vertical dotted line represents the point at which the 5% probability level is breached. The value at risk is -3.9 in this example.

2.6.3 The picture also shows the CVaR, the conditional or tail value at risk, also sometimes known as expected shortfall (ES), although this terminology is used elsewhere with a slightly different meaning. This represents the average loss in the shaded area shown: first define the quantile function:

$$q_X(p) = F_X^{-1}(p) = \inf\{x : F_X(x) \geq p\} \in \mathbb{R} \cup \{-\infty, \infty\},$$

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

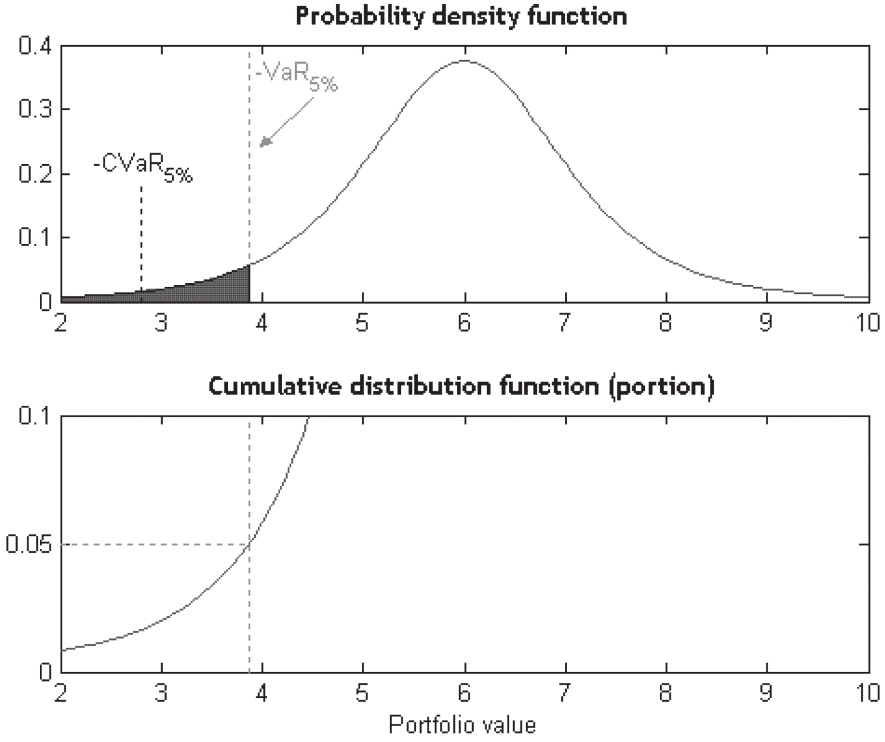


Figure 3. Relationship between value at risk (VaR) and tail value at risk (CVaR)

and then:

$$TailVaR_{\varepsilon}(X) = \alpha^{-1}\mathbb{E}\left[\left(F_X^{-1}(\varepsilon) - X\right)^{+}\right] - F_X^{-1}(\varepsilon).$$

2.6.4 In the intuitive case when there are no jumps in the cumulative distribution function, the quantile function is, up to a sign, the same as the value at risk, and the CVaR is the expected loss in the tail bounded by the VaR. That is, informally, TailVaR is the average over the worst alpha% of outcomes (e.g. alpha = 5%). Where the portfolio can take particular values with positive probability (e.g. this will be the case for typical option payoffs; or when working with simulations) slightly more care is required.

2.6.5 Interest in tail-risk measures has increased in the wake of Artzner *et al.* (1999). A substantial corpus of papers examines measures of risk, with particular focus on those measures that are coherent.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

2.6.6 An alternative definition of TailVar/CVaR is as the solution of an optimisation problem:

$$CVaR_{5\%}(X) = \inf\left\{\frac{\left((s-X)^{+}\right)}{5\%} - s : s \in \mathbb{R}\right\}.$$

2.6.7 This formulation is due to Rockafellar & Uryasev (1999); see also Pflug (2000). Crucially, the functional (of s and X) being optimised here is a convex function. This means that very often — for example, when seeking the minimum CVaR for a given expected return — any local minimum is guaranteed to be a global minimum, and convex optimisation techniques can be used to steadily locate the desired optimum.

2.6.8 A similar feature is that the nature of the optimisation is often changed from being one of a 'global' nature to one that is 'local'. By fixing s, small changes in one part of the distribution of X can be considered, and the effect on the functional analysed near where these changes are made. In the more usual form (as in the definition of tail mean above), the percentile (the Value at Risk) has to be recalculated for any small change — and this is a global calculation (i.e. it depends on the whole distribution). Features of the optimum can therefore be derived before optimising s, and this makes the search for the optimum considerably easier.

3. Techniques for Finding an Optimal Investment Strategy

3.1 *Introduction*

3.1.1 In the previous section, we introduced two possible measures that could be used to help determine good investment strategies: TailVaR and utility functions. In this section, we turn to some techniques that have been developed to solve or simplify these problems. We have made no attempt to be encyclopaedic — instead the focus is on techniques that we happen to have found useful. These are:

— *Dynamic programming*. Bellman's approach to multi-period optimisation is very powerful. It leads to a recursive approach that can be applied either numerically or in some cases analytically.
— *Cox–Huang approach*. This builds on dynamic asset pricing theory developed in connection with option pricing. In a complete market setting this often allows a significant shortcut to be made and analytical solutions to be developed.
— *Numerical approaches*. We outline some approaches that we have found to be useful.

3.1.2 We also discuss constraints at the end of this section — although their effect is often a feature in any situation where the above techniques are

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

used, we are seeking here to make general comments on multi-period optimisation and the techniques that need to be applied in practice. It therefore seems to make sense to discuss the impact of constraints separately.

3.2 *Dynamic Stochastic Programming and Bellman's Principle of Optimality*

3.2.1 The aim of dynamic programming is to find a multi-period strategy that maximises an objective function (a function of the strategy). Dynamic programming originated, in many respects, from original work by Richard Bellman. His insight was to see how a given multi-period problem could be embedded within a wider set of problems, which can then be broken down into simpler (single-period) problems. Often this arises because the objective is a function not only of the strategy but also of the initial state. Oddly, it is often simpler to consider the larger class of problems corresponding to different initial states and different time horizons. This is because one can then take a path-wise perspective: at any point on an optimal path, the *remaining* path should be optimal for the appropriate problem started at that point.

3.2.2 This 'principle of optimality' immediately suggests a recursive approach: suppose that the optimal policies and resulting value of the objective function V_{t+1} has already been derived for time $t+1$. (At the terminal horizon, this value function is just the original objective function.) The value V_{t+1} depends on the state at time $t+1$. The decision problem at time t then amounts to finding a policy that maximises the expected value of V_{t+1} — crucially, this is a single-period problem, with much lower dimensionality than the original multi-period problem.

3.2.3 In addition to the obvious computational gain that this procedure offers, a major benefit is in fact to clarify the *structure* of an optimal solution. It has become clearer what each individual decision is seeking to achieve, and how this relates to other decisions made along the path or nearby paths. This structure is in accordance with many real world situations — decisions must be taken one at a time, and the objective is usually to maximise some objective prospectively, independently of how the portfolio or other system has happened to arrive in its current state. At the same time, it is easy (and we will discuss an example below) to construct examples that do not exhibit this structure.

3.2.4 Bellman's principle of optimality is somewhat heuristic, and making it into a more precise statement to which optimisation procedures can be applied requires a more explicit statement of the objective function in a particular situation. Additive-utility is the classic financial example (Samuelson (1969), Merton (1969): a strategy consists of an investment policy α_t and consumption (drawdown) of c_t from the portfolio at each time t, leading to a terminal asset portfolio with value w_T; the objective is to maximise

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

$$\mathbb{E}_0\left(\sum_{t=0}^{T-1} u_t(c_t) + u_T(w_T)\right). \tag{4}$$

3.2.5 Here u_t are utility functions, whose structure is unimportant for this discussion, and where the sum is replaced by an integral in the continuous-time version. The potential value that can be still achieved at time t is

$$V(t, w_t) = \max_{\{c_s, \alpha_s\}_{s=t}^{T-1}} \mathbb{E}_t\left(\sum_{s=t}^{T-1} u_s(c_s) + u_T(w_T)\right). \tag{5}$$

3.2.6 The optimality principle then leads to the Bellman equation which is the recursive procedure

$$V(t, w_t) = \max_{c_t, \alpha_t} \mathbb{E}_t\big(u_t(c_t) + V\big(t+1, w_{t+1}(\alpha_t)\big)\big). \tag{6}$$

3.2.7 This recursion is started with the initial condition $V(T, w) = u_T(w)$, and — at least in principle! — the functions $V(t, w)$ can be derived one by one.

3.2.8 This discussion is evidently based on intuition that depends upon discrete time. In a continuous time setting, similar but more technical considerations continue to apply and a similar equation, now known as the Hamilton Jacobi Bellman or HJB equation, can be derived.

3.2.9 The procedure described here is variously called dynamic stochastic programming or more simply dynamic programming, or stochastic control. (Arising from engineering contexts, the strategy that should be followed at each time is known as a 'control'.)

3.3 *Cox–Huang Approach*

3.3.1 While very well-suited to numerical optimisation, the value function V may be difficult to derive analytically except in special cases. Merton's solution of the continuous-time version of (6), where the utility is of CRRA form, is such a special case and was a significant milestone.

3.3.2 As mentioned above, even where it is difficult to solve, the structure of the Bellman equation, in identifying what each decision must do in order to be part of a dynamically optimal strategy, is itself a major advance on the apparently daunting original form of the problem.

3.3.3 A second approach which also gives insight into the structure is the Cox–Huang, or martingale, approach, developed in Cox *et al.* (1989) and Karatzas *et al.* (1987). Here, rather than focussing on each decision, at each time, attention is paid to the final result, i.e. the set of consumption amounts $\{c_t\}_0^{T-1}$ and the final wealth w_T. These must satisfy a budget constraint —

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

that the value at time zero of these 'payoffs' equals the current wealth. The payoffs that maximise the expected utility (4) can therefore be sought, subject to this constraint. If such payoffs exist, the question of course arises as to whether they are achievable, i.e. does there exist a strategy $\{\alpha_t\}_0^{T-1}$ which, together with the consumption policies, produces these payoffs? The answer is 'yes', provided that the market is dynamically complete — this is simply a fancy way of saying that any payoff is achievable, and it is a rather restrictive condition, but this does hold in many situations where the Bellman equation can be explicitly solved and in many simple models.

3.3.4 The construction of a strategy that leads to a particular payoff is a familiar topic in option pricing theory. From a probabilistic perspective, the appropriately discounted value of the payoff is a martingale and the martingale representation theorem guarantees the existence of a strategy supporting this value. From an economic perspective, this strategy is simply one that delta hedges the discounted value at each point in time. Bellman's value function can thus, in this case, be replaced by the value of the future part of the optimal payoff and the strategy is secondary to this payoff.

3.4 *Numerical Approaches — Introduction*

3.4.1 As mentioned above, the recursive structure of the Bellman equation makes it rather well suited for numerical approximation and optimisation. An estimate of the value function can be made at each time step, so that a solution is evolved by moving backwards in time, taking into account portfolio constraints along the way.

3.4.2 In constructing estimates of functions, an approximation methodology is required. The most common varieties used in finance are trees and lattices, and although these are well known we provide a brief summary here, paying particular attention to the PDE solution techniques that have been developed for lattice approximations. The PDE to which we have applied these ideas below is the Fokker–Planck equation, which we also discuss here.

3.4.3 *Tree approaches*

3.4.3.1 A classic approach that has been used in the option pricing literature (following Cox *et al.* (1979) for example) is to construct a tree. The tree encodes a simplified set of future paths for asset prices, starting with the current set of prices. A simple example of a recombining binomial tree is illustrated in Figure 4.

3.4.3.2 There is a single node at time 0, and the price of the (single) instrument being modelled can go either up or down. After one time step there are two possible future states. Over n time steps, there are 2^n possible paths for asset prices and either 2^n states at the end of these steps for a non-recombining tree, or, in a recombining tree (i.e. a tree where going up then down produces an identical situation as going down then up), $n + 1$ states.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

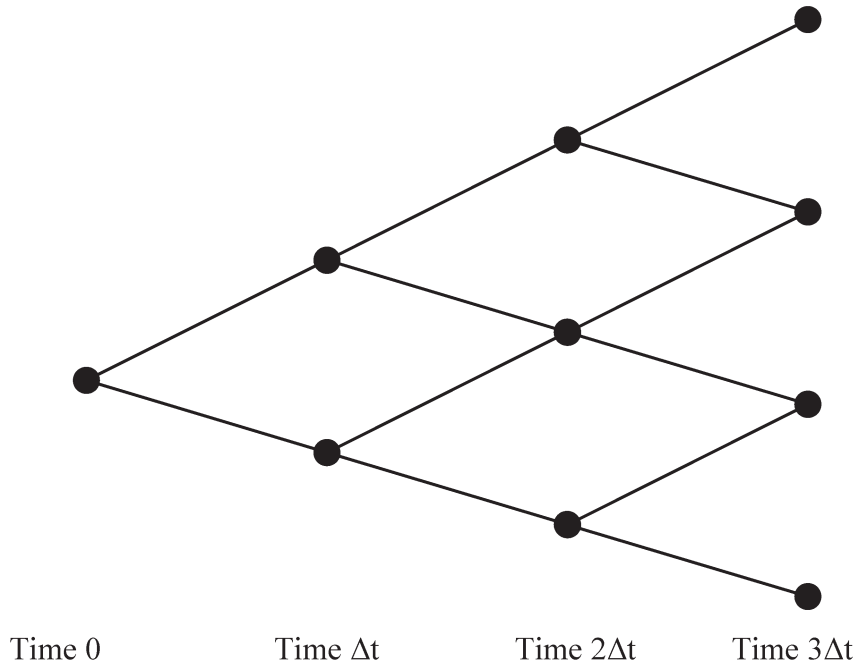


Figure 4. Binomial tree, with recombination

3.4.3.3 Standard arbitrage arguments can be used to price any stochastic payoff payable at time $n\Delta t$ — the price is given by working backwards. That is, the price at the final set of nodes is known; the price one step back is given by constructing an appropriate hedge portfolio for the payoff at each node and then valuing this hedge portfolio. This gives the prices at each of the nodes at time $(n-1)\Delta t$. Iterating this procedure then produces a price at the initial node, time 0.

3.4.3.4 This approach works well if the final payoff is known and if the tree provides a good approximation to the set of future paths. For many standard models and options, this is indeed the case, and it can even be shown that the binomial tree approaches the continuous-time model as the step size tends to zero. This is not surprising: it is well understood how the binomial distribution approaches the normal distribution in an example of the central limit theorem in action. In a simple model, e.g. where the mean and variance are constant over time, the calibration is straightforward: the size of the up and down moves can be set so that the distribution of outcomes at $n\Delta t$ matches the expected distribution.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

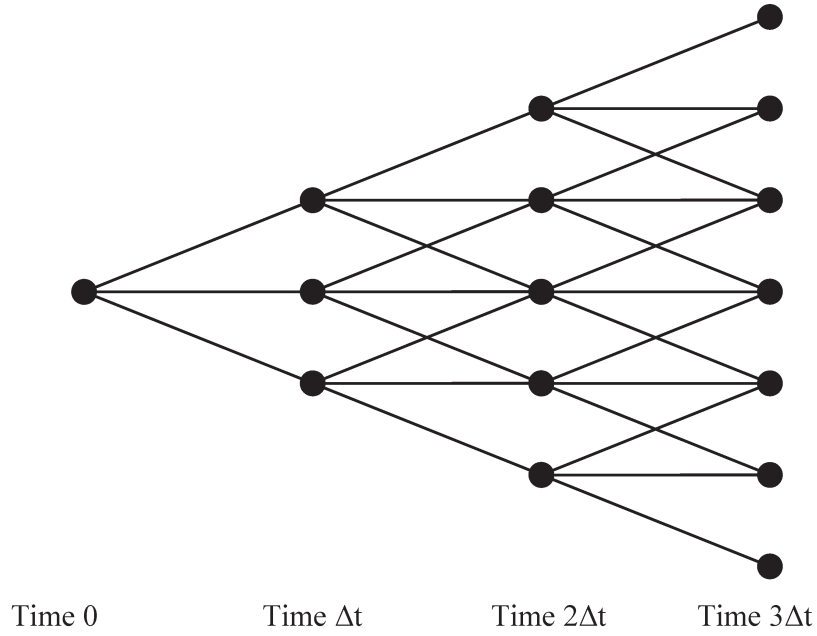


Figure 5. Trinomial tree, with recombination

3.4.3.5 In some circumstances, a trinomial tree approach may be preferred (e.g. this helps the calibration when fitting 'implied volatility' surfaces) — in this case, there are obviously three possible branches at each node. This is illustrated schematically in Figure 5.

3.4.3.6 In this case, with a recombining tree, after n steps there are $2n + 1$ nodes. The total number of nodes is $(n + 1^2)$.

3.4.4 *Lattice approaches*

3.4.4.1 To value an option, a more general alternative to building a tree involves using a lattice. Here, a subtly different perspective is taken. Given stochastic differential equations for a set of asset classes, a partial differential equation (PDE) can be written down which governs the evolution of the price of an option over time. Now, rather than building a tree and arguing that the tree is a good approximation to the continuous limit, the goal is to find an approximation (in discrete time) to the solution to this PDE (in continuous time). The PDE has the generic form

$$\frac{\partial f}{\partial t} = a + b\frac{\partial f}{\partial x} + c\frac{\partial^2 f}{\partial x^2},$$

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

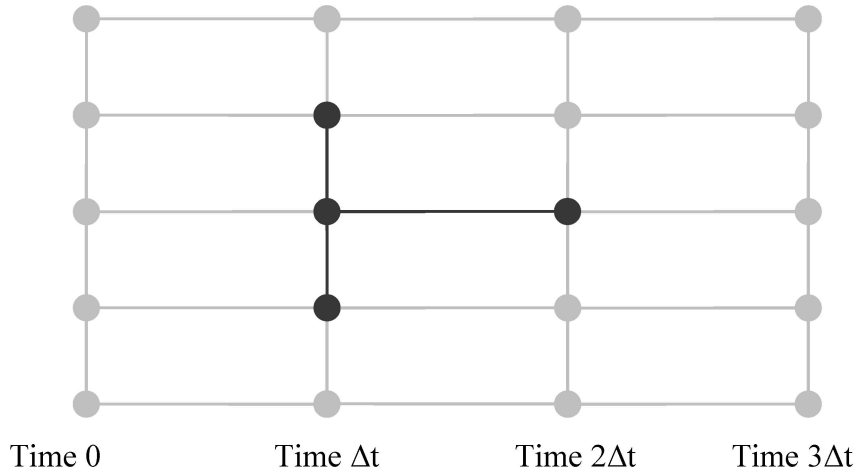


Figure 6. Finite difference lattice — explicit method

(here a, b, c are usually functions rather than constant parameters).

3.4.4.2 The numerical solution of PDEs, particularly PDEs of the parabolic form that arise in this context, is a subject that has been well studied beyond finance — for example in engineering contexts. The simplest approach, the finite difference (FD) approach, approximates the region over which the PDE is to be solved by a lattice (see Figure 6). The differentials that appear in the PDE are approximated by finite differences, for example:

$$\frac{\partial f}{\partial t}(x, t) \approx \frac{f(x, t + \Delta t) - f(x, t)}{\Delta t}.$$

3.4.4.3 The differential equation then can be approximated by a set of difference equations and these equations can be readily solved. A set of general results then provide confirmation that, under certain conditions, as the step sizes of the lattice $\Delta x, \Delta t \to 0$, the lattice values tend to the solution to the continuous PDE.

3.4.4.4 A number of choices can be made about the approximations used to define these derivatives. These are illustrated in Figures 6 to 8. If the derivatives in the x direction are estimated using values of the function at time t then the method is called 'explicit': the estimate of $f(x + \Delta x, t + \Delta t)$ depends on the 3 nearest values at time t. If the derivatives are estimated using time $t + 1$ values, then this is the implicit method. Finally, using a weighted average of these leads is known as a 'theta' method (with theta being the relative weight).

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

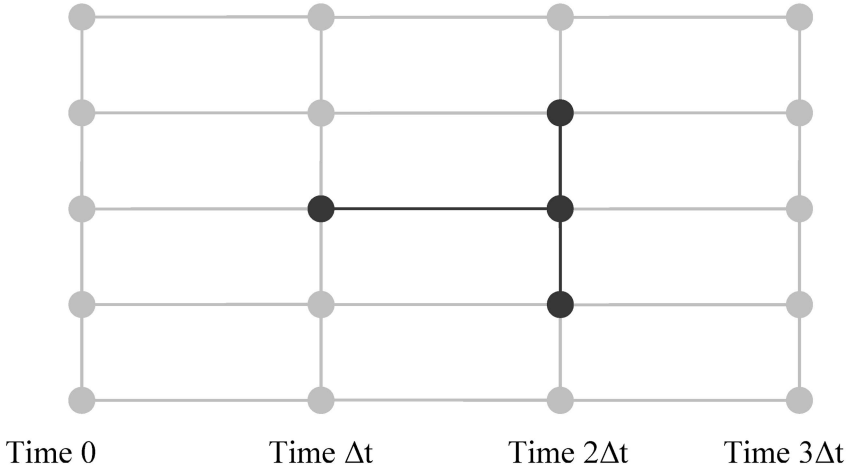


Figure 7. Finite difference lattice — implicit method

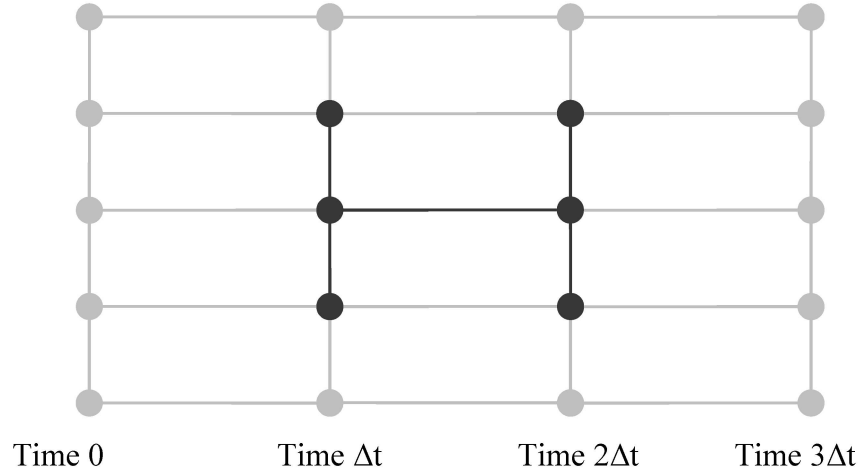


Figure 8. Finite difference lattice — theta method (e.g. Crank-Nicolson)

3.4.4.5 The conditions for convergence in the limit as the step sizes go to zero for an implicit or theta method are less severe than in the explicit method, which runs the risk of magnifying any rounding errors that arise in the numerical approximation. The main example of a theta method is to set theta equal to a half: this is usually the most accurate approximation that can

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

be made (the order of the approximation is two rather than one). This method is due to Crank & Nicolson (1947).

3.4.4.6 These finite-difference approximation schemes can be applied to the well known PDE arising in option pricing. The 'forward direction' in time for the PDE is actually backwards in time for the option. Just as in the tree approach, we start with a set of values at the terminal horizon and derive values for prior times by moving backwards one step at a time.

3.4.4.7 For more detailed discussions, we refer to the reader to Duffy (2006), Clewlow *et al.* (1998) or Seydel (2009), for example.

3.4.5 *Combined lattice/tree approach*

3.4.5.1 We have used a combined approach in Section 6 below. This seeks to combine the advantage of a smaller number of nodes in the tree approach with the PDE-friendly structure of a lattice. We would like to optimise a strategy $\alpha(x, t)$ defined on a tree, but calculate a functional of this strategy by solving a PDE on a lattice. Fortunately this is straightforward: all that is required is a mapping that produces a strategy on a lattice given one defined on a tree — this is illustrated by Figure 9.

3.4.5.2 At each time, the x direction (the vertical direction in Figure 9) can be broken into a series of intervals, with the interval boundaries set to the lattice points midway between each tree node. The number of intervals thus equals the number of nodes, and a function on the tree-nodes is mapped directly to a step function on the whole real line.

3.4.5.3 More explicitly, but less clearly, suppose that there are M lattice points per node ($M-1$ lattice points between each node: $M=2$ in the picture) and $2N+1$ lattice points in total (N above and N below the initial node: $N=7$ in the picture). Then given a function $f(x)$ on the nodes $-t \le x \le t$ at time t, we can define a function $f(l)$ on the lattice points $-N \le l \le N$ via

$$f(l) = \begin{cases} f\left(\text{int}\left(\dfrac{l}{2M}+\dfrac{1}{2}\right)\right) & \text{if} \quad x-\dfrac{1}{2} < \dfrac{l}{2M} < x+\dfrac{1}{2} \\ & \text{for some integer } x, -t < x < t \\ \dfrac{1}{2}f\left(\dfrac{l}{2M}+\dfrac{1}{2}\right)+\dfrac{1}{2}f\left(\dfrac{l}{2M}-\dfrac{1}{2}\right) & \text{if} \quad \dfrac{l}{2M} = x+\dfrac{1}{2} \\ & \text{for some integer } x, -t < x < t \\ f(t) & \text{if} \quad t-\dfrac{1}{2} < \dfrac{l}{2M} \\ f(-t) & \text{if} \quad \dfrac{l}{2M} < -t+\dfrac{1}{2}. \end{cases}$$

3.4.5.4 Here 'int' denotes 'nearest integer'. The idea is that node x gets mapped to the lattice points $(2M-1)x$ to $(2M+1)x$. It is worth pointing out

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

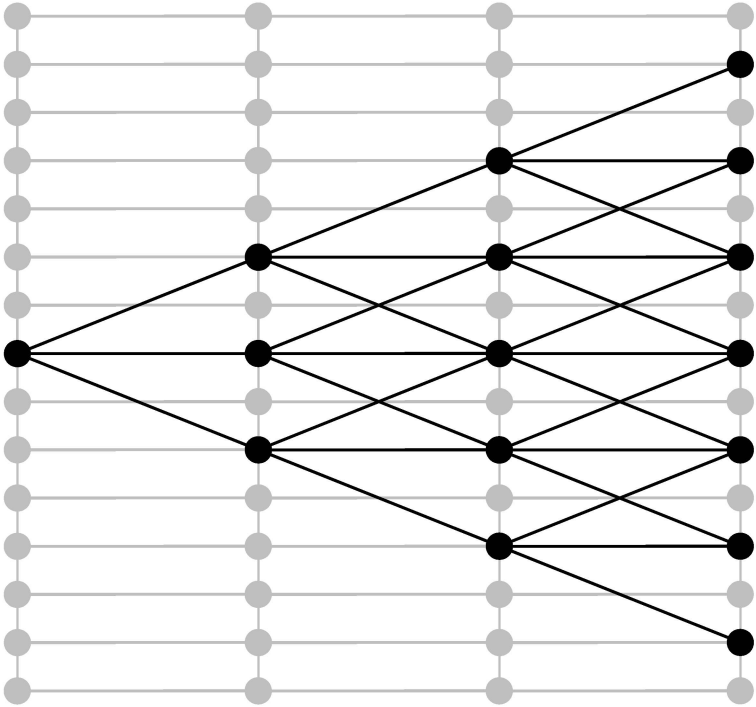


Figure 9. Combined tree and lattice

that although the number of lattice points, $2N + 1$, should satisfy at least $N \geq 2MT$, where T is the number of time steps, in practice a larger N is often useful so that a boundary condition on $p(\pm N)$ can be imposed without hindering numerical stability.

3.4.6 *The Fokker–Planck equation*

3.4.6.1 In the absence of an optimal substructure to the dynamic optimisation problem, i.e. where the problem can no longer be broken into single-decision pieces, the problem must be attacked 'all at once'. This unfortunately sets the stage for the 'curse of dimensionality' to enter — the explosion in the number of states and parameters can become a serious problem, leading to impractically long computer run times. The aim is to optimise a functional of a strategy that consists of asset allocations at each future time, along each possible path. Estimating this functional for a given strategy appears to require a costly Monte Carlo procedure, and this must happen at each iteration in the search for an optimum.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

3.4.6.2 The Fokker–Planck equation allows us to manage some of this dimensionality. It allows the evolution of the probability density of the portfolio value to be modelled easily without resorting to Monte Carlo simulation. The Fokker–Planck equation is a PDE of a similar type to that which arises in option pricing. Whereas option prices satisfy a PDE — essentially the backward Kolmogorov equation — it turns out that the probability distribution of outcomes satisfies a forward Kolmogorov equation, also known as the Fokker–Planck equation. The coefficients in this equation depend on the asset strategy — solving this PDE approximately for a discrete version of the asset strategy thus provides an estimate of the probability distribution at the time horizon, and from this estimate the functional to be optimised can be readily derived. A nonlinear optimisation routine can then be applied to find the optimal asset strategy on a lattice, which should be close to the optimal asset strategy defined in a continuous-space domain.

3.4.6.3 A strategy for a dynamic optimisation problem consists of a set of asset weights w at each time and portfolio value. In principle, the weights could depend upon the path taken up to that time. However, as the objective functions we consider are defined in terms of statistics defined at the terminal time, the optimal weight at time t will depend only on this path via the portfolio value at time t. This assumes that the future opportunity set is homogeneous (but if it is not, the inclusion of state variables that determine the regime at time t allows the process to work) and that transaction costs are negligible. This is necessary to allow the tree to recombine.

3.4.6.4 If :

(i) the weights are functions of time t and portfolio value Y alone,
(ii) the vector of drifts of the various assets modelled is constant or dependent on time only', and
(iii) the variance covariance matrix of the various assets is either constant or dependent on time only,

then the drift and volatility of the strategy depend only on time and portfolio value.

3.4.6.5 Thus, the portfolio value satisfies a stochastic differential equation of the form

$$dY = Y\mu'(w(t, Y))dt + Y\sigma(w(t, Y))dW_t.$$

The probability density function, $p_Y(y, t)$ (informally, $p_Y(y, t)dy = \mathbb{P}(Y(t) \in [y, y + dy])$) then satisfies:

$$\frac{\partial p(y, t)}{\partial t} = -\frac{\partial}{\partial y}[x\mu'(y, t)p(y, t)] + \frac{1}{2}\frac{\partial^2}{\partial y^2}\left[x^2\sigma^2(y, t)p(y, t)\right].$$

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

This is the Fokker–Planck or forward-Kolmogorov equation.

3.4.6.6 In fact, it's simpler, and numerically more stable, to do a log-transformation and express the equation in terms of $p_X(x,t) = \exp(x)p_Y \times (\exp(x), t)$, the pdf of $X = \ln(Y)$. With a slight abuse of notation, Ito's formula leads us to set $\mu = \mu' - \sigma^2/2$, then giving:

$$dX = \mu(w(t, X))dt + \sigma(w(t, X))dW_t,$$

and so

$$\frac{\partial p_X(x,t)}{\partial t} = -\frac{\partial}{\partial x}[\mu(x,t)p_X(x,t)] + \frac{1}{2}\frac{\partial^2}{\partial x^2}\left[\sigma^2(x,t)p_X(x,t)\right]. \tag{7}$$

3.4.6.7 This is the equation to which we apply the methods such as Crank–Nicolson discussed above. It is worth mentioning that, although we have stated the forward Kolmogorov equation for diffusions, i.e. solutions of stochastic differential equations based on Brownian motion, a similar equation is valid in much more general contexts, the main requirement being the Markov property that the future depends on the past only via the current state (no path dependence). Thus, for example, a similar approach can be taken for regime models that depend on finite-state Markov chains, for jump process models, etc.

3.4.6.8 The Kolmogorov forward equation is also used in option pricing — when a model that captures the 'volatility smile' is required. Option prices can be used to infer a risk-neutral evolution of the probability density, and the Fokker–Planck equation is then inverted to determine a volatility surface consistent with that evolution.

3.5 *Portfolio Constraints*

3.5.1 In many analytically-tractable situations, constraints are few. In practice, there are often multiple constraints. A classic constraint is a 'long only' one, corresponding to an investor who is unable to create short, or leveraged exposures to a basic set of asset classes. These constraints may well be binding for problems which are otherwise rather benign. There are at least three approaches to these kinds of problem, following the pattern outlined already in this section:

— *Dynamic programming*. The recursive Bellman equation approach described above can handle constraints at each stage rather easily. This is most generic, and thus requires most creativity, to solve. This may work analytically in special situations but is most useful in numerical work.

— *Duality approach*. As described in Karatzas & Shreve (1998) and Rogers (2003), Lagrange multipliers (in fact, a Lagrange multiplier process) can

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

be introduced to produce a modified optimisation problem. A solution to the unconstrained modified problem is often a solution to the unmodified constrained problem, in the usual way, given appropriate Lagrange multipliers. Again, in some situations, the dual problem (effectively, finding the appropriate Lagrange multipliers) can be more readily solved and a solution then found for the original ('primal') problem.

— *Numerical approach*. Where there is no optimal substructure allowing a Bellman-style decomposition into sub-problems, a brute force approach, using a discretisation of the problem on a suitable tree, is required. As usual, care and attention is required to set up the algorithm to ensure that a solution can be found in reasonable time. Analytical work can sometimes be undertaken to simplify the problem before resorting to the computer. For example, at each point the chosen portfolio will usually be mean-variance optimal and a parameterisation of the mean-variance efficient frontier means that only a single parameter rather than n portfolio holdings need be found at each point of the tree. The parameter used can be thought of as a Lagrange multiplier, and thus this approach can be linked into the duality approach above.

3.5.2 In Section 6 below, we show how a numerical approach can be used to examine a problem with portfolio constraints.

4. Optimal Utility Approach

4.1 *Introduction*

4.1.1 There is a strong academic tradition that quantifies an investor's multi-period investment problem as the maximisation of the expected utility to be gained from an income/wealth stream:

$$\text{Maximise } \mathbb{E}\left(\sum_{t=0}^{T-1} u_t(c_t) + u_T(w_T)\right).$$

4.1.2 Here c_t is the income drawn down from the portfolio at time t and the portfolio value at the terminal time T is w_T. As a model for individual decision-making under uncertainty, this has strong arguments in favour, and has been studied in a wide range of contexts, mostly exercises in establishing extensions of the pioneering work of Merton (1969). In the context of a homogeneous cash-and-stock model (the 'Black Scholes' model) and a particular form of the utility functions, Merton established simple formulations of the optimal strategy and drawdown policies.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

4.2 *Whose Utility?*

4.2.1 However, one of the problems that motivates our interest in multi-period investment problems is the question of how a pension plan's investment strategy should evolve over time. While the utility framework is well-grounded for individuals, it is much less so in the context of a pension plan. Indeed, in a pensions context, the main difficulty to answer is: who is the investor? Who is taking risks & benefiting from returns in a pension plan? Plan members, acting through the trustees as their agents, have a defined benefit, so that, provided the sponsoring employer does not default, they have no exposure to the funding level of the plan. Of course, the proviso that the employer remains solvent is a strong one: companies regularly go insolvent. On the other hand, the sponsoring employer — or rather its shareholders — is exposed to the funding level of the plan: as the funding level goes up or down, less or more future contributions will be required. Black (1976) and Exley *et al*. (1997), for example, argue that the shareholders, as the bearers of the risk of the asset performance, are the relevant parties. But shareholders will measure utility at the level of their entire portfolio — optimisation at the level of a particular piece of the portfolio would be sub-optimal. Rather, the proper objective for shareholders to ask of a company should be to maximise the value of the company rather than trading off risk and return.

4.2.2 These arguments are forceful (and we do not dispute them) but largely ignored in practice — pension plan trustees do take risk in order to improve funding levels and reduce contributions, and the problem of how to vary risk over time as the funding level changes is an important one. In studying this problem, we therefore choose to proceed regardless: at its simplest, this section could therefore simply be regarded as studying the important problem of an individual investor who has a target liability and/or minimum income level; but more controversially this analysis helps to identify multi-period strategies that are consistent with goals of this kind in an institutional context. This could perhaps be framed in terms of the utility of an agent (e.g. a trustee or finance director) rather than the ultimate risk taker themselves. For a fuller discussion of the perspective of these different stakeholders, we refer the reader to Barnes *et al*. (2008).

4.3 *Overview*

4.3.1 In much of the academic literature, the fashion is for the utility functions to be of the 'constant relative risk aversion' (CRRA) type — essentially, the utility function is a power-law function of wealth. In this case, the classic Merton result (in what would now be called the homogeneous Black–Scholes model case) is that the investment strategy is a constant mix of risky assets and bonds; the rate of drawdown (usually called 'consumption') is a deterministic function depending on the relative attractiveness of consumption and the bequest motive (the desire to have assets remaining at the terminal time).

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

4.3.2 We are interested in the following adjustments to the standard CRRA function:

— *Liability benchmark*. Rather than being concerned with the value of the assets in cash terms, it is the surplus or deficit relative to a stochastic (i.e., variable) benchmark, which might be the value of a target income or target liability depending on the investor;
— *Target funding level.* The typical objective is to reach a particular 'funding level' relative to the liability benchmark (likely to be 100%) — the additional utility for exceeding this level is likely to be minimal; and
— *Minimum funding level.* We consider several ways in which a minimum funding level might be incorporated, inspired by different ways in which a pension plan's economic context could be interpreted.

4.3.3 For each of these variations to standard utility functions, an optimal solution can be developed mathematically — this is described in Appendix C, using the Cox–Huang approach. Here in the main text we focus instead on the properties of these solutions. Several perspectives can be usefully taken:

— *The dynamic strategy*. The portfolio strategy is a function of (a) time to the investment horizon and (b) the funding level. This can often be characterised in terms of the risk of the portfolio relative to the liability benchmark;
— *Payoff profile.* The net result of following the dynamic strategy through time is a final asset level that is a function of the cumulative return of a static risky portfolio over the liability benchmark. The shape of this profile is a simple consequence of the shape of the utility function (see earlier comments on the Cox–Huang approach), and the dynamic strategy can be thought of as the 'delta hedging' strategy that leads to this payoff;
— *Outcome distribution.* By combining the probability distribution of the risky portfolio with the payoff profile, the distribution of achieving particular funding level outcomes can be determined. The shape of these distributions, and in particular the tails, emphasises the potential benefits of dynamic strategies to control future outcomes for investors.

4.3.4 One consequence of these perspectives is that the utility framework is not strictly necessary — in client conversations it can be easier to focus instead on either the payoff profile or the distribution of outcomes. The dynamic strategy then follows naturally from these. It is often easier for a client to consider the trade off between potential outcomes rather than quantifying an abstract utility function. The utility function therefore only serves to give confidence to the adviser that the resulting strategy provides a time-consistent means of managing risk.

4.3.5 In the main part of this section we have restricted ourselves to

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

models that are analytically tractable. This enables clean results to be derived, and intuition to be gained, but ultimately more realistic models are needed to make decisions in practical situations. We seek to identify general principles (e.g. on the risk budgeting process over time) that are expected to hold when more realistic asset pricing, discrete rebalancing, a reduced investment set and transaction costs are introduced. A more complex addition is to include non-investment strategy (e.g. contribution rate strategy; benefit improvement strategy) into the optimisation problem. In the final subsection we therefore turn to the consideration some of these important practical aspects. To what extent does discrete (rather than continuous) rebalancing impact the investor's outcomes? Can the use of options help to mitigate this impact? And can simplified strategies lead to similar outcomes?

4.4 *Model*

4.4.1 The formal derivation of results described in Appendix C requires a lognormal model of asset prices. This framework is sufficiently rich to allow a wide range of behaviours to be captured — for example, multi-factor stochastic interest rates following a Hull–White or similar form can be readily incorporated. Robustness in the face of stochastic interest rates is clearly crucial for any analysis of a liability-relative situation to be taken seriously. But lognormality and the assumption that volatility is deterministic are at variance with the reality that markets exhibit both fat tails and volatility clustering. These results should therefore be treated as indicative of general principles rather than specific advice — the existence of fully worked out analytic solutions in a rich stochastic environment provides a high degree of comfort and intuition in designing practical dynamic strategies, but the precise strategies should not be followed slavishly. In this main part of the paper we therefore seek to develop themes and identify features that apply across the wide class of solutions and appear to use to be relevant more generally — we hope to have provided sufficient detail to allow readers to incorporate and test these themes within their own model environments.

4.4.2 To keep things transparent in the main part of the paper, we have used an extremely simple background model: fixed interest rates and a risk asset portfolio with a Sharpe ratio (market price of risk) of 0.33. We have assumed a core relative risk aversion of 0.8. The assumptions are summarised in Table 2.

4.5 *Introducing a Minimum Solvency Level*

4.5.1 The standard CRRA utility function (20) is a power law at all levels of wealth. The 'constant relative risk aversion' name reflects this power law structure: the risk/reward trade off looks identical, in percentage terms, at all levels of wealth. It is therefore reasonable that Merton (1969) was able

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

Table 2. Parameters used to generate graphs in this section. See equation (22) in Appendix C for an explanation of the parameters

Parameter	Value
μ	3.8%
σ	15%
γ	0.8
r	5%

to show that the optimal portfolio strategy, in a continuous-time framework, is independent, in percentage terms, at all wealth levels.

4.5.2 It is reasonable to dispute this constancy of relative risk aversion, and as a first adjustment we suppose that the investor has a minimum income level M. We have in mind an investor planning for retirement income, so this could be considered as a stream of payments payable after the investor's investment horizon time T. (This is different to Merton's context, where the horizon is thought of as the investor's death, and the income is payable as portfolio drawdown in the period up to the time horizon.) To keep the mathematics manageable, it is easier to model the income stream as a lump sum payable at $T' > T$, where $T' - T$ is an approximate duration of this income stream. This then leads to a minimum asset value $MP(t, T')$ at any time $t < T$.

4.5.3 We have investigated two obvious mechanisms for enforcing this minimum asset level. The first is to amend the utility function by 'shifting it to the right' by an amount $FP(T, T')$ (it is not a problem that this shift is stochastic). The second is to simply truncate the utility function at $FP(T, T')$ and set utility equal to $-\infty$ for asset values below this amount:

$$\begin{aligned} &\text{Shifted utility} && u_1(w_T) = u[w_T - MP(T, T')] \\ &\text{Truncated utility} && u_2(w_T) = \begin{cases} u(w_T) \text{ if } w_T \geq MP(T, T') \\ -\infty \text{ if } w_T < MP(T, T'). \end{cases} \end{aligned} \tag{8}$$

For clarity, these versions are illustrated in Figure 10.

4.6 *Impact of Minimum Solvency Level on Dynamic Strategy*

4.6.1 Given such a utility function, how would an investor, looking to maximise their expected utility at a chosen time horizon, behave? In broad terms, we would expect that for asset levels well in excess of the minimum, the minimum plays no role and so the optimal dynamic strategy is close to Merton's result of being independent of the asset level in percentage terms. But if the asset level is close to the desired floor, we would expect less risk to be taken, so as to avoid falling through the floor. This is indeed the case, but the ways that risk reduces in the two cases are quite different.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

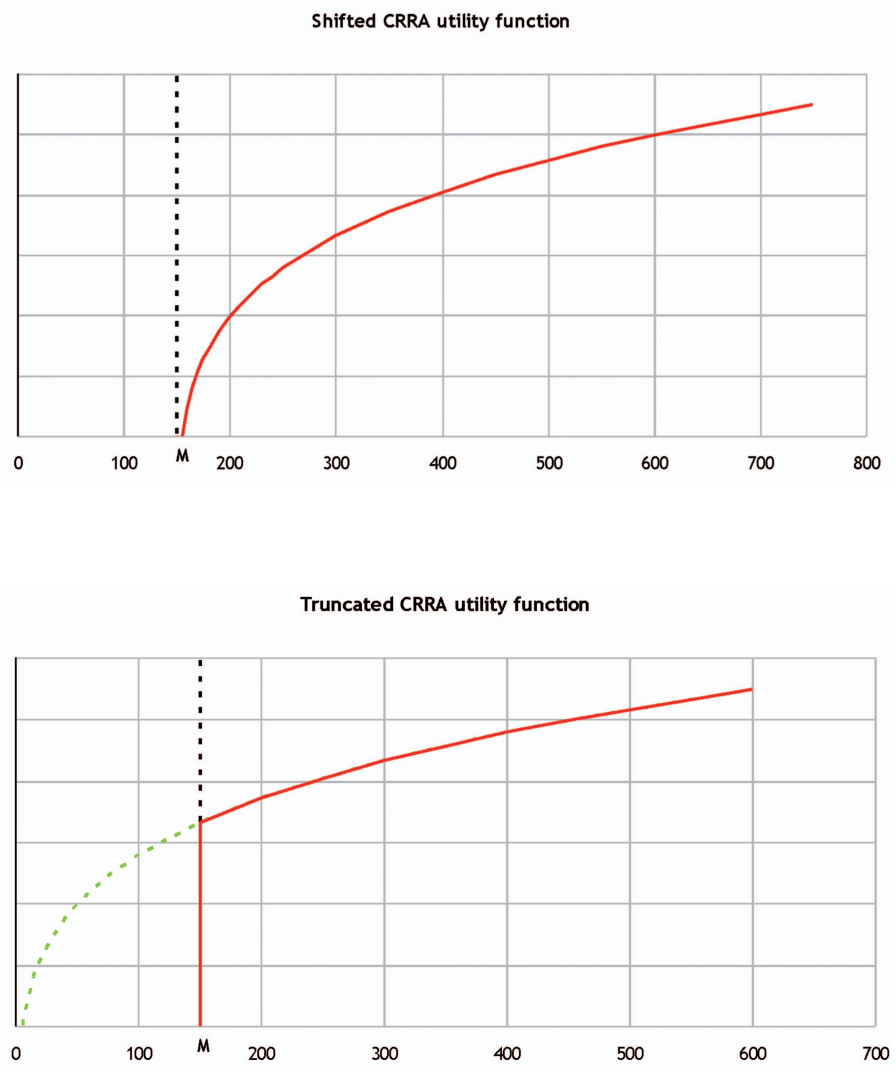


Figure 10. Two ways to enforce a minimum solvency level: shift and truncation for a constant minimum level $MP(T, T') = 150$. The green dotted line in the truncated case represents the 'base case' of a CRRA utility function

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

As described above, we can investigate the dynamic strategy from a number of perspectives.

4.6.2 In all the cases we are considering, the asset allocation strategy can be expressed as a dynamic mix of two portfolios: a risky portfolio and a liability hedge portfolio:

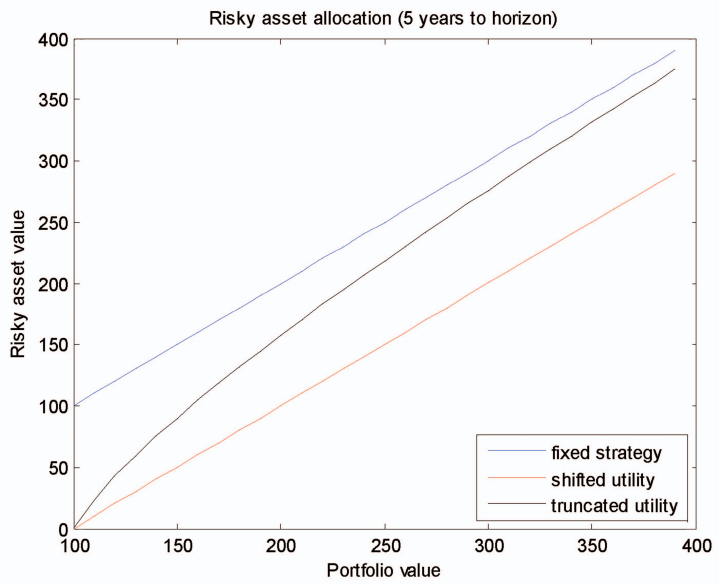
— The risky portfolio is the 'Merton portfolio' which an unconstrained CRRA investor would hold — it has fixed percentage allocations to its component assets, assumed to be continuously rebalanced. The portfolio contains a bond that matures at the investor's chosen time horizon but is otherwise independent of the time horizon. This portfolio can be thought of as a fixed mix of a 'growth-optimal' portfolio (a portfolio with highest possible Sharpe ratio) and this bond.
— The liability hedge portfolio, i.e. the bond$P(t, T')$, does not arise in the base CRRA case, but only appears due to adjustments such as (8).

4.6.3 The dynamic strategy is a function of the portfolio value and the time horizon — the first panel in Figure 11 shows a time-slice through these surfaces at a point 5 years from the time horizon. The value of the minimum income floor is 100 at that time, in this example. The three lines show:

— The base case. The short-dashed blue line shows the strategy for a CRRA investor. They hold all of their assets in the risky portfolio, and so this is a 45-degree line that goes through the origin;
— The shifted-utility case. It turns out that the strategy for this investor is also a straight line in this picture: the amount invested in the risky asset is equal to the buffer that the investor has over the floor. This is in fact a well-known strategy. The proportion in the risky asset portfolio, and thus the proportion in the growth-optimal portfolio, is a fixed multiple of the buffer. This strategy is known as 'Constant Proportion Portfolio Insurance' (CPPI) and is a classic method of portfolio insurance.
— The truncated-utility case. The strategy here is intermediate between the first two cases: the tolerance for risk is higher than in the CPPI strategy, but this falls rapidly to zero as the floor is approached. This strategy can be more easily understood by examining the second panel in Figure 11, which shows the final payoff from the 3 strategies, compared to the base case. The picture is a classic 'hockey stick' picture of an option payoff: this is a covered put. An alternative to holding a dynamic mix of the risky portfolio and the liability-hedge portfolio is to hold the risky portfolio together with a relative-performance put option, which protects the portfolio from falling below the value of the minimum income floor. Put option strategies are the main competitor to CPPI strategies in portfolio insurance and it is therefore pleasing to see this strategy emerge so naturally.

4.6.4 Finally, our third suggested method of investigating strategies is to

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press



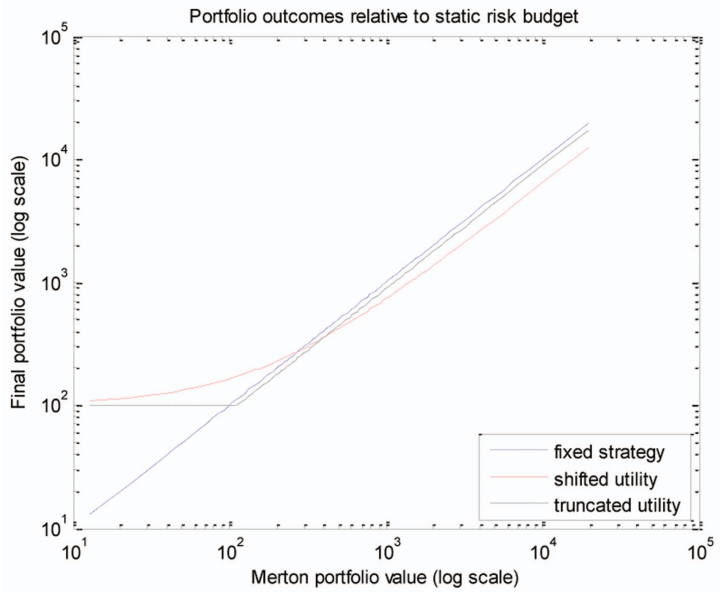


Figure 11. The impact on a dynamic strategy of a minimum income level

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

consider the distribution of outcomes. It is tempting to attempt to plot a probability density function of outcomes. However, as the truncated-utility function leads to a put option strategy, and thus a positive probability of exactly hitting the income floor. Figure 12 shows two ways to avoid this pictorial problem: the cumulative probability distribution function (where a point mass simply appears as a vertical segment in the chart) and a histogram, which is just a discrete version of the density function.

4.6.5 These plots demonstrate the same features that were present in Figure 11, overlaid with the probability of particular outcomes. The positive probability of hitting the floor is clearly seen in the truncated-utility function: the distribution of outcomes is truncated as well as the utility function. Similarly, the shifted utility function leads to the left-hand tail of the outcome distribution being shifted to the right, beyond the desired floor. The necessary penalty for this is a lower likelihood of very large positive outcomes.

4.7 *Introducing a Target Solvency Level*

4.7.1 In addition to a minimum income floor, it is also important to incorporate the idea of a target income level. This is done by introducing satiation: more wealth is good up until a target is reached, but beyond this, additional wealth generates no additional utility.

4.7.2 This is illustrated in Figure 13, which can be compared to the methods used to incorporate a minimum (see (8) and Figure 10 above).

4.7.3 In an equivalent way to the discussion of a minimum solvency level, here we have in mind a stylised target liability $LP(t, T')$ payable at $T' > T$. The adjustment to the utility function is then straightforward:

$$\text{Target utility} \quad u_3(w_T) = \begin{cases} u(w_T) \text{ if } w_T \leq LP(T, T') \\ u(LP(T, T')) \text{ if } w_T > LP(T, T') \end{cases}. \tag{9}$$

4.7.4 Again, we can usefully investigate the resulting dynamic strategy from a number of interlocking perspectives.

4.7.5 Figure 14 is equivalent to Figure 2 above in the minimum-income case. The upper graph shows the amount invested in the risky asset portfolio as a function of the total portfolio value — it can be seen that the amount invested is always less than in the CRRA base case. This is intuitive: the potential reward for risk-taking is limited on the upside, so risk-taking is less attractive. This effect becomes more pronounced the nearer the investor is to the target liability level. Finally, when the asset level is at (or above) the target level, there is no benefit from risk-taking and so the investor simply holds the liability hedge portfolio.

4.7.6 The lower graph of Figure 14 shows the final outcome from following the dynamic strategy of which the upper graph is only a slice. The

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

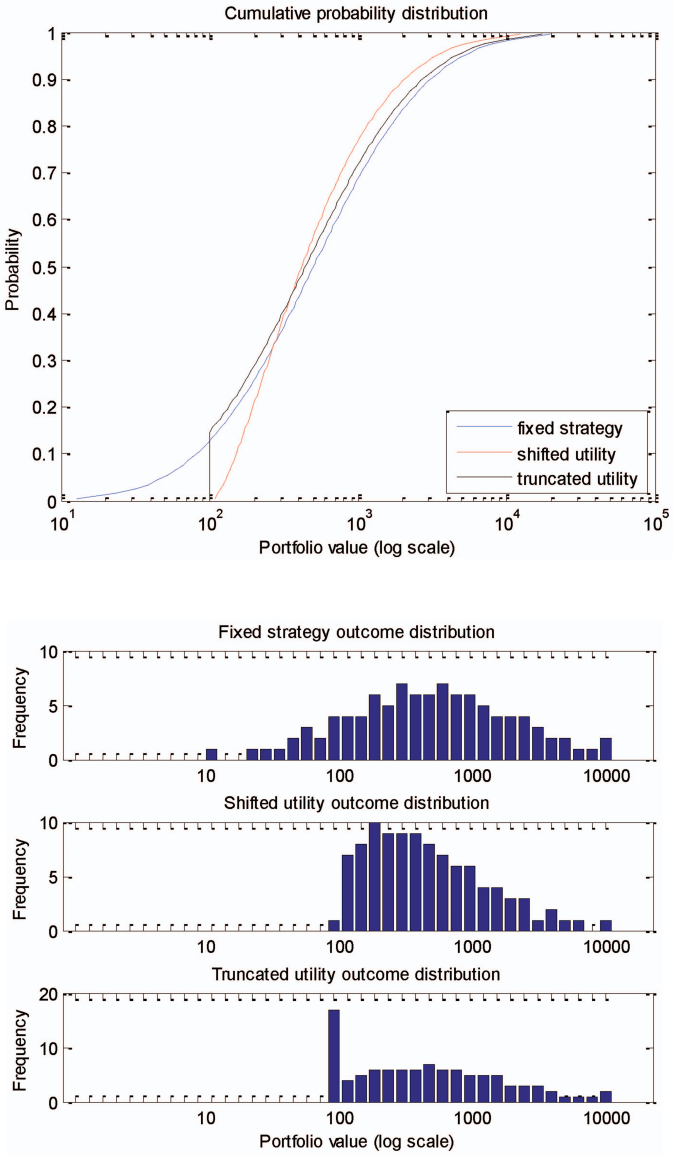


Figure 12. Distribution of outcomes for the minimum-solvency utility functions compared to the base case ('fixed strategy'). Plots assume an income floor of with a value of 100 at the time horizon, a horizon of 12 years and an initial asset level equal to 3 times the value of the floor

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

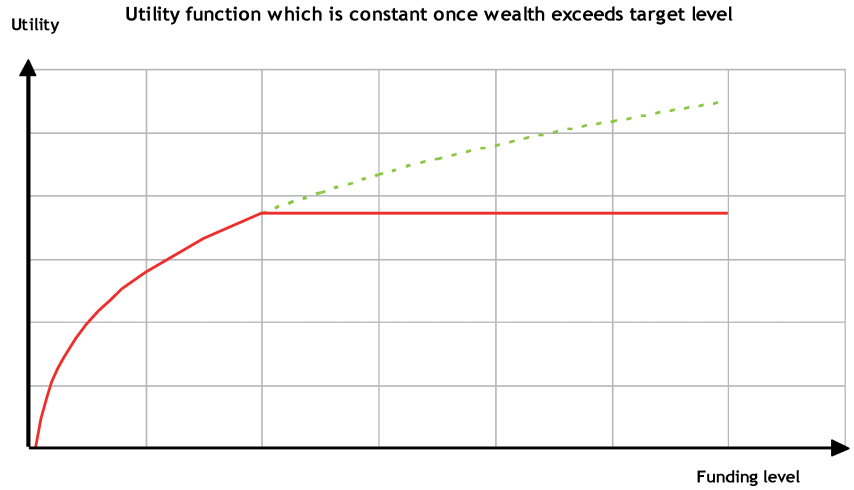


Figure 13. Utility function with a target solvency level

structure is that of a covered call: an equivalent static strategy is to hold the risky portfolio together and write calls that sell away the upside for which there is no additional utility.

4.7.7 Figure 15 shows the dynamic strategy in more detail. The upper graph recasts the left-hand panel of Figure 14 in percentage terms: the percentage falls steadily to zero as the funding level improves. This is just the 'five years from the time horizon' slice of the full dynamic strategy, which is then shown in the lower graph. It is worth noting a few features:

— For a given funding level, it can be seen that the proportion in risky asset increases as the time horizon shortens: the risk-reducing impact of the utility cap becomes less, since the likelihood of reaching the liability target falls over time.
— However, risk does not increase in an unbounded way: the 100% risky-asset proportion remains as a ceiling. The explanation for this is that there is a level of risk that the investor is generally happy to take, expressed in the relative risk aversion parameter of the (base) utility function.

4.7.8 This shape of dynamic strategy surface is a familiar consequence of the covered-call structure of the payoff. As mentioned earlier, it is simpler in practice for clients to focus on this payoff rather than on the utility function. Figure 16 shows the distribution of outcomes as a probability distribution — again, this is a simple consequence of the structure but can be an easier point of focus for an investor attempting to trade off different parts of this distribution.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

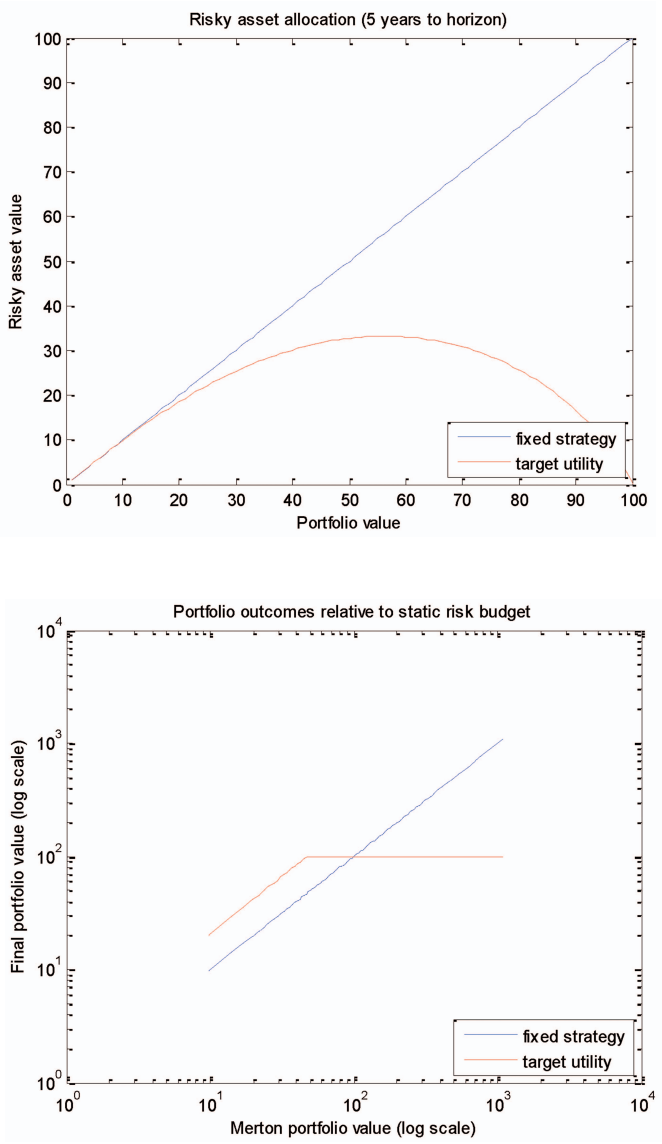
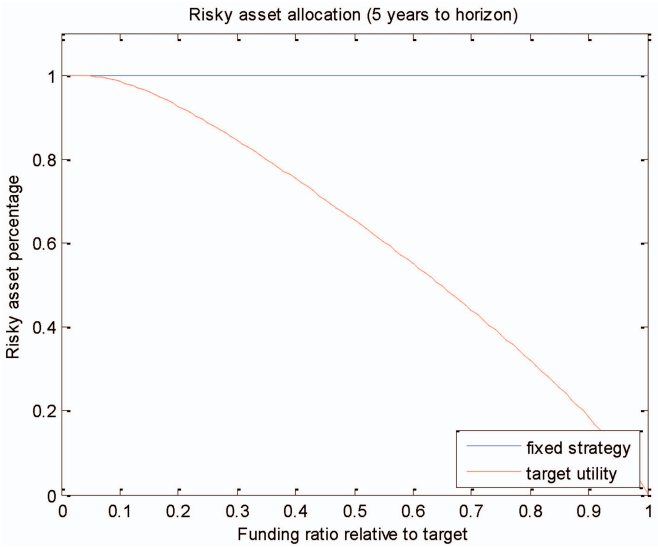


Figure 14. The impact on a dynamic strategy of a target income level. Five years from the time horizon, the liability value is 100. The right hand panel shows the outcomes over a 5 year horizon assuming an initial funding level of 80% (i.e. asset value of 80)

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press



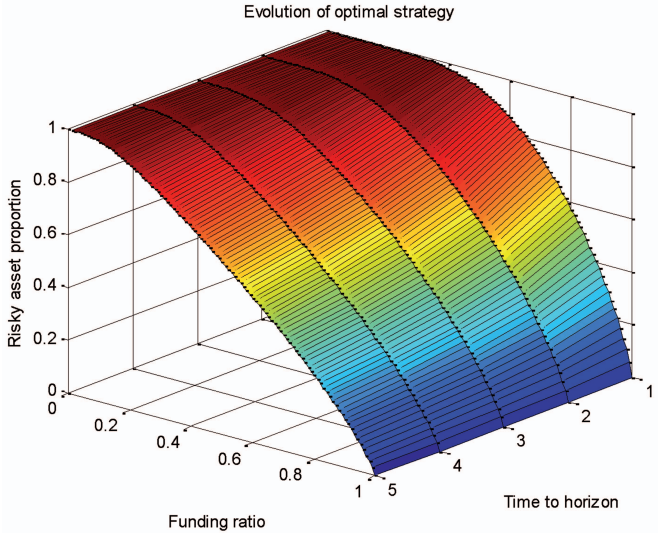


Figure 15. The dynamic strategy as a proportion of assets. Same assumptions as Figure 14

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

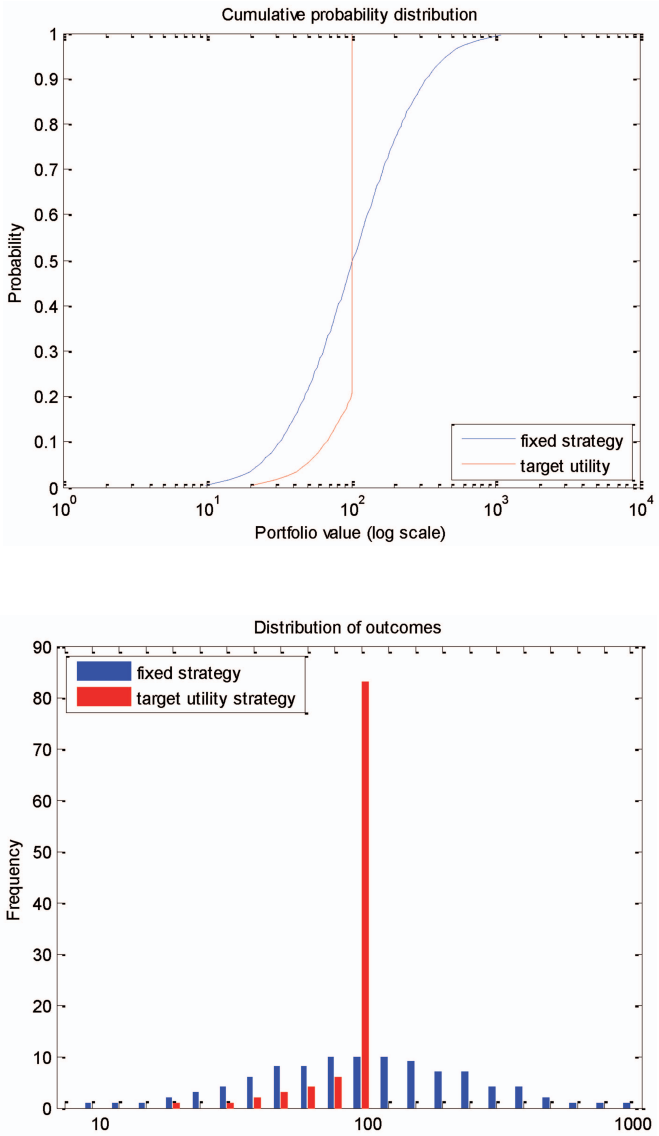


Figure 16. Distribution of outcomes for the target-income utility functions compared to the base case ('fixed strategy'). Plots assume a liability target of with a value of 100 at the time horizon, a horizon of 5 years and an initial funding ratio of 80%

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

4.7.9 The trade off is very clear in this case: potential upside is foregone in return for increased certainty of meeting the liability target. If the utility function correctly expresses the investor's objectives, then this is a simple decision: the high returns foregone would not be appreciated while the pain of not meeting the target would be keenly felt.

4.7.10 The results for a target-liability objective, and in particular the option-like structure of the optimal strategy, appear superficially very similar to the minimum-solvency problem discussed above. However, as a final point we mention that these are different in character. In the earlier case, the investor starts above the minimum and then seeks a strategy that will fully ensure that this remains the case. In this case, the target cannot be achieved with certainty and therefore a strategy is sought that maximises the likelihood of meeting the objective, while controlling the risk of falling short, not removing the risk.

4.8 *Combined Utility Functions*

4.8.1 In practice, there may be both a target income level and a minimum required income for the investor. This would suggest a combination of the two approaches, and this subsection briefly outlines two examples of such combined approaches. Appendix C shows that such combinations introduce no new essential complexity to the problem.

The first example is shown in Figure 17. This is a combination of the truncated utility function and the target utility functions discussed above.

4.8.2 Figure 18 summarises the optimal dynamic strategy and resulting

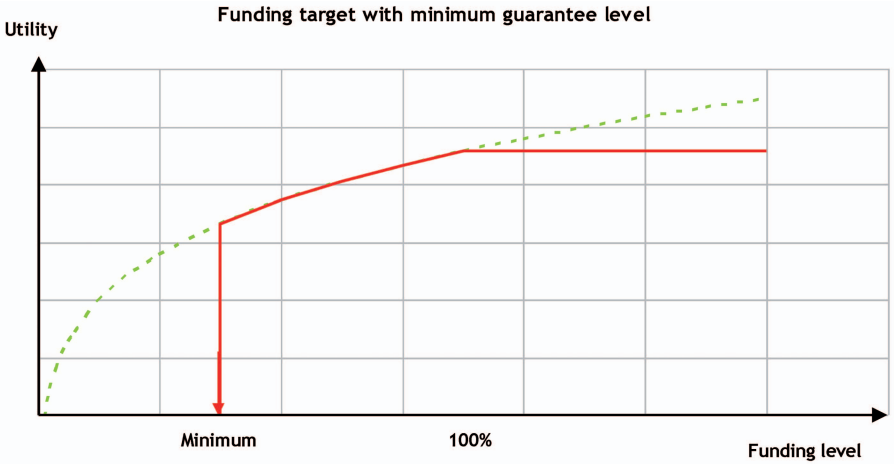
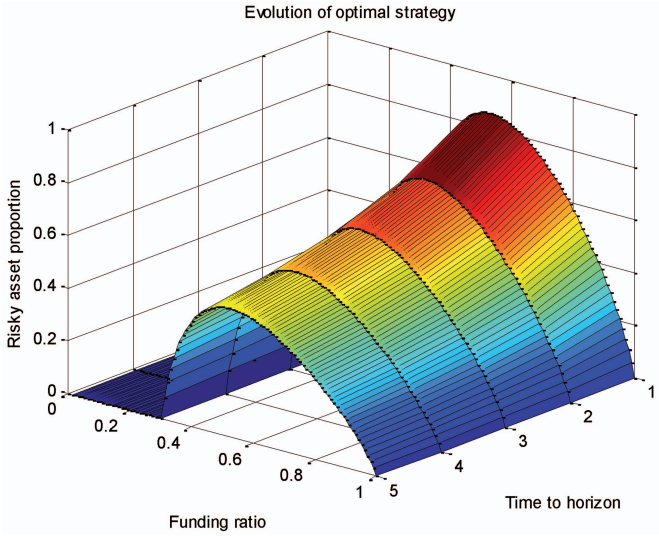


Figure 17. A utility function expressing a minimum income level and a target income level

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press



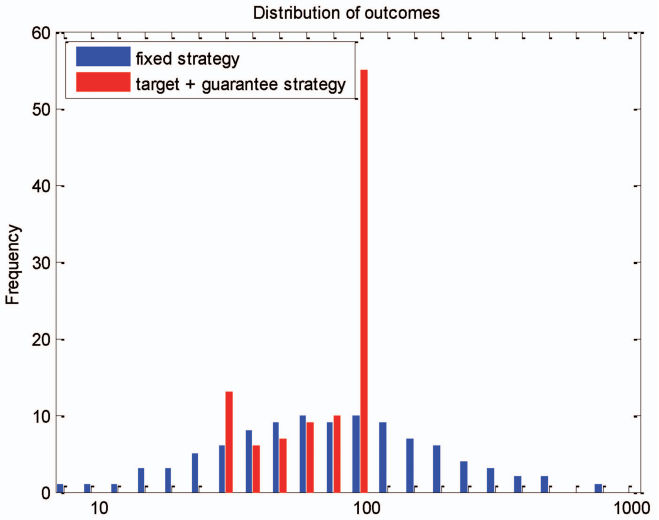


Figure 18. Dynamic strategy and outcomes for an investor with a utility function of the form illustrated in Figure 17. The minimum income level is 30% of the target liability. The outcomes are shown for an investor with assets equal to 60% of the target initially, and a five year horizon

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

distribution of outcomes. The payoff (not shown) has a collar strategy: equivalent to a holding in the risky asset portfolio combined with puts to protect the downside and calls to give away the unwanted upside. The dynamic strategy involves taking less risk as either of the two boundaries is approached, with more risk taken between these two zones. Based on this structure, the distribution of final outcomes is immediate: it is bimodal, with positive probabilities of hitting either boundary, and a straightforward density in between these limits.

4.8.3 The second example is somewhat different, and illustrated in Figure 19. This incorporates a target income level, as above but now there is no loss in utility if income falls below a lower level.

4.8.4 The idea behind this structure is that it might reflect a means-tested environment, where if income falls below a particular level, then external (government) assistance becomes payable, bringing income back up to the minimum level. In a pension environment, the existence of an insurer such as the PBGC in the US or PPF in the UK may mean that members' benefits are similarly protected from asset performance below a certain level.

4.8.5 This utility function has the somewhat worrying feature that it is not convex, as standard economic theory would normally suggest. (Such economic theory doesn't normally allow for means testing.) However, from a

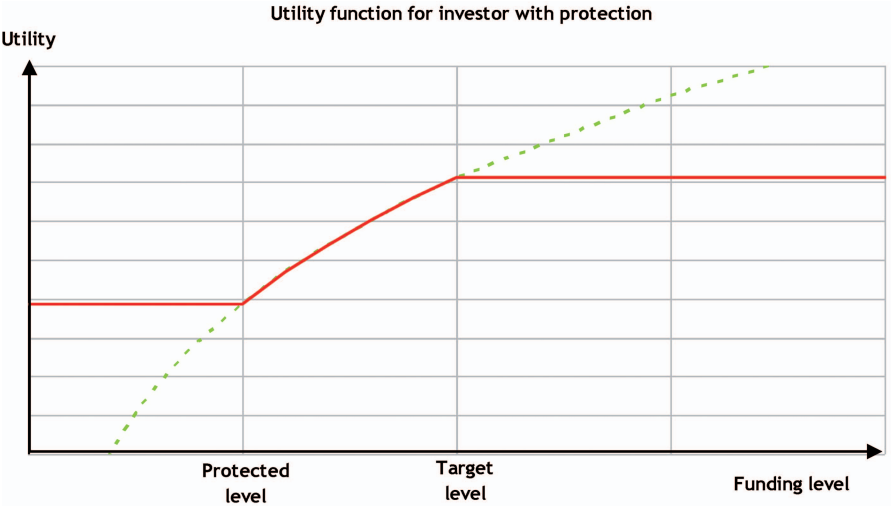


Figure 19. A utility function expressing a protected-income level and a target income level

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

mathematical point of view, the problem of maximising expected utility can still be stated, and again, Appendix C outlines the solution. The lack of convexity is important here too — and the 'convexified' form of the utility function (the surface of the convex hull of the points below the graph in Figure 17) is all that is required to allow the solution to proceed.

4.8.6 The optimal strategy is now also worrying — if neither surprising nor unfamiliar for those who work with pension plans. The optimal strategy is now akin to 'double or quits'. For low funding levels, more and more risk is taken — much more than in the base case of an investor with a constant level of relative risk aversion. In direct contrast to the behaviour near the target, where potential upside was unrewarded, now potential downside is painless and so the level of risk is unbounded. This strategy of 'gaming' whoever is providing the protection is unappealing from a systemic perspective, but does help to explain some more extreme forms of investor behaviour.

4.8.7 Of course, for many pension plans, the floor does involve pain: the level of benefits may be materially less than those intended under the Plan rules; trustees may be nervous of legal action in following such a strategy; or they may feel morally bound to secure Plan benefits without recourse to the State. In these circumstances, the floor may 'harden' and a utility function such as those discussed in Section 4.5 above would be more appropriate. As we have seen this leads to less and less risk being taken as the floor is approached, rather than more and more.

4.8.8 The outcomes under the optimal strategy (shown in Figure 20) are almost binary: either the target is hit (with fairly high probability) or all the assets are lost. Indeed the optimal payoff incorporates a binary option of this kind (see Appendix C).

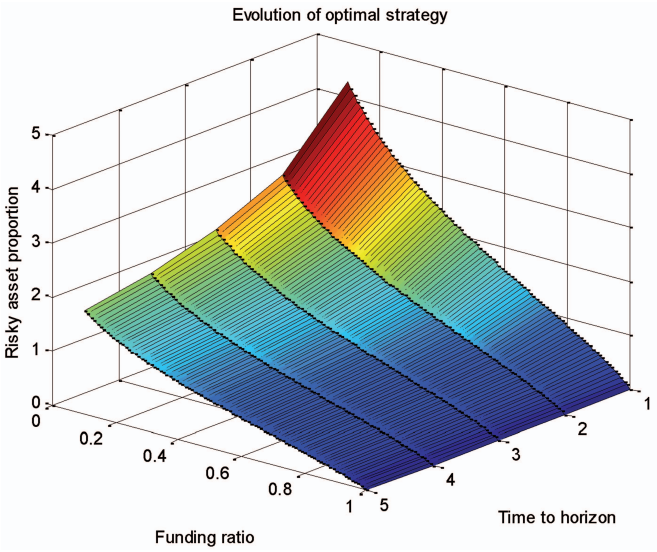
4.9 *Practical Aspects*

The discussion so far has relied upon an idealised situation, where assets are assumed to be rebalanced continuously and without transaction cost. Of course we don't believe these assumptions, any more than we believed the assumed market dynamics accurately reflected real-world asset moves. Making these assumptions allowed us to make progress and derive potential strategies — but we must remember that a key challenge now is to assess how robust these strategies are to real world conditions. We will not undertake this fully here, but instead indicate alterations that can be useful in practice. We focus on the 'target liability' model described in Section 4.7 but similar considerations apply elsewhere.

4.9.1 *Discrete rebalancing*

4.9.1.1 We can follow the Bellman approach to establish the utility-optimal strategy to follow if we reflect the fact that rebalancing is only possible at discrete points rather than continuously. The idea here is that the

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press



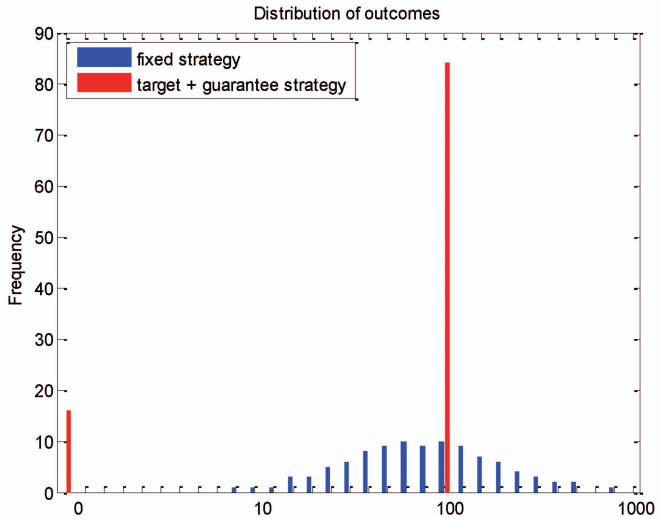


Figure 20. Dynamic strategy and outcomes for an investor with a utility function of the form illustrated in Figure 19. The protected income level is 30% of the target liability. The outcomes are shown for an investor with assets equal to 60% of the target initially, and a five year horizon

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

funding level relative to the target income level could be monitored at regular intervals (e.g. daily, monthly, annually) and the strategy rebalanced accordingly at these intervals. Figure 21 shows the result of reviewing the strategic portfolio weights annually over a five year period.

4.9.1.2 The upper graph of Figure 21 shows the expected utility that would result from this annual rebalancing strategy. It can be seen that the utility is (of course!) less than would be achievable with a continuously-rebalanced strategy, but that it is already a significant improvement on not rebalancing the strategic portfolio weights at all over this five year period. The expected utility score can not be used directly for quantifying this gain, as the same preferences are expressed if we multiply the utility function by a positive constant or add a constant to it. One common approach is to express the result of following a strategy as a certainty-equivalent. Alternatively different strategies can be compared by seeing what funding level would be required to achieve the same average-utility outcome. Thus, from the graph in Figure 21 it can be seen that for an investor with an initial funding ratio of 70%, an annual-review policy does about as well as a high-frequency-review investor would achieve with an initial funding level of 67% and a no-review investor with an initial funding ratio of about 78%.

4.9.1.3 As described earlier, utility provides a useful metric for expressing the investor's preferences for potential outcomes, but looking at the full distribution of outcomes is also useful and avoids the messy issue of calibrating the utility function. The lower graph of Figure 21 seeks to do this. It shows the result of 1000 simulated outcomes from following four strategies: a high-frequency (continuous) review, an annual review where expected utility is optimised, an annual review where the portfolio weights are set so that the average outcome is expected to exactly hit the liability target, and a no-review comparison. In addition, all strategies except for the (somewhat theoretical) high-frequency review strategy are further constrained by not allowing portfolio leverage (a reasonable and common imposition for an investor not able to rebalance very frequently).

4.9.1.4 As expected the continuously-rebalanced strategy is extremely focused on the target outcome and does not overshoot. Conversely, failure to rebalance leads to a very wide distribution of possible outcomes. However, it is striking — but reassuring — how well the simple target-return rebalancing rule does. It is slightly more likely to over- or under-shoot the target compared to the rule based on optimal utility. One feature which is less visible is that the left-hand tail of the continuously-rebalanced strategy is slightly fatter than the annual rules. This is due to the freedom allowed to this rule to use leverage — the tails of both the annual-review strategies are more controlled than the no-review strategy.

4.9.2 *One-way versus two-way rebalancing*

4.9.2.1 If the funding level falls, then the dynamic strategies discussed

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

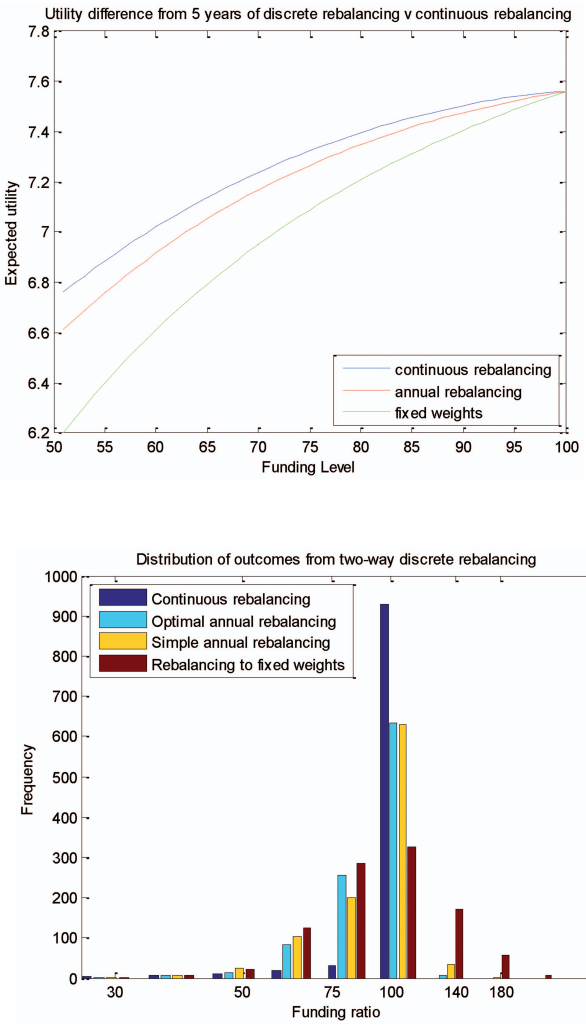


Figure 21. Impact of following an annual rather than continuous rebalancing strategy over a 5 year period. Model assumptions are as in Section 4.4. The upper graph shows expected utility and the lower graph the distribution of 1000 simulated outcomes with an initial funding ratio of 90%. The result of following a fixed-weight strategy is also shown as well as (on the right) the impact of following a simple target-return strategy rather than a utility-optimised strategy. See further description in the text

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

above will increase strategic risk. However, this may feel unattractive for an investor seeking to meet a target income. In practice there may be other levers available — the key ones being additional contributions and moving the target horizon. The determination of a dynamic strategy therefore needs to cover the extent to which contributions and timescale will be reviewed in light of realised performance in addition to the asset allocation. In our experience, the typical approach is to change these additional controls relatively infrequently (e.g. annually or on a triennial cycle) whereas more frequent rebalancing in line with the investor's goals is both possible and desirable. Such strategies can be simulated in the same way, and the volatility of contributions observed to be lower in many cases.

4.9.2.2 However, one aspect of these additional levers is that it can become less desirable for an investor to increase risk if the investment return outcomes are poor. The objective for these investors is to have a dynamic derisking strategy — to take advantage of opportunities to lock in gains if they arise, but otherwise use a combination of the other levers in poor outcomes.

4.9.2.3 The impact of following such a strategy is shown in Figure 22. The resulting distribution is similar in the one-way and two-way rebalancing

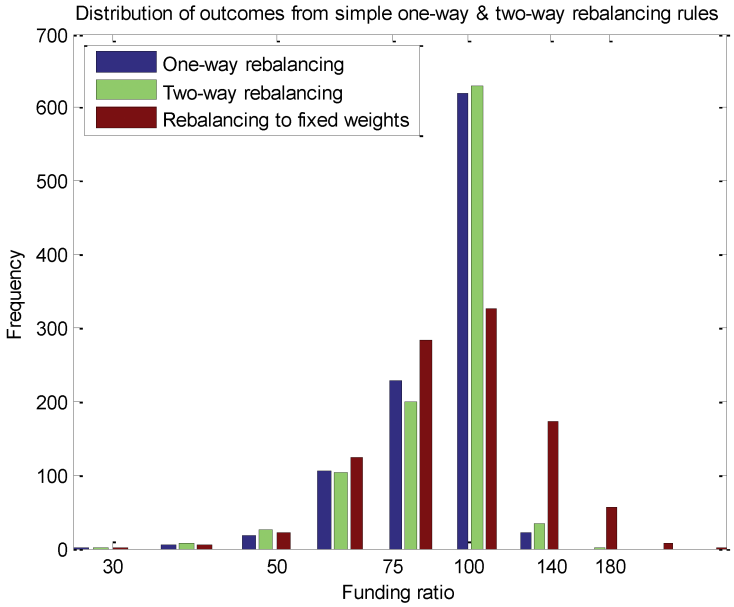


Figure 22. Impact of one-way (derisking only) rebalancing strategy compared to a strategy that both de-risks and re-risks

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

alternatives: the likelihood of hitting the target remains high, and the likelihood of poor outcomes remains low, but there is a slightly higher risk of missing the target by a modest amount. This is consistent with the aims of this kind of strategy.

4.9.2.4 There is of course a lot of potential variety in strategies that use a more realistic combination of levers to meet an investor's goals. Although a formal trade off between contribution volatility, time horizon slippage and missing the goal can be set up, we have found it more productive to pick strategies that are intuitive and take sensible account of the amount of time an investor is able to invest in review as the target nears. The strategies discussed here provide a cookbook for developing internally-coherent strategies — some judgement is useful to determine which ones should be applied in a given situation.

4.9.3 *Use of options*

4.9.3.1 As described above, the utility-optimal dynamic strategy in the target-liability case amounts to delta-hedging a relative performance option. The option is based on the performance of a growth-optimal (and so liability-independent) portfolio relative to a liability hedge portfolio. This naturally suggests an alternative to the dynamic strategy, namely the static strategy of just holding the option. This raises a number of issues, for example:

4.9.3.2 Who should rebalance? In theory, buying an option and following a dynamic strategy should lead to identical outcomes. In one case, the end-investor performs the dynamic rebalancing; in the other, the option-writer undertakes dynamic hedge management. This is therefore just an outsourcing question: if an investor lacks the capability or scale to rebalance efficiently, it should be relatively more attractive to use options.

4.9.3.3 Is the option price reasonable? An option premium will normally exceed a model price, particularly if the model ignores transaction costs and the possibility of extreme price moves. In addition, the pricing is dependent on market liquidity and some option markets are more liquid than others. An option on two underlying portfolios, where one is a diversified basket and the other the value of a long-dated stream of payments, and where the tenor of the option is several years long, is unlikely to be liquid. Nonetheless, other options can still be useful and it therefore makes sense to focus on liquid markets where competitive pricing can be realised.

4.9.3.4 How would an option interact with other investor needs? It was described in the previous section how changes in contribution level and timescale adjustments are often made at infrequent reviews set some time ahead. It may make sense to use options that expire at a time close to these review dates, in order to manage event risks associated with these dates and thus enable a smoother strategy.

4.9.3.5 Figure 23 shows the potential impact of using options in a very simple example. Here, as before, rebalancing is assumed to take place

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

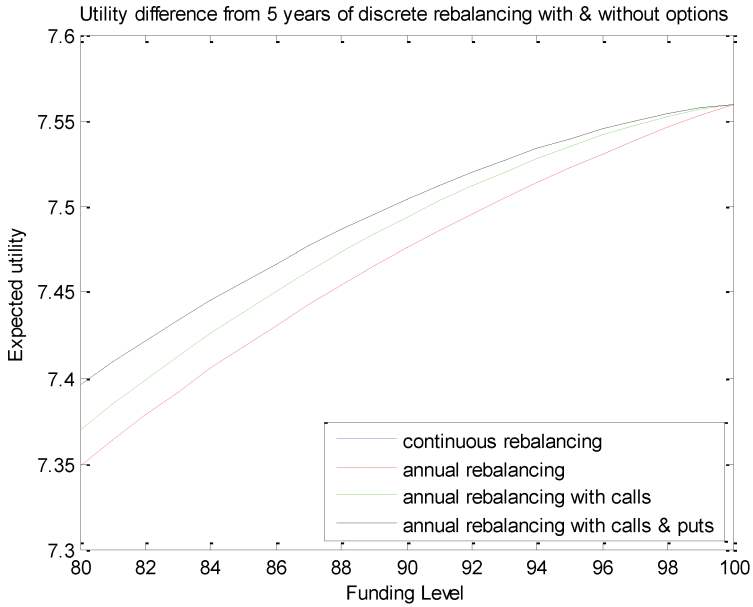


Figure 23. Impact of including call options on expected utility. Same assumptions as Figure 21, except that calls struck at the liability target. The impact of including puts is also shown (but the line is coincident with the continuously-rebalanced strategy at this scale)

annually — rather than using options that expire at the chosen time horizon, options are chosen that expire at each annual review. Call options are sold to minimise the possibility of the funding level exceeding the target liability at the review date. These calls are 'covered' in the sense the option underlying is already held within the portfolio. The additional premium then acts to increase the funding level in all scenarios (whether or not the option is exercised). It can be seen that a further significant increase in utility is achieved. This is most significant when the funding level is already close to 100% — essentially this allows more targeted risk to be taken in these cases. When the funding level is some way below 100%, the likelihood of exceeding 100% by the end of the year is much lower and so the premium & benefit from the option strategy are small.

4.9.3.6 Similar results hold when a minimum income level is sought as in Section 4.6. Without put options, the presence of a floor would force an investor who is only able to review relatively infrequently to scale back their risk budget significantly. Using put options, again with annual expiries for example, allows much more upside risk to be taken without putting the floor

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

in peril. Indeed, at any funding level, using put options allows a material increase in the risk budget to be achieved (effectively, leveraging the portfolio). Figure 23 also shows that in the target-liability case, the effect of this leverage is to achieve a distribution of outcomes that is very close to the theoretical optimum.

4.9.3.7 Overall, therefore, the targeted use of options can help an investor to meet their goals, as expressed by the utility function.

5. Risk-Constrained Utility Optimisation

5.1 *Introduction*

5.1.1 As mentioned above, tail risk constraints can be used in addition to, or instead of, utility functions as a way to control the losses in a portfolio. In the next section, we look at a mean-tail-risk optimisation problem, and a numerical approach to its solution. As an intermediate example, we first look at a utility-tail-risk optimisation problem. This turns out to be a straightforward addition to the utility approach of the previous section.

5.1.2 Further details on the results outlined in this section can be found in Gandy (2005) and Gabih *et al.* (2004).

5.2 *Statement of the Problem*

5.2.1 The objective is to find a strategy that maximises utility subject to the CVaR risk being below a limit:

$$\begin{array}{c} \text{Max } \mathbb{E}(U(X_T)) \\ CVaR_\varepsilon(X_T) \leq -c \end{array}. \tag{10}$$

5.2.2 We assume a geometric Brownian motion model for returns on assets, i.e. the total return of the i'th asset evolves according to a stochastic differential equation of the form

$$S_i(t) = S_i(0)\exp\left(\left(\mu_i - \frac{1}{2}\Sigma_{ii}\right)t + A_i' W(t)\right).$$

5.2.3 Here $W(t)$ is a multivariate Brownian motion; A_i is a vector containing the (constant) exposures of asset i to each driver; and $\Sigma_{ij} = A_i' A_j$ is the covariance of assets i and j. To avoid arbitrage opportunities, the expected returns μ_i must be of the form

$$\mu_i = r + A_i'\lambda. \tag{11}$$

5.2.4 Here, $\lambda 0$ is a vector each of whose components is the market price of risk for the underlying drivers; and r is the (again, assumed constant) rate

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

of interest on cash. The pricing kernel $m(t) = \exp\left(-rt - \lambda'\lambda t/2 - \lambda' W\right)$ is a key object of interest: the product of this kernel and any total return index will be a martingale (under the real-world measure). The existence of this pricing kernel, i.e. the condition (11), is thus a familiar one on the existence of an equivalent martingale measure.

5.3 *Optimal Payoff*

5.3.1 The model as described is dynamically complete, i.e. any payoff at time T can be replicated using a suitable dynamic hedging strategy. This therefore allows us again to use the martingale approach (rather than the more complicated Bellman approach) to solving the optimisation problem (10). We can first derive the optimal payoff X_T and subsequently determine the investment strategy that supports it.

5.3.2 This is straightforward using standard variational calculus: the optimal payoff is of the form

$$X_T = \begin{cases} (\lambda_1 m_T)^{-1/\gamma} & \text{if} \quad h < (\lambda_1 m_T)^{-1/\gamma} \\ h & \text{if} \quad (\lambda_1 m_T)^{-1/\gamma} < h < (\lambda_1 m_T - \lambda_2)^{-1/\gamma} \\ (\lambda_1 m_T - \lambda_2)^{-1/\gamma} & \text{if} \quad h > (\lambda_1 m_T - \lambda_2)^{-1/\gamma} \end{cases}.$$

5.3.3 Here the portfolio value h is attained with positive probability — this acts as a threshold value, with different behaviour of the payoff above and below this value.

5.3.4 For clarity, in the absence of the risk constraint, the optimal payoff profile would simply be $X_T = (\xi m_T)^{-1/\gamma}$ for some constant ξ. These payoffs are all functions of the so-called 'growth-optimal' portfolio (GOP) $G_t = m_t^{-1}$ at time T. The typical structure of the payoff from the constrained problem (10) is illustrated in Figures 24 and 25. (Special cases can arise from some constraints: if the constraint does not bite then $\lambda_2 = 0$ and the unconstrained-optimum applies; in other cases we can have $X_T \geq h$ everywhere or even $X_T \in \{0, h\}$ everywhere. The optimal X_T is always non-decreasing with the value of the GOP.)

5.3.5 The constants λ_1, λ_2, h must satisfy the budget constraint $\mathbb{E}(m_T X_T) = X_0$ and the shortfall constraint $\mathbb{E}\left((h - X_T)^+\right) = \varepsilon(h - c)$. The remaining degree of freedom can be found by maximising the expected utility. Each of these constraints and operations are straightforward to carry out numerically.

5.4 *Investment Strategy Example*

5.4.1 The weights in each risky asset are always in the same relative proportions: these can for example be identified with the weights in the growth-optimal portfolio, or equivalently the weights in the unconstrained optimal portfolio. The allocations thus vary over time along a single

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

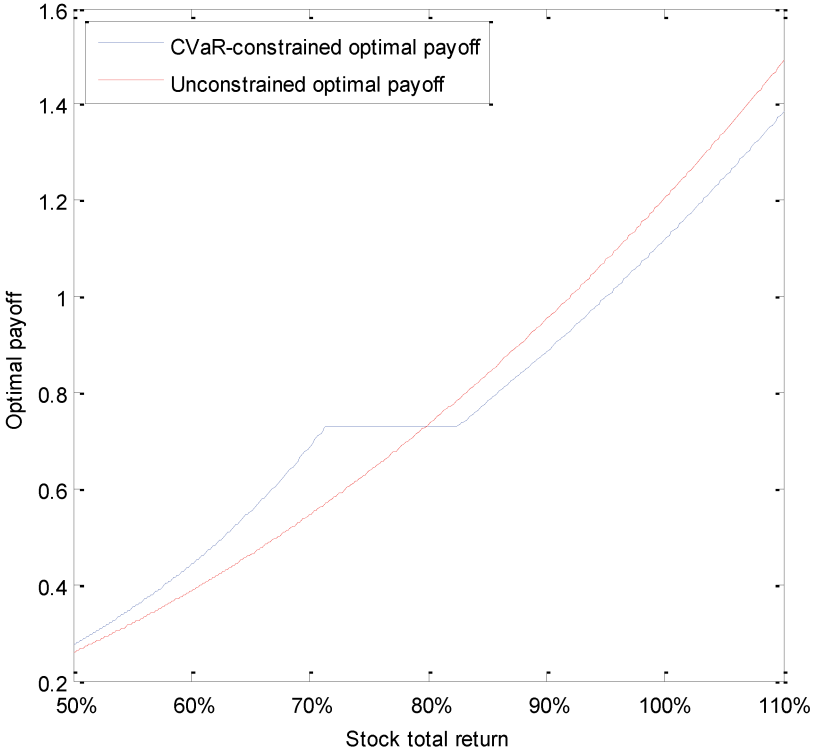


Figure 24. Impact of conditional value at risk constraint on optimal payoff at time horizon, one-asset case

dimension, namely the total exposure to risky assets can vary (with the balance made up with cash or bonds). The evolution of this exposure could equally be described as the variation of the risk budget, or return target of the portfolio over time. In the pictures below, we choose to show the total risky asset allocation as a percentage of the unconstrained asset allocation: where this percentage is above (below) one we are therefore taking more (less) risk than we would in the absence of a risk constraint.

5.4.2 There is therefore no additional complexity here with having more than one risky asset; so we keep things simple in the following example, which has one asset, with $\Sigma_{11} = (15\%)^2, \lambda = 1/3, r = 5\%$. The CVaR constraint we impose is $CVaR_{5\%} \leq -70\%$, i.e. the average loss for the worst 5% of scenarios can be at most 30% over a 1-year horizon. This is a fairly mild restriction, which allows the graphs below to be easily readable. (The tail VaR for the unconstrained optimal portfolio is $CVaR_{5\%} = -57\%$.)

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

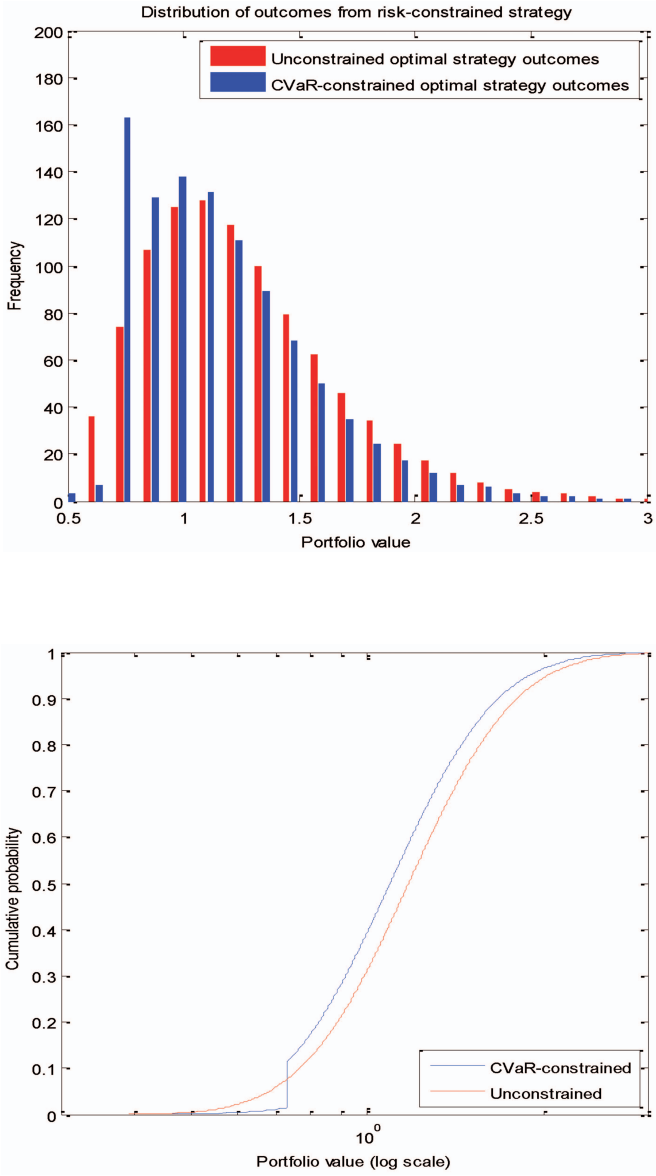


Figure 25. Impact of conditional value at risk constraint on distribution of outcomes: discrete form of probability density function (upper graph) and cumulative distribution function (lower graph)

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

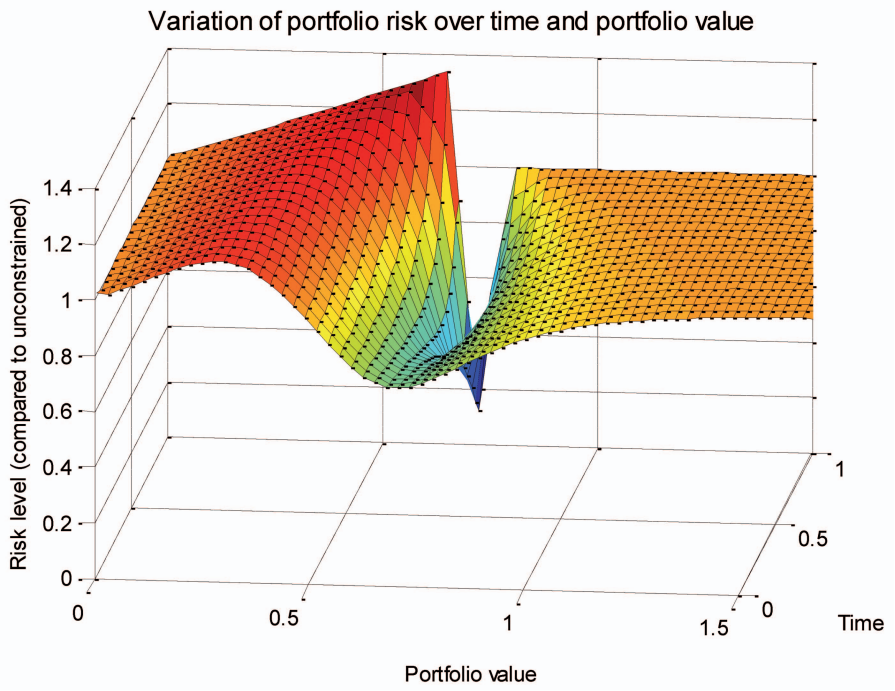


Figure 26. Optimal portfolio risk for problem (10), as function of portfolio value

5.4.3 The evolution of the strategic allocation over time is illustrated in the following charts.

5.4.4 Figure 26 shows how the risk level, as a proportion of that which would apply in the absence of a risk constraint, varies over time. Initially, when the portfolio value is 1, the optimal level of risk is below the unconstrained version. If the portfolio value stays around or above a threshold value, the risk remains below 1; but if the value falls too far below then it is optimal to take more risk. The highest possible level of risk is obtained if we are just below the threshold level very close to the time horizon. Figure 27 illustrates this from a different perspective: the time zero and time 1 slice is given at the bottom, to make the structure of the surface in Figure 26 clearer; and on the right we also give the equivalent surface if we plot against the unconstrained portfolio value rather than the constrained portfolio value. As the optimal constrained portfolio attains the threshold value h with a strictly positive probability (see Figure 24), this gives a slightly smoother picture.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

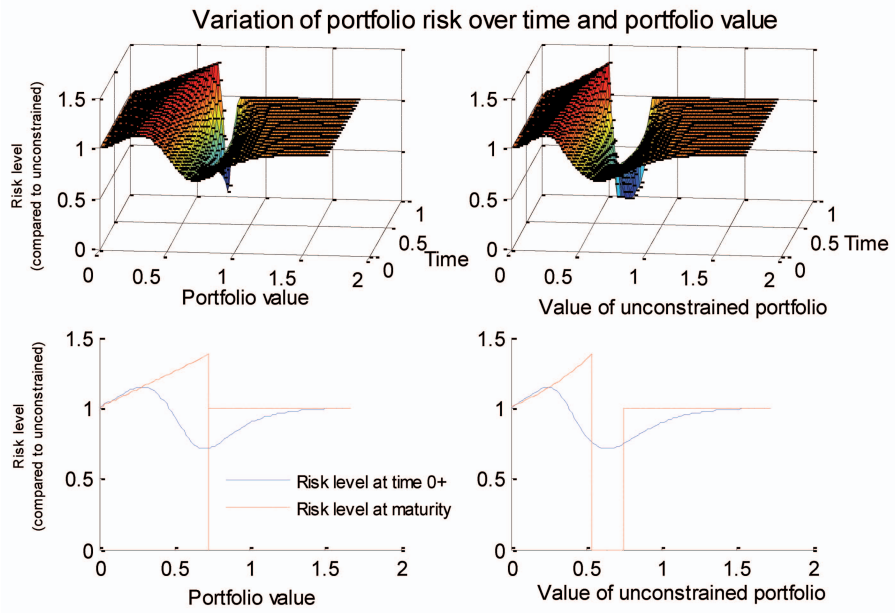


Figure 27. Optimal portfolio risk for problem (10), as function of unconstrained portfolio value

5.5 *Comments*

5.5.1 Clearly a TailVaR constraint is very different from a hard constraint that such as led to CPPI or put option strategies in the previous section. Controlling, but not eliminating, tail risk, is clearly cheaper, enabling better average outcomes, and may be attractive for most investors. Using a VaR constraint also exerts some control, but although this has an impact on the size of the tail, it is much less effective within the tail itself.

5.5.2 On the other hand, it is questionable whether the TailVaR constraint, as formulated here, is appropriate. It focuses explicitly on the outcomes as seen at the outset of the strategy. In contrast to the expected utility approach, which has a consistent interpretation through time, that is not the case here: if very poor returns are experienced, the strategy will tend to raise risk, in contrast to the strategies considered above. While this is sensible when viewed at the outset (this is likely to happen rarely), it may not be once the returns have occurred. A risk constraint which moves appropriately through time is required — remaining solvent with a 199 in 200 confidence level *at all times* is rather different from the requirement we have analysed here. We leave this to further research.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

6. Dynamic Stochastic Programming Approach

6.1 *Introduction*

6.1.1 In this section we turn to the problem of constructing a multi-period portfolio strategy that maximises return subject to a TailVaR constraint, long-only constraints and discrete times at which the portfolio is reviewed. The dynamic programming technique — the recursive stage-by-stage approach described in Section 3.2 — can in principle still be applied, but the state space would need to be expanded to take account of the TailVaR constraint, and the problem becomes somewhat unwieldy. Rather we choose to attack the problem more directly, using techniques set out in Sections 3.4 to 3.8. The approach falls under the general heading of stochastic programming, which at its simplest is merely a framework for looking at multi-period optimisation problems in the presence of uncertainty.

6.1.2 A general introduction to stochastic programming can be found in, for example, Rockafellar (2001); we have drawn much inspiration from Dempster *et al.* (2003).

6.2 *Overview of the Approach Used*

6.2.1 The earlier parts of the paper have relied on the existence of a utility function to describe the preferences of the investor. However, as individuals, we rarely know what our utility function is and as described in Section 4.2, in a collective investment such as a pension scheme or a company it is not always clear how the utility functions of individual members or shareholders should be aggregated to form the utility preferences of the institutional investor. Therefore it is often useful to formulate an investment approach that is not dependent on the existence of a utility function.

6.2.2 In contrast, institutional investors have become increasingly familiar with the use of various statistical measures of risk such as tracking error, VaR and TailVaR. We have therefore set out to optimise the trade off between portfolio growth and such risk measures rather than relying on the existence of a utility function. In the following we have used TailVaR (whose analytical properties are slightly safer from the perspective of an optimiser, as described in Section 2.6), but similar ideas could be applied to any risk measure (e.g. VaR) defined in terms of the distribution of portfolio outcomes. For ease of exposition, we have focussed our optimisation on the terminal wealth distribution but in principle this method could be extended to include risk metrics that also refer to intermediate distributions.

6.2.3 The other key distinction compared to earlier parts of the paper is that we proceed down a numerical rather than analytical route. The original Merton (1969) paper and the problems considered above were sufficiently stylised that an analytical solution was possible. The nature of many real-world problems, in particular the inclusion of constraints such as long only

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

investment usually force one to use a numerical algorithm, so it is important to have efficient procedures able to cope with these situations.

6.2.4 The discussion so far is likely to lead us to follow an approach similar to Dempster *et al.* (2003). They used a tree based approach which we significantly simplify as follows:

— Write down stochastic differential equations for the various asset classes.
— Divide the time interval up into a number of decision points when the portfolio would be rebalanced.
— For each asset class/economic variable used, discretise the possible outcomes.

So far we are following a similar approach to the Cox Rubinstein tree approach described in Section 3.5, with the difference that the tree is non-recombining and there are multiple asset/economic variables being modelled so the tree is multidimensional.

— Assume that various portfolio weights are applied to each asset at each decision point.
— Combine the portfolio weights and the asset class behaviour to calculate the portfolio value at any point on the tree. Together with the probability of each node being passed through, this allows the portfolio probability distribution at any point in time to be calculated.
— Finally, optimise the weights to produce a portfolio probability distribution that is optimal for the chosen risk and return objective.

6.2.5 We initially attempted to build a model with similar characteristics and found that the number of variables tends to increase exponentially and make the method intractable for the types of problem we wished to solve. More detail of the reasons for this are to be found in Appendix D.

6.2.6 We were therefore is a position of wishing to use a tree approach such as was set out in Dempster *et al*. (2003) while avoiding the exponential proliferation of variables. We used a number of techniques to control the numerical growth of the problem:

(i) use asset level SDEs that were consistent with tree recombination;
(ii) map the high dimensional tree in all the assets to a tree in the portfolio value (although the behaviour at each node still is dependant on all the assets);
(iii) introduced a 'local efficient frontier' to make decision at each node single-dimensional; and
(iv) replaced the tree-based approach that was focused on the evolution of asset classes with a lattice approach focused on the evolution of the probability density of the resulting portfolio values.

6.2.7 We briefly describe these in the next two sections (but refer to

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

Section 3 for a fuller overview of tree and lattice approaches commonly used in similar applications).

6.3 *Portfolio Value Tree and Local Efficient Frontier*

6.3.1 At first glance, for a non-recombining tree in all the assets, the number of dimensions of the tree grows exponentially with the number of asset classes — we need to separately model movements in each asset class, appropriately weighted to take account of co-dependency. However there are two observations that allow us to simplify the problem:

(1) A recombining tree will have fewer nodes than a non-recombining tree.
(2) Certain sets of economic events that will have the equivalent effects on the portfolio and so may be grouped together. For example, if the portfolio is equally weighted between US and European equities then, in the absence of transaction costs, we will be indifferent between US equities rising 10% while European equities stay flat or vice versa.

6.3.2 Since the strategy should be a function of the portfolio value at any stage, i.e. independent of which combination of asset returns led to this portfolio value, these two observations allow us to significantly reduce the dimensionality of the problem. We can replace a non-recombining tree in all assets classes by a recombining tree in the whole portfolio value. There is a single dimension to deal with in each time slice, and the number of nodes at a given point in time increases linearly rather than exponentially. This makes the problem much more tractable as we have reduced the problem to a two dimensional one in portfolio size and time. The price that we have pad for this is to require all assets to follow Geometric Brownian motion. The model therefore does not cope well with complex yield curve models.

6.3.3 This leaves us with something similar to a Cox–Rubinstein American option valuation on a recombining tree. The difference is that at each node we have a set of portfolio weights. The weights at that node enable us to calculate the local volatility and growth rate at that node. This is very similar to the local volatility models that have been used to provide option prices consistent with a volatility smile (e.g. Dupire (1994)).

6.3.4 Although this dramatically reduces the number of nodes, the size of the decision to be made at each node will still remain multi-dimensional (i.e. what weights do we invest in each asset class?). However, if the underlying drivers of return are normally distributed, we can make a further dimensional reduction. This is because in this case, the optimal weights along each path will be locally efficient, i.e. they will be on the mean-variance efficient frontier — if there is a portfolio with a higher mean and the same risk, this will be preferred. It therefore makes sense to parameterise the efficient frontier, by a parameter α say, and seek to find the optimal $\alpha(x, t)$ rather than the weight-vector $w(x, t)$. In the time-homogeneous set-up, we can determine global functions $\mu(\alpha)$, $\sigma(\alpha)$ for the corresponding (log) drift and volatility.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

6.4 *PDE and Lattice Approach*

6.4.1 Finally, we needed to solve some technical problems with the approach. As the portfolio weights vary with portfolio size and time, so does the portfolio volatility, producing a so called local volatility model. Such models have been used in option pricing in order to create an option pricing tree that is consistent with the well known volatility smile and skew. Such local volatility models, have the potential when used with a recombining tree to produce negative probabilities at various nodes. Clearly this aspect is aesthetically unappealing because we are trying to model the real world where probabilities are positive. There is also a more serious concern here in that negative node probabilities can destabilise the tree.

6.4.2 We were therefore in a position where we sought a method that retained the good properties of the recombining trees (discrete time, controlled growth in complexity as number of assets and time periods increase) but that also had more stable numerical properties. We decided to remove the negative probabilities problem by recharacterising the recombining tree as a lattice style solution to a partial differential equation (see Section 3.4.5). Lattice PDE methods are well known in finance — within option pricing theory they are known to provide a more controlled approach than trees. (See Duffy (2006) for a fuller discussion.) Characterising the problem as a lattice PDE problem allowed us to improve the stability properties by switching from explicit to implicit PDE solution methods.

6.4.3 The PDE in question is the Fokker–Planck equation (see Section 3.4.6 above) — this governs the evolution of the probability density of the portfolio value process. As described in Section 3.4.5, we approximate the various functions of the portfolio value (the strategy risk and the probability density function) by step functions on a tree within the lattice. We justify the use of the Fokker–Planck equation in Section 6.6.

6.5 *Summary of Method*

6.5.1 After these simplifications, we therefore end up with the following algorithm, starting from a set of parameters α or asset weights w on the nodes of a tree:

— approximate $\alpha(x, t)$ or $w(x, t)$ by a step function, defined by values on the nodes of a tree;
— this determines the corresponding drift and volatility functions, $\mu(x, t), \sigma(x, t)$;
— approximate the probability density function $p_X(x, t)$ from the Fokker–Planck equation (7) using a numerical scheme such as Crank–Nicolson, starting from an initial delta function $p_X(x, 0) = \delta_0(x)$ (corresponding to an initial portfolio value of 1); and
— determine the mean and TailVaR of the portfolio at time T using $p(x, T)$.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

6.5.2 So to summarise the work so far, we now have a PDE for the portfolio value. The coefficients of the PDE vary for each point on the lattice because we vary the portfolio weights at each point. We then number the nodes so that they can index a weight matrix.

6.5.3 We now define a weight vector to be (for any given node) the asset allocation at that node expressed as a vector (so all the components will add to 1). Further we define a strategy to be a matrix of weight vectors such that each node has a weight vector. Therefore a strategy can be thought of as an object with three indices being time, portfolio value and asset.

6.5.4 Therefore if we have any given strategy we can combine this with the lattice PDE to calculate the terminal probability distribution of the portfolio and hence its risk statistics such as TailVaR. If we then combine this with an optimisation algorithm we are able to find the optimum strategy. Those who are familiar with the theory of PDE Constrained Optimisation may then see this algorithm as an application of that theory.

6.5.5 This therefore enables an objective function to be defined in terms of a set of tree-node parameters. An optimisation routine (we have used Matlab and NAG) can then be applied to find the set of parameters that maximises this objective function.

6.5.6 A final aspect of this is to note that the discussion so far has been in terms of assets that follow some kind of Brownian motion. Such models are often criticised because they produce insufficiently fat tails. One suggestion for further work is to note that although fat tails do not fit well into the methodology of partial differential equations they can be expressed using Partial Integro-Differential Equations (PIDEs). Jump processes can be designed that are consistent with fat tails and finite difference methods have been developed that enable the lattice solution to PIDEs thus we could envisage that our methodology could be modified to allow for a fat tail aware dynamic portfolio optimisation strategy.

6.5.7 Having given an overview of our rationale we now specify the method.

6.6 *Method*

This method seeks to solve the following problem: given a universe of possible assets that evolve under Geometric Brownian Motion with a vector of known expected growth rates, a covariance matrix, a measure of risk and a given tolerance for that risk what is the optimum way to (re)allocate the portfolio over time. The method does not require use of a utility function but instead requires use of well known statistical measures of risk (variance, Tail VaR, etc.). Importantly we require that the volatility of each asset be constant or at least deterministic with a given term structure.

6.6.1 *Algorithm*

In outline, the method works as follows:

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

(i) Discretise the portfolio value between a minimum and maximum level and the time between inception and maturity by dividing both into segments so that taken together they form a grid. Each intersection of gridlines is referred to as a node and these nodes are enumerated using a node number.
(ii) Enumerate the universe of possible assets.
(iii) Define a (2 dimensional) weight matrix as the percentage holding of the portfolio in a given asset at a given node.
(iv) Make an initial guess at what the weight matrix might be (e.g. by doing a single period Markowitz optimisation then assuming that the weight matrix holds those weights at every node).
(v) Use the weight matrix to calculate using matrix multiplication the portfolio mean growth rate and volatility separately at every node.
(vi) The portfolio is considered as a probability density function over the space of future time and portfolio values. It therefore follows the Fokker–Planck equations with volatility and drift varying by time and portfolio value.
(vii) The initial condition for the Fokker–Planck equation is that the portfolio has infinite point density at its initial value (as the Dirac delta function).
(viii) Solve the difference equations that arise from discretising the Fokker–Planck equations — this produces a terminal probability distribution.
(ix) Calculate the risk measure from the terminal probability distribution.
(x) Use an optimisation routine at step (iv) together with the risk calculation in step (ix) to iteratively estimate the optimal weight matrix.

6.6.2 *Detail*

6.6.2.1 Filling in the detail, we assume that the economics are that each asset follows Geometric Brownian Motion, possibly with term dependence in the drift and volatility. Our strategy is to define the Stochastic Differential Equation (SDE) then deduce the Partial Differential Equation (PDE) and hence the difference equations.

6.6.2.2 Thus, we assume that the price of a unit of asset i follows the SDE

$$dS^i = \mu^i(t)S^i dt + S^i\sigma^i(t)dZ^i.$$

6.6.2.3 Suppose that we hold $n^i(P, t)$ units of asset i then the value of the portfolio P at time t is $P = \sum n^i S^i$. The multidimensional Ito formula then gives:

$$dP = \sum \mu^i(P, t)S^i dt + \sum n^i S^i \sigma^i(P, t)dZ^i$$

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

or,

$$dP = \mu^P(P, t)Pdt + \sigma^P(P, t)PdZ. \tag{12}$$

6.6.2.4 This is the portfolio-level SDE where we define
— $n^i S^i = w^i P$,
— $\mu^P(t)P = \Sigma\mu^i(t)S^i$, and
— $\sigma^P(t)^2 = w^T Cw$.

6.6.2.5 These formulae can be interpreted as saying that the portfolio drift is the mean of the asset drifts, and that portfolio volatility is derived from C, the variance-covariance matrix in the normal way. Z is standard Brownian motion.

6.6.2.6 We now wish to transform the SDE equation 12 into a partial differential equation (PDE). This is common in derivative pricing although usually in that case one is working backwards from known values. However we are not trying to value a derivative backwards but to find the paths of the portfolio total value from a known initial value. Therefore we use the Kolmogorov forward equation or Fokker–Planck equation to give us a PDE for the probability density of P. Let the log portfolio value probability density be $p(x, t)$ then:

$$\partial/\partial t\ p(x, t) = -\partial/\partial x[\mu(x, t)p(x, t)] + (1/2)\partial^2/\partial x^2[\sigma^2(x, t)p(x, t)]. \tag{13}$$

6.6.2.7 The above problem has stochastic volatility. It is well known that such problems can have stability problems induced by negative probabilities. We avoid this by using an implicit PDE finite difference method.

6.6.2.8 First then we need to set up a grid. We have time go from 0 to T in units of Δt (chosen to generate an integer number of time steps) denoted by index i so we consider times:

$$0, \Delta t, 2\Delta t, \ldots, i\Delta t, \ldots, T.$$

6.6.2.9 We also let the log portfolio range in value between $P_{\min}$ and $P_{\max}$ in steps of ΔP with grid points denoted by index j so the portfolio can take values:

$$P_{\min}, P_{\min} + \Delta P, P_{\min} + 2\Delta P, \ldots, P_{\min} + j\Delta P, \ldots, P_{\max}.$$

6.6.2.10 Therefore we denote positions on the grid by (i, j). We will be setting boundary conditions that $P_{\min}$ and $P_{\max}$ have probability zero so they need to be sufficiently far away from the starting point for this. We also need to choose them so that P start lies on a grid line.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

6.6.2.11 We will also need initial conditions and boundary conditions. Now we know the portfolio value at time zero. Thus the probability density function at time zero is the Dirac delta function, i.e. infinitely dense at P_{start} and 0 density elsewhere. On a grid this means that:

$$p(0,j) = \begin{cases} 1 & \text{if } P_{\min} + j\Delta P = P_{\text{start}} \\ 0 & \text{otherwise.} \end{cases}$$

6.6.2.12 We also need boundary conditions and these are:

$$P(i,0) = 0$$
$$P(i,(P_{\max} - P_{\min})/\Delta P) = 0.$$

6.6.2.13 Finally, we need to express the PDE as a recurrence relationship, noting that we have boundary conditions at initial rather than terminal time as is normally the case in a derivative pricing algorithm and therefore what would normally be an implicit method is an explicit method and vice versa:

$$\partial p/\partial t = [p(i+1,j) - p(i,j)]/\Delta t$$
$$\partial/\partial x[\mu p] = [\mu(i,j+1)p(i,j+1) - \mu(i,j)p(i,j)]/\Delta x$$
$$\partial^2/\partial x^2[\sigma^2 p] = [\mu(i,j+1)p(i,j+1) - 2\mu(i,j)p(i,j) + \mu(i,j-1)p(i,j-1)]/\Delta x^2.$$

6.6.2.14 Having set up the PDE grid we now have a mapping from strategies to our objective function. We are therefore able to impose any constraints (e.g. long only weights) and optimise the objective function. The objective function would be

$$f(w(i,j) = \text{total return} - \lambda^* \text{ Tail VaR}.$$

6.6.2.15 The optimisation is non-linear and we have used MatLab and NAG algorithms to perform this.

6.7 *Results*

6.7.1 We now show the results of this method which are typically displayed graphically. The types of output that we have developed are as follows:

(i) A Markowitz style efficient frontier. As in traditional mean variance portfolio optimisation we show for each level of risk the maximum return that can be obtained. However we show 2 differences from the traditional model. Firstly we show the x axis as Tail VAR rather than portfolio volatility. Secondly, we show, in addition to the single period optimisation, the optimisation that is produced by breaking the period down into several sub-periods and allowing rebalancing at each sub-

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

period. As might be expected, the dynamically rebalanced portfolio always shows at least as great return for a given level of risk. Also the two curves meet at the right hand side because that corresponds to both strategies choosing maximum risk maximum return assets throughout as would be expected.

(ii) A terminal probability distribution showing at the end of the portfolio projection period the probability distribution of the portfolio. This allows one to visually confirm that the TailVaR is optimised.

(iii) A contour map showing the evolution of the portfolios probability distribution. Thus the last time period of this is equivalent to the distribution in (ii). At time zero the probability density will be infinite at one point (the initial portfolio value). This initial probability distribution can be thought of as the Dirac Delta Function.

(iv) A separate contour map for each asset showing its weight within the portfolio for each future state. This can be used to check that the algorithm is reallocating into lower risk assets following falls in portfolio value (see Figure 28).

6.7.2 The efficient frontier shows the same kind of shape that we would expect given what we have discussed. The terminal probability distribution shows an interesting skewness to the right. We interpret this as follows. The

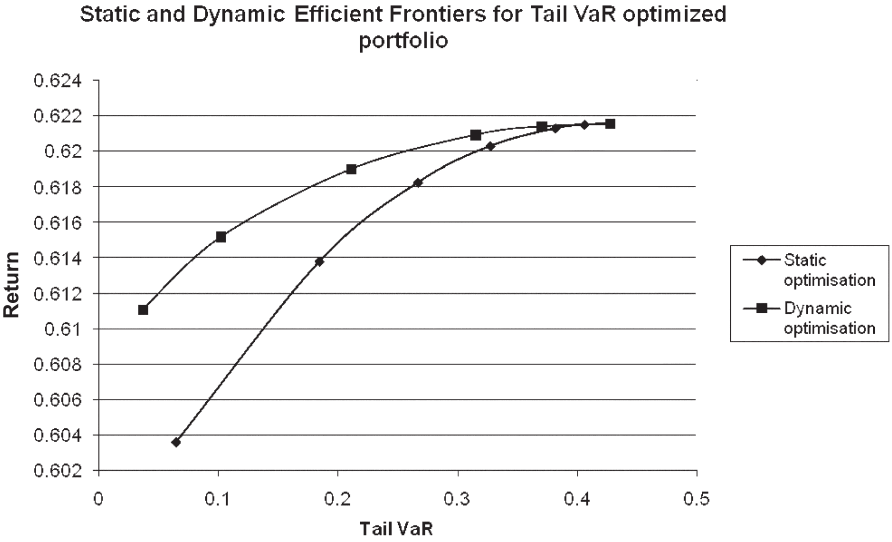


Figure 28. Comparison of efficient frontiers with and without dynamic rebalancing

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

algorithm is punished heavily for any probability that the portfolio falls below a certain level. Therefore as that level is approached it reallocates assets into low risk assets. This produces a peak in probability close to but above the level that we use to define TailVaR. However on paths where financial markets perform well, the portfolio can afford to reallocate towards riskier higher returning assets, hence the skew to the right.

7. Final Thoughts

7.1 *Summary*

7.1.1 In this paper we have demonstrated how a focus on an investor's goals leads to strategies that vary over time as those goals become more or less easily achievable. The objective of strategic asset allocation should not be to arrive at a single portfolio allocation once every three years, but also to determine a set of (preferably simple) rules that will allow the portfolio allocation to respond to events in a way that is consistent with the investor's goals.

7.1.2 There are a number of issues that this paper has highlighted in passing but to which we have so far devoted insufficient attention. We now correct these lacunae, pointing out where possible work done elsewhere and as an encouragement for others to conduct further research in this area.

7.2 *Time Consistency*

7.2.1 Although tail value at risk is an appealing risk measure in a single-period framework, it has theoretical shortcomings in the multi-period (i.e. dynamic) context that we have been analysing here. This seems to have been first highlighted in Artzner (2002). Artzner *et al.* (2004) discuss the example illustrated in Figure 29.

7.2.2 A coherent measure $\rho(X)$ of the risk of a strategy's outcome X can be thought of as a capital requirement: the amount of capital required to make the portfolio 'acceptable'. Indeed, defining the set of outcomes X which would be acceptable (the set of X for which $\rho(X) \leq 0$) rather than the risk metric itself is a common approach to analysing risk measures in the substantial recent literature on risk measurement. To do this uses the apparently-reasonable assumption that the capital required for a set of outcomes $X + m$, for any constant m, should be $\rho(X + m) = \rho(X) - m$; it is then easy to see that a risk measure can be defined as the least capital that would need to be added to make the set of outcomes 'acceptable'.

7.2.3 In a multi-period setting, we need to make capital requirements not just at some initial time, but at subsequent times too. These should be consistent in some sense. For example, some authors (e.g. Kupper *et al.* (2009), Boyle *et al.* (2005)) frame this in terms of a recursion property, such as:

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

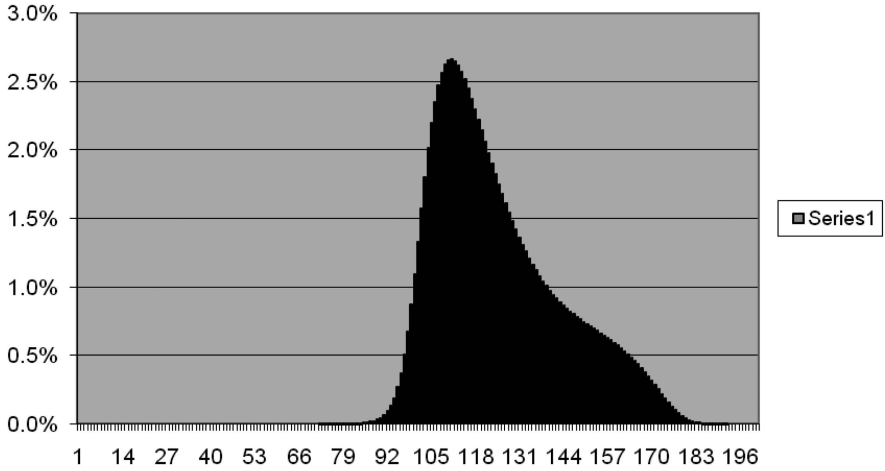


Figure 29. Terminal probability distribution of portfolio values

$$\rho_t(X) = \rho_t(-\rho_{t+1}(X)). \tag{14}$$

7.2.4 This says that the capital requirement at time t can be determined from the potential capital requirements at time $t+1$.

7.2.5 To assess tail value at risk against this time-consistency requirement, consider the example illustrated in Figure 30 and tail VaR assessed at a 2/3 probability level. In this example, there is a binary split at time zero followed by a three-way split at either branch at time 1. At time zero, the tail value of risk is +1 (minus the average of −10, 12, 14 and −20), while at time 1 it is −1 *whichever outcome* (u or d) *occurs*. In capital terms, this would mean that 1 unit of capital would be required at time 0 but 1 unit of capital could be released at time 1. This is clearly inconsistent with the recursion property (14) as well as with many weaker forms of time consistency.

7.2.6 Unfortunately, time consistency in a strong sense is difficult to achieve. For example a useful property of value at risk and tail value at risk are that they are 'law invariant' in the sense that the risk measure depends only on the set of possible outcomes and their probability distribution. It might be hoped that some suitably minor adjustment to one of these risk measures could be made to make them dynamically consistent. Kupper & Schachermayer (2009) show that, under apparently mild technical assumptions, this hope is misplaced — the only law-invariant dynamically consistent risk measures are of the form

$$\rho_t(X) = \frac{1}{\gamma}\log \mathbf{E}_t(\exp(-\gamma X)). \tag{15}$$

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

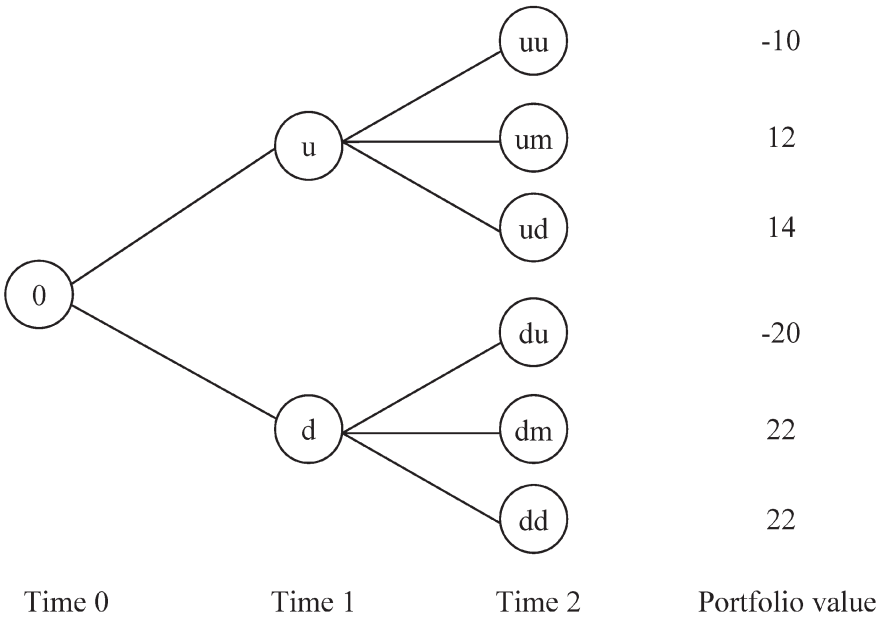


Figure 30. Multi-period TailVaR strategy problem from Artzner *et al.* (2004)

7.2.7 These are the so-called 'entropic' risk measures. If the 'translation invariance' requirement $\rho(X+m)=\rho(X)-m$ is not required, then the slightly more general class

$$\rho_t(X)=u^{-1}\mathbf{E}_t(u(X)) \tag{16}$$

of certainty-equivalent expected utility measures is obtained. This leads to the more recent investigation as to whether weaker notions of time consistency could be used. Roorda & Schumacher (2005) propose two alternatives for dynamically-consistent versions of tail VaR.

7.2.8 How could these theories be applied in practice? We have outlined two separate methods, a utility-based and a statistics-based method. The utility method allowed us to construct optimal strategies analytically and produced strategies that make intuitive sense.

7.2.9 In the specific optimisation problems discussed in Sections 5 and 6, the problem of dynamic inconsistency can be clearly seen: initially, a risk budget is set which offers a reasonable trade off between risk and return. But this trade off does not remain unchanged along each path — if things go

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

'well' then the risk budget can be increased, and the risk metric can become more and more irrelevant. In contrast, if things go sufficiently 'badly', then the ability to meet a risk constraint becomes more and more difficult — eventually, taking more risk becomes attractive, because de-risking offers no return and no help towards meeting the risk constraint. In both cases, the risk constraint is relative to the amount of risk that is deemed reasonable at the initial time — there is no opportunity or requirement to update this at subsequent times, in light of experience.

7.2.10 A related criticism can also be made of the utility measures discussed in Section 4: the time horizon plays an essential role in helping to determine the strategy. In the applications we have in mind, the choice of a target time horizon is an important part of the governance process. However. as Henderson & Hobson (2007) describe, if strategies that are independent of the time horizon are required, conditions must be placed on the kinds of utility function that can be used.

7.3 *Tail Behaviour*

7.3.1 Partly for the sake of ease of exposition, but also in order to allow us to make rapid progress, we have used fairly simple models of investment markets in this paper. It is important to test the intuition we have gained in more realistic models. In particular, for strategies that are aiming to control a measure of downside risk such as TailVaR, the weight of the tails being modelled should be expected to have a material impact on the strategy. Compared to a static procedure, dynamic optimisation will be able to spot, and exploit, a wider range of opportunities in a mis-specified model, so this model risk is important to understand and control. On the other hand, the dynamic procedure may identify unintended consequences of the specification of the investor's goals — e.g. the 'double or quits' strategy that emerged when a protected level of benefits was included in the utility function in Section 4. Accepting or modifying these features should be an important educational aspect of the strategic asset allocation process.

7.3.2 Using techniques such as copulas or extreme value theory to adequately model tail behaviour prior to dynamic optimisation should be important in practice. One example we have seen is Martellini & Milhau (2009), who include some tests of dynamic strategies such as those discussed here against models which include jump behaviour. There is certainly a need for further work in this direction.

7.4 *Further Reading*

We have tried to describe a reasonably diverse range of methods for tackling dynamic optimisation problems, and to touch on key areas of the large finance literature on these and related topics. We hope that the references within the bibliography, to which we have tried to refer where

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

appropriate, should offer the reader interested in more detail on multi-period risk measurement, PDE methods, optimal investment theory, etc.

Acknowledgements

We have benefitted from fruitful conversations with numerous clients and colleagues. The authors would particularly like to thank Matthew Barnes, Alvar Chambers, Anthony Earnshaw, Richard McMahon & Peter Ridges, the members of a working party in 2007/2008 which shaped the agenda for the utility theory based part of this work and Rob Vakil, Con Keating and Paul Bourdon for helpful comments on the PDE components. We are also grateful to the attendees at the Institute sessional meeting, and an anonymous referee, for feedback which helped to identify areas where further work or explanation was required. All errors remain, of course, our own responsibility.

References

Acerbi, C. & Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking and Finance*, **26**, 1487-1503.

Artzner, P., Delbaen, F., Eber, J.M. & Heath, D. (1997). Thinking coherently. *Risk*, **10**, November, 68-71.

Artzner, P., Delbaen, F., Eber, J.M. & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, **9**(3), 203-228.

Artzner, P. (2002). *Conditional value at risk: is it good in the multiperiod case?* IIR Conference on volatility and risk, London, February 2002.

Artzner, P., Delbaen, F., Eber, J.M., Heath, D. & Ku, H. (2007). Coherent multiperiod risk adjusted values and Bellman's principle. *Annals of Operations Research*, **152**, 5-22.

Barnes, M., Chambers, A., Earnshaw, A., Jarvis, S.J., McMahon, R. & Ridges, P. (2008). *Controlling the growth of your hedge. Presented to the Finance Investment and Risk Management Conference*, Manchester, June 2008.

Bellman, R. (1957). *Dynamic programming.* Princeton University Press.

Black, F. & Perold, A. (1992). Theory of constant proportion portfolio insurance. *Journal of Economic Dynamics and Control*, **16**, 403-426.

Box, G.E.P. & Draper, N.R. (1987). *Empirical model-building and response surfaces*. Wiley. ISBN 0471810339

Boyle, P., Hardy, M. & Vorst, T. (2005). Life after VaR. *Journal of Derivatives*, **13**, 48-55.

Cairns, A.J.G. (2000). *Some notes on the dynamics and optimal control of stochastic pension fund models in continuous time* ASTIN Bulletin **30:** 19-55. Available, with erratum, at http://www.ma.hw.ac.uk/~andrewc/papers/

Chapman, R.J., Gordon, T.J. & Speed, C.A. (1999). Pensions, funding and risk. *British Actuarial Journal*, **7**(4), 605-662.

Cheridito, P., Delbaen, F. & Kupper, M. (2006). Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability*, **11**, 57-106.

Clewlow, L. & Strickland, C. (1998). *Implementing derivatives models*. Wiley.

Cox, J. & Huang, C.F. (1989). Optimal consumption and portfolio policies when asset prices follow a diffusion process. *Journal of Economic Theory*, **49**, 33-83.

Cox, J.C., Ross, S.A. & Rubinstein, M. (1979). Option pricing: a simplified approach. *Journal of Financial Economics*, **7**, 229-263.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

Crank, J. & Nicolson, P. (1947). A practical method for numerical evaluation of solutions of partial differential equations of the heat conduction type. *Proceedings of the Cambridge Philosophical Society*, **43**, 50-67.
Dempster, M.A.H., Germano, M., Medova, E.A. & Villaverde, M. (2003). Global asset-liability management. *British Actuarial Journal*, **9**, 137-216.
Derman, E. & Kani, I. (1994). *The volatility smile and its implied tree*. Goldman Sachs, January 1994.
Duffie, D. (2001). *Dynamic asset pricing theory*, 3rd edition. Princeton University Press.
Duffy, D.J. (2006). *Finite difference methods in financial engineering*. Wiley.
Dupire, B. (1994). *Pricing with a smile*. Risk, January 1994.
Exley, C.J., Mehta, S.J.B. & Smith, A.D. (1997). The financial theory of defined benefit pension schemes. *British Actuarial Journal*, **3**, 835-966.
Gabih, A. & Wunderlich, R. (2004). Optimal portfolios with bounded shortfall risks. In J. vom Scheidt & M. Richter (eds). *Tagungsband zum Workshop 'Stochastische Analysis'* (21-42). TU Chemnitz, Fakultät für Mathematik.
Gandy, R. (2005). *Portfolio optimization with risk constraints*. PhD-Dissertation, Universität Ulm.
Goetzmann, W., Ingersoll, J., Spiegel, M. & Welch, I. (2002). *Sharpening Sharpe ratios.* NBER Working Paper No. 9116.
Henderson, V. & Hobson, D. (2007). Horizon-unbiased utility functions. *Stochastic processes and their applications*, **117**, 1621-1641.
Jarvis, S.J., Southhall, F. & Varnell, E. (2001). *Modern valuation techniques.* Presented to Staple Inn Actuarial Society.
Karatzas, I., Lehoczky, J. & Shreve, S. (1987). Optimal portfolio and consumption decisions for a 'small investor' on a finite horizon. SIAM *Journal of Control and Optimisation*, **25**, 1157-1186.
Karatzas, I. & Shreve, S.E. (1998). *Methods of mathematical finance*, Springer, New York.
Kemp, M.H.D. (1996). *Asset/liability modelling for pension funds*. Staple Inn Actuarial Society.
Kupper, M. & Schachermeyer, W. (2009). Representation results for law-invariant time consistent functions. *Mathematics and Financial Economics*, **2**, 189-210.
Leland, H.E. (1980). Who should buy portfolio insurance? *Journal of Finance*, **35**(2), 581-594.
Martellini, L. & Milhau, V. (2009). *Measuring the benefits of dynamic asset allocation strategies in the presence of liability constraints*. EDHEC, March 2009.
Merton, R. (1969). Lifetime portfolio selection under uncertainty: the continuous-time case. *Review of Economics and Statistics*, **51**(3), 247-257.
Merton, R. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, **3**, 373-413.
Nielsen, P.H. (2006). *Optimal investment strategies: a short survey of classical and recent results.* Presented to the International Congress of Actuaries.
Pflug, G.Ch. (2000). Some remarks on the Value-at-Risk and the conditional Value-at-Risk. In S. Uryasev (ed.). *Probabilistic constrained optimization: methodology and applications*. Kluwer Academic Publishers, 2000.
Rockafellar, R.T. (2001). *Optimisation under uncertainty*. University of Washington lecture notes. Available at http://www.math.washington.edu/~rtr/uncertainty.pdf
Rockafellar, R.T. & Uryasev, S. (2000). Optimization of conditional Value-At-Risk. *The Journal of Risk*, **2**(3), 21-41.
Rogers, L.C.G. (2003). *Duality in constrained optimal investment and consumption problems: a synthesis*. Paris-Princeton Lectures on Mathematical Finance 2002, (Springer Lecture Notes in Mathematics 1814), 95-131.
Rom, B.M. & Ferguson, K. (1994). *Post-modern portfolio theory comes of age*. 4th AFIR International Colloquium, 349-364.
Roorda, B. & Schumacher, J.M. (2007). Time consistency conditions for acceptability measures, with an application to tail value at risk. *Insurance: Mathematics and Economics*, **40**, 209-230.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

Roorda, B. Schumacher, J.M. (2010). *When can a risk measure be updated consistently?* Preprint.

Ross, S. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, **13**(3), 341-360.

Samuelson, P.A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, **51**, 239-246.

Seydel, R. (2009). Tools for Computational Finance (4th edition). Springer.

Sharpe, W.F. (1964). Capital asset prices: a theory of market equilibrium under conditions of risk, *Journal of Finance*, **19**(3), 425-442.

Smith, A.D. (1996). How actuaries can use financial economics. *British Actuarial Journal*, **2**, 1057-1193.

Uryasev, S. (2000). Conditional Value-at-Risk: optimization algorithms and applications. *Financial Engineering News*, **14** (February), 1-5.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

APPENDIX A

COMMON UTILITY FUNCTIONS

We provide some more detail on the utility functions that are regularly used in the finance literature.

A.1 *Definition*

The basic idea of a utility function is to express mathematically the personal value that will arise from different levels of wealth. Generally, utility functions have two reasonable properties:

— *Increasing*. More wealth is always better.
— *Convexity*. The increase in utility for each additional dollar of wealth diminishes as the initial level of wealth increase.

A.2 *Measures of Risk Aversion*

A2.1 The Arrow–Pratt measure of *absolute risk aversion* (ARA) at a particular level of wealth w is:

$$ARA(w) = -\frac{U''(w)}{U'(w)}. \tag{17}$$

A2.2 This effectively expresses the trade off required between additional return and variance when starting from a given level of wealth.

The *relative* risk aversion is:

$$RRA(w) = w.ARA(w) = -\frac{w.U''(w)}{U'(w)}. \tag{18}$$

A2.3 This measures the aversion to risks that are proportional to the initial level of wealth.

A2.4 Arrow argued that investors' aversion should decrease with wealth (i.e. ARA is a decreasing function of w), i.e. that willingness to face a risk of a given size increases with wealth. But if the risk rises in the same proportion as the wealth then the willingness falls (i.e. RRA is an *increasing* function of w).

A.3 *Constant Relative Risk Aversion (CRRA) Utility Functions*

A3.1 As noted in the main text, the CRRA family of utility functions can be defined via:

$$U(w) = \frac{w^{1-\gamma} - 1}{1 - \gamma}. \tag{19}$$

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

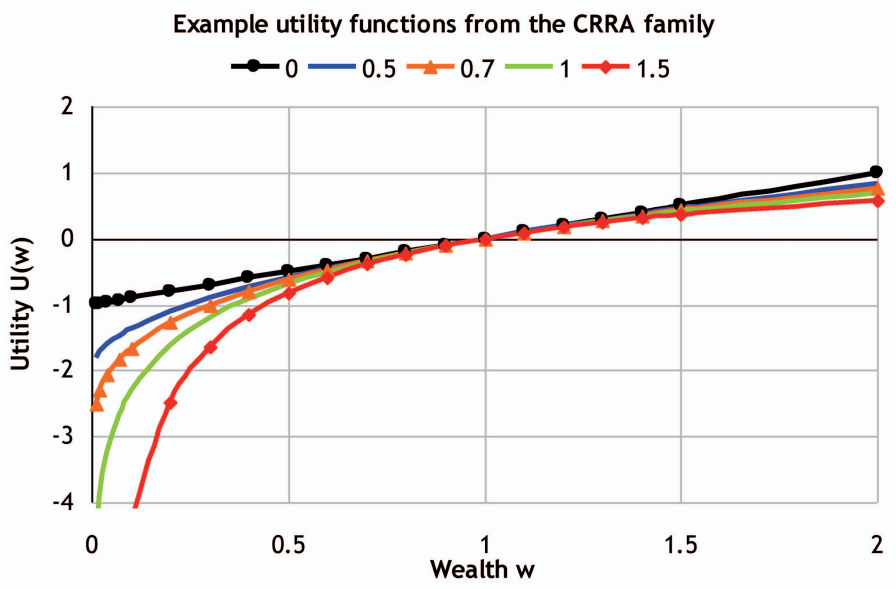


Figure 31. Example CRRA utility functions

A3.2 It is easy to check that this family does indeed have constant relative risk aversion γ.

A3.3 Since it is the expected value of a utility function that is maximised, changing a utility function by the addition of, or multiplication with, a constant has no effect on the investment strategy that is followed. A utility function $U(w) = w^{1-\gamma}$ would therefore be equivalent. The form of (6) is used so that the $\gamma = 1$ limit exists: it is simply $U(w) = \ln w$.

A3.4 Note that:

— The $\gamma = 0$ case is linear.
— The cases $0 < \gamma < 1$ are bounded below.
— The $\gamma = 1$ case is logarithmic.
— The cases $\gamma > 1$ are bounded above.

A3.5 This family is sketched in Figure 31 for different levels of the relative risk aversion γ.

A.4 *Hyperbolic Absolute Risk Aversion (HARA) Utility Functions*

The absolute risk aversion of the CRRA family is γ/w, a function whose graph is a hyperbola. Other hyperbolic absolute risk aversion functions, of the form $ARA(w) = a + (bw + c)^{-1}$, define the HARA family of utility

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

functions. These utility functions are also commonly used in academic papers.

A.5 *Further Reading*

For a fuller introduction to the use of utility functions within the context of dynamic investment strategy problems, we recommend Nielsen (2006).

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

APPENDIX B

OPTIMAL SHARPE RATIO PORTFOLIOS

B.1 *The Model*

B1.1 This Appendix summarises results in Goetzmann *et al.* (2002). They investigate models which allow rebalancing or which contain assets such as options whose returns are not normally distributed. They show how to construct portfolios that have superior Sharpe ratios. This Appendix looks at just one of their examples (continuous-time rebalancing) to illustrate their point.

B1.2 The model consists of the usual 'Black Scholes' set up of lognormally distributed asset prices:

— Deterministic bond prices $P_{tT} = P(t, T) = \exp(-r(T-t))$;
— Stock price $S_t = \exp\big((\mu - \sigma^2/2)t + \sigma W_t\big)$ driven by a Brownian motion (Wiener) process W_t;
— Stochastic discount process $m_t = \exp\big(-(r + \lambda^2/2)t - \lambda W_t\big)$, where $\lambda = (\mu - r)/\sigma$ is the instantaneous Sharpe ratio for the stock.

The Sharpe ratio of investing a static stock investment to time T is

$$\frac{\exp(\mu T) - \exp(rT)}{\exp(\mu T)\sqrt{\exp(\sigma^2 T) - 1}} \approx \frac{\mu - r}{\sigma}\sqrt{T} = \lambda\sqrt{T} \text{ for small } T.$$

B.2 *Maximum Sharpe Ratio Portfolio*

Suppose a portfolio has value (payoff) X_T at time T. This has value $\mathbb{E}_0(m_T X_T)/m_0$ at time zero, and so the expected Sharpe ratio is:

$$\frac{\mathbb{E}_0(X_T) - \mathbb{E}_0(m_T X_T)/(m_0 P(0, T))}{\sqrt{Var(X_T)}}.$$

This can be optimised by maximising a Lagrangian:

$$\mathbb{E}_0(X_T) - \lambda\mathbb{E}_0\big(X_T^2\big) - \mu\mathbb{E}_0(m_T X_T).$$

This then gives $X_T = \alpha - \beta m_T$ which has value $X_t = \alpha P(t, T) - \beta m_t \exp\big((\lambda^2 - 2r)(T-t)\big)$ and expected return and variance $\mathbb{E}_0(X_T) = \alpha - \beta P(0, T)$, $Var(X_T) = \beta^2 \exp(-2rT)\big(\exp\big(\lambda^2 T\big) - 1\big)$. Therefore the expected Sharpe ratio is:

$$sign(\beta).\sqrt{\exp(\lambda^2 T) - 1}.$$

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

B.3 *Return Distribution*

The probability density function for the stock and the maximum Sharpe ratio portfolio follow directly from the pdf for the normal distribution (set $\beta = 1$ for simplicity here):

$$f_{S_T}(x) = \frac{1}{\sqrt{2\pi}\sigma x}\exp\left(-\frac{\left(\ln(x) - \left(\mu - \sigma^2/2\right)T\right)^2}{2\sigma^2 T}\right)$$

$$f_{X_T}(x) = \frac{1}{\sqrt{2\pi}\lambda(\alpha - x)}\exp\left(-\frac{\left(\ln(\alpha - x) + \left(r - \lambda^2/2\right)T\right)^2}{2\lambda^2 T}\right).$$

These are plotted in the graph shown in Section 2.3 above and repeated in Figure 32.

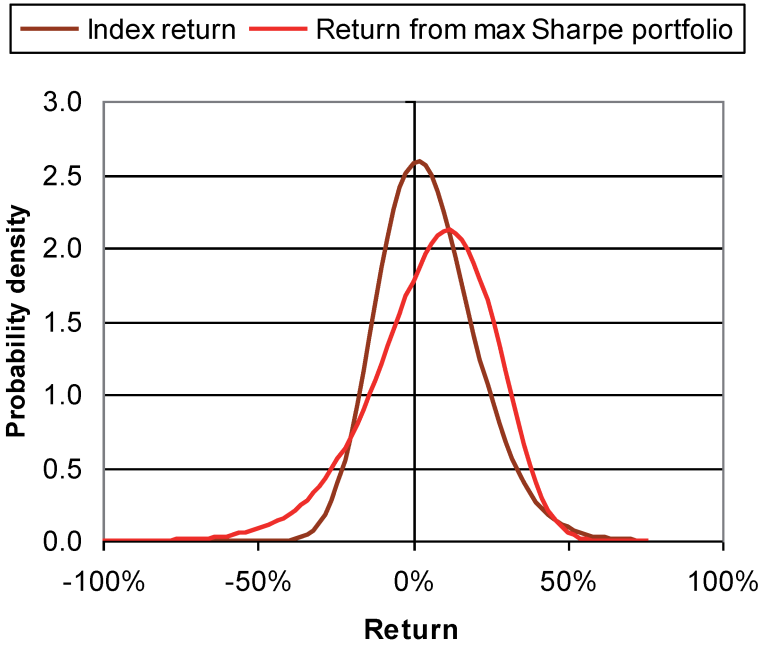


Figure 32. Probability density of a mean-variance optimal dynamic strategy

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

APPENDIX C

DERIVATION OF EXPECTED UTILITY RESULTS

C.1 *Modified Family of Utility Functions*

C1.1 The basic CRRA utility function is:

$$U(w;\gamma)=\frac{w^{1-\gamma}-1}{1-\gamma}. \tag{20}$$

C1.2 In the case $\gamma=1$, this is to be interpreted as $U(w;1)=\ln w$. It is easy to check that the family is then continuous at $\gamma=1$.

C1.3 We define a family of extended utility functions for $F\le M<L\le\infty$:

$$U(w_T;\gamma,T',F,M,L)=\frac{\left([w_T-FP(T,T')]_{(M-F)\cdot P(T,T')}^{(L-F)\cdot P(T,T')}\right)^{1-\gamma}-1}{1-\gamma}.$$

C1.4 Again this is interpreted as a logarithm in the case $\gamma=1$. Here:

— $[x]_A^B=\max(A,\min(B,x))$ denotes a value x limited to lie in the interval $[A,B]$;
— $P(s,t)$ is the price at time s of a zero-coupon bond paying 1 unit at time $t>s$; and
— γ is a relative risk aversion parameter.

The idea behind the three modifications is:

— F is a floor relative to which risk is measured;
— M is a minimum payment level, below which some kind of guarantee may apply; and
— L is a target liability payment due to be paid at time $T'>T$.
— In addition, we may seek to impose a hard constraint $w_T\le GP(T,T')$. Here $G\le M$.

Figure 33 may help to illustrate the idea here ($G=25$, $F=50$, $M=150$, $L=350$).

Figure 33 shows the typical situation in the case that $F\le M$. When $F\ge M$, then F becomes a hard floor to the wealth level, and M becomes irrelevant.

C.2 *Unmodified Optimal Strategy: the Merton Portfolio*

C2.1 The optimisation problem in the unmodified case is to maximise

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

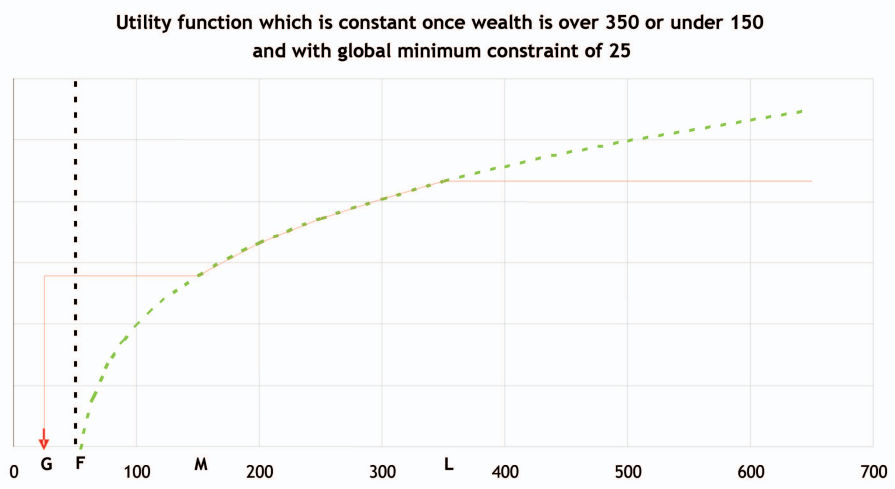


Figure 33. An example of the modified family of utility functions

$$U(w_T) = \frac{w_T^{1-\gamma} - 1}{1-\gamma}. \tag{21}$$

C2.2 Subject to a budget constraint $\mathbb{E}_0(m_T w_T) = w_0$. Here m_t is a stochastic discount process. We work in a deterministic diffusion setting, so that this discount process satisfies an SDE of the form:

$$dm_t = -r_t dt - \lambda(t) dW_t.$$

C2.3 The market price of risk process $\lambda(t)$ may be vector-valued if there is more than one Brownian motion driving the model. [For background on stochastic discount factors, see for example Duffie (1992), Cochrane (2001) or for a more elementary introduction see Jarvis *et al*. (2001). As an example, the 'Black–Scholes' model consists of two basic assets: a money market account B_t and a stock S_t — these are of the form

$$\begin{aligned} dS_t &= S_t(\mu dt + \sigma dW_t) & S_t &= \exp\left(\mu t - \tfrac{1}{2}\sigma^2 t + \sigma W_t\right) \\ dB_t &= rB_t dt & B_t &= \exp(rt) \\ dm_t &= m_t(-r dt - \lambda dW_t) & m_t &= \exp\left(-rt - \tfrac{1}{2}\lambda^2 t - \lambda W_t\right). \end{aligned} \tag{22}$$

These are consistent (i.e. the model is free of arbitrage) if $\mu = \lambda\sigma$ — from which comes the terminology 'market price of risk'.]

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

C2.4 The inverse of the stochastic discount factor is a portfolio process h_t, known as the 'growth optimal portfolio'. This is therefore:

$$h_t = m_t^{-1} = \exp\left(\frac{1}{2}\int^t \lambda(u)^T\lambda(u)du + \int^t \lambda(u)dW_u\right). \tag{23}$$

C2.5 The Cox–Huang approach then shows that the solution to (21) is a portfolio process which we denote as $X(\gamma, T)_t$ and refer to, for convenience, as the 'Merton portfolio' or 'T-optimal portfolio' even though the context is slightly different to Merton's:

$$\frac{dX(\gamma, T)_t}{X(\gamma, T)_t} = \gamma^{-1}\frac{dh_t}{h_t} + (1-\gamma^{-1})\frac{dP_{tT}}{P_{tT}} \tag{24}$$

$$\begin{aligned} X(\gamma, T)_t &= (h_t)^{1/\gamma}(P_{tT})^{1-1/\gamma}\exp\left(\int^t c_X(u)du\right) \\ c_X(u) &= \frac{\gamma-1}{2\gamma^2}\left|\sigma(u, T) - \lambda(u)\right|^2. \end{aligned} \tag{25}$$

C2.6 The SDE shows that the optimal *strategy* is to invest a proportion γ^{-1} of wealth at any time in the growth-optimal portfolio; and the remainder in a bond expiring at the chosen time horizon. The second equation describes the strategy *payoff* in terms of the growth-optimal portfolio. When $\gamma = 1$, the (logarithmic) utility-optimal portfolio is exactly the growth-optimal portfolio.

C2.7 In the Black–Scholes case, we can be more explicit and express the growth-optimal portfolio as a constantly-rebalanced portfolio — with a proportion λ/σ in stocks and the remainder in cash. However, in this Appendix we will not express portfolios as stock/cash mixes, rather expressing the risky (non-bond) elements of portfolios in terms of the Merton portfolio $X(\gamma, T)_t$.

C.3 *Convex Form of Modified Utility Function*

C3.1 The optimisation problem we want to solve is

$$Max\, U(w_T; \gamma, T', F, M, L) \text{ subject to } w_t = \mathbb{E}_t(m_T w_T)\big/m_t, \quad w_T \geq G \cdot P(T, T'). \tag{26}$$

C3.2 This involves a utility function that is not convex if $G < M$, due to the horizontal section below M. In order to state the form of the solution of the dynamic optimisation problem for the modified utility function, it is helpful to define an 'effective minimum' liability level. This amounts to

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

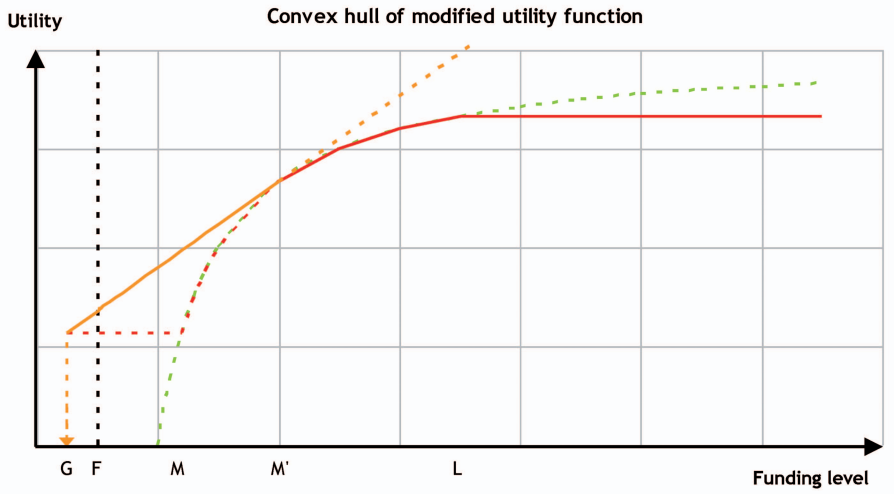


Figure 34. Convexified version of modified utility function

establishing the point M' at which a line drawn from $(G, U(M))$ touches the utility function, as illustrated in Figure 34.

C3.3 Note that $M' \leq L$, and equality is possible. It will turn out that the solution to the original problem is also a solution to the problem with this convex utility function, and in particular the optimal wealth does not take values on the (open) interval (G, M').

C.4 *Modified Optimal Strategy: Statement*

C4.1 The optimal strategy is defined implicitly as follows. For a given level of assets w_t at time t, there is a unique ξ such that:

$$w_t = F \cdot P(t, T') + (G - F) \cdot P(t, T')\Phi(-d_2) + \xi \cdot X(\gamma, T)_t[\Phi(-d_3) - \Phi(-d_1)] + (L - F) \cdot P(t, T')\Phi(d_4) \quad (27)$$

where:

$$d_1 = \frac{1}{\sigma}\left[\ln\left(\frac{\xi X(\gamma, T)_t}{(M' - F)P(t, T')}\right) + \frac{1}{2}\sigma^2\right], \quad d_3 = \frac{1}{\sigma}\left[\ln\left(\frac{\xi X(\gamma, T)_t}{(L - F)P(t, T')}\right) + \frac{1}{2}\sigma^2\right]$$
$$d_2 = d_1 - \sigma, d_4 = d_3 - \sigma.$$

C4.2 The optimal strategy is then such that (27) holds for all time (with

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

an unchanging value of ξ). The strategy that supports this is given by the SDE:

$$\begin{aligned} dw_t = &[\xi\Phi(-d_3) - \xi\Phi(-d_1)]dX(\gamma, T)_t \\ &+ [F + (G - F)\Phi(-d_2) + (L - F)\Phi(d_4)]dP(t, T') \\ &- \frac{(G - M)P(t, T')\phi(-d_2)}{\left[\int_t^T \sigma_u^2 du\right]^{1/2}}\left(\frac{dX(\gamma, T)_t}{X(\gamma, T)_t} - \frac{dP(t, T')}{P(t, T')}\right). \end{aligned} \tag{28}$$

C4.3 Where σ_u is the log-volatility of $X(\gamma, T)_t / P(t, T')$ at time u. This formula says that the dynamic strategy can be considered as a dynamic mix of two portfolios: the Merton portfolio $X(\gamma, T)_t$ and a liability-hedging portfolio $P(t, T')$. The amounts invested in each portfolio over time are just the coefficients of $dX(\gamma, T)_t$ and $dP(t, T')$ in this equation.

C.5 *Modified Optimal Strategy: Sketch Proof*

C5.1 To derive the formulae (27), (28), the Cox–Huang approach can be used. The first step is to derive the optimal payoff, i.e. which function w_T maximises the utility in (26) subject to the budget constraint. Multiplying the budget constraint by a Lagrange multiplier ξ and using the calculus of variations then quickly shows that w_T must be of the form:

$$w_T^* = \begin{cases} G \cdot P(T, T') & \text{if} \quad F \cdot P(T, T') + \xi X(\gamma, T)_T < M' \cdot P(T, T') \\ F \cdot P(T, T') + \xi X(\gamma, T)_T & \text{if} \quad M' \cdot P(T, T') \le F \cdot P(T, T') + \xi X(\gamma, T)_T \le L \cdot P(T, T') \\ L \cdot P(T, T') & \text{if} \quad L \cdot P(T, T') < F \cdot P(T, T') + \xi X(\gamma, T)_T. \end{cases}$$

C5.2 This is illustrated in Figure 35 — it is a piecewise-linear function of the Merton portfolio value.

C5.3 Although this looks complex, this is simply a combination of the bond $P(t, T')$ or the T-optimal portfolio $X(\gamma, T)_t$ and familiar option payoffs: a binary option (giving the jump from G to M' shown above), and two call options (or, due to put-call parity, this can be expressed in terms of put options).

C5.4 The value of this payoff is readily calculated using standard option formulae, to be (27). Furthermore, either using the well-known deltas of these options or by direct differentiation, the formula (28) follows.

C.6 *Optimal Modified Utility: Statement*

C6.1 As a further exercise in option pricing mathematics, it is straightforward to calculate the expected utility. In the case $\gamma \neq 1$ this gives:

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

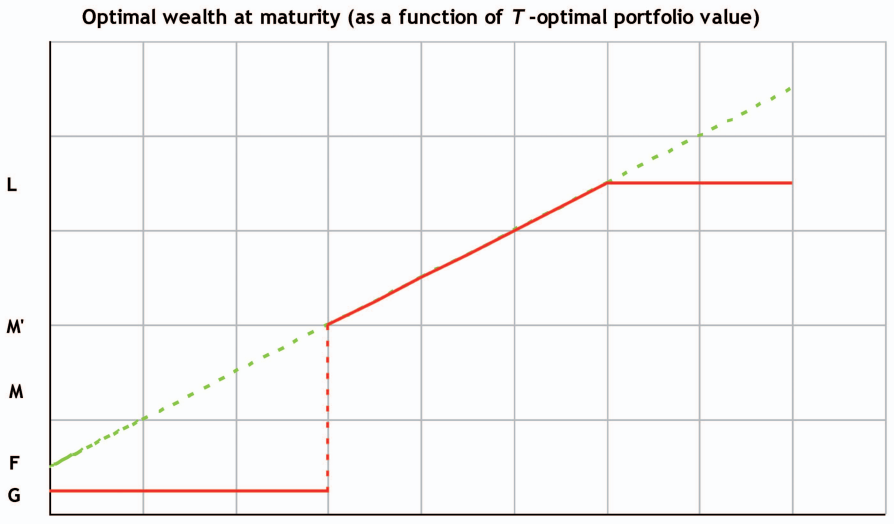


Figure 35. Piecewise-linear function of the Merton portfolio value

$$
\begin{aligned}
&(1-\gamma)\mathbb{E}_t(U(w_T^*;\gamma,T',F,M,L))+1\\
&\quad=\left(\frac{(M-F)P(t,T')}{P(t,T)}\right)^{1-\gamma}\exp\left(-\int_t^T c_Y(u)du\right)\Phi\\
&\qquad\times\left(\frac{1}{\sigma}\left[-\ln\left(\frac{\xi X(\gamma,T)_t}{(M'-F)P(t,T')}\right)-\frac{2\gamma-1}{2}\sigma^2\right]\right)\\
&\qquad+\left(\frac{\xi X(\gamma,T)_t}{P(t,T)}\right)^{1-\gamma}\exp\left(-\gamma\int_t^T c_X(u)du\right)(\Phi(-d_3)-\Phi(-d_1))\\
&\qquad+\left(\frac{(L-F)P(t,T')}{P(t,T)}\right)^{1-\gamma}\exp\left(-\int_t^T c_Y(u)du\right)\Phi\\
&\qquad\times\left(\frac{1}{\sigma}\left[\ln\left(\frac{\xi X(\gamma,T)_t}{(L-F)P(t,T')}\right)+\frac{2\gamma-1}{2}\sigma^2\right]\right).
\end{aligned}
$$

Here, $c_X(u)$ is as in (25) above, and:

$$
c_Y(u)=\frac{\gamma-1}{2}\left[\left|\lambda(u)-\sigma(u,T)\right|^2-\left|\lambda(u)-\sigma(u,T')\right|^2\right]-\frac{(\gamma-1)^2}{2}\left|\sigma(u,T')-\sigma(u,T)\right|^2.
$$

C6.2 This expression for optimal utility can be useful in numerical exercises where the Bellman value function needs to be calculated: this

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

function acts as an upper bound; and can be used as an initial approximation, in discrete contexts. A detailed proof of the formulae in this Appendix offer little insight but can be obtained from the authors on request.

C.7 *Special Cases*

C7.1 Although it was easier to state the result in the full generality of having G, F, M, L all playing a role, it is instructive to look at particular cases. When $F = G = M = 0$, i.e. there is a liability target but no floor, the wealth process reduces to a 'covered call':

$$w_t^* = \xi \cdot X(\gamma, T)_t \Phi(-d_3) + L \cdot P(t, T') \Phi(d_4)$$
$$d_3 = \frac{1}{\sigma}\left[\ln\left(\frac{\xi X(\gamma, T)_t}{L P(t, T')}\right) + \frac{1}{2}\sigma^2\right], \quad d_4 = d_3 - \sigma.$$

C7.2 When $M = F = G$ and $L = \infty$, i.e. there is no target but the CRRA function is 'shifted to the right' the formula becomes a simple sum:

$$w_t^* = F \cdot P(t, T') + \xi \cdot X(\gamma, T)_t.$$

C7.3 Finally, when $M' = M = G > 0 = F$, i.e. the CRRA function is truncated at M, the formula becomes a 'covered put':

$$w_t^* = \xi \cdot X(\gamma, T)_t \Phi(d_1) + M \cdot P(t, T') \Phi(-d_2)$$
$$d_1 = \frac{1}{\sigma}\left[\ln\left(\frac{\xi X(\gamma, T)_t}{M P(t, T')}\right) + \frac{1}{2}\sigma^2\right], \quad d_2 = d_1 - \sigma.$$

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

APPENDIX D

ON THE CURSE OF DIMENSIONALITY

D.1 *The Non-Recombining Tree*

The non-recombining tree approach is as follows:

— One defines a set of investment factors (stock market indices, FX rates, yield curve parameters).
— One then writes down a set of simultaneous SDEs for these variables.
— One constructs a multidimensional tree for them.
— One assigns an unknown set of portfolio weights to each node.
— One then runs an optimisation to find the global set of weights that optimises the risk return trade off for that portfolio.

D.2 *Advantages and Disadvantages*

D2.1 This non-recombining tree approach that we started with has a number of distinct advantages being:

— Because (apart from the covariance matrix) each asset can be modelled independently, it is possible to incorporate arbitrarily complex asset price dynamics such as complex yield curve models.
— Because the model is non-recombining it is possible to easily ensure that one never suffers from negative probabilities of a node being realised. (To see this consider a binary node. It is always possible to assign each branch emanating from it the probability half and solve for the position of the subsequent nodes in asset price space using the local drift and volatility of that asset at that node. In contrast because a recombining tree is more constrained we need to be more flexible in setting transition probabilities.

D2.2 However, we found that when we tried to implement the non-recombining approach in practice we very quickly had exponentially increasing number of variables in our optimisation problem rendering it intractable for the type of problem we were solving. The reason for this is as follows:

D2.3 Suppose that we have 10 asset factors, and 10 time steps at which decisions are made. The methods then require construction of a non-recombining tree. For a decision to be made at a node, the node has to split in the dimension of each decision variable. Thus with 10 assets we need to split each node into 2^{10} paths (approximately 1,000). We then construct a (non-recombining) tree with 10 time steps. Therefore this has $1{,}000^{10}$ (or 10^{30} final nodes). At each node there are 10 weights therefore we have to solve an optimisation problem with 10^{31} variables.

D2.4 In contrast, the recombining portfolio lattice/tree method described in the text can also optimise 10 assets for 10 time steps. We use

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press

only 1 SDE being the total size of the portfolio. We construct a finite difference mesh with 10 time steps and 10 space steps giving us 100 nodes. At each node we have a set of 10 portfolio weights to choose giving us a problem with 1,000 variables to solve. Therefore recombining portfolio lattice/tree method has reduced the scale of the problem by a factor of the order of 10^{28}, making tractable an otherwise intractable problem.

https://doi.org/10.1017/S1357321700005742 Published online by Cambridge University Press