

OPTIMAL PORTFOLIOS THROUGH BELLMAN NUMERICS

ERIK BØLVIKEN AND PÅL NICOLAI HENRIKSEN

ABSTRACT. A numerical strategy for solving low-dimensional Bellman equations through the traditional backwards recursion is formulated. A simple error analysis suggests that the approach handles many multi-period portfolio selection problems, and a number of examples confirm this experimentally. Minimum downside risk procedures are studied and it is demonstrated how multi-period efficient frontiers can be calculated for such criteria. A closing example examines the impact of heavy-tailed distributions on optimal, multi-period risk.

1. INTRODUCTION

Financial portfolios that are optimal over many periods are solutions of stochastic control problems and Hamilton Jacobi Bellman equations. There is an enormous literature on their theoretical and numerical properties. Continuous-time summaries are, for example, Fleming and Soner [14], Kushner and Dupuis [23], Pham [26] and in insurance Schmidli [31]. The present paper is concerned with the mathematically less demanding discrete-time view, but practical computation is still a major problem. One approach is the fairly recent suggestion in Rogers [29], but the traditional answers are stochastic, dynamic programming (Wallace and Ziemba [34]) and the Bellman principle that goes back to Bellman [3]. High-dimensional portfolio selection through stochastic, dynamic programming has been reported for example in Hilli, Koivu, Pennanen and Ranne [17] and Krokmal, Uryasev and Palmquist [22], and it is in Krokmal and Uryasev [21] and Rockafeller and Uryasev [28] shown how this approach reduces conditional-value-at-risk criteria to convex programming. The Bellman line works through recursions running backwards in time, and optimal strategies are inferred from the preceding one one time step ahead. This technique is used to determine optimal financial portfolios in Balduzzi and Lynch [1] and Barberis [2] with more algorithmic contributions in Brandt, Goyal, Santa-Clara and Stroud [9], Garlappi and Skoulakis [16] and van Binsbergen and Brandt [33] while Dickson and Waters [12] and Korn and Wiese [20] are contributions in insurance.

Numerical solutions of the Bellman recursion go far back, see Tapiero and Sulem [32] and references therein. Time is discrete, but the state space is not, and the traditional way is to

Date: June 2, 2010.

Key words and phrases. Bellman equation, downside risk, efficient frontier, heavy-tailed, multi-period, numerical error, optimal portfolio.

replace continuous models and distributions by discrete analogues. This is known as the Markov chain approximating method, related to, but not quite the same as the idea advocated here. Bellman schemes run on a grid of say N points in the state space, and each time step demands N optimizations of functions that are multi-dimensional integrals. The integrands depend on the solution from the preceding calculation, but the latter is only available on the grid that has been selected for it, and this is not enough to proceed further. The Markov chain approximation stays clear of this problem by working on a discrete state space all along, but another way to break the deadlock is to use function approximation where values we need are determined numerically from those we have. A modern application of that idea is the much cited paper by Longstaff and Schwartz [24] on financial options, but it goes far back, almost half a century (at least) to Bellman, Kalaba and Kotkin [4]. The discretization of the state space is now a purely numerical issue and not something that takes place in the model sphere as with the Markov chain approximation. An advantage of this viewpoint is that we draw on the high accuracy of function approximation. Many criteria in finance vary fairly slowly which creates a good basis for interpolation or even regression methodology.

The purpose of this paper is to formulate a general numerical strategy and examine the solutions that is obtained with it when it is put to work on several interesting and not quite common low-dimensional problems. A simple error analysis of the basic scheme is developed in the next section. It will suggest that N may not have to be overwhelmingly large, and practical experience reported in Sections 4 and 5 offer support. This would be particularly important when the state space has higher dimension than in our examples. All grids in this paper are cartesian products. With n points for each variable $N = n^d$ which grows rapidly with the dimension d , but it is quite possible that more sophisticated grid design could reduce it. If we can't afford to make N high enough, the solutions become sub-optimal, but they might still represent sensible investment strategies. Successful algorithms require for each time step optimization, numerical integration and function approximation. Quasi Monte Carlo (or even ordinary Monte Carlo) may be used to evaluate the integrals (it's done in Sections 5 and 6), but care must be exercised to make criteria smooth functions so that optimization methodology works as intended. The approach handles criteria like probable or expected shortfall. Those are not differentiable, but often this is only in the beginning as they from the second round on have entered integrals that make them smooth functions with derivatives of any order.

2. COMPUTATIONAL APPROACH

2.1. The Bellman equation. Let X_0, X_1, \dots be a vectorial Markov process influenced by decision vectors π_0, π_1, \dots under our control. The problem addressed is how $\pi_k = \pi_k(X_k)$ should be selected to steer X_K at some terminal point in time towards a result to our advantage. This will here mean a strategy leading to the minimum of $E\{H(X_K)\}$ for some criterion $H(x)$, for example one of those in 2.4 below. The target is in all our examples in terms of accumulated capital, a single variable, but nothing is gained by bringing that in at this stage. Optimal decisions may clearly depend on other variables as well, see Section 3. The formal condition behind everything is

that history and future (X_0, \dots, X_{k-1}) and (X_{k+1}, \dots, X_K) are for all k stochastically independent given $X_k = x$ and the decision $\pi_k = \pi_k(x)$.

This is the set-up of Markovian decision processes and leads to the classical equation in Bellman [3]. Introduce the function

$$(2.1) \quad C_k(x, \pi) = E\{H(X_K)|X_k = x, \pi_k = \pi\}$$

and suppose π for given x is selected as the decision vector $\hat{\pi}_k = \hat{\pi}_k(x)$ that minimizes $C_k(x, \pi)$. This is done for all x and k , and buried inside the function $C_k(x, \pi)$ is the condition that all future decisions are optimal in this sense. There is a joint recursion that applies to $C_k(x, \pi)$ and their minima $\hat{C}_k(x)$. If $p_{k+1}(x'|x, \pi)$ is the conditional density function for X_{k+1} given $X_k = x$ and $\pi_k = \pi$ at time k , then

$$(2.2) \quad C_k(x, \pi) = \int_{-\infty}^{\infty} \hat{C}_{k+1}(x')p_{k+1}(x'|x, \pi)dx';$$

where

$$(2.3) \quad \hat{C}_k(x) = \inf_{\pi} C_k(x, \pi);$$

see Bellman [3]. The scheme works backwards, starting at expiry where $\hat{C}_K(x) = H(x)$. Next $C_{K-1}(x, \pi)$ is obtained by (2.2) and $\hat{C}_{K-1}(x)$ by (2.3), and we go on until the initial $X_0 = x_0$ and its optimal decision vector $\hat{\pi}_0(x_0)$ has been reached. The sequence of vectors $\hat{\pi}_k(x)$ so computed can then be used to control the process whatever happens after $X_0 = x_0$ in the beginning.

2.2. Numerical solutions. Let $C_k^\dagger(x, \pi)$ and $\hat{C}_k^\dagger(x)$ be numerical approximations of the exact functions $C_k(x, \pi)$ and $\hat{C}_k(x)$ defined by the Bellman scheme (2.2) and (2.3). Suppose they are run on a grid of N points, say x_k^i for $i = 1, \dots, N$. How we choose it is important, but N does not necessarily have to be huge; see Sections 4 and 5. Details of grid selection are presented among the examples. With the grid taken care of the analogy of (2.2) and (2.3) becomes

$$(2.4) \quad C_k^\dagger(x_k^i, \omega) = \int_{-\infty}^{\infty} \hat{C}_{k+1}^\dagger(x)p_{k+1}(x|x_k^i, \pi)dx, \quad i = 1, \dots, N$$

and

$$(2.5) \quad \hat{C}_k^\dagger(x_k^i) = \inf_{\pi} C_k^\dagger(x_k^i, \pi), \quad i = 1, \dots, N.$$

As before, the scheme runs backwards, starting at $\hat{C}_K^\dagger(x) = H(x)$.

Implementing such a scheme demands the integral (2.4) and the minimum (2.5) which take numerical methods. There are for the minimization highly developed optimization procedures that are globally convergent and guaranteed to locate a local minimum; see Fletcher [15]. They will not work well with Monte Carlo inaccuracies in the functions $\hat{C}_k^\dagger(x_k^i, \pi)$, and care must be exercised to avoid that. The integration step raises a number of points. Fastest by far is Gaussian quadrature; see Press, Teukolsky, Vetterling and Flannery [27]. Those requires the integrand to have derivatives of high order which may seem incompatible with expected shortfall or value-at-risk, but the problem only occurs at the start of the recursion where $\hat{C}_K^\dagger(x) = H(x)$, and $\hat{C}_k^\dagger(x)$ is typically smooth further down. Sometimes the density function $p_{k+1}(x|x_k^i, \pi)$ is not available in

closed form and can only be sampled. Integration must now be carried out by ordinary or quasi Monte Carlo.

Whatever method employed there is the problem from the second round on that the integrand is typically in demand where it has not been computed. Gaussian quadrature grids sometimes take us around this obstacle so that $\widehat{C}_k^\dagger(x)$ is being computed on exactly those abscissas x_k^i needed for the evaluation of the next integral. Such techniques have been successfully used in stochastic particle filtering, see Bølviken and Storvik [11] and references therein, and they did work when tried on some of the examples of this paper. Yet the range of application seems too narrow, and Monte Carlo evaluation of integrals wouldn't be covered at all. Extrapolation of $\widehat{C}_k^\dagger(x)$ beyond the grid x_k^i where it has been computed therefore seems unavoidable. Almost half a century has passed since Bellman, Kalaba and Kotkin [4] suggested function approximation as a solution. Their motivation was problems in physics, but the idea is all the more attractive in finance where the underlying criterion often change slowly from one x to another. Lagrange interpolation and two-dimensional splines are used with the examples below, but there are many other possibilities, and with Monte Carlo integration regression smoothing as in Longstaff and Schwarz [24] may be more attractive than strict interpolation. There is also a choice between local approximations as in Section 4 and global ones as in the Longstaff-Schwarz article.

2.3. Numerical error. The question is how numerical error propagates as the recursion evolves. A partial answer is provided by the following lemma:

Lemma *Suppose function approximation is the only source of error. The approximate and exact values of the objective functions then satisfy*

$$(2.6) \quad \max_i |\widehat{C}_k^\dagger(x_k^i) - \widehat{C}_k(x_k^i)| \leq \sup_x |\widehat{C}_{k+1}^\dagger(x) - \widehat{C}_{k+1}(x)|$$

for $k = 0, \dots, K - 1$.

This is an error diminishing property. The maximum error at time k among the grid points is smaller than its maximum at $k + 1$, but of course inaccuracy goes further up when the function $\widehat{C}_k^\dagger(x)$ during the next iteration is extrapolated beyond the points x_k^i where it has been computed. It is possible to run a similar argument when integration error is included. Again there is in the set-up an contraction of old error as the recursion proceeds, but when Monte Carlo and quasi Monte Carlo are used as in Sections 5 and 6 numerical inaccuracy is larger. By contrast the optimization step may not bring much more error if $C_k^\dagger(x, \pi)$ are smooth functions of π .

Some control of the aggregate is obtained by introducing ϵ as the maximum error of all function approximations involved. With integration and optimization error disregarded it follows from the lemma that $(K - k)\epsilon$ is an upper bound on the error in $\widehat{C}_k^\dagger(x)$. Numerical methods may guarantee a very small ϵ . With Lagrange interpolation in the next section we may for the smooth and slowly varying value functions in finance organize things so that ϵ is of order say 10^{-6} and smaller. The algorithm may under such circumstances run for a long time without being much disturbed by numerical error.

The lemma is proved by subtracting the exact $C(x_k^i, \pi)$ from its approximation $C_k^\dagger(x_k^i, \pi)$. It then follows by (2.2) and (2.4) that

$$C_k^\dagger(x_k^i, \pi) - C_k(x_k^i, \pi) = \int_{-\infty}^{\infty} \{\widehat{C}_{k+1}^\dagger(x') - \widehat{C}_{k+1}(x')\} p_{k+1}(x'|x_k^i, \pi) dx'.$$

Hence

$$|C_k^\dagger(x_k^i, \pi) - C_k(x_k^i, \pi)| \leq \int_{-\infty}^{\infty} |\widehat{C}_{k+1}^\dagger(x') - \widehat{C}_{k+1}(x')| p_{k+1}(x'|x_k^i, \pi) dx'$$

so that

$$(2.7) \quad \sup_{x_k^i, \pi} |C_k^\dagger(x_k^i, \pi) - C_k(x_k^i, \pi)| \leq \sup_x |\widehat{C}_{k+1}^\dagger(x) - \widehat{C}_{k+1}(x)|$$

since the integral over the density function $p_{k+1}(x'|x_k^i, \pi)$ is one. Fix j , let $\delta > 0$ and select π_δ so that $C_k(x_k^j, \pi_\delta) \leq \inf_\pi C_k(x_k^j, \pi) + \delta$. Then

$$\inf_\pi C_k^\dagger(x_k^j, \pi) - \inf_\pi C_k(x_k^j, \pi) \leq C_k^\dagger(x_k^j, \pi_\delta) - C_k(x_k^j, \pi_\delta) + \delta \leq \sup_{x_k^i, \pi} |C_k^\dagger(x_k^i, \pi) - C_k(x_k^i, \pi)| + \delta.$$

Since this holds for all $\delta > 0$, we must have

$$\inf_\pi C_k^\dagger(x_k^j, \pi) - \inf_\pi C_k(x_k^j, \pi) \leq \sup_{x_k^i, \pi} |C_k^\dagger(x_k^i, \pi) - C_k(x_k^i, \pi)|,$$

and this must be equally valid if $C_k^\dagger(x_k^j, \pi)$ and $C_k(x_k^j, \pi)$ change places so that

$$|\inf_\pi C_k^\dagger(x_k^j, \pi) - \inf_\pi C_k(x_k^j, \pi)| \leq \sup_{x_k^i, \pi} |C_k^\dagger(x_k^i, \pi) - C_k(x_k^i, \pi)|.$$

The right hand side is the left hand side of (2.7) so that for all j

$$|\inf_\pi C_k^\dagger(x_k^j, \pi) - \inf_\pi C_k(x_k^j, \pi)| \leq \sup_x |\widehat{C}_{k+1}^\dagger(x) - \widehat{C}_{k+1}(x)|,$$

and the lemma follows.

2.4. Criteria. Optimal solutions depend heavily on what optimal means. In the examples below the target is capital Y_K at some terminal date, but the specification of $H(y)$ still remains. One possibility is $H(y) = -U(y)$ where U is one of the utility functions in economic literature. This leads to the maximization of expected utility, but the thrust of the present paper is more towards downside criteria such as

$$(2.8) \quad H(Y_K) = I(Y_K < b) \quad \text{and} \quad H(Y_K) = (b - Y_K)_+$$

where $I(Y_K < b) = 1$ and $(b - Y_K)_+ = b - Y_K$ if $Y_K < b$ with both being zero otherwise. Their expected terminal values $E\{H(Y_K)\}$ are then probable and expected shortfall, $P(Y_K < b)$ and $E(b - Y_K)_+$. Schemes that minimize them will be developed below.

It is also possible to introduce Lagrange set-ups and through those get hold of minimum risk procedures with expectations $E(Y_K)$ assigned a given level. For example, let

$$(2.9) \quad H(Y_K) = -Y_K + \lambda Y_K^2$$

where $\lambda > 0$ is a fixed coefficient. Suppose we have found decision functions $\hat{\pi}_0, \dots, \hat{\pi}_{K-1}$ that minimize the expectation of (2.9). This means that

$$-\hat{E}(Y_K) + \lambda \hat{E}(Y_K^2) \leq -E(Y_K) + \lambda E(Y_K^2)$$

where \hat{E} operates under the optimal strategies $\hat{\pi}_0, \dots, \hat{\pi}_{K-1}$ and E under an arbitrary different one π_0, \dots, π_{K-1} . But if expectations are equal so that $\hat{E}(Y_K) = E(Y_K)$, then $\hat{E}(Y_K^2) \leq E(Y_K^2)$ which implies that $\widehat{\text{var}}(Y_K^2) \leq \text{var}(Y_K^2)$ as well. All these optimum strategies are therefore minimum variance ones, and the efficient K -period mean/variance frontier is computed by varying λ . This can be done for all models and in all situations where numerical Bellman recursions work.

The same trick may be attempted when probable and expected shortfall replace variance as risk measure. Now similar to (2.9)

$$(2.10) \quad H(Y_K) = -Y_K + \lambda I(Y_K < b) \quad \text{and} \quad H(Y_K) = -Y_K + \lambda(b - Y_K)_+.$$

Suppose $E\{H(Y_K)\}$ has been minimized for one of them, say the first one. Then the shortfall probability $\hat{P}(Y_K < b)$ for the optimum strategies is less than $P(Y_K < b)$ for any other strategy for which $E(Y_K) = \hat{E}(Y_K)$, and an efficient frontier is found when λ is varied. If a value-at-risk calculation rather than a shortfall one is sought, we must adjust b so that $\hat{P}(Y_K < b) = \epsilon$ for a given ϵ . The solution \hat{b}_ϵ of this equation is then a maximum value-at-risk value, and as λ is varied, $\hat{E}(Y_K)$ and \hat{b}_ϵ define an efficient frontier. When the argument is applied to expected instead of probable shortfall, we obtain an efficient conditional value-at-risk frontier of the same type. Examples of such calculations are presented in Section 5.

2.5. Numerical illustration. A simple example where numerical approximations can be compared to an exact solution is provided by the Merton problem where the investor chooses between one risky asset following a geometric Brownian motion and another one earning fixed return, see Merton [25]. It is convenient to write the value process S_t of the risky asset as

$$(2.11) \quad \frac{dS_t}{S_t} = \xi_R dt + \sigma_R dW_t$$

where ξ_R and σ_R are the same parameters as in the discrete time analogue (4.1) and (5.1) below, and W_t is the standardized Wiener process. With π_t being the weight on the risky asset at time t the portfolio value Y_t grows according to $dY_t = (1 - \pi_t)r_0 dt + \pi_t dS_t$ where r_0 is a risk-free spot rate of interest.

The optimal choice of π_t reduces to a strikingly simple form if the goal of the investor is to maximize at some terminal time T the expected CRRA utility $E\{U(Y_T)\}$ where $U(y) = y^c/c$ for a risk aversion parameter $c < 1$. This is a popular choice in academic literature, perhaps because it makes optimal solutions independent of wealth, see Brandt, Goyal, Santa-Clara and Stroud [9] for an application in our context. It can be shown (Boyle, Imai and Tan [8]) that the optimal value of π_t is

$$(2.12) \quad \hat{\pi}_t = \frac{\xi_R - r_0}{\sigma_R^2(1 - c)}$$

which defines the fixed mix strategy where the weight depends neither on t nor on what happens in the market.

K	$U(y) = -y^{-2}/2$				$U(y) = -y^{-5}/5$				$U(y) = -y^{-10}/10$			
	4	12	100	Merton	4	12	100	Merton	4	12	100	Merton
$\hat{\pi}_0$.3330	.3332	.3333	.3333	.1662	.1665	.1666	.1667	.0906	.0908	.0909	.0909

TABLE 1. *Optimal weight $\hat{\pi}_0$ for the discrete time Merton problem when the number of rebalances K is varied.*

This situation is a continuous time special case of the model in Section 5, and we have tested whether one of our programs there (the quadrature version) is able to reconstruct this solution. Investments were followed over one year with $\xi_R = 0.08$, $\sigma_R = 0.20$ and $r_0 = 0.04$ with the number of rebalances K varied from 4 (quarterly) to 100 (twice a week) as shown in Table 1 for three utility functions. The exact Merton weight (corresponding to infinite K) is reproduced quite closely even when $K = 4$.

3. MARKOV MODELLING

How examples are entered the framework defined above is now illustrated when the target is the capital Y_K accumulated at some termination date. Processes must be Markovian, and this means additional variables in the state vectors X_k . The dimension should be kept low to make algorithms work as effectively as possible. Detailed models are not needed at this stage.

3.1. Cash and equity. How investments should be divided between cash and equity is one of the classics and will here serve as a core example to be extended in several directions. With a single asset of equity the capital Y_k of the investor evolves according to the recursion

$$(3.1) \quad Y_{k+1} = \{1 + (1 - \pi_k)r_{k+1} + \pi_k R_{k+1}\}Y_k, \quad k = 0, \dots, K - 1$$

where π_k is the weight on equity at time k , R_{k+1} its return during the ensuing period and r_{k+1} the interest rate earned by the cash account. The problem is to choose the weights π_0, \dots, π_{K-1} dynamically as the process unfolds. Standard assumptions are interest rates following a Markov process (for example Cox-Ingersol-Ross or Vasiček) and equity returns that are stochastically independent. Random terms driving the two processes may well be correlated. Under these circumstances all information about the future resides in Y_k and r_k and the state process is the two-dimensional $X_k = (Y_k, r_k)$ with optimum weights of the form $\hat{\pi}_k = \hat{\pi}_k(y, r)$.

3.2. Many equity classes. Suppose there are several sub-classes of equity as in Brandt, Goyal, Santa-Clara and Stroud [9]. With J such classes the recursion (3.1) becomes

$$(3.2) \quad Y_{k+1} = \left(1 + (1 - \sum_{j=1}^J \pi_{jk})r_{k+1} + \sum_{j=1}^J \pi_{jk}R_{jk+1} \right) Y_k,$$

with J equity returns $R_{1k+1}, \dots, R_{Jk+1}$ and J weights $\pi_{1k}, \dots, \pi_{Jk}$. This seemingly more complicated problem might not be *that* much harder computationally. A joint model such as a dependent log-normal is needed for the equity returns, but in the absence of transaction costs the state vector can still be $X_k = (Y_k, r_k)$, and the optimum weights $\hat{\pi}_{jk} = \hat{\pi}_{jk}(y, r)$ are cast in the same

mould as before. What will be a bit more laborious is the optimization step (2.3) which is now with respect to J weights instead of one, but modern optimization software is equal to the task. Another problem is the density function $p_{k+1}(x'|x, \omega)$ of transitions from X_k to X_{k+1} now being unavailable in closed form since a sum of dependent log-normal variables is involved, but that can be overcome by Monte Carlo.

3.3. Bonds added. Suppose a bond earning risk premium over bank accounts is added. Writing R_{ek} and R_{bk} for the returns on equity and bond and π_{ek} and π_{bk} for their weights the recursion (3.1) now is extended to

$$(3.3) \quad Y_{k+1} = \{1 + (1 - \pi_{bk} - \pi_{ek})r_{k+1} + \pi_{bk}R_{bk+1} + \pi_{ek}R_{ek+1}\}Y_k.$$

The state vector depends on the model for the bond. A simple version is the one in Korn and Kraft [19] where the risk premium is fixed and the distribution of R_{bk+1} determined by the interest rate r_k . This would leave the same state vector $X_k = (Y_k, r_k)$ as before.

3.4. Transaction costs. Expenses due to rebalancing expand the state space. Consider again the cash and equity example (3.1), and suppose selling or buying stock entails cost that must be subtracted. Let w_k be the weight on equity at time k before it has been changed to the value we want. The recursion can then be written

$$(3.4) \quad Y_{k+1} = \{1 + (1 - \pi_k)r_{k+1} + \pi_k R_{k+1}\}(1 - \eta|\pi_k - w_k|)Y_k$$

$$(3.5) \quad w_{k+1} = \pi_k R_{k+1} / (1 + (1 - \pi_k)r_{k+1} + \pi_k R_{k+1})$$

where $\eta > 0$ defines rebalance cost. To make the future Markovian we now need the three-dimensional process $X_k = (Y_k, r_k, w_k)$. If there are J equity classes, we must keep track on the weights of all of them, and the state vector becomes $J + 2$ -dimensional.

4. MINIMUM EXPECTED SHORTFALL

The first example is strategies minimizing expected shortfall when there is a choice between one stock index and one money market investment. Their returns R_k and r_k run over an equidistant time sequence with increment $h > 0$ according to the recursions

$$(4.1) \quad \log(1 + R_k) = \xi_R h + \sigma_R \sqrt{h} \varepsilon_{Rk},$$

$$(4.2) \quad \log\left(\frac{r_k}{\xi_r}\right) = -\frac{\sigma_r^2 h}{4 - 2ah} + (1 - ah) \log\left(\frac{r_{k-1}}{\xi_r}\right) + \sigma_r \sqrt{h} \varepsilon_{rk}$$

for $k = 0, 1, \dots$. Here ξ_R , σ_R , ξ_r , σ_r and a are parameters and ε_{rk} and ε_{Rk} independent $N(0, 1)$ variables. Our implementation allows pairwise correlated error terms ε_{rk} and ε_{Rk} at the same point in time, but the experiments below assume independence. Equity log-returns $\log(1 + R_k)$ has mean ξ_R , and the model tends to a geometric Brownian motion as $h \rightarrow 0$. Interest rates r_k evolve through a first order auto-regression on log-scale, sometimes called the Black-Karinsinsky model, see James and Webber [18]. The first term on the right makes ξ_r mean rate of interest.

There are no transaction costs. It then follows as in Section 3 that $X_k = (Y_k, r_k)$ is a Markov process, and optimal weights on equity are of the form $\pi_k = \pi_k(y, r)$ where y is accumulated capital and r the interest rate at time k . The scheme in Section 2.2 was implemented with an

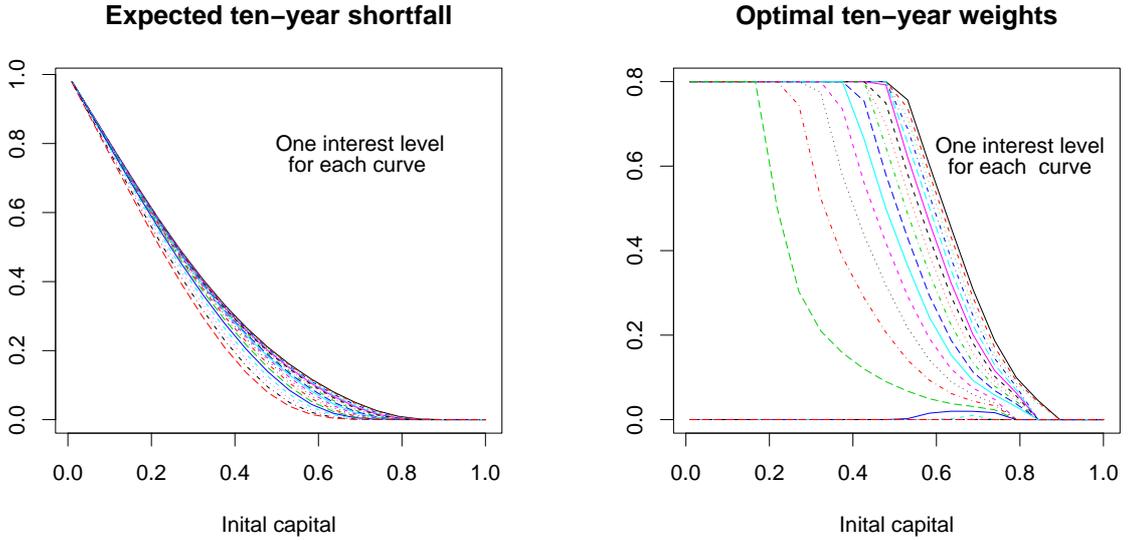


FIGURE 1. *Optimal ten-year expected shortfall (left) and optimal ten-year weights on equity (right) as a function of initial capital for different levels of the initial interest.*

equidistant grid in the state space and local four-point Lagrange interpolation for the function approximation. Gaussian quadrature was used to evaluate integrals (the Legendre type with twelve weights/abscissas) and Brent's inverse parabola method to find the optima; consult Press, Teukolsky, Vetterling and Flannery [27] for descriptions of these methods. Distributions were truncated at the $2.5 \cdot 10^{-7}$ percentile. Both Lagrange interpolation and Gaussian quadrature are one-dimensional methods that had to be applied in the y and r directions successively. The first integral required special care since the integrand has discontinuities in the first derivative, but this problem disappears thereafter, as remarked earlier. The grid was constructed as follows. If b is the threshold of the expected shortfall, then the optimum $\hat{C}_k(y) = 0$ and $\hat{\pi}_k(y, r) = 0$ if $y > b$ since the investor is certain to reach the goal by placing everything in the money market. We may therefore limit y to the interval $(0, b)$, and similarly r to a finite interval that ignores variation below or above extreme percentiles ($2.5 \cdot 10^{-7}$ with the experiments).

Results shown are for a period of 10 years with annual portfolio rebalances. Equity weights had to be between 0 and 0.8. Parameters were selected as $h = 1$, $\xi_R = 5.1\%$, $\sigma_R = 22.85\%$, $\xi_r = 4\%$, $\sigma_r = 0.357$ and $a = 0.3$. Annual volatilities of R_k and r_k are then 18.5% and 2.1% respectively, and interest rate may go up to 10 – 12% due to the heavy skewness of log-normal distributions. Figure 1 shows the solutions for planning ten years ahead when there are nine future possibilities for adjusting the portfolio. The threshold was $b = 1$, and $n = 20$ grid points was used in each direction (which took around 10 seconds to compute). On the left the minimum expected shortfall $\hat{C}_0(y, r)$ is plotted as functions of y for the 20 different values of the interest rates. The higher the interest rate the lower the curve. Note the smoothness that is so important for the iterative scheme to work well. The optimum weights $\hat{\pi}_0(y, r)$ on the right are downwards

n	percentiles (%)							
	0.1	1	5	25	50	75	90	99
5	-.609	-.389	-.076	.472	.682	.779	.892	1.243
10	-.600	-.379	-.061	.499	.632	.715	.838	1.323
20	-.600	-.381	-.063	.503	.623	.708	.825	1.335
50	-.602	-.384	-.066	.504	.624	.710	.828	1.336

TABLE 2. *Percentiles for ten-year returns when the number of grid points n per variable is varied.*

monotone functions of y ; i.e the less capital accumulated the more risk a downside risk criterion forces us to take on. There is again an ordering of the curves according to the interest rate. When money market earnings are moderate, more is placed in the stock market and quite sensitively so. The strange pattern in the extreme lower right corner is the combined effect of a discrete grid *and* lack of uniqueness. When zero expected shortfall is reachable, the optimum is an interval of weights.

The optimal solution was simulated to examine sensitivity of the size of the grid and to compare the solution with the fixed mix strategy. It was then assumed that the investor wanted 5% return of his investment per year or 62.9% over ten years which was used as the benchmark for the expected shortfall. Interest rate was 2% in the beginning. Monte Carlo simulations were run for ten years with annual rebalances, taking the weights from the optimal functions that had been computed (which required interpolation from the table of values available). Table 4 compares percentiles of ten-year returns for grids based on $n = 5, 10, 20$ and 50 points per variable. Results are quite stable and *not* heavily influenced by n . One million simulations were used (recycled for the four experiments), and random error is small. We might go down to 10 and perhaps even lower which is an important consideration when addressing problems of higher dimension.

Figure 2 shows ten-year returns. Those from the optimum expected shortfall strategy are plotted on the left against returns from equity only. The conservative line when the benchmark is likely to be reached is evident as simulations are kept back by the benchmark line once it has been exceeded. Comparisons with results from fixed mix strategy (Section 2.5) is presented on the right. The procedures were calibrated so that the medians were approximately equal which lead to much higher equity weights for fixed mix ($w = 0.76$ against the average 0.25 for minimum shortfall). This accounts for the much flatter percentile curves for expected shortfall corresponding to a much higher downside and much lower upside.

5. COMPARISONS

This section compares different criteria and integration methods. It will be demonstrated that quasi-Monte Carlo (slower computationally, but more widely applicable) may replace quadrature. There are as in the previous section two risky assets with models

$$(5.1) \quad \log(1 + R_k) = \xi_R h + \sigma_R \sqrt{h} \varepsilon_{Rk},$$

$$(5.2) \quad r_k = a \xi_r h + (1 - ah)r_{k-1} + \sigma_r \sqrt{r_k} \sqrt{h} \varepsilon_{rk}.$$

for $k = 1, 2, \dots$, using the same mathematical notation as before. Interest rates r_k now follow the popular Cox-Ingersoll-Ross model, see James and Webber [18], but optimal investments are still of the form $\pi_k = \pi_k(y, r)$.

The experiments below are run with a second implementation of the algorithm in Section 2.2 with interpolation in the state space carried out by two-dimensional splines. There is a number of variants here, see Duchon [13]. We have used univariate splines of the cubic type made two-dimensional through tensor products. Grids were equidistant with the same mesh everywhere, but dynamic in the number of points n used per variable. The spread of wealth (and to some degree interest rate) is highest in the beginning of the recursion. A possible strategy is therefore to reduce n gradually after specifying an initial maximum, and this line was followed in Sections 5 and 6 with $n = 5$ acting as a floor. Integrals were evaluated by Gaussian quadrature (the Hermite version with 20 abscissas/weights) and quasi Monte Carlo (one of the Sobols with 10000 simulations). Gauss-Hermite quadrature works excellently with the smooth CRRA criterion, but it could also be used with the downside risk provided some care was exercised with the first, non-smooth integrand, as remarked earlier. Optimization was carried out by Brent's method as in Section 4.

Evaluations below are annual and based on 20 rebalances of the portfolio free of charge, about two weeks and a half between each. Parameters were $\xi_R = 0.08$, $\sigma_R = 0.20$, $\xi_r = 0.04$, $\sigma_r = 0.0364$ and $a = 0.0824$ with $h = 0.05$. Initial values of capital and interest rate were $Y_0 = 1$ and $r_0 = 0.04$. The optimal weight $\hat{\pi}_0$ on equity for the coming year (with 19 additional rebalances planned) is shown in Table 3 under variation of technical features of the implementation. Its value is quite

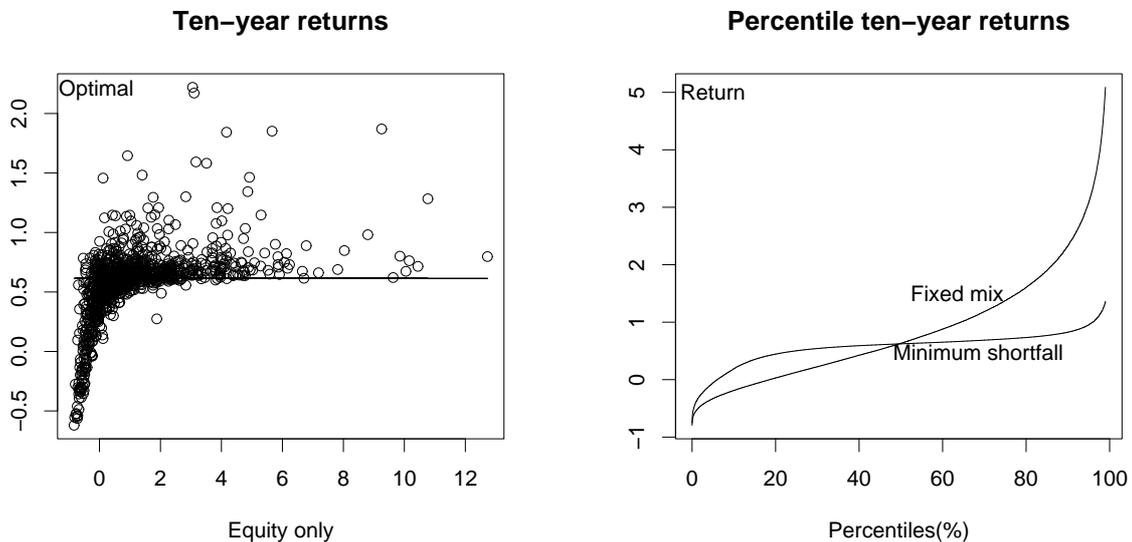


FIGURE 2. **Left:** Simulated ten-year returns for the the minimum shortfall strategy against those for equity with the benchmark as the horizontal line. **Right:** Percentiles curves for minimum shortfall and fixed mix.

Maximum n		$U(y) = -y^{-2}/2$				$H(y) = -y + 1.35(1.025 - y)_+$			
y	r	Quadrature		quasi Monte Carlo		Quadrature		Quasi Monte Carlo	
		$\hat{\pi}_0$	CPU	$\hat{\pi}_0$	CPU	$\hat{\pi}_0$	CPU	$\hat{\pi}_0$	CPU
49	21	0.3333	435	0.3326	3118	0.1882	4541	0.1889	1015
49	12	0.3333	242	0.3326	1809	0.1890	2569	0.1897	567
32	12	0.3336	148	0.3329	1223	0.1991	1633	0.1998	340
49	6	0.3333	123	0.3326	1025	0.1891	1353	0.1898	287
32	6	0.3336	81	0.3329	697	0.1991	950	0.1998	185
19	6	0.3341	49	0.3333	426	0.2700	560	0.2710	106
9	6	0.3396	29	0.3389	249	0.3274	470	0.3284	54
6	6	0.3229	23	0.3222	201	0.4230	344	0.4245	40

TABLE 3. *The initial optimal weight $\hat{\pi}_0$ and computer time (seconds) for different number of points n in the y and r -space.*

robust towards n , and the initial maxima (first two columns) may be lowered a good deal before much change can be observed. Nor does the choice of integration method amount to much, though quasi-Monte Carlo is distinctly slower computationally. The weights vary with the criterion as they must.

How strongly optimal risk depends on the criterion chosen is illustrated in Figure 3 where density functions of annual returns have been computed from 10000 simulated scenarios. Implementation was as in Section 4 with optimal weights for the twenty rebalances throughout the year taken from those calculated (interpolation necessary). The CRRA utilities on the left of Figure 3 yields allmost Gaussian risk with spread depending on the risk aversion factor c . Losses are punished harder as c is raised which leads to less use of the stock market and less uncertainty. The Gaussian form is entirely understandable through the central limit theorem since the optimal strategy is close to the fixed weight one. With downside criteria given by Equation (2.9) this is different. Minimizing expected shortfall as on the right of Figure 3 lead to spike-like risk profiles in a fairly narrow band with occasional large losses or gains. The ten-year returns in Figure 2 were similar. Shortfall at 0% or 4% means that losses are proclaimed for earnings less than 0% and 4%.

The last round of experiments was an attempt to calculate multi-period efficient frontiers in the Markowitz sense, but with value-at-risk instead of standard deviation. Value-at-risk is linked to probable shortfall as explained in Section 2.4. The risk profile under the latter has the same spiky behaviour as under expected shortfall and is shown in Figure 4 left. Efficient frontiers are plotted on the right. They contain some Monte Carlo error (10000 simulations used), but they do convey a general picture. Curves are quite steep which means that raising value-at-risk means serious losses in expected value. Indeed, they become progressively larger as we move from 90% value-at-risk to 99% since the curves become steeper and steeper.

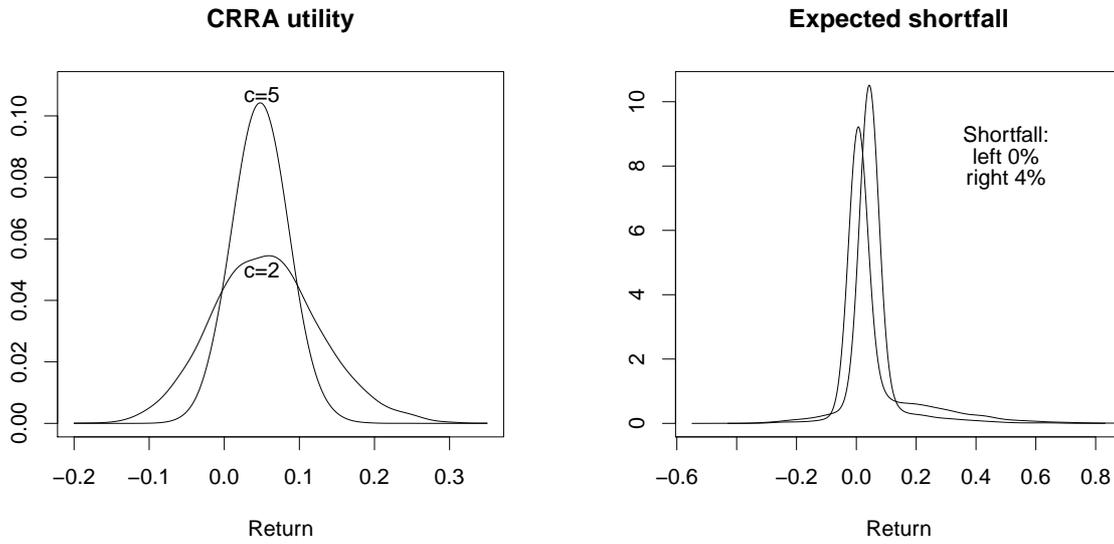


FIGURE 3. Density functions for annual returns under optimal strategies under CRRA utility (left) and expected shortfall right (right).

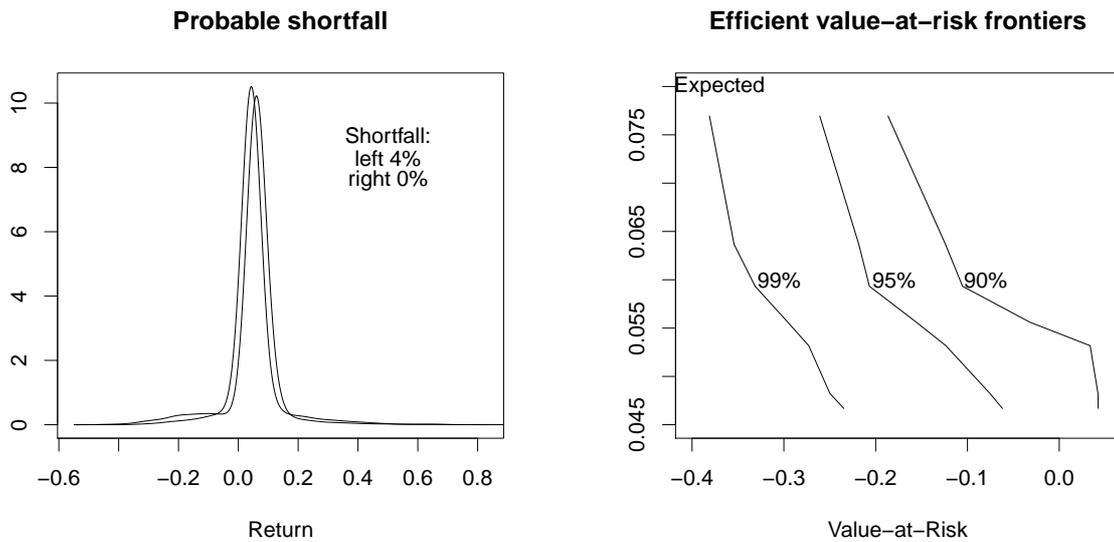


FIGURE 4. Density functions for annual returns under optimal strategies for probable shortfall (left) and mean/value-at-risk efficient frontiers for the same criterion (right).

6. HEAVY-TAILED DISTRIBUTIONS

What is the impact of heavy-tailed distributions on optimal portfolio strategies and risk? The problem is conveniently addressed through the Normal Inverse Gaussian (NIG) distributions for

	$U(y) = -y^{-2}/2$		$H(y) = -y + (0.975 - y)_+$	
	week	year	week	year
GBM	0.294	0.251	0.170	0.962
NIG	0.303	0.254	0.075	0.579

TABLE 4. *Optimal weights $\hat{\pi}_0$ under GBM and NIG models for weekly and annual time horizons.*

which the density function is

$$f(x) = \frac{\delta\alpha}{\pi} \exp\left(\delta\sqrt{\alpha^2 - \beta^2} + \beta(x - \mu)\right) \frac{K_1(\alpha\sqrt{\delta^2 + (x - \mu)^2})}{\sqrt{\delta^2 + (x - \mu)^2}}, \quad -\infty < x < \infty$$

where K_1 is the first-order modified Bessel function of the second kind. Among the four parameters, μ , δ , α and β , the second and third are positive and $\alpha > |\beta|$. The model has been popular in finance, and optimal investment under it has been studied in Benth, Karlsen and Reikvam[5, 6]. It becomes Gaussian if $\beta = 0$, $\delta = 1/\sqrt{\alpha}$ and $\alpha \rightarrow \infty$ and is heavy-tailed when α is smaller. There is also a convolution property which implies that NIG-distributed log-returns remain NIG on all time scales with the same α and β and with μ and δ proportional to the time increment h ; for these results and also sampling consult Rydberg [30].

The NIG-distribution is in this section used for equity log-returns whereas the interest rate model is the same as in Section 5. Its parameters were taken from the main index of the Oslo Stock Exchange (OSEBX). Daily log-returns from March 2005 to March 2010 gave moment estimates $\mu = 0.0032$, $\delta = 0.0167$, $\alpha = 42.7954$ and $\beta = -7.2915$. This corresponds to skewness -0.6093 and kurtosis 7.7573 . The large kurtosis bears evidence of heavier tails and a much better fit than the log-normal would have produced (as has often been observed). Consequences for financial risk will now be examined through optimal portfolios under both models. The parameters must then be selected to make mean and volatility of equity returns under the log-normal equal to those under NIG which requires $\xi_R = 0.1270$ and $\sigma_R = 0.3205$ in (5.1). Discrepancies between results are then due to the unequal shapes of the two distributions only. Technical implementation was as described in Section 5 except that numerical integration was carried out by ordinary Monte Carlo (100000 simulations) to escape the intricacies of quasi Monte Carlo sampling with NIG; but consult Benth, Groth and Kettler [7] to see how this could have been done.

The first round of experiments (in Table 4) illustrates how the optimum weight $\hat{\pi}_0$ in the beginning may depend on equity model, criterion and time to expiry. The log-normal corresponds to an ordinary geometric Brownian motion for equity and is referred to as GBM. Initial capital and rate of interest were $Y_0 = 1$ and $r_0 = 0.04$, and there are four rebalances for the weekly computation and 100 for the annual one. Note how little the distribution matters when we are using the utility function on the left and how much more sensitive it becomes with the other criterion on the right where losses are punished much harder.

Density functions of weekly returns following daily rebalancing are shown in Figure 5 when the optimal strategy was expected shortfall with $b = 0.975$, $\lambda = 1$ and $Y_0 = 1$ in the beginning. Investment rules were adapted to both the log-normal and the NIG equity model. There were

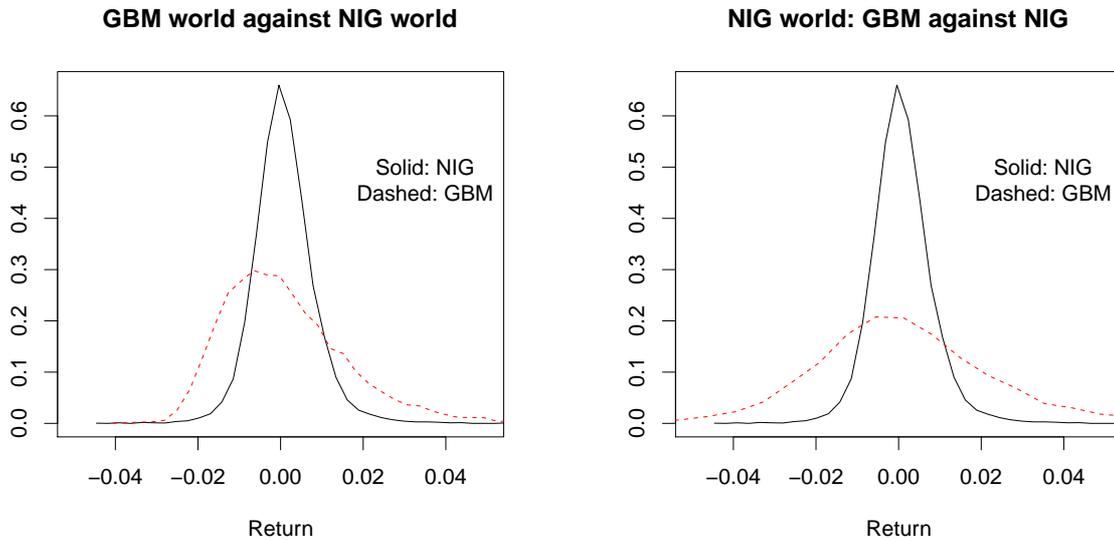


FIGURE 5. Weekly returns after daily rebalancing with a minimum shortfall strategy. **Left:** Results under GBM and NIG models, **right:** GBM and NIG strategies compared under the NIG model.

10000 Monte Carlo realisations of the portfolio with weights computed from the optimal ones dependent on how the simulations evolve. How terminal wealth is distributed under the two models are shown in Figure 5 left. The higher risk under the seemingly low-risk log-normal may appear paradoxical, but it is caused by the minimum shortfall procedure now tolerating much higher equity weights than under NIG. Another issue is performance when the model used to plan investments isn't the true one. That is illustrated in Figure 5 right where the optimal weights calculated under the log-normal has been applied to NIG-distributed scenarios which is the situation when the heavy-tailedness of daily equity returns is ignored. Risk now becomes much larger since much more money is invested in equity than the shortfall criterion deems prudent under NIG. The investor would be left with much less control than he or she imagines.

7. CONCLUSIONS AND FURTHER WORK

Practical experimentation and a theoretical argument have shown the power of Bellman recursions as a tool for the construction of optimal, multi-period financial strategies. Optimization, numerical integration and function approximation are required. For optimization software to work properly criteria functions must be kept smooth during the recursion. Only one-dimensional minimization whas been tried, but extensions to portfolio weights π_k that are vectors are straightforward. The dimension of the state space is much more critical. Not only does it blow up the size of the numerical grid, but it also limits the numerical tools available. Some of the most efficient ones like Lagrange or Chebyshev interpolation and Gaussian quadrature are inherently one-dimensional, and the curse of the dimensionality limits the number of directions they can be

applied. Quasi and even ordinary Monte Carlo (with common random numbers to impose as much smoothness as possible) may take care of the integration step at higher cost computationally.

The biggest challenge when dealing with problems of higher dimension may be grid construction and function approximation. Two-dimensional splines were used in Sections 5 and 6, and three-dimensional ones are available as well, but perhaps efforts based on non-cartesian grids are more promising. This may lead to computation on irregular patterns of points, but fast function approximation is still possible through radial basis functions (Buhmann [10]) or even regression (as in Longstaff and Schwartz [24]). It may also be that the criteria functions in finance that tend to vary fairly slowly as strategies are varied, allow a good deal of sub-optimality before the quality of the solution is degraded substantially. A lot remains before the potential of these possibilities has been explored.

REFERENCES

- [1] Balduzzi, P. and Lynch, A. W. (1999). Transaction costs and predictability: some utility cost calculations. *Journal of Financial Economics* Vol. 52, pp. 47–78.
- [2] Barberis, N. (2000). Investing for the Long Run when Returns Are Predictable. *Journal of Finance* Vol. 55, pp. 225–264.
- [3] Bellman, R.E. (1957). *Dynamic Programming*. Princeton University Press.
- [4] Bellman, R, Kalaba, R. and Kotkin, B. (1963). Polynomial approximation. A new Technique in Dynamic Programming: Allocation Processes. *Mathematics of Computation*, Vol. 17, pp. 155-161.
- [5] Benth, F.E., Karlsen, K.H. and Reikvam, K. (2001). A Note on Portfolio Management under non-Gaussian log-returns. *International Journal of Theoretical and Applied Finance*, 4, 711-731.
- [6] Benth, F.E., Karlsen, K.H. and Reikvam, K. (2001). Optimum Portfolio Management rules in a non-Gaussian Market with Durability and Intertemporal Substitution. *Finance and Stochastics*, 5, 447-467.
- [7] Benth, F.E., Groth, M. and Kettler, P. (2006). A quasi-Monte Carlo algorithm for the the normal inverse Gaussian distribution and valuation of financial derivatives. *International Journal of Theoretical and Applied Finance* Vol 9, No.5, pp 843-867.
- [8] Boyle, P., Imai, J. and Tan, K. S. (2008). Computation of optimal portfolios using simulation-based dimension reduction. *Insurance: Mathematics and Economics* Vol. 43, pp. 327–338.
- [9] Brandt, M. W., Goyal, A., Santa-Clara, P. and Stroud, J. R. (2005). A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Return Predictability. *Review of Financial Studies*, Vol. 18, No. 3, pp. 831–873.
- [10] Buhmann, M. (2003). *Radial Basis Functions: Theory and Implementations*. Cambridge University Press, Cambridge.
- [11] Bølviken, E. and Storvik, G. (2001). Deterministic and Stochastic Filters in State Space Models. In *Sequential Monte Carlo Methods in Practice*, Doucet, A., de Freitas, N. and Gordon, N (eds), Springer Verlag, New York, pp. 97-116.
- [12] Dickson, D. C. M. and Waters, H. R. (2006). Optimal Dynamic Reinsurance. *Astin Bulletin*, Vol. 36, No. 2, pp. 415–432.
- [13] Duchon, J. (1976). Splines minimizing rotation invariant seminorms in sobolev spaces, Constructive theory of functions of several variables, *Springer: Berlin / Heidelberg* pp. 85–100.
- [14] Fleming, W.H. and Soner, H.M. (2006). *Controlled Markov Processes and Viscosity Solutions*, 2nd edition Springer-Verlag, Berlin Heidelberg.

- [15] Fletcher, R. (1987). *Practical Methods of Optimization*. John Wiley & Sons, Chichester.
- [16] Garlappi, L. and Skoulakis, G. (2009). Numerical Solutions to Dynamic Portfolio Problems: The Case for Value Function Iteration using Taylor Approximation. *Computational Economics*, Vol. 33, pp. 193–207.
- [17] Hilli, P., Koivu, M., Pennanen, T. and Ranne, A. (2007). A Stochastic Programming Model for Asset Liability Management of a Finnish Pension Company. *Annals of Operations Research*, 152, pp. 115–139.
- [18] James, J. and Webber, N. (2000). *Interest Rate Modelling*. John Wiley & Sons, Chichester
- [19] Korn, R. and Kraft, H. (2001). A Stochastic Control Approach to Portfolio Problems with Stochastic Interest Rates. *Siam Journal on Control and Optimization*, Vol. 40, No. 4, pp. 1250–1269.
- [20] Korn, R. and Wiese, A. (2008). Optimal Investment and Bounded Ruin Probability: Constant Portfolio Strategies and Mean-Variance Analysis. *Astin Bulletin*, Vol. 38, No. 2, pp. 423–440.
- [21] Krokmal, P. and Uryasev, S. (2007). A Sample-Path Approach to Optimal Position Liquidation. *Annals of Operations Research*, 152, pp 193–225.
- [22] Krokmal, P., Uryasev, S. and Palmquist, J. (2002). Portfolio Optimization with Conditional Value-at-Risk Objective and Constraints. *Journal of Risk*, 4, pp. 43–68.
- [23] Kushner, H.J. and Dupuis, P.G. (1992). *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York.
- [24] Longstaff, F. and Schwartz, E. (2001). Valuing American Options by Simulation: A Simple Least-Squares Approach. *The Review of Financial Studies*, Vol. 14, pp. 113–147.
- [25] Merton, R.C. (1971). Optimum consumption and portfolio rules in a continuous time model. *Journal of Economic Theory* Vol. 3, pp. 273–413.
- [26] Pham, H. (2009). *Continuous-time Stochastic Control and Optimization with Financial Applications*. Springer-Verlag, Berlin Heidelberg.
- [27] Press, W.H., Teukolsky, S.A., Vetterling, W.T and Flannery, B. (2007). *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, Cambridge.
- [28] Rockafeller, R.T. and Uryasev, S. (2000). Optimization of Conditional Value at Risk. *Journal of Risk*, 2, pp. 21–41.
- [29] Rogers, L.C.G. (2007). Pathwise Stochastic Optimal Control, *SIAM Journal on Control and Optimization*, 46, 1116–1132.
- [30] Rydberg, T. H. (1997). The normal inverse Gaussian Levy process: simulation and approximation. *Communications in Statistics: Stochastic Models*, Vol. 13, pp. 887–910.
- [31] Schmidli, H. (2008). *Stochastic Control in Insurance*. Springer-Verlag, Berlin Heidelberg.
- [32] Tapiero, C. S. and Sulem, A. (1994). Computational Aspects in Applied Stochastic Control. *Computational Economics*, Vol. 7, pp. 109–146.
- [33] van Binsbergen, J. H. and Brandt, M. W. (2007). Solving dynamic portfolio choice problems by recursing on optimized portfolio weights or on the value function? *Computational Economics*, Vol. 29, pp. 355–367.
- [34] Wallace, S. and Ziemba, W.T. (2005). *Applications of Stochastic Programming*. MPS-Siam series on Optimization.

(Erik Bølviken), UNIVERSITY OF OSLO AND GABLERWASSUM, DEPARTMENT OF MATHEMATICS, P.O. BOX 1053, BLINDERN, N-0316 OSLO, NORWAY

E-mail address: erikb@math.uio.no

(Pål Nicolai Henriksen), UNIVERSITY OF OSLO, DEPARTMENT OF MATHEMATICS, P.O. Box 1053, BLINDERN, N-0316 OSLO, NORWAY

E-mail address: paalnhe@math.uio.no