Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of Banking & Finance 29 (2005) 2821–2848

Journal of BANKING & FINANCE

www.elsevier.com/locate/jbf

Intertemporal asset allocation: A comparison of methods ☆

Jérôme Detemple [a,d,*], René Garcia [b,d,e], Marcel Rindisbacher [c,d]

[a] *Boston University School of Management, Boston, MA 02215, United States*
[b] *Université de Montréal, Montreal, Que., Canada*
[c] *Rotman School of Management, University of Toronto, Toronto, Ont., Canada*
[d] *CIRANO, Montreal, Que., Canada*
[e] *CIREQ, Montreal, Que., Canada*

Available online 25 March 2005

Abstract

This paper compares two recent Monte Carlo methods advocated for the computation of optimal portfolio rules. The candidate methods are the approach based on Monte Carlo with Malliavin Derivatives (MCMD) proposed by Detemple, Garcia and Rindisbacher [Detemple et al., 2003. A Monte-Carlo method for optimal portfolios. Journal of Finance 58, 401–406] and the approach based on Monte Carlo with regression (MCR) of Brandt, Goyal, Santa-Clara and Stroud [Brandt et al., 2003. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. Working paper, Wharton School]. Our comparisons are carried out in the context of various intertemporal portfolio choice problems with two assets, a risky asset and a riskless asset, and different configurations of the state variables. The specifications studied include a linear model with a single state variable admitting an exact solution and a non-linear model with two state variables that requires a purely numerical resolution. The accuracies of the candidate methods are compared. We provide, in particular, efficiency plots displaying the speed–accuracy trade-off for various selections

☆ Excellent research assistance was provided by Marc Boucher. We thank MITACS for financial support. The second author is grateful to Hydro-Quebec and the Bank of Canada for financial support.
* Corresponding author. Boston University School of Management 546E, 595 Commonwealth ave., Boston, MA 02215, United States. Tel.: +1 617 353 4297; fax: +1 617 353 6667.
E-mail address: detemple@bu.edu (J. Detemple).

doi:10.1016/j.jbankfin.2005.02.004

of the relevant simulation and discretization parameters. MCMD is shown to dominate in all the settings considered.

JEL Classification: G11

Keywords: Asset allocation; Monte Carlo methods; Malliavin derivatives; Regression method

1. Introduction

Asset allocation, in recent years, has become a topic of intense interest for academics and practitioners of finance. Partly fueling this interest has been the development of theoretical and numerical methods bringing elegant, but abstract models, in the realm of practical applications. The availability of this technology has made it possible to implement realistic models with large numbers of assets and state variables. This paper provides a detailed review of the most versatile methodology available, based on the Malliavin calculus. The implementation of this methodology in the context of asset allocation problems was initially carried out by Detemple et al. (2003). This paper will provide further insights into their Monte Carlo with Malliavin derivatives (MCMD) method and, in particular, will provide additional results relating to its performance. For this purpose, we review another recent approach, the Monte Carlo regression (MCR) approach of Brandt et al. (2003). A comparative study will be conducted to document the performances of the two candidate methods. This comparison will provide further evidence of the dominance of the MCMD method.

The traditional approach to asset allocation relies on the mean–variance analysis of Markowitz (1952). The central idea is the notion that higher risk warrants a higher average return. The optimal portfolio trades off these aspects and results in an allocation which is ex ante mean–variance optimal. This static approach is simple and intuitive and has had enormous practical and professional impact. Even today, it is still driving part of the academic research. It also underlies a vast majority of the products offered by financial intermediaries, as well as the advice provided by many financial planners.

Modern asset allocation theory, however, has long recognized the shortcomings of the mean–variance approach. One obvious element is that the static formulation of this theory does not capture the essence of the problem as investors are often faced with long horizons and changing financial markets. That is, it fails to capture the dynamic aspect of the problem. The breakthrough contribution of Merton (1971) lays the foundation for dynamic portfolio choice. The central idea is based on the observation that means and variances of asset returns do not stay constant over time, but change in response to economic and business conditions. This simple remark suggests the need to design investment strategies that take advantage of these fluctuations by limiting their welfare-reducing effects. Practically, investors should modify the mean–variance allocation so as to hedge against fluctuations in the opportunity set. Optimal portfolios are therefore composed of a standard, mean–variance component, as well as of intertemporal hedging components that provide insurance against shocks to return moments.

Yet, in spite of the appeal and the elegance of Merton's ideas, implementation has lagged. One has to wait the mid-1990s for an application, by Brennan et al. (1997), in a non-trivial context with four assets and four state variables. Their implementation applies numerical methods to solve the second-order partial differential equation (PDE) that characterizes the value function for the dynamic portfolio problem. In spite of the relatively simple nature of their setting the resulting portfolio rule displays striking variations, as it jumps from 0% to 100% of wealth in stocks, following small changes in the underlying variables.[1]

In order to bypass the difficulties encountered in PDE-based implementations of the model, Detemple et al. (2003)—hereafter DGR—propose a Monte Carlo method that relies on several advances in financial economics achieved during the last two decades.[2] The first of these is the application of the risk neutralization procedure (see Cox and Ross, 1976; Harrison and Kreps, 1979; Harrison and Pliska, 1981) to consumption–portfolio choice models. This development, pioneered by Pliska (1986), Karatzas et al. (1987) and Cox and Huang (1989), led to the resolution of consumption–portfolio problems in settings with complete markets and von Neumann–Morgenstern time separable preferences. Specifically, the approach permits a closed-form solution for the optimal consumption rate and its associated wealth process.[3] The second innovation is the application of the Malliavin calculus, i.e. the stochastic calculus of variations, in order to derive an explicit formula for the financing portfolio. This contribution, due to Ocone and Karatzas (1991), expresses the optimal portfolio as the expected value of a functional that depends on the Malliavin derivatives of the coefficients of the model, namely the interest rate and the market price of risk. To the non-expert the formula appears somewhat abstract as it is derived in the context of models with Ito price processes and unspecified adapted coefficients. For implementation DGR (2003) focus on a more parsimonious structure with a finite, but arbitrarily large, number of diffusion processes. In this context Malliavin derivatives can be characterized as solutions of stochastic differential equations. This, along with the representation of portfolios as expected values, suggests the use of Monte Carlo simulation for computation purposes. DGR develop this simulation-based approach and apply it to various market and preference structures. They also evaluate its performance against

[1] This feature may be due to instabilities associated with the PDE method employed. The non-linearity of the equation, the infinite natural domain of the state variables and the absence of natural boundary conditions when the domain is truncated may be factors contributing to the unstable behavior of the solution.

[2] Numerical methods for PDEs also suffer from the curse of dimensionality and can only deal with low dimensional problems. Monte Carlo methods, on the other hand, can be employed for large dimensional systems.

[3] As Cox and Huang (1989) point out, the optimal portfolio can also be characterized as the solution to a "simplified" PDE depending of the state price density. Tests carried out by DGR (2003) show that finite difference methods applied to this equation are dominated by MCMD.

finite-difference PDE methods and against the Monte Carlo covariation method of Cvitanic et al. (2003). Their results show that MCMD fares better in terms of accuracy and speed of computation.

This paper reviews the fundamentals of the approach advocated by DGR and provides complementary results about its performance. Another method, proposed recently for portfolio computation in dynamic settings, is the simulation-regression approach of Brandt et al. (2003). The Monte Carlo regression approach replaces the value function satisfying the Bellman equation by a Taylor approximation. The optimal portfolio for the optimization problem based on this approximation can be expressed in terms of conditional moments of the future value function and of the returns. A backward resolution procedure, combined with a regression approach, is then used to compute these moments. The resulting portfolio value is an approximation to the optimal portfolio in the original problem.

A comparison of the candidate methods is performed in the context of two models. The first model has a single state variable, the market price of risk, that follows an Ornstein–Uhlenbeck process. A closed-form solution is provided by Wachter (2002) for the case of von Neumann–Morgenstern (vNM) preferences with constant relative risk aversion. The availability of an explicit solution provides a benchmark true value that is easy to compute and renders the model particularly attractive for comparison of different numerical approaches. The second model is the non-linear model with two state variables developed by DGR (2003). Although the true portfolio value is not available it can be calculated by any convergent method, such as MCMD, using a sufficiently large number of replications and discretization points. Computation, in this case, is time intensive. Both models have von Neumann–Morgenstern preferences with constant relative risk aversion.

Our results document the respective properties of the candidate methods. Broadly speaking these results establish the dominance of MCMD over MCR. This latter method exhibits a varying degree of performance depending on the context. In the Ornstein–Uhlenbeck model of Wachter, calibrated to the data, it has substantial relative errors for certain parameter configurations. Accuracy seems to improve in models, such as the non-linear model with two state variables of DGR (2003), in which the distribution of future state variables is not too disperse. Nevertheless, efficiency studies in both of these settings show that it is dominated by MCMD by a factor of 10 or better (i.e. MCMD provides estimates at least 10 times more accurate for a given budget of computation time).

Section 2 describes the consumption–portfolio choice problem in a general model with complete markets and diffusive prices, and solves for optimal consumption. Section 3 details the optimal portfolio policy. The first part of that section provides a simple review of principles of Malliavin calculus geared towards applications in finance. The second part details the portfolio policy and discusses its structure and properties. Section 4 addresses computational matters and explains the numerical implementation of MCMD and MCR. Section 5 performs several comparative studies to evaluate the two candidate methods. Comparisons are carried out in the context of a model with a single state variable admitting a closed-form solution and a model with a pair of state variables requiring a numerical resolution. Conclusions are formulated in Section 6. An Appendix A contains the proofs of some of the results.

2. The economic model

We formulate a consumption–portfolio model in manner of Merton (1971). The investor operates in a frictionless financial market in which asset prices and state variables follow a joint diffusion process. The investor has a finite planning horizon $[0, T]$.

2.1. The financial market

The financial market has d risky assets (stocks and bonds) and one riskless asset. The price of risky asset i, $i = 1, \ldots, d$, is given by

$$\mathrm{d}S_{it} = S_{it}[(\mu_i(t, Y_t) - \delta_i(t, Y_t))\mathrm{d}t + \sigma_i(t, Y_t)\mathrm{d}W_t], \tag{1}$$

where μ_i is the drift, δ_i the dividend yield and σ_i the $1 \times d$ vector of volatility coefficients. These coefficients depend on a $k \times 1$ vector of state variables $Y = (Y_1, \ldots, Y_k)'$. The riskless asset pays interest at the rate $r(t, Y_t)$, which also depends on the state variables. For notational convenience we will write μ_t for the $d \times 1$ vector of expected risky asset returns at date t and σ_t for the $d \times d$ matrix of return volatilities. Similarly we will write r_t for the interest rate. We assume that σ is invertible at all times (i.e. the market is complete).

The price system described above uniquely induces the d-dimensional vector of market prices of risk $\theta_t = (\theta_{1t}, \ldots, \theta_{\mathrm{d}t})'$ defined by $\theta_t \equiv \sigma_t^{-1}(\mu_t - r_t\mathbf{1})$ where $\mathbf{1} = (1, \ldots, 1)'$ is the d-dimensional vector of ones. This vector captures the premia implicitly assigned by the financial market to the various sources of uncertainty affecting the economy, i.e. the Brownian motions. Thus, W_i carries a premium θ_{it} at time t. The associated state price density (SPD) is $\xi_{t,v} \equiv \exp\left(-\int_t^v (r_s + \frac{1}{2}\theta_s'\theta_s)\mathrm{d}s - \int_t^v \theta_s'\,\mathrm{d}W_s\right)$. This state price density is the stochastic discount factor used to value any asset with cash flows contingent on the sources of uncertainty W.

2.2. State variables

The k state variables $Y = (Y_1, \ldots, Y_k)'$ affect the opportunity set, i.e. the means and variances of asset returns and the riskfree rate. The market prices of risk and the interest rate can be chosen as state variables (by setting $Y_1 = r$ and $Y_j = \theta_j$, $j = 1, \ldots, d$). Additional state variables may include dividend–price ratios, firm sizes, revenues, and other factors needed to describe returns. The evolution of state variables is given by

$$\mathrm{d}Y_t = \mu^Y(t, Y_t)\mathrm{d}t + \sigma^Y(t, Y_t)\mathrm{d}W_t, \tag{2}$$

where $\mu^Y(t, Y_t)$ is the $k \times 1$ vector of drift coefficients and $\sigma^Y(t, Y_t)$ is a $k \times d$ matrix of volatility coefficients.

2.3. Consumption, portfolios and wealth

In this setting the investor consumes and invests in assets. Wealth is the value of assets, i.e. of the portfolio held. Let X_t denote wealth at time t. If π_t is the $d \times 1$ vector of proportions of wealth invested in the risky assets (hence $1 - \pi_t'\mathbf{1}$ is the proportion invested in the riskless asset) and c_t the amount withdrawn from the portfolio for consumption then wealth evolves according to

$$\mathrm{d}X_t = (X_t r_t - c_t)\mathrm{d}t + X_t \pi_t'[(\mu_t - r_t\mathbf{1})\mathrm{d}t + \sigma_t \mathrm{d}W_t] \tag{3}$$

subject to some initial value x.

2.4. Preferences

We consider investors with time-separable von Neumann–Morgenstern preferences who care about intermediate consumption as well as terminal wealth. A consumption-terminal wealth plan (c, X_T) is ranked according to the expected utility criterion

$$\max E\left[\int_0^T u(c_v, v)\mathrm{d}v + U(X_T, T)\right]. \tag{4}$$

The utility functions u and U are assumed to be strictly increasing, strictly concave and differentiable. To simplify matters we assume that the utility functions satisfy the Inada conditions $\lim_{c\to 0} u'(c,t) = \lim_{X\to 0} U'(X,T) = \infty$ and $\lim_{c\to\infty} u'(c,t) = \lim_{X\to\infty} U'(X,T) = 0$, where $u'(c,t) \equiv \partial u(c,t)/\partial c$ and $U'(X,T) \equiv \partial U(X,T)/\partial X$.

2.5. The dynamic consumption–portfolio choice problem

The investor's objective is to maximize preferences

$$\max E\left[\int_0^T u(c_v, v)\mathrm{d}v + U(X_T, T)\right] \tag{5}$$

with respect to (c, π, X_T) subject to the following constraints:

$$\mathrm{d}X_t = (r_t X_t - c_t)\mathrm{d}t + X_t \pi_t'[(\mu_t - r_t\mathbf{1})\mathrm{d}t + \sigma \mathrm{d}W_t], \tag{6}$$

$$c_t \geqslant 0, \tag{7}$$

$$X_t \geqslant 0 \tag{8}$$

for all $t \in [0, T]$. The first constraint (6) is the dynamic budget constraint (i.e. the evolution of wealth given a consumption–portfolio policy (c, π)). The second constraint, (7), captures the physical restriction that consumption cannot become negative (i.e.

$c_t \geqslant 0$ for all $t \in [0, T]$). The last one, (8), stipulates that wealth cannot become negative. This restriction has the nature of a no-default condition.

It is well known that the assumption of infinite marginal utility at zero (the Inada condition) ensures that optimal consumption is non-negative and this guarantees a non-negative wealth (because wealth is the present value of future consumption). The constraints, in this case, can be ignored in the optimization problem. In the absence of Inada conditions, the constraints will bind and need to be incorporated in the optimization problem.

2.6. The equivalent static choice problem

In order to solve the dynamic consumption–portfolio choice problem described above it is useful to reformulate it as a static optimization problem. Cox and Huang (1989) and Karatzas et al. (1987) have shown that the equivalent static problem is to maximize

$$\max_{c, X_T} E\left[\int_0^T u(c_v, v)\mathrm{d}v + U(X_T)\right] \tag{9}$$

subject to the static budget constraint

$$E\left[\int_0^T \xi_v c_v \mathrm{d}v + \xi_T X_T\right] \leqslant x, \tag{10}$$

the inequality constraint (7) and the terminal consumption constraint $X_T \geqslant 0$. In this formulation the investor selects the consumption plan (c, X_T) subject to a static budget constraint, (10), and non-negativity constraints on consumption, (7) and $X_T \geqslant 0$. The budget constraint (10) stipulates that the present value of the plan must be bounded above by initial wealth. In other words the present value of expenditures cannot exceed the value of initial resources.

2.7. Necessary and sufficient conditions for optimality

The static problem is a constrained optimization problem which can be approached by forming the Lagrangian

$$\mathscr{L}(c, X_T, y) \equiv E\left[\int_0^T u(c_v, v)\mathrm{d}v + U(X_T)\right] + y\left(x - E\left[\int_0^T \xi_v c_v \mathrm{d}v + \xi_T X_T\right]\right),$$

where y is the positive multiplier for the static budget constraint (10), and solving $\min_{y>0}\max_{c\geqslant 0, X_T\geqslant 0}\mathscr{L}(c, X_T, y)$. The multiplier represents the shadow price of the budget constraint. The Lagrange optimization problem can be solved in two steps, first ignoring the non-negativity constraints $c \geqslant 0$, $X_T \geqslant 0$, then checking that the solution derived satisfies the required conditions.

The first order conditions for the Lagrange optimization problem (hence for the static problem (9), (10)),

$$u'(c_t, t) = y\xi_t, \tag{11}$$

$$U'(X_T, T) = y\xi_T, \tag{12}$$

$$E\left[\int_0^T \xi_v c_v \mathrm{d}v + \xi_T X_T\right] \leqslant x, \tag{13}$$

show that optimal consumption is selected so as to equate the marginal utility of an additional unit of consumption to the marginal cost. The latter is given by the state price density adjusted by the shadow price of the budget constraint. Conditions (11)–(13) are also sufficient due to the strict concavity of the utility function.

2.8. Optimal consumption and wealth

In order to solve the system of optimality conditions (11)–(13) let $I(y, t)$ and $J(y, T)$ be the respective inverses of the marginal utility functions $u'(c_t, t)$ and $U'(X_T, T)$. The strict concavity of the utility functions ensures that these inverse functions exist and that they are unique. Since marginal utilities range from 0 to ∞ the inverse functions are defined over the positive reals. Our candidate optimal consumption and terminal wealth are then given by the functions

$$c_t = I(y\xi_t, t), \tag{14}$$

$$X_T = J(y\xi_T, T). \tag{15}$$

Substituting these expressions in (13) shows that y satisfies the associated budget constraint $E[\int_0^T \xi_v I(y\xi_v, v)\mathrm{d}v + \xi_T J(y\xi_T, T)] = x$. Given that the inverse marginal utility functions $I(\cdot, \cdot)$ and $J(\cdot, \cdot)$ are strictly decreasing and range from $+\infty$ to 0, there is a unique solution y^* and this solution is strictly positive. The candidate optimal policies are then given by the functions (14), (15) evaluated at y^*. As the functions $I(\cdot, \cdot)$ and $J(\cdot,\cdot)$ are strictly positive these policies are indeed optimal for the constrained problem.

With the expressions for optimal consumption and terminal wealth it is easy to derive a formula for optimal wealth. Indeed, wealth represents the present value of future consumption and is therefore given by

$$X_t^* = E_t\left[\int_t^T \xi_{t,v} I(y^*\xi_v, v)\mathrm{d}v + \xi_{t,T} J(y^*\xi_T, T)\right] \tag{16}$$

for all $t \in [0, T]$. Given that the consumption functions $c_t^* = I(y^*\xi_t, t)$ and $X_T^* = J(y^*\xi_T, T)$ are non-negative, it is immediate to see that optimal wealth satisfies the constraint $X_t^* \geqslant 0$ at all times.

3. The optimal portfolio policy

The optimal portfolio policy is the portfolio policy that finances consumption and terminal wealth. Stated differently, it is the portfolio generating the optimal wealth process in (16).

The optimal portfolio corresponds to the integrand in the Martingale representation of the optimal discounted wealth process (see Section 3.1.5). The Clark–Ocone formula from Malliavin calculus gives an explicit expression for this integrand and therefore for the optimal portfolio. Our next section presents basic results from Malliavin calculus that are needed to derive this integrand.

3.1. An introduction to Malliavin calculus

The Malliavin calculus is a calculus of variations for stochastic processes. It applies to *Wiener (or Brownian) functionals*, i.e. random variables and stochastic processes that depend on the trajectories of Brownian motions. The Malliavin derivative, which is one element of this calculus of variations, measures the effect of a small variation in the trajectory of an underlying Brownian motion on the value of a Wiener functional.

3.1.1. Smooth Brownian functionals

To set the stage consider a Wiener space generated by the d-dimensional Brownian motion process $W=(W_1,\dots,W_d)'$. As is well known we can associate each state of nature with a trajectory of the Brownian motion (the set of states of nature is the space of trajectories). Let $(t_1,\dots,t_n)$ be a partition of the time interval $[0,T]$ and let $F(W)$ be a random variable of the form

$$F(W)\equiv f(W_{t_1},\dots,W_{t_n}),$$

where f is a continuously differentiable function. The random variable $F(W)$ depends (smoothly) on the d-dimensional Brownian motion W at a finite number of points along its trajectory; it is called a *smooth Brownian functional*.

3.1.2. The Malliavin derivative of a smooth Brownian functional

The Malliavin derivative of F is the change in F due to a change in the path of W. To simplify matters assume first that $d=1$, i.e. there is a unique Brownian motion. Consider shifting the trajectory of W by ε starting at time t. Suppose $t_k\leqslant t<t_{k+1}$ for some $k=1,\dots,n$. The Malliavin derivative of F at t is defined by

$$\begin{aligned}\mathscr{D}_tF(W)&\equiv\left.\frac{\partial f(W_{t_1}+\varepsilon\mathbf{1}_{[t,\infty[}(t_1),\dots,W_{t_k}+\varepsilon\mathbf{1}_{[t,\infty[}(t_k),\dots,W_{t_n}+\varepsilon\mathbf{1}_{[t,\infty[}(t_n))}{\partial\varepsilon}\right|_{\varepsilon=0}\\&=\lim_{\varepsilon\to0}\frac{F(W+\varepsilon\mathbf{1}_{[t,\infty[})-F(W)}{\varepsilon},\end{aligned}\tag{17}$$

where $\mathbf{1}_{[t,\infty[}$ is the indicator of the set $[t,\infty[$ (that is $\mathbf{1}_{[t,\infty[}(s)=1$ for $s\in[t,\infty[$; $=0$ otherwise). In more compact form we can write

$$\mathscr{D}_tF(\omega)=\sum_{j=k}^{n}\partial_jf(W_{t_1},\dots,W_{t_k},\dots,W_{t_n})\mathbf{1}_{[t,\infty[}(t_j),\tag{18}$$

where ∂_jf is the derivative with respect to the jth argument of f.

A simple example will illustrate the notion. Consider the price of the stock in the Black–Scholes model. Its value at date T is given by

$$S_T = S_0 \exp\left(\left(\mu - \frac{1}{2}\sigma^2\right)T + \sigma W_T\right),$$

where W_T is the terminal value of the univariate Brownian motion process defining the uncertainty in this model. Since $S_T = f(W_T)$ with $f(x) = S_0 \exp((\mu - \frac{1}{2}\sigma^2)T + \sigma x)$ it is clear that S_T is a smooth Brownian functional. A direct application of the definition gives

$$\mathscr{D}_t S_T = \partial f(W_T)\mathbf{1}_{[t,\infty[}(T) = \sigma S_0 \exp\left(\left(\mu - \frac{1}{2}\sigma^2\right)T + \sigma W_T\right) = \sigma S_T.$$

In this example the stock price depends only on the Brownian motion at time T. The Malliavin derivative is then the derivative with respect to W_T. This reflects the fact that a perturbation of the path of the Brownian motion from t onward, affects S_T only through the terminal value W_T.

Suppose next that $d > 1$, i.e. the underlying Brownian motion is multi-dimensional. The Malliavin derivative of F at t is now a $1 \times d$-dimensional vector denoted by $\mathscr{D}_t F = (\mathscr{D}_{1t}F, \ldots, \mathscr{D}_{dt}F)$. The ith coordinate of this vector, $\mathscr{D}_{it}F$, captures the impact of a perturbation in W_i by ε starting at some time t. If $t_k \leqslant t < t_{k+1}$ we have

$$\mathscr{D}_{it}F = \sum_{j=k}^{n} \frac{\partial f}{\partial x_{ij}}(W_{t_1}, \ldots, W_{t_k}, \ldots, W_{t_n})\mathbf{1}_{[t,\infty[}(t_j), \tag{19}$$

where $\partial f/\partial x_{ij}$ is the derivative with respect to the ith component of the jth argument of f (i.e. the derivative with respect to W_{it_j}).

3.1.3. The domain of the Malliavin derivative operator

The definition above can be extended to random variables that depend on the path of the Brownian motion over a continuous interval $[0, T]$. This extension uses the fact that a path-dependent functional can be approximated by a suitable sequence of smooth Brownian functionals. The Malliavin derivative of the path-dependent functional is then given by the limit of the Malliavin derivatives of the smooth Brownian functionals in the approximating sequence. The space of random variables for which Malliavin derivatives are defined is called $\mathbb{D}^{1,2}$. This space is the completion of the set of smooth Brownian functionals with respect to the norm $\|F\|_{1,2} = (E(F^2) + E(\int_0^T \|\mathscr{D}_t F\|^2 \mathrm{d}t))^{\frac{1}{2}}$ where $\|\mathscr{D}_t F\|^2 = \sum_i (\mathscr{D}_{it}F)^2$.

3.1.4. Malliavin derivatives of Riemann, Wiener and Ito integrals

This extension enables us to handle stochastic integrals, that depend on the path of the Brownian motion over a continuous interval, in a very natural manner. Consider, for instance, the stochastic Wiener integral $F(W) = \int_0^T h(t)\mathrm{d}W_t$, where $h(t)$ is a function of time and W is one-dimensional. Integration by parts shows that $F(W) = h(T)W_T - \int_0^T W_s \mathrm{d}h(s)$. Straightforward calculations give, for $t \in [0, T]$,

$$F(W + \varepsilon \mathbf{1}_{[t,\infty[}) - F(W) = h(T)(W_T + \varepsilon \mathbf{1}_{[t,\infty[}(T)) - \int_0^T (W_s + \varepsilon \mathbf{1}_{[t,\infty[}(s))\mathrm{d}h(s)$$
$$- \left(h(T)W_T - \int_0^T W_s \,\mathrm{d}h(s) \right)$$
$$= h(T)\varepsilon - \int_0^T \varepsilon \mathbf{1}_{[t,\infty[}(s)\mathrm{d}h(s) = \varepsilon h(t).$$

It then follows, from the definition (17), that $\mathscr{D}_t F = h(t)$, for $t \in [0, T]$. The Malliavin derivative of F at t is the volatility $h(t)$ of the stochastic integral at t: this volatility measures the sensitivity of the random variable F to the Brownian innovation at t.

Next, let us consider a random Riemann integral with integrand that depends on the path of the Brownian motion. This Brownian functional takes the form $F(W) \equiv \int_0^T h_s \mathrm{d}s$ where h_s is a progressively measurable process (i.e. a process which depends on time and the past trajectory of the Brownian motion) such that the integral exists (i.e. $\int_0^T |h_s|\mathrm{d}s < \infty$ with probability one). We now have

$$F(W + \varepsilon \mathbf{1}_{[t,\infty[}) - F(W) = \int_0^T (h_s(W + \varepsilon \mathbf{1}_{[t,\infty[}) - h_s(W))\mathrm{d}s.$$

As $\lim_{\varepsilon\to 0}(h_s(W + \varepsilon \mathbf{1}_{[t,\infty[}) - h_s(W))/\varepsilon = \mathscr{D}_t h_s(W)$ it follows that $\mathscr{D}_t F = \int_t^T \mathscr{D}_t h_s \mathrm{d}s$.

Finally, consider the Ito integral $F(\omega) = \int_0^T h_s(W)\mathrm{d}W_s$. To simplify the notation write $h^\varepsilon \equiv h(W + \varepsilon \mathbf{1}_{[t,\infty[})$ and $W^\varepsilon \equiv W + \varepsilon \mathbf{1}_{[t,\infty[}$. Integration by parts then gives, for $t \in [0, T]$,

$$F^\varepsilon - F = \int_0^T (h_s^\varepsilon - h_s)\mathrm{d}W_s + \int_0^T h_s^\varepsilon d(W_s^\varepsilon - W_s)$$
$$= \int_t^T (h_s^\varepsilon - h_s)\mathrm{d}W_s + h_T^\varepsilon (W_T^\varepsilon - W_T) - \int_0^T (W_s^\varepsilon - W_s)\mathrm{d}h_s^\varepsilon - \int_0^T \mathrm{d}[W^\varepsilon - W, h^\varepsilon]_s$$
$$= \int_t^T (h_s^\varepsilon - h_s)\mathrm{d}W_s + h_T^\varepsilon \varepsilon - \varepsilon \int_t^T \mathrm{d}h_s^\varepsilon$$
$$= \int_t^T (h_s^\varepsilon - h_s)\mathrm{d}W_s + \varepsilon h_t^\varepsilon.$$

The second equality above uses $h_s^\varepsilon = h_s$ for $s < t$ to simplify the first integral and the integration by parts formula to expand the second integral. The third equality is based on the fact that the cross-variation is null (i.e. $[W^\varepsilon - W, h^\varepsilon]_T = 0$) because $W_s^\varepsilon - W_s = \varepsilon \mathbf{1}_{[t,\infty[}(s)$ and $\mathbf{1}_{[t,\infty[}(s)$ is of bounded total variation.[4] The last equality uses, again, the integration by parts formula to simplify the last two terms. Since $\lim_{\varepsilon\to 0}(h_s^\varepsilon - h_s)/\varepsilon = \mathscr{D}_t h_s$ we obtain $\mathscr{D}_t F = h_t + \int_t^T \mathscr{D}_t h_s \,\mathrm{d}W_s$, for $t \in [0, T]$.

Malliavin derivatives of Wiener, Riemann and Ito integrals depending on multidimensional Brownian motions can be defined in a similar manner. As in Section

[4] The *total variation* of a function f is $\lim_{N\to\infty} \sum_{t_n \in \Pi^N([0,t])} | f(t_{n+1}) - f(t_n) |$ where $\Pi^N([0, t])$ is a partition with N points of the interval $[0, t]$.

3.1.2 the Malliavin derivative is a d-dimensional process which can be defined component-by-component, by the operations described above.

3.1.5. Martingale representation and the Clark–Ocone formula

In Wiener spaces martingales with finite variances can be written as sums of Brownian increments.[5] That is, $M_t = M_0 + \int_0^t \phi_s \mathrm{d}W_s$ for some progressively measurable process ϕ, which represents the volatility coefficient of the martingale. This result is known as the Martingale representation theorem. One of the most important benefits of Malliavin calculus is to identify the integrand ϕ in this representation. This is the content of the Clark–Ocone formula.

The Clark–Ocone formula states that any random variable $F \in \mathbb{D}^{1,2}$ can be decomposed as

$$F = \mathbf{E}[F] + \int_0^T \mathbf{E}_t[\mathscr{D}_t F]\mathrm{d}W_t \tag{20}$$

where $\mathbf{E}_t[\cdot]$ is the conditional expectation at t given the information generated by the Brownian motion W. For a martingale closed by $F \in \mathbb{D}^{1,2}$ (i.e. $M_t = \mathbf{E}_t[F]$) conditional expectations can be applied to (20) to obtain $M_t = \mathbf{E}[F] + \int_0^t \mathbf{E}_s[\mathscr{D}_s F]\mathrm{d}W_s$.

An intuitive derivation of this formula can be provided along the following lines. Assume that $F \in \mathbb{D}^{1,2}$. From the martingale representation theorem we have $F = \mathbf{E}[F] + \int_0^T \phi_s \mathrm{d}W_s$. Taking the Malliavin derivative on each side, and applying the rules of Malliavin calculus described above, gives $\mathscr{D}_t F = \phi_t + \int_t^T \mathscr{D}_t \phi_s \mathrm{d}W_s$. Taking conditional expectations on each side now produces $E_t[\mathscr{D}_t F] = \phi_t$ (since $\mathbf{E}_t[\int_t^T \mathscr{D}_t \phi_s \mathrm{d}W_s] = 0$ and ϕ_t is known at t). Substituting this expression in the representation of F leads to (20).

The results above also show that the Malliavin derivative and the conditional expectation operator commute. Indeed, let $v \geqslant t$ and consider the martingale M closed by $F \in \mathbb{D}^{1,2}$. From the representations for M and F above we obtain

$$\mathscr{D}_t M_v = \int_t^v \mathscr{D}_t \mathbf{E}_s[\mathscr{D}_s F]\mathrm{d}W_s + \mathscr{D}_t \mathbf{E}_t[F],$$

$$\mathscr{D}_t F = \int_t^T \mathscr{D}_t \mathbf{E}_s[\mathscr{D}_s F]\mathrm{d}W_s + \mathscr{D}_t \mathbf{E}_t[F].$$

Taking the conditional expectation at time $v \geqslant t$ of the second expression gives $E_v[\mathscr{D}_t F] = \int_t^v \mathscr{D}_t \mathbf{E}_s[\mathscr{D}_s F]\mathrm{d}W_s + \mathscr{D}_t \mathbf{E}_t[F]$. Since the formulas on the right hand sides of $\mathscr{D}_t M_v$ and $E_v[\mathscr{D}_t F]$ coincide we conclude that $\mathscr{D}_t M_v = E_v[\mathscr{D}_t F]$. Using the definition of M_v we can also write $\mathscr{D}_t E_v[F] = E_v[\mathscr{D}_t F]$: the Malliavin derivative operator and the conditional expectation operator commute.

[5] A *Wiener space* is the canonical probability space $(\mathscr{C}_0(\mathbb{R}_+;\mathbb{R}^d), \mathscr{B}(\mathscr{C}_0(\mathbb{R}_+;\mathbb{R}^d)), \mathbf{P})$ of nowhere differentiable functions $\mathscr{C}_0$, endowed with its Borel sigma field and the Wiener measure. The *Wiener measure* is the measure such that the d-dimensional coordinate mapping process is a Brownian motion.

3.1.6. The chain rule of Malliavin calculus

In applications one often needs to compute the Malliavin derivative of a function of a path-dependent random variable. As in ordinary calculus, a chain rule also applies in the Malliavin calculus. Let $F = (F_1, \ldots, F_n)$ be a vector of random variables in $\mathbb{D}^{1,2}$ and suppose that ϕ is a differentiable function of F with bounded derivatives. The Malliavin derivative of $\phi(F)$ is then

$$\mathscr{D}_t\phi(F) = \sum_{i=1}^{n} \frac{\partial\phi}{\partial x_i}(F)\mathscr{D}_t F_i,$$

where $\frac{\partial\phi}{\partial x_i}(F)$ represents the derivative relative to the ith argument of ϕ.

3.1.7. Malliavin derivatives of stochastic differential equations

For applications to portfolio allocation it is essential to be able to calculate the Malliavin derivative of the solution of a stochastic differential equation (i.e. the Malliavin derivative of a diffusion process). The rules of Malliavin calculus presented above can be used to that effect.

Suppose that a state variable Y_t follows the diffusion process $\mathrm{d}Y_t = \mu^Y(Y_t)\mathrm{d}t + \sigma^Y(Y_t)\mathrm{d}W_t$ where Y_0 is given and $\sigma^Y(Y_t)$ is a scalar (W is single dimensional). Equivalently, we can write the process Y in integral form as

$$Y_t = Y_0 + \int_0^t \mu^Y(Y_s)\mathrm{d}s + \int_0^t \sigma^Y(Y_s)\mathrm{d}W_s.$$

Using the results presented above, it is easy to verify that the Malliavin derivative $\mathscr{D}_tY_s$ satisfies

$$\mathscr{D}_tY_s = D_tY_0 + \int_t^s \partial\mu^Y \mathscr{D}_tY_v\,\mathrm{d}v + \int_t^s \partial\sigma^Y \mathscr{D}_tY_v\,\mathrm{d}W_v + \sigma(Y_t).$$

Since $\mathscr{D}_tY_0 = 0$, the Malliavin derivative obeys the following linear SDE:

$$d(\mathscr{D}_tY_s) = [\partial\mu^Y(Y_s)\mathrm{d}s + \partial\sigma^Y(Y_s)\mathrm{d}W_s](\mathscr{D}_tY_s) \tag{21}$$

subject to the initial condition $\lim_{s\to t}\mathscr{D}_tY_s = \sigma^Y(Y_t)$.

If Y is a $k \times 1$ vector and $\sigma^Y(Y_t)$ is a $k \times d$ matrix (W is a d-dimensional Brownian motion) the same arguments apply to yield

$$d(\mathscr{D}_tY_s) = \left[\partial\mu^Y(Y_s)\mathrm{d}s + \sum_{j=1}^{d}\partial\sigma_j^Y(Y_s)\mathrm{d}W_{js}\right](\mathscr{D}_tY_s) \tag{22}$$

subject to the initial condition $\lim_{s\to t}\mathscr{D}_tY_s = \sigma^Y(Y_t)$. In this multi-dimensional setting $\sigma_j^Y(Y)$ is the jth column of $\sigma^Y(Y)$. The Malliavin derivative $\mathscr{D}_tY_s$ is the $k \times d$ matrix $\mathscr{D}_tY_s = (\mathscr{D}_{1t}Y_s, \ldots, \mathscr{D}_{dt}Y_s)$.

3.2. The optimal portfolio

Let us now apply these notions to the derivation of the optimal portfolio. As mentioned before the optimal portfolio finances the wealth process (16). Equivalently, it represents the volatility of (16) (see the dynamic budget constraint (3)). The Clark–Ocone formula, which provides an expression for this volatility coefficient, can be applied to conclude that $X_t^*\pi_t^{*\prime}\sigma_t = \mathscr{D}_t X_t^*$. It then suffices to compute $\mathscr{D}_t X_t^*$ from (16) in order to derive the portfolio formula.

A renormalization of the optimal wealth process simplifies this computation. Consider the discounted wealth process $\xi_t X_t^*$. By Ito's lemma its volatility is $\xi_t X_t^*(\pi_t^{*\prime}\sigma_t - \theta_t')$. It follows immediately that $\xi_t X_t^*(\pi_t^{*\prime}\sigma_t - \theta_t') = \mathscr{D}_t(\xi_t X_t^*)$. Solving for π^* gives the following expression for the optimal portfolio:

$$\pi_t^* = (\sigma_t')^{-1}\theta_t + (\xi_t X_t^*)^{-1}(\sigma_t')^{-1}(\mathscr{D}_t(\xi_t X_t^*))', \tag{23}$$

where

$$\xi_t X_t^* = E_t\left[\int_t^T \xi_v I(y^*\xi_v, v)\mathrm{d}v + \xi_T J(y^*\xi_T, T)\right].$$

In this expression $c_v^* = I(y^*\xi_v, v)$ represents optimal consumption and $X_T^* = J(y^*\xi_T, T)$ is optimal terminal wealth. The basic rules of Malliavin calculus yield the expression

$$\mathscr{D}_t(\xi_t X_t^*) = E_t\left[\int_t^T Z_1(y^*\xi_v, v)\mathscr{D}_t\xi_v\,\mathrm{d}v + Z_2(y^*\xi_T, T)\mathscr{D}_t\xi_T\right], \tag{24}$$

where

$$Z_1(y^*\xi_v, v) = I(y^*\xi_v, v) + y^*\xi_v I'(y^*\xi_v, v) = c_v^*\left(1 - \frac{1}{R_u(c_v^*, v)}\right), \tag{25}$$

$$Z_2(y^*\xi_T, T) = J(y^*\xi_T, T) + y^*\xi_T J'(y^*\xi_T, T) = X_T^*\left(1 - \frac{1}{R_U(X_T^*, T)}\right), \tag{26}$$

where $I'(y^*\xi_v, v)$, $J'(y^*\xi_T, T)$ are the derivatives with respect to the first argument $y^*\xi$ of the inverse marginal utility functions and $R_u(x,v) = -u''(x,v)x/u'(x,v)$, $R_U(X,T) = -U''(X,T)X/U'(X,T)$ are the coefficients of relative risk aversion ($u''(c,t) \equiv \partial^2 u(c,t)/\partial c^2$ and $U''(X,T) \equiv \partial^2 U(X,T)/\partial X^2$). Similarly, from the definition of the stochastic discount factor

$$\xi_v \equiv \exp\left(-\int_0^v \left(r_s + \frac{1}{2}\theta_s'\theta_s\right)\mathrm{d}s - \int_0^v \theta_s'\mathrm{d}W_s\right)$$

we obtain

$$\mathscr{D}_t\xi_v \equiv -\xi_v\left(\int_t^v (\mathscr{D}_t r_s + \theta_s'\mathscr{D}_t\theta_s)\mathrm{d}s + \int_t^v \mathrm{d}W_s' \cdot \mathscr{D}_t\theta_s + \theta_t'\right).$$

The chain rule can then be used to write $\mathscr{D}_t\xi_v = -\xi_v(H_{t,v}' + \theta_t')$ with

$$H_{t,v}' = \int_t^v (\partial r(Y_s, s) + \theta_s'\partial\theta(Y_s, s))\mathscr{D}_t Y_s\mathrm{d}s + \int_t^v \mathrm{d}W_s' \cdot \partial\theta(Y_s, s)\mathscr{D}_t Y_s \tag{27}$$

and where $\mathscr{D}_t Y_s$ satisfies the stochastic differential equation

$$\mathrm{d}\mathscr{D}_t Y_s = \left[\partial\mu^Y(s, Y_s)\mathrm{d}s + \sum_{j=1}^{d} \partial\sigma_j^Y(s, Y_s)\mathrm{d}W_{js}\right]\mathscr{D}_t Y_s; \quad \mathscr{D}_t Y_t = \sigma^Y(t, Y_t). \tag{28}$$

Substituting (24)–(28) back in (23), collecting terms and simplifying produces our portfolio formula in Proposition 1 below (for details see Proof of Proposition 1 in Appendix A).

Proposition 1. *Consider the consumption–portfolio problem described in Section 2. The optimal consumption policy is* $c_v^* = I(y^*\xi_v, v)$ *and optimal terminal wealth is* $X_T^* = J(y^*\xi_T, T)$. *The optimal portfolio policy has the decomposition* $X_t^*\pi_t^* = X_t^*[\pi_{1t}^* + \pi_{2t}^*]$ *where* π_{1t}^* *is the mean–variance demand and* π_{2t}^* *the intertemporal hedging demand. The two parts are given by*

$$X_t^*\pi_{1t}^* = -E_t\left[\int_t^T \xi_{t,v}(y^*\xi_v)I'(y^*\xi_v, v)\mathrm{d}v + \xi_{t,T}(y^*\xi_T)J'(y^*\xi_T, T)\right](\sigma_t')^{-1}\theta_t,$$

$$X_t^*\pi_{2t}^* = -(\sigma_t')^{-1}E_t\left[\int_t^T \xi_{t,v}Z_1(y^*\xi_v, v)H_{t,v}\mathrm{d}v + \xi_{t,T}Z_2(y^*\xi_T, T)H_{t,T}\right],$$

where $Z_1(y^*\xi_v, v)$ *and* $Z_2(y^*\xi_T, T)$ *are defined in (25) and (26), the random variable* $H_{t,v}'$ *is defined in (27) and the Malliavin derivative of the state variables,* $\mathscr{D}_t Y_s$, *satisfies the stochastic differential equation (28). Optimal wealth is given by* $X_t^* = E_t[\int_t^T \xi_{t,v}I(y^*\xi_v, v)\mathrm{d}v + \xi_{t,T}J(y^*\xi_T, T)]$.

Proposition 1 shows that the optimal portfolio decomposes in two parts. The first one, π_1^*, is a mean–variance component motivated by the local properties of asset returns (which are embodied in the first two instantaneous moments). The second part, π_2^*, is an intertemporal hedging demand as defined in Merton (1971). This component is comprised of an interest rate hedge and a market price of risk hedge. The appearance of these two hedges follows from the fact that the stochastic discount factor (the SPD) is determined by r and θ. Since optimal consumption and terminal wealth are functions of the SPD the investor will naturally hedge against fluctuations in these quantities.

For constant relative risk aversion the portfolio simplifies as follows:

Proposition 2. *Suppose that the investor displays constant relative risk aversion R and has subjective discount factor* $\eta_t \equiv exp(-\beta t)$ *where* β *is a constant discount rate. The optimal consumption policy and the optimal terminal wealth are respectively given by* $c_v^* = (y^*\xi_v/\eta_v)^{-1/R}$ *and* $X_T^* = (y^*\xi_T/\eta_T)^{-1/R}$. *The optimal portfolio policy is given by* $X_t^*\pi_t^* = X_t^*[\pi_{1t}^* + \pi_{2t}^*]$ *where*

$$X_t^*\pi_{1t}^* = \frac{X_t^*}{R}(\sigma_t')^{-1}\theta_t,$$

$$X_t^* \pi_{2t}^* = -X_t^* \rho (\sigma_t')^{-1} \frac{E_t\left[\int_t^T \xi_{t,v}^\rho \eta_{t,v}^{1/R} H_{t,v} \mathrm{d}v + \xi_{t,T}^\rho \eta_{t,T}^{1/R} H_{t,T}\right]}{E_t\left[\int_t^T \xi_{t,v}^\rho \eta_{t,v}^{1/R} \mathrm{d}v + \xi_{t,T}^\rho \eta_{t,T}^{1/R}\right]}$$

with $\rho = 1 - 1/R$.

The portfolio formula of Proposition 2 serves as the basis for the models in which the different candidate methods will be evaluated.

4. Computing the optimal portfolio

In this section we review two recent approaches that have been proposed for the computation of asset allocation rules.

4.1. Monte Carlo with Malliavin derivatives (DGR, 2003)

In order to implement the portfolio formula in Proposition 2 we need to calculate the conditional expectations that appear in the hedging terms. These conditional expectations are taken over the random variables $\xi_{t,v}, H_{t,v}$ which in general will be path-dependent functionals of the state variables Y_s and their Malliavin derivatives $\mathscr{D}_t Y_s$. This complexity in the structure of the hedges naturally suggests the use of Monte Carlo simulation for computation purposes.

Our simulation approach proceeds as follows. First, note that the hedging demand can be written in the form

$$\pi_{2t}^* = -\rho (\sigma_t')^{-1} \frac{E_t[F_{t,T}]}{E_t[G_{t,T}]},$$

where $F_{t,T} \equiv F_{t,T}^c + F_{t,T}^x$ and $G_{t,T} \equiv G_{t,T}^c + G_{t,T}^x$, with

$$F_{t,v}^c \equiv \int_t^v \xi_{t,s}^\rho \eta_{t,s}^{1/R} H_{t,s} \mathrm{d}s \quad \text{and} \quad F_{t,T}^x \equiv \xi_{t,T}^\rho \eta_{t,T}^{1/R} H_{t,T},$$

$$G_{t,v}^c \equiv \int_t^v \xi_{t,s}^\rho \eta_{t,s}^{1/R} \mathrm{d}s \quad \text{and} \quad F_{t,T}^x \equiv \xi_{t,T}^\rho \eta_{t,T}^{1/R}.$$

To calculate π_2^* write the random variables appearing in the hedges as a joint system $(Y_v, \mathscr{D}_t Y_v, K_{t,v}, H_{t,v}, F_v^c)$, where

$$K_{t,v} \equiv \int_t^v \left(r_s + \frac{1}{2}\theta_s'\theta_s\right) \mathrm{d}s + \int_t^v \theta_s' \mathrm{d}W_s,$$

$$H_{t,v}' \equiv \int_t^v \partial r(Y_s, s) \mathscr{D}_t Y_s \mathrm{d}s + \int_t^v \theta_s' \partial\theta(Y_s, s) \mathscr{D}_t Y_s \mathrm{d}s + \int_t^v \mathrm{d}W_s' \cdot \partial\theta(Y_s, s) \mathscr{D}_t Y_s$$

and $\xi_{t,v} = \exp(-K_{t,v})$. By Ito's lemma we can write the dynamics of this system as

$$\mathrm{d}K_{t,s} = \left(r_s + \frac{1}{2}\theta_s'\theta_s\right)\mathrm{d}s + \theta_s'\mathrm{d}W_s,$$

$$\mathrm{d}H_{t,s}' = \partial r(Y_s, s)\mathscr{D}_t Y_s \mathrm{d}s + (\mathrm{d}W_s + \theta(Y_s, s)\mathrm{d}s)'\partial\theta(Y_s, s)\mathscr{D}_t Y_s,$$

$$\mathrm{d}F_{t,s}^c = \xi_{t,s}^{\rho}\eta_{t,s}^{1/R} H_{t,s}\mathrm{d}s,$$

$$\mathrm{d}G_{t,s}^c = \xi_{t,s}^{\rho}\eta_{t,s}^{1/R}\mathrm{d}s$$

with initial conditions $K_{t,t} = 0$, $H_{t,t}' = \mathbf{0}$, $F_{t,t}^c = \mathbf{0}$ and $G_{t,t}^c = 0$, along with Eqs. (2) and (28) for $(Y_s, \mathscr{D}_t Y_s)$.

Next, simulate M trajectories of the solutions of these equations. This can be performed using various discretization schemes, such as the Euler scheme or the Milshtein scheme. Let N be the number of (time) discretization points in the scheme selected. This simulation yields M estimates $\{(Y_s^{N,i}, \mathscr{D}_t^{N,i} Y_s, K_{t,s}^{N,i}, H_{t,s}^{N,i}, F_{t,s}^{c,N,i}, G_{t,s}^{c,N,i}) : s \in [t, T]\}$ of the trajectories $\{(Y_s, \mathscr{D}_t Y_s, K_{t,s}, H_{t,s}, F_{t,s}^c, G_{t,s}^c) : s \in [t, T]\}$. From the terminal values of the simulated processes construct M estimates of the random variables $F_{t,T}$ and $G_{t,T}$ in the hedging demand. Averaging over these M values yields the following estimate of the hedging demand:

$$\widehat{\pi_{2t}^*} = -\rho(\sigma_t')^{-1}\frac{\frac{1}{M}\sum_{i=1}^M F_{t,T}^{N,i}}{\frac{1}{M}\sum_{i=1}^M G_{t,T}^{N,i}} = -\rho(\sigma_t')^{-1}\frac{\sum_{i=1}^M F_{t,T}^{N,i}}{\sum_{i=1}^M G_{t,T}^{N,i}}.$$

The properties of this estimator are studied in Detemple et al. (2003). As shown in their paper the accuracy of the estimate depends on the number of Monte Carlo replications M and on the number of time discretization points N. Convergence to the true value is at the rate $1/\sqrt{M}$ as long as the ratio $\sqrt{M}/N$ is held constant (i.e. when M and N are simultaneously increased so as to leave the ratio $\sqrt{M}/N$ unchanged). For a single Monte Carlo replication, the Euler scheme converges weakly at the rate $1/\sqrt{N}$.

4.2. Monte Carlo with regression (BGSS, 2003)

Brandt et al. (2003) also present a simulation-based method for solving portfolio choice problems, but set in discrete time. The approach is based on the standard recursive dynamic programming algorithm, but instead of seeking closed form solutions, it simulates returns and state variables and uses a regression approach to calculate expectations. It also uses an approximation of the value function in order to calculate approximate "optimal" policies. This simulation-based approximation procedure is very general. As MCMD, it applies to large-scale problems with path-dependent and non-stationary dynamics as well as non-standard preferences. It therefore represents a prime candidate for comparison with MCMD. We summarize the main steps of the approach next, in the context of the pure portfolio problem (with utility of terminal wealth).

The algorithm is recursive in nature. The starting point is the Bellman equation for the value function, V, associated with the dynamic portfolio problem

$$V_t(X_t, Z_t) = \max_{\pi_t} E_t[V_{t+1}(X_t(\pi_t' R_{t+1}^e + R^f), Z_{t+1})]. \tag{29}$$

In this equation, X_t is the endogenous wealth at time t, Z_t is a vector of exogenous state variables at t, R_{t+1}^e is the vector of risky assets' excess returns from t to $t+1$, R^f is the return on the risk-free asset and finally, π_t is the portfolio over which the optimization is conducted. The first-order condition (FOC) for this optimization problem is

$$E_t[\partial_1 V_{t+1}(X_t(\pi_t' R_{t+1}^e + R^f), Z_{t+1}) R_{t+1}^e] = 0, \tag{30}$$

where $\partial_1 V_{t+1}$ is the derivative of the value function with respect to the first argument (future wealth associated with the portfolio policy π).

The method proposed by BGSS consists of three steps. The first step simplifies the initial optimization problem (29) by expanding the value function in a Taylor series around $X_t R^f$, the future value (at $t+1$) of current wealth. To account for departures from quadratic utility and Gaussian returns, BGSS propose the fourth-order expansion [6]

$$\begin{aligned} V_t^a(X_t, Z_t) = \max_{\pi_t} E_t \Bigg[& V_{t+1}^a(X_t R^f, Z_{t+1}) + \partial_1 V_{t+1}^a(X_t R^f, Z_{t+1})(X_t \pi_t' R_{t+1}^e) \\ & + \frac{1}{2} \partial_1^2 V_{t+1}^a(X_t R^f, Z_{t+1})(X_t \pi_t' R_{t+1}^e)^2 + \frac{1}{6} \partial_1^3 V_{t+1}^a(X_t R^f, Z_{t+1})(X_t \pi_t' R_{t+1}^e)^3 \\ & + \frac{1}{24} \partial_1^4 V_{t+1}^a(X_t R^f, Z_{t+1})(X_t \pi_t' R_{t+1}^e)^4 \Bigg], \end{aligned}$$

where V^a represents the value function for this new (approximate) problem. Let π_t^a be the solution of the approximate problem. The FOC leads to the following implicit expression for π^a:

$$\begin{aligned} \pi_t^a = & -\{E_t[\partial_1^2 V_{t+1}^a(X_t R^f, Z_{t+1}) R_{t+1}^e (R_{t+1}^e)'] X_t^2\}^{-1} \\ & \times \Bigg\{ E_t[\partial_1 V_{t+1}^a(X_t R^f, Z_{t+1}) R_{t+1}^e] X_t + \frac{1}{2} E_t[\partial_1^3 V_{t+1}^a(X_t R^f, Z_{t+1})((\pi_t^a)' R_{t+1}^e)^2 R_{t+1}^e] X_t^3 \\ & + \frac{1}{6} E_t[\partial_1^4 V_{t+1}^a(X_t R^f, Z_{t+1})((\pi_t^a)' R_{t+1}^e)^3 R_{t+1}^e] X_t^4 \Bigg\} \\ \equiv & -\{E_t[B_{t+1}] X_t\}^{-1} \{E_t[A_{t+1}] + E_t[C_{t+1}(\pi_t^a)] X_t^2 + E_t[D_{t+1}(\pi_t^a)] X_t^3\}. \end{aligned} \tag{31}$$

The structure of this equation shows that the solution π_t^a depends on conditional moments involving the derivatives of the value function and powers of the returns. Assume for the time being that these moments can be calculated by some procedure. The solution of (31) is then computed as follows. First, compute the solution of the quadratic problem corresponding to the second-order expansion of the value

[6] Brandt et al. (2003) report that a fourth-order expansion around $X_t R^f$ gives very accurate results for the particular problems they considered.

function. This gives an explicit expression which can be used as an initial guess for solving (31). Second, substitute this initial guess into the right-hand side of (31) to produce a new estimate of π^a on the left-hand side. Finally, iterate by repeating the previous step until the distance between consecutive estimates falls below some pre-selected tolerance level.

The second step of the method involves the forward simulation of a large number M of sample paths of the vector $Y_t = [R_t^e, Z_t]$ where Z_t is a vector of state variables which may include a finite number of past returns. This set of paths serves as the underlying tree for the application of a recursive procedure where the portfolio is approximated at each step and along each trajectory m by the solution of (31).

The third step consists in computing the expectations appearing in (31) and solving for the portfolio at date t. Suppose that approximate weights π_s^a have been found for $s = t+1, \dots, T-1$. The corresponding terminal wealth is

$$X_T^a = X_t^a R^f \prod_{s=t+1}^{T-1} (\pi_s^a R_{s+1}^e + R^f) \tag{32}$$

assuming that the riskfree rate applies between t and $t+1$. The coefficients in (31) are then replaced by

$$A_{t+1} = E_{t+1}\left[\partial u(X_T^a) \prod_{s=t+1}^{T-1} (\pi_s^a R_{s+1}^e + R^f)\right] R_{t+1}^e \tag{33}$$

in the case of A_{t+1}, and similar expressions for B_{t+1}, C_{t+1} and D_{t+1}. Letting $a_{t+1} \equiv \partial u(X_T^a)\prod_{s=t+1}^{T-1}(\pi_s^a R_{s+1}^e + R^f)R_{t+1}^e$ be the random variable inside the expectation in (33) and defining b_{t+1}, c_{t+1} and d_{t+1} in a similar manner leads to

$$\pi_t^a = -\{E_t[b_{t+1}]X_t^a\}^{-1}\{E_t[a_{t+1}] + E_t[c_{t+1}(\pi_t^a)](X_t^a)^2 + E_t[d_{t+1}(\pi_t^a)](X_t^a)^3\}. \tag{34}$$

This relation serves as an approximation of the optimal portfolio policy.

To find this approximation the expectations of a, b, c, d must be computed. To do this Brandt et al. (2003) rely on the regression method suggested by Longstaff and Schwartz (2001) in the context of American option pricing. This simple approach uses regressions across the simulated paths to evaluate conditional expectations. Let y be a typical element of the vector $[a, b, c, d]$. The expectation of y_{t+1} is computed by regressing y_{t+1} on a vector of polynomial bases in the state variables Z_t. That is

$$E_t[y_{t+1}] = \varphi(Z_t)' k_t$$

where k_t is the vector of regression parameters to be estimated. The fitted values of this regression are used to construct estimates of the time t-conditional expectations of a_{t+1}, b_{t+1}, c_{t+1} and d_{t+1}, along each path m. Solving (34) produces the approximate portfolio $\pi_t^{a,m}$.

The backward construction described above can be implemented for all indices t running from $T-1$ to 0. For the computation of the approximate portfolio π_0^a at the initial date it suffices to solve (34) with time index $t = 0$. At that date the expectations $E_0[a_1], E_0[b_1], E_0[c_1(\pi_0^a)]$ and $E_0[d_1(\pi_0^a)]$ on the right hand side of the equation are projections on a set of constants as there is a unique set of initial values for the state variables (i.e. the vector of independent variables Z_0 is single valued). This is equivalent to estimation by simple Monte Carlo averaging, i.e. for $x_1 = a_1, b_1, c_1(\pi_0^a), d_1(\pi_0^a)$ the estimate $\widehat{x}_1$ of the conditional expectation $E_0[x_1]$ is $\widehat{x}_1 = \frac{1}{M}\sum_{j=1}^{M} x_1^j$. Substituting in (34) and solving gives the portfolio estimate π_0^a.

5. A comparison of methods

In order to compare the candidate methods we focus on two specifications of the consumption–portfolio model. The first one is the linear specification with closed form solution proposed by Wachter (2002). The second is the non-linear model with two state variables, developed by DGR (2003). Since BGSS (2003) provide a detailed presentation of MCR for the maximization of the utility of terminal wealth we compare MCR and MCMD in this setting.

5.1. Two frameworks for comparison

Let us first describe the two models used for comparison.

5.1.1. A single state variable model with explicit solution (WA)

In the first model the investor has constant relative risk aversion R and maximizes utility in a financial market with a single risky stock and the riskless asset. The underlying source of uncertainty is a (single) Brownian motion W. The interest rate r is constant and the market price of risk θ follows the Ornstein–Uhlenbeck (OU) process

$$d\theta_t = A(\overline{\theta} - \theta_t)dt + \Sigma dW_t, \quad \theta_0 \text{ given}, \tag{35}$$

where A, $\overline{\theta}$, Σ are positive constants. The stock return has constant volatility σ. In the version of this model studied the investor cares about the expected utility of terminal wealth. We refer to this setting as WA.

The closed form solution for WA can be found in Wachter (2002).[7] Assume that the determinant condition

$$\Sigma^{-2}A^2 + \rho(1 + 2\Sigma^{-1}A) \geqslant 0, \tag{36}$$

holds, where $\rho = 1 - 1/R$, and define the constants

$$G = -\Sigma^{-1}A - \sqrt{\Sigma^{-2}A^2 + \rho(1 + 2\Sigma^{-1}A)}$$

and $\alpha = 2(A + \Sigma G)$. In the pure portfolio problem, the optimal demand for the stock is $\pi_t^* = \pi_{1t}^* + \pi_{2t}^*$ where $\pi_{1t}^* = (1/R)(\sigma_t)^{-1}\theta_t$ is the mean–variance demand and

[7] For this model, Wachter shows that the problem reduces to a system of Riccati ordinary differential equations. Liu (1998) and Schroder and Skiadas (1999) show that the same reduction applies when state variables follow affine processes.

$$\pi_{2t}^{*}=-\frac{\rho}{R}[B(t,T)+C(t,T)\theta_t]\Sigma\sigma^{-1}$$

with

$$B(t,T)=\frac{2\left(1-\exp\left(-\frac{1}{2}\alpha(T-t)\right)\right)^2}{\alpha(\alpha+(\rho-G)\Sigma(1-\exp(-\alpha(T-t))))}A\overline{\theta}, \tag{37}$$

$$C(t,T)=\frac{1-\exp(-\alpha(T-t))}{\alpha+(\rho-G)\Sigma(1-\exp(-\alpha(T-t)))}, \tag{38}$$

represents the intertemporal hedging demand.

5.1.2. A nonlinear model with two state variables (DGR1)

The second model is the benchmark non-linear model developed by DGR (2003), which we call DGR1. The state variables, in this setting, correspond to the interest rate and the market price of risk (r,θ) which evolve according to

$$\mathrm{d}r_t=\kappa_r(\bar{r}-r_t)(1+\phi_r(\bar{r}-r_t)^{2\eta_r})\mathrm{d}t-\sigma_r r_t^{\gamma_r}\mathrm{d}W_t,\quad r_0 \text{ given}, \tag{39}$$

$$\mathrm{d}\theta_t=(\kappa_\theta(\overline{\theta}-\theta_t)+\mu_\theta^r(r_t,\theta_t))\mathrm{d}t+\sigma_\theta(\theta_t)\mathrm{d}W_t,\quad \theta_0 \text{ given}, \tag{40}$$

where

$$\mu_\theta^r(r_t,\theta_t)\equiv\delta_r(\bar{r}-r_t)(\theta_l+\theta_t)\left(1-\left(\frac{\theta_l+\theta_t}{\theta_l+\theta_u}\right)\right), \tag{41}$$

$$\sigma_\theta(\theta_t)=\sigma_\theta(\theta_l+\theta_t)^{\gamma_{1\theta}}\left(1-\left(\frac{\theta_l+\theta_t}{\theta_l+\theta_u}\right)^{1-\gamma_{1\theta}}\right)^{\gamma_{2\theta}}. \tag{42}$$

The coefficients $(\kappa_r,\bar{r},\phi_r,\eta_r,\sigma_r,\gamma_r,\kappa_\theta,\overline{\theta},\eta_\theta,\sigma_\theta,\theta_l,\theta_u,\gamma_{1\theta},\gamma_{2\theta})$ are constants. Moreover, $(\kappa_r,\bar{r},\kappa_\theta,\theta_l,\theta_u)$ are positive, and $\overline{\theta}\in(-\theta_l,\theta_u)$. The Brownian motion W is unidimensional.

The interest rate process (39) is mean reverting with constant elasticity of variance (NMRCEV) given by $2\gamma_r$. The speed of mean-reversion is non-linear and modelled by the function $\kappa_r(1+\phi_r(\bar{r}-r_t)^{2\eta_r})$. When $\phi_r>0$, this specification implies an increased pull towards the long run mean as the deviation $|\bar{r}-r_t|$ increases. Empirical motivation for this structure is provided by DGR (2003). The presence of non-linearities in the mean is also documented by Ait-Sahalia (1996) and Ahn and Gao (1999). The market price of risk process exhibits linear mean reversion and has an elasticity of variance with hyperbolic structure. The drift also depends on the interest rate, which has been shown to be a good predictor of the market price of risk. The formulation adopted ensures that the process stays between two reflecting bounds, given by the parameters $-\theta_l$ and θ_u. The process is said to exhibit mean reversion with hyperbolic elasticity of variance and interest rate dependence in the drift (MRHEVID).

5.2. Comparing MCMD and MCR

This section reports comparison results for MCMD and MCR. Five versions of MCR are tested: three of these involve regressions on linear terms, the last two regress on power series of the returns. The MCR methods are labelled:

- *MCR-lin-1*. Regression on the market price of risk for the current date.
- *MCR-lin-2*. Regression on the market prices of risk for the current and previous dates.
- *MCR-lin-3*. Regression on the market prices of risk for the current and the past three dates.
- *MCR-poly-1*. Regression on a second-order power series of the market price of risk for the current date.
- *MCR-poly-2*. Regression on a third-order power series of the market price of risk for the current date.

Our first comparison is carried out in the context of WA. The availability of a closed-form solution for this model enables us to calculate the exact error associated with each method. Table 1 shows the results of a simulation exercise based on $N = 20$ time points per year and $M = 20{,}000$ trajectories.[8] The approximate portfolio in MCR is calculated using 200 iterations for the resolution of (34). This framework enabled us to replicate the results reported by BGSS (2003) in their Table 1. Computations are performed for risk aversions $R = 2, 3, 4, 5$ and investment horizons $T = 1, 2, 5$; parameter values are those described in Table 1. The experiment was carried out for several states of the random number generator. As the results obtained were roughly similar across states we report results for a particular state.

Several conclusions emerge from this simple exercise. The first one is that MCMD produces more accurate values, in this particular example across all 12 pairs (R, T) studied. The gain in accuracy, relative to the other methods tested, is often large and sometimes exceeds a factor of 10. The systematic pattern of dominance across all parameter values tested is of course extreme and one may naturally wonder about the relative performance of the methods in a larger experiment. Our efficiency study below will examine this issue and confirm the broad conclusions emerging from the analysis of this limited experiment.

The second conclusion is that there are variations in performance across the 5 versions of MCR tested. For some values of risk aversion and horizon all 5 versions produce closely related relative error (RE) values (e.g. for $T = 1$). In other cases there are more important differences (e.g. for $T = 5$).

Lastly, it should be noted that MCR often fails to produce portfolio values and, instead, returns infinite values (NaN). In the table, this happens for the polynomial approximations MCR-poly-1 and MCR-poly-2 (MCR-poly-2 fails in over 66% of the cases considered in Table 1). The problem can be traced back to the fact that the algorithm employed for the resolution of (34) diverges for certain parameter values. In a broader set of experiments that we conducted, cases of divergence were recorded for almost every variant of MCR. The occurrence of problematic cases appeared to increase with the length of the investment period.

These results in Table 1 give a preliminary feel for the relative properties of the candidate methods based on a given number of trajectories of the underlying Brown-

[8] For MCR the duration of a time period is taken to be 1/20. The total number of periods is $T \times 20$.

Table 1
Portfolio demands using MCMD and MCR in the Ornstein–Uhlenbeck model

T	Risk aversion (RA)							
	2		3		4		5	
	π	RE	π	RE	π	RE	π	RE
1								
Exact	20.53		12.83		9.31		7.31	
MCMD	20.05	2.36	12.18	5.02	8.59	7.73	6.54	10.47
MCR-lin-1	24.06	17.18	15.02	17.11	10.90	17.05	8.55	17.00
MCR-lin-2	24.11	17.41	15.07	17.48	10.94	17.50	8.59	17.50
MCR-lin-3	24.34	18.55	15.28	19.11	11.12	19.41	8.74	19.59
MCR-poly-1	24.01	16.96	14.97	16.68	10.85	16.47	8.50	16.33
MCR-poly-2	24.13	17.54	15.07	17.47	10.93	17.38	8.57	17.32
2								
Exact	15.14		8.42		5.74		4.33	
MCMD	14.97	1.08	8.27	1.81	5.60	2.43	4.20	3.01
MCR-lin-1	18.62	22.99	10.40	23.49	7.08	23.33	5.33	23.07
MCR-lin-2	18.49	22.18	10.28	22.08	6.98	21.52	5.24	20.99
MCR-lin-3	18.71	23.61	10.46	24.26	7.13	24.15	5.37	23.92
MCR-poly-1	18.73	23.77	10.50	24.66	7.16	24.75	5.40	24.67
MCR-poly-2	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN
5								
Exact	4.18		0.18		−0.68		−0.90	
MCMD	4.92	17.55	1.13	535.40	0.37	154.49	0.21	122.78
MCR-lin-1	9.31	122.71	2.99	1575.45	1.16	269.53	0.43	147.81
MCR-lin-2	9.28	121.96	2.94	1545.61	1.11	262.53	0.39	143.34
MCR-lin-3	9.75	133.21	3.38	1794.63	1.50	318.93	0.73	180.26
MCR-poly-1	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN
MCR-poly-2	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN

The fraction of wealth invested in the stock market is given in terms of relative risk aversion (RA) and investment horizon (T) in the model with Ornstein–Uhlenbeck market price of risk process (model WA). Parameter values are $A = 0.2712$, $\overline{\theta} = 0.9456$, $\Sigma = 0.2268$, $\sigma = 0.2$, $r = 0.06$ and $\theta_0 = 0.10$. The exact solution is calculated using the formulas of Wachter (2002) in Section 5.1.1. Each computation is performed using MCMD, MCR-lin-1, MCR-lin-2, MCR-lin-3, MCR-poly-1 and MCR-poly-2. For computations we use $M = 20{,}000$ trajectories and $N = 20$ discretization points per year. Risk aversion varies from 2 to 5 in unit increment. Investment horizon takes values 1, 2 and 5. Portfolio shares (π) and relative errors (RE) are expressed as percentages. The relative errors associated with the different methods are shown for state 0 of the random number generator.

ian motion and a limited number of risk aversion and horizon combinations. In order to provide more conclusive evidence we conduct a large-scale study involving a large number of draws (here 10,000) for the parameters of the model and the initial values of the state variables. For each draw relative errors and computation times are recorded, for each method. A measure of accuracy, root mean square relative error (RMSRE), and a measure of speed, inverse average time (IAT), are computed from this sample, again for each method. This experiment is repeated for different discretization values N and different numbers of trajectories M. The speed–accuracy trade-off can then be graphed to evaluate the relative performance of the candidate methods. Given the difficulties experienced with the polynomial-regression methods in the initial experiment, we focus on the linear approximations MCR-lin1, MC-lin-2, MCR-lin-3 and on MCMD.

To conduct our experiment we draw the parameters R, T, θ_0, r_0 from independent uniform distributions (the values of other parameters remain as specified in Table 1). Specifically, R is uniform over the interval $[0.5, 5]$, T is uniform over the discrete set $\{1,2, \ldots ,5\}$, θ_0 is uniform over $[0.30, 1.50]$ and r_0 is uniform over $[0.01, 0.10]$. Each draw consists of a vector $[R, T, \theta_0, r_0]$. Errors and computation times are recorded, for each method, for the pairs $(M,N) = \{(1000, 10), (4000, 20), (9000, 30), (16{,}000, 40)\}$. These combinations of M, N are chosen so as to quadruple M when N is doubled, leaving the ratio $\sqrt{M}/N$ constant.[9] Out of the 10,000 parameter draws 1826 violated the determinant condition (36) for the closed form solution of Wachter (2002), or failed to produce finite values for at least one of the regression methods. Eliminating these problematic values left a sample of 8164 "good" draws which constitute our sample. Statistics, such as RMSRE and IAT, are computed over this sample.

Fig. 1 displays the results from this experiment. The first observation is that MCMD converges faster as M and N increase. Although somewhat difficult to infer from the graph convergence is, in fact, not assured for MCR.[10] The second observation is that all three regression methods have a very similar performance. Regressing on additional lagged returns does not appear to improve performance in a significant way. The last observation emerging from the experiment is that MCMD improves on MCR by a factor in excess of 10. For instance, for a speed in the neighborhood of 4 the RMSRE of MCMD nears 8×10^{-1} while that of MCR is about 2×10.

Our second experiment is a large-scale study performed in the context of the nonlinear model DGR1. Parameter values are those reported in Appendix C of DGR (2003, pp. 441–442). In this instance 600 vectors R, T, θ_0, r_0 are drawn from independent uniform distributions: R is uniform over $[0.5,5]$, T is uniform over the discrete set $\{1,2, \ldots ,5\}$, θ_0 is uniform over $[-0.10, 0.40]$ and r_0 is uniform over $[0.01, 0.10]$. For each draw the benchmark true value of the portfolio is computed using the convergent method of DGR with $M = 250{,}000$, $N = 200$ and variance reduction by antithetics. The long computation times needed to calculate the benchmark values explains the more modest size of the sample in this experiment. As before errors relative to the benchmark and computation times are recorded for the pairs $(M,N) = \{(1000, 10), (4000, 20), (9000, 30), (16{,}000, 40)\}$. Out of the 600 parameter draws 65 produced infinite values for at least one of the regression methods. This left a sample of 535 "good" draws over which statistics were computed.

[9] The ratio $\sqrt{M}/N$ is the efficiency ratio for MCMD. Increasing M and N while maintaining this ratio constant ensures convergence to the true value without modifying the structure of the second order bias (see Detemple et al. (2003)).

[10] Suppose that the order of the polynomials used to approximate conditional expectations is fixed. In this case MCR converges to the *projection* of the portfolio on the linear space spanned by the approximating polynomials and not to the portfolio policy itself (see Clément et al., 2002). If the projection residual is large the approximation error of MCR will also be large, independently of the number of replications used to estimate the regression coefficients.

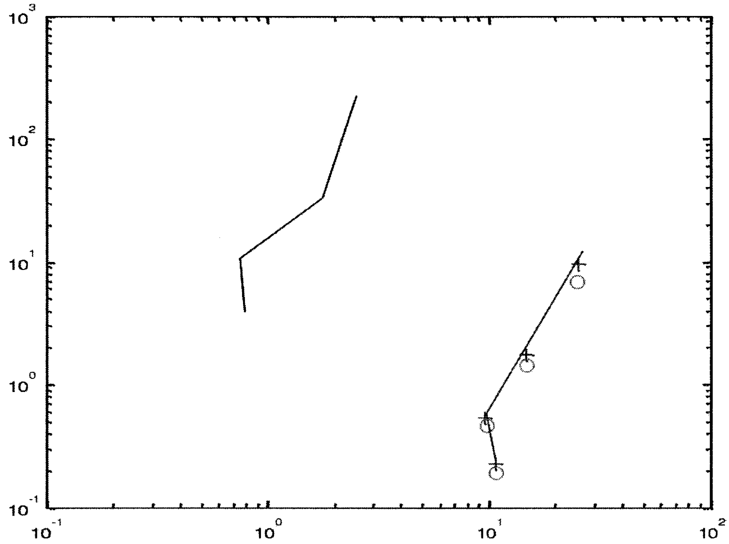

Fig. 1. This figure shows the speed–accuracy trade-off for MCMD (plain line top left corner), MCR-lin-1 (plain line bottom right corner), MCR-lin-2 (plus '+') and MCR-lin-3 (circle 'o'), in the context of the linear model WA. Speed is measured by the inverse of the average computation time over the sample (y-axis). Accuracy is measured by root mean square relative error (x-axis). Four points, corresponding to the pairs $(M,N) = \{(1000,10),(4000,20),(9000,30),(16{,}000,40)\}$, are graphed for each method.

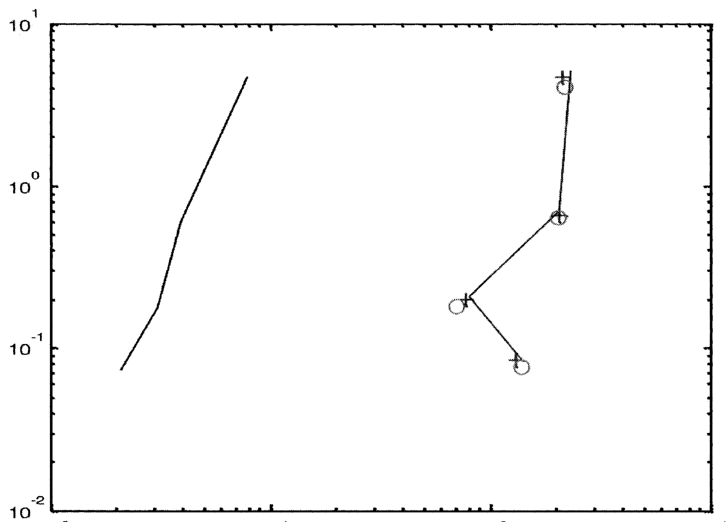

Fig. 2. This figure shows the speed–accuracy trade-off for MCMD (plain line left side), MCR-lin-1 (plain line right side), MCR-lin-2 (plus '+') and MCR-lin-3 (circle 'o'), in the context of the nonlinear model DGR1. Speed is measured by the inverse of the average computation time over the sample (y-axis). Accuracy is measured by root mean square relative error (x-axis). Four points, corresponding to the pairs $(M,N) = \{(1000,10),(4000,20),(9000,30),(16{,}000,40)\}$, are graphed for each method.

The results obtained confirm the conclusions reached in the prior experiment. Fig. 2 shows that MCMD converges fairly smoothly, as M and N increase, while MCR does not. The three linear versions of MCR exhibit, again, a very similar trade-off between speed and accuracy. The performance advantage of MCMD over MCR appears even more clearly in this model. The gain in accuracy approaches a factor of 10^2 for a speed in the neighborhood of 10^{-1} (i.e. for the pair $(M,N) = (16{,}000, 40)$).

6. Conclusion

In this paper we have provided a review of the MCMD method developed by Detemple et al. (2003) for the computation of optimal portfolio policies. This approach, which combines Monte Carlo simulation and Malliavin derivatives, is extremely flexible and accurate. One of its benefits is to permit the implementation of realistic models, featuring multiple assets and state variables, non-linear dynamics of the state variables and general time-separable von Neumann–Morgenstern preferences. It also identifies the primitive components of the optimal portfolio, including the intertemporal hedging terms whose importance was first highlighted by Merton (1971).

The comparative studies that we performed show that MCMD dominates the Monte Carlo Regression approach of BGSS (2003). In large-scale studies conducted in the context of two models, MCMD proved to be significantly more efficient. Efficiency graphs show that it provides at least a tenfold increase in accuracy for a given budget of computation time.

MCMD essentially resolves the implementation question for general diffusion models with complete financial markets. It also holds much promise for the analysis of more general consumption–portfolio problems. Directions for future research include its extension to settings with incomplete markets and constraints, discontinuous price processes and non-separable preferences. Preliminary results (see for instance Detemple and Rindisbacher (2003) for a specialized model with incomplete markets) have already made inroads in some of these areas and provide a glimpse of the potential benefits associated with a full blown extension of the method.

Appendix A: Proofs

Proof of Proposition 1. Since $\xi_t X_t^* = E_t\left[\int_t^T \xi_v I(y^*\xi_v, v)\mathrm{d}v + \xi_T J(y^*\xi_T, T)\right]$ an application of the Clark–Ocone formula shows that

$$\xi_t X_t^* \pi_t^{*\prime}\sigma_t - \xi_t X_t^*\theta_t' = -E_t\left[\int_t^T \xi_v Z_1(y^*\xi_v, v)\mathrm{d}v + \xi_T Z_2(y^*\xi_T, T)\right]\theta_t' \\ - E_t\left[\int_t^T \xi_v Z_1(y^*\xi_v, v)H_{t,v}'\mathrm{d}v + \xi_T Z_2(y^*\xi_T, T)H_{t,T}'\right],$$

where

$$Z_1(y^*\xi_v, v) = I(y^*\xi_v, v) + y^*\xi_v I'(y^*\xi_v, v),$$
$$Z_2(y^*\xi_T, T) = J(y^*\xi_T, T) + y^*\xi_T J'(y^*\xi_T, T),$$
$$H'_{t,v} = \int_t^v (\mathscr{D}_t r_s + \theta'_s \mathscr{D}_t \theta_s)\mathrm{d}s + \int_t^v \mathrm{d}W'_s \cdot \mathscr{D}_t \theta_s$$

and $\mathscr{D}_t r_s$, $\mathscr{D}_t \theta_s$ and $\mathscr{D}_t Y_s$ are Malliavin derivatives that satisfy the relevant equations (see (27) and (28)).

Using the definition of X_t^*, we can write

$$X_t^* - E_t\left[\int_t^T \xi_{t,v} Z_1(y^*\xi_v, v)\mathrm{d}v + \xi_{t,T} Z_2(y^*\xi_T, T)\right]$$
$$= -E_t\left[\int_t^T \xi_{t,v}(y^*\xi_v) I'(y^*\xi_v, v)\mathrm{d}v + \xi_{t,T}(y^*\xi_T) J'(y^*\xi_T, T)\right]$$

and therefore,

$$X_t^* \pi_t^{*\prime} \sigma_t = -E_t\left[\int_t^T \xi_{t,v}(y^*\xi_v) I'(y^*\xi_v, v)\mathrm{d}v + \xi_{t,T}(y^*\xi_T) J'(y^*\xi_T, T)\right]\theta'_t$$
$$- E_t\left[\int_t^T \xi_{t,v} Z_1(y^*\xi_v, v) H'_{t,v}\mathrm{d}v + \xi_{t,T} Z_2(y^*\xi_T, T) H'_{t,T}\right].$$

Transposing this expression and identifying the first term with π_1^* and the second with π_2^* gives the formula stated. □

Proof of Proposition 2. For power utility function $u(c,t) = \eta_t c^{1-R}/(1-R)$ and $U(X,T) = \eta_T X^{1-R}/(1-R)$, with $\eta_t \equiv \exp(-\beta t)$, we obtain

$$I(y\xi_v, v) = (y\xi_v/\eta_v)^{-1/R}, \quad J(y\xi_T, T) = (y\xi_T/\eta_T)^{-1/R},$$
$$y\xi_v I'(y\xi_v, v) = -(1/R)(y\xi_v/\eta_v)^{-1/R} = -(1/R) I(y\xi_v, v),$$
$$y\xi_T J'(y\xi_T, T) = -(1/R)(y\xi_T/\eta_T)^{-1/R} = -(1/R) J(y\xi_T, T),$$
$$Z_1(y\xi_v, v) = (1-1/R) I(y\xi_v, v), \quad Z_2(y\xi_T, T) = (1-1/R) J(y\xi_T, T).$$

Substituting these expressions in the policies of Proposition 1 gives the formulas stated. □

References

Ahn, D., Gao, B., 1999. A parametric nonlinear model of term structure dynamics. Review of Financial Studies 12, 721–762.

Ait-Sahalia, Y., 1996. Testing continuous-time models of the spot interest rate. Review of Financial Studies 9, 385–426.

Brandt, M.W., Goyal, A., Santa-Clara, P., Stroud, J.R., 2003. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. Working paper, Wharton School, 2003, Review of Financial Studies.

Brennan, M., Schwartz, E., Lagnado, R., 1997. Strategic asset allocation. Journal of Economic Dynamics and Control 21, 1377–1403.

Clément, E., Lamberton, D., Protter, P., 2002. An analysis of a least squares regression method for American option pricing. Finance and Stochastics 5, 449–471.

Cox, J.C., 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., 1976. The valuation of options for alternative stochastic processes. Journal of Financial Economics 3, 145–166.

Cvitanic, J., Goukasian, L., Zapatero, F., 2003. Monte Carlo computation of optimal portfolio in complete markets. Journal of Economic Dynamics and Control 27, 971–986.

Detemple, J.B., Garcia, R., Rindisbacher, M., 2003. A Monte-Carlo method for optimal portfolios. Journal of Finance 58, 401–446.

Detemple, J.B., Rindisbacher, M., forthcoming. Explicit solutions for a portfolio problem with incomplete markets and investment constraints. Working paper, Boston University, 2003, Mathematical Finance.

Detemple, J.B., Garcia, R., Rindisbacher, M., 2003. Asymptotic properties of Monte Carlo estimators of diffusion processes. Working paper, CIRANO 2004, Journal of Econometrics.

Harrison, J.M., Kreps, D.M., 1979. Martingales and arbitrage in multiperiod securities markets. Journal of Economic Theory 20, 381–408.

Harrison, J.M., Pliska, S.R., 1981. Martingales and Stochastic integrals in the theory of continuous trading. Stochastic Processes and Applications 11, 215–260.

Karatzas, I., Lehoczky, J.P., Shreve, S.E., 1987. Optimal portfolio and consumption decisions for a "small investor on a finite horizon". SIAM Journal of Control and Optimization 25, 1557–1586.

Liu, J., 1998. Portfolio selection in stochastic environments. Working paper, Stanford University.

Longstaff, F., Schwartz, E., 2001. Valuing American options by simulation: A simple least-squares approach. Review of Financial Studies 14, 113–147.

Markowitz, H., 1952. Portfolio selection. Journal of Finance 7, 77–91.

Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous time model. Journal of Economic Theory 3, 273–413.

Ocone, D., Karatzas, I., 1991. A generalized Clark representation formula, with application to optimal portfolios. Stochastics and Stochastics Reports 34, 187–220.

Pliska, S., 1986. A stochastic calculus model of continuous trading: Optimal portfolios. Mathematics of Operations Research 11, 371–382.

Schroder, M., Skiadas, C., 1999. Optimal consumption and portfolio selection with stochastic differential utility. Journal of Economic Theory 89, 68–126.

Wachter, J., 2002. Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets. Journal of Financial and Quantitative Analysis 37, 63–91.