Quantitative Finance, 2014
Vol. 14, No. 10, 1739–1751, http://dx.doi.org/10.1080/14697688.2011.649780

Routledge
Taylor & Francis Group

Dynamic option hedging via stochastic model predictive control based on scenario simulation

ALBERTO BEMPORAD*†, LEONARDO BELLUCCI‡ and TOMMASO GABBRIELLINI‡

†IMT Institute for Advanced Studies Lucca, Piazza San Ponziano 6, 55100 Lucca, Italy
‡MPS Capital Services, Modelling & Quantitative Analysis, Viale Mazzini 23, Siena, Italy

(*Received 10 December 2009; revised 22 March 2011; in final form 8 December 2011*)

Derivative contracts require the replication of the product by means of a dynamic portfolio composed of simpler, more liquid securities. For a broad class of options encountered in financial engineering we propose a solution to the problem of finding a hedging portfolio using a discrete-time stochastic model predictive control and receding horizon optimization. By employing existing option pricing engines for estimating future option prices (possibly in an approximate way, to increase computation speed), in the absence of transaction costs the resulting stochastic optimization problem is easily solved at each trading date as a least-squares problem with as many variables as the number of traded assets and as many constraints as the number of predicted scenarios. As shown through numerical examples, the approach is particularly useful and numerically viable for exotic options where closed-form results are not available, as well as relatively long expiration dates where tree-based stochastic approaches are excessively complex.

Keywords: Financial options; Hedging techniques; Stochastic control; Stochastic programming; Dynamical systems; Exotic options

JEL Classification: C1, C6, C15, C61

1. Introduction

Modern pricing theory, based on the seminal works of Black and Scholes (1973) and Merton (1973), approaches the valuation of a derivative contract as a replication problem: the price of the financial contract is the cost to trade a self-financing portfolio of more liquid and simpler securities so as to match the option payoff in every state of the world. The allocating strategy is generally dynamic and chosen to locally eliminate any (modelled) risk from a portfolio composed of the derivative security and the replication strategy.

A popular and well-studied class of synthetic financial securities are European vanilla options: a *call* (*put*) *option* gives the holder the right to buy (sell) the underlying asset at a given *expiration date* in the future for a predetermined *strike price* (see, e.g., Hull (2006) for the basics of financial options). From the point of view of an investment firm, the problem of writing an option amounts to jointly determining (i) the *price* the customer must pay to obtain the right to exercise the option, and (ii) the dynamic strategy for managing this money by creating a portfolio and periodically changing its composition during the life of the option. The strategy should make the value of the portfolio equal to the payoff amount to be paid to the customer at the expiration date, *regardless of the realized price evolution* of the assets underlying the option and composing the portfolio. In other words, the strategy should *hedge* the option as much as possible against the indeterminism associated with those price evolutions, or, in control engineering terms, the strategy should *reject* the effects of the stochastic variables affecting the evolution of the market through a proper *feedback* action.

Several approaches have been proposed for option pricing and dynamic hedging. The common goal of any approach is to eliminate (or, at least, reduce as much as possible) the risk that the value of the portfolio at the expiration date does not match the payoff, i.e. that the hedging error is non-zero. Such approaches can be roughly divided into two main categories: *local* approaches,

*Corresponding author. Email: alberto.bemporad@imtlucca.it

where the aim is to attempt to eliminate risk instantaneously, and *global* approaches, where the entire life span of the option is considered and the objective is to minimize the variance of the hedging error at the expiration date.

The seminal works of Merton (1973) and Black and Scholes (1973), a real milestone in finance and still a reference point, and their extensions to models with stochastic volatility (Heston 1993; Heston and Nandi 2000), belong to the first category. In general, they aim at perfect hedging by eliminating the risk at each time instant through a proper rebalancing of assets in the portfolio, usually continuously in time. Solutions can be obtained in different ways. *Analytical* approaches provide invaluable insight into understanding and operating with options, although their main operational limitation is that the pricing formulas and the hedging strategies can only be obtained under restrictive assumptions on the pricing model, such as log-normal or GARCH (Generalized AutoRegressive Conditional Heteroscedasticity) price models, and specific payoff functions (such as European call/put options). When closed-form solutions are not available, *numerical methods* (Duffie 1996) can be employed to solve the Kolmogorov backward equation associated with the model. *Simulation* is another method often used by investment firms to price options (Boyle 1977, Hull and White 1987). A (large) set of scenarios is generated by Monte Carlo simulation for the future prices of the underlying assets; the final value of the asset price of each scenario is used to compute the payoff value; the average of such payoff values, discounted by the interest rate, provides the option price. In view of such a current practice for option pricing, we focus our attention only on the hedging problem.

Approaches that instead look at the entire life of the option aim at minimizing risk at the expiration date. The problem can be cast as a stochastic optimal control problem and relies on the Hamilton–Jacobi–Bellman partial differential equation. This category includes *multi-stage stochastic programming* approaches, where the probability space of the option price realization is discretized in an *n*-ary tree (binary and trinomial are the most used), and then the pricing and hedging problem is solved as a stochastic linear programming problem (Edirisinghe *et al.* 1993, Klaassen 1998, Kouwenberg and Vorst 1998, Zhao and Ziemba 2001, Gondzio *et al.* 2003). Although special structures can be exploited to solve the associated large-scale linear program (Birge and Louveaux 1997), the approach is often limited for numerical reasons. In fact, the number of nodes in the tree (which is proportional to the number of optimization variables) is exponential in the number of trading periods, and heavily depends on the number of branches at each node. Nonetheless, not only is the approach useful for relatively infrequent portfolio rebalancing and/or short expiration dates, but subtle issues such as lack of arbitrage conditions when discretizing the probability space (Høyland and Wallace 2001) provide interesting insights (Klaassen 2002). Stochastic *dynamic programming* (DP) approaches (Fedotov 1999, Bertsimas *et al.* 2001) also discretize the probability space and solve the pricing and hedging problem backwards in time. While the method is appealing, its main limitation is due to the numerical explosion when the number of trading periods is large and several assets are traded.

This paper attacks the hedging problem from the viewpoint of feedback control systems. Rephrased using the jargon of systems theory, the hedging problem can be solved by a feedback control strategy that, based on *measurements* such as the current value of the portfolio and the state of the market (e.g., price and variance of the underlying assets), determines the *actuation signals*, i.e. the trading moves with a given sampling frequency (ideally, continuously in time). The control law, together with the stochastic pricing models of the underlying assets it is based upon, determines the financial properties of the derivative security. Hence, in control systems terms, the option hedging problem can be stated as follows: given the initial condition for the wealth of the portfolio (=option price), find a feedback control law (=hedging strategy) that brings the error between the final value of the wealth of the portfolio and the payoff function to zero, rejecting the effects of the stochastic uncertainty affecting the portfolio dynamics.

Based on such a control theoretical approach, in this paper we use stochastic model predictive control (MPC) concepts (Kolmanovsky *et al.* 2002, Muñoz de la Peña *et al.* 2005, Couchman *et al.* 2006, Wang and Boyd 2008, Primbs and Sung 2009) to develop a dynamic hedging strategy. MPC is an optimization-based policy that solves a finite-horizon open-loop optimal control problem at each sampling instant. Each optimization yields a sequence of optimal control moves, but only the first move is applied to the process: at the next time step, the computation is repeated over a shifted time-horizon by taking the most recently available state information as the new initial condition of the optimal control problem. For this reason, MPC is also called *receding* or *rolling horizon control* (Bemporad and Morari 1999, Rawlings 2000, Bemporad *et al.* 2002, 2004, Bemporad 2006, Mayne and Rawlings 2009). *Stochastic* model predictive control (SMPC) can be seen as a suboptimal way of solving a stochastic multi-stage dynamic programming problem: rather than solving the problem for the whole option-life horizon, a smaller problem is solved repeatedly from the current time-step t up to a certain number N of time steps in the future by suitably re-mapping the condition at the expiration date into a value at time $t+N$. SMPC has been proposed for financial applications only very recently, for example by Herzog *et al.* (2006) for portfolio optimization and by Meindl (2006), Meindl and Primbs (2008) and Primbs (2009) for option pricing and hedging. In particular, Meindl (2006) and Meindl and Primbs (2008) use clustering ideas to reduce the computational complexity and at the same time mitigate arbitrage opportunities, which are intrinsic to the optimization over multiple independent scenarios (the so-called 'fan'), rather than on scenario trees.

This paper proposes a novel SMPC approach to the dynamic hedging of a broad class of options. At each trading date, given the current state of the market, the SMPC algorithm computes the optimal asset quantities to compose the portfolio by minimizing the variance of the future hedging error, i.e. the difference between the price of the option at the next trading date and the wealth of the replicating portfolio at the same date. The minimization only requires the solution of a simple *least-squares optimization* problem at each trading date. To be able to handle very general stock price models and exotic payoffs, for which no analytic hedging policy exists, the variance of the predicted hedging error is formulated on-line using a pricing engine that generates a finite number of future scenarios of option prices, rather than analytically deriving expected values from pricing models as in Herzog *et al.* (2006) and Primbs (2009). To evaluate each option price, the pricing engine employs either Monte Carlo simulation (*on-line* computations) or *off-line* function approximation (still based on Monte Carlo simulation) to approximate the option value as a function of the state of the market (such as the price of the underlying stock), so that on-line evaluation is very fast. The approach proved to perform well not only for plain vanilla European call options, for which analytical methods also perform well, but also for an exotic 'Napoleon cliquet' option, for both log-normal Black and Scholes (1973) and Heston (1993) stock price models.

The approach of this paper is related to the Hedged Monte Carlo (HMC) approach (Potters *et al.* 2001, Pochart and Bouchaud 2004, Petrelli *et al.* 2008), where function approximation is used not only for pricing, but also to determine (off-line) a hedging policy as a function of the market state, using backward-in-time computations similarly to DP. The main differences between the SMPC and HMC approaches are that (i) SMPC can very easily adapt to changes of the market model, being the optimal hedging problem constructed on-line rather than pre-solved off-line for a given market model, (ii) the use of basis functions in HMC might lead to larger approximation errors, especially for a large number of components of the market state (and possibly of other model parameters that one may want to change on-line), and (iii) on the other hand, as in DP, HMC only requires the evaluation on-line of the predetermined optimal hedging function, while in SMPC a least-square problem needs to be solved based on the predicted possible future option prices.

The paper is organized as follows. In section 2 we introduce the dynamical models that we adopt in the paper for asset prices, for the synthetic option and its payoff, and for the wealth of the portfolio. In section 3 we state the hedging problem as a stochastic feedback control problem and show that minimum variance objectives are best candidates. In section 4 we address the posed option hedging problem as a stochastic model predictive control problem and present the proposed receding horizon optimization strategy. Finally, in section 5 we report numerical examples on European call and 'Napoleon cliquet' options, and draw some conclusions in section 6.

2. Model formulation

2.1. Asset price dynamics

Consider the problem of dynamically hedging a European option† $\mathcal{O}$ defined over n spot prices x_i of underlying assets, $i=1,\dots,n$, satisfying the stochastic differential equations in the real-world probability measure

$$\mathrm{d}x_i(\tau) = \mu_i^x(x_i(\tau), y_i(\tau))\mathrm{d}\tau + \sigma_i^x(x(\tau), y(\tau))\mathrm{d}z_i^x, \tag{1a}$$

$$\mathrm{d}y_i(\tau) = \mu_i^y(y_i(\tau))\mathrm{d}\tau + \sigma_i^y(y(\tau))\mathrm{d}z_i^y, \tag{1b}$$

where $z_i^x(\tau)$ and $z_i^y(\tau)$ are Wiener processes, namely $\mathrm{d}z_i^x$ and $\mathrm{d}z_i^y$ are correlated Gaussian variables with zero mean and variance $\mathrm{d}\tau$. In equation (1) we assume that $x_i \geq 0$, $\forall i=1,\dots,n$, $\forall \tau \geq 0$. Model (4) is a rather general form that covers several popular models, including the log-normal stock price model

$$\mathrm{d}x_i(\tau) = (\mu\,\mathrm{d}\tau + \sigma\,\mathrm{d}z_i^x)x_i(\tau) \tag{2}$$

($y_i(\tau) \equiv 0$, $z_i^y(\tau) \equiv 0$), and the model of Heston (1993)

$$\mathrm{d}x_i(\tau) = (\mu_i^x\,\mathrm{d}\tau + \sqrt{y_i(\tau)}\mathrm{d}z_i^x)x_i(\tau), \tag{3a}$$

$$\mathrm{d}y_i(\tau) = \theta_i(k_i - y_i(\tau))\mathrm{d}\tau + \omega_i\sqrt{y_i(\tau)}\mathrm{d}z_i^y, \tag{3b}$$

where (3b) is the Cox *et al.* (1985) process for the variance $y_i(\tau)$, and $\mathrm{d}z_i^x$ has correlation ρ_i with $\mathrm{d}z_i^y$.

In this paper we are interested in evaluating $x_i(\tau)$ and $y_i(\tau)$ at certain trading dates $\tau = t\Delta_T$, where $t \in \mathbb{Z}$ and $t \geq 0$ denotes a discrete-time index.‡ To this end, we discretize (1) into the difference equations

$$x_i(t+1) = f_i(x_i(t), y_i(t), z_i^x(t)), \tag{4a}$$

$$y_i(t+1) = g_i(y_i(t), z_i^y(t)), \tag{4b}$$

where, with a slight abuse of notation, $[\cdot](t)$ denotes the value of $[\cdot]$ at time $\tau = t\Delta_T$, $z^x(t)$, $z^y(t) \in \mathbb{R}^n$ are random Gaussian vectors with zero mean $E[z_i^x(t)] = 0$, $E[z_i^y(t)] = 0$, $\forall i=1,\dots,n$, $\forall t \geq 0$, and covariance matrix $\Phi^x(t) = E[z^x(t)z^x(t)']$, $\Phi^y = E[z^y(t)z^y(t)']$, with $\Phi^x(t)$, $\Phi^y(t)$ positive semidefinite $\forall t \geq 0$. In (4), functions f_i and g_i are either expressed analytically in an explicit way through the exact integration of (1), or simply represent a numerical integration engine providing $x_i(t+1)$, $y_i(t+1)$ as an implicit function of $x_i(t)$, $y_i(t)$. For example,

$$x_i(t+1) = \mathrm{e}^{(\mu - (1/2)\sigma^2)\Delta t + \sigma\sqrt{\Delta t}z_i^x(t)}x_i(t) \tag{5}$$

is the analytical solution of (2). In the sequel we denote by $x(t) = [x_1(t) \ \ \dots \ \ x_n(t)]' \in \mathbb{R}^n$ the vector of asset prices, and by $y(t) = [y_1(t) \ \ \dots \ \ y_n(t)]' \in \mathbb{R}^n$ the associated vector of additional state variables of the asset price models.

†The approach can easily be extended to other options as described by Bertsimas *et al.* (2001), such as Asian options.
‡The results presented in this paper can immediately be extended to non-uniform trading intervals Δ_T.

2.2. Option price and payoff function

We assume that the portfolio associated with option $\mathcal{O}$ is updated every Δ_T units of time, and denote by T the maturity of $\mathcal{O}$ expressed in terms of the number of sampling steps. The payoff $p(T)$ of $\mathcal{O}$ is described by the function

$$p(T) = \mathcal{P}(m(T)) \tag{6}$$

of the state $m(T)$ of the considered asset market at the expiration date, for example $m(T) = x(T)$. We denote by $p(t)$ the price of the hedged option at a generic intermediate time $t\Delta_T$,

$$p(t) = (1+r)^{t-N}\tilde{E}[\mathcal{P}(m(T)) \mid m(t)] \tag{7}$$

(see, for instance, Hull (2006, p. 589)), where $m(t) = \{x(0), \ldots, x(t), y(t)\}$ (or, more generally, $m(t) = \{x(0), \ldots, x(t), y(0), \ldots, y(t)\}$) and $\tilde{E}[p(T)]$ is the expected value of the payoff in the risk-neutral measure, given the asset prices realized up to time t and the current state $y(t)$ of the asset price model. In (7), $r = e^{r_a \Delta T} - 1$ is the return of the risk-free investment over Δ_T, and r_a is the annualized continuously compounded interest rate, which we assume to be constant.

Equation (7) can also be restated recursively as

$$p(t) = (1+r)^{-1}\tilde{E}[p(t+1) \mid m(t)]. \tag{8}$$

For instance, for a European call option on a single stock x with strike price K, we have

$$p(T) = \max\{x(T) - K, 0\}, \tag{9}$$

$m(t) = \{x(t), y(t)\}$, and

$$p(t) = e^{-r(N-t)}\tilde{E}[\max\{x(T) - K, 0\} \mid x(t), y(t)].$$

In particular, for log-normal price models, $m(t) = x(t)$. For 'Napoleon cliquet' path-dependent exotic options

$$p(T) = \max\left\{0, C + \min_{i \in \{1,\ldots,N_{\text{fix}}\}} \frac{x(t_i) - x(t_{i-1})}{x(t_{i-1})}\right\}, \tag{10}$$

where t_i, $i = 1, \ldots, N_{\text{fix}}$, are the fixing dates, and C is a fixed value. In this case, $m(t) = \{x(t_0), \ldots, x(t_k), x(t), y(t)\}$, where k is the fixing index such that $t_k \le t < t_{k+1}$.

2.3. Portfolio dynamics

Assume that a portfolio $\mathcal{W}$ consisting of assets x_i, $i = 1, \ldots, n$, and risk-free investments is dynamically managed by the option writer. Let $u_i(t)$ denote the number of assets i, $i = 1, \ldots, n$, contained in the portfolio during the time interval $[t\Delta_T, (t+1)\Delta_T)$, $t = 0, \ldots, T$, and let $u_0(t)$ be the amount of wealth allocated to risk-free investments. The trading moves $u_i(t)$, $i = 0, \ldots, n$, are decided at time $t\Delta_T$. The total wealth $w(t)$ of $\mathcal{W}$ in money units invested at time $k\Delta_T$ is

$$w(t) = u_0(t) + \sum_{i=1}^{n} x_i(t)u_i(t), \tag{11}$$

where $x_i(t)$ is the spot price of asset i at the trading time instant (we assume that the value of x_i is continuous across the time instant the asset is traded, and therefore is the same immediately before and immediately after trading). After a trading period Δ_T, the portfolio $\mathcal{W}$ is worth

$$w(t+1) = (1+r)u_0(t) + \sum_{i=1}^{n} x_i(t+1)u_i(t). \tag{12}$$

By assuming the standard *self-financing* condition (that is, the wealth of the portfolio is always totally reinvested), we obtain the following dynamical equation for the wealth of $\mathcal{W}$:

$$w(t+1) = (1+r)w(t) + \sum_{i=1}^{n} b_i(t)u_i(t), \tag{13}$$

where $b_i(t) \triangleq x_i(t+1) - (1+r)x_i(t)$. The initial condition $w(0)$ is set equal to the price paid by the customer of option $\mathcal{O}$, $w(0) = (1+r)^{-N}\tilde{E}[p(T) \mid x(0), y(0)]$.

We remark on a few features enjoyed by the stated model. (i) The assets' dynamics (1) do not depend on trading decision $u_i(t)$, a reasonable assumption if the volumes traded in $\mathcal{W}$ are negligible with respect to the volumes exchanged on the entire market. (ii) As a consequence, also the option price $\mathcal{P}$, and therefore its expected value $\tilde{E}[\mathcal{P}]$, do not depend on $u_i(t)$. (iii) Dynamics (13) is a first-order linear stochastic discrete-time system.

Although in this paper we do not consider transaction costs, taking the viewpoint of a financial institution for which they can be considered negligible, the approach can be extended to deal with transaction costs using quadratic programming.

3. Stochastic control problem

Based on the models developed in section 2, the dynamic option problem can be reformulated in system theoretical terms as a stochastic control problem. The wealth $w(t) \in \mathbb{R}$ represents the state and output of the regulated process, the traded asset quantities $u(t) \in \mathbb{R}^n$ are the manipulated variables, and the option price $p(t)$ is the target reference for $w(t)$. In particular, the control objective is to make $w(T)$ as close as possible to $p(T)$, for any possible realization of the asset prices $x(t)$. This can be labeled as a reference tracking and disturbance rejection problem.

By defining the tracking error $e(t) \triangleq w(t) - p(t)$, the objective can be restated as minimizing $e(t)$ for all possible asset price realizations. This can be achieved by minimizing the trade-off between the variance and mean of $e(T)$

$$\begin{aligned} J(e(T), \alpha) &= \mathrm{Var}[e(T)] + \alpha E[e(T)]^2 \\ &= E[(e(T) - E[e(T)])^2] + \alpha E[e(T)]^2, \end{aligned} \tag{14}$$

where $E[\cdot]$ ($\mathrm{Var}[\cdot]$) is the expectation (variance) in the real-world probability measure, and $\alpha \in [0, +\infty)$. Clearly, minimizing the squared mean $E[e(T)]^2$ ($\alpha \to \infty$) would be a very risky approach, as $E[e(T)]$ would be small but

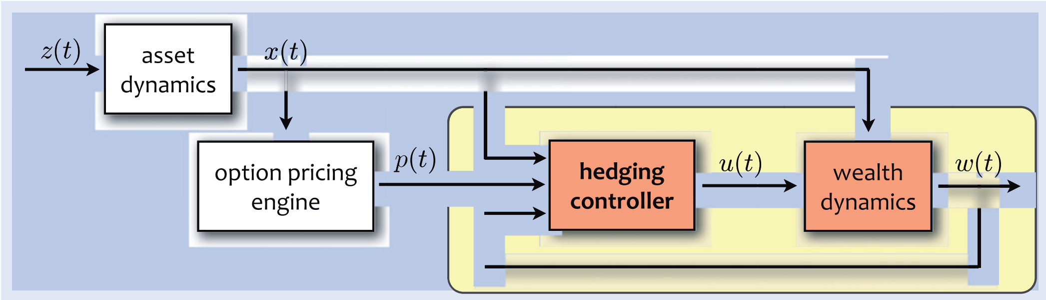

Figure 1. The dynamic option hedging problem as a feedback control problem.

the variance $\mathrm{Var}[e(T)] = E[(e(T) - E[e(T)])^2]$ could be very large. A special case would be to minimize the expected value of the squared mean $E[e^2(T)]$, since

$$E[e(T)^2] = E[(e(T) - E[e(T)])^2] + \alpha E[e(T)]^2,$$

for $\alpha = 1$. The following proposition shows that, under non-arbitrage conditions, if the minimum variance criterion

$$J(e(T), \alpha) = E[(e(T) - E[e(T)])^2] \tag{15}$$

is minimized and the minimum is zero, then $E[e(T)]$ will also be zero.

Proposition 3.1: *Consider an ideal minimum variance hedging strategy choosing $u(t)$, $t = 0, \ldots, T-1$, such that $\mathrm{Var}[w(T) - p(T)] = 0$ and let $w(0) = p(0) = (1+r)^{-T}\tilde{E}[p(T)]$. Then $w(t) - p(t) \equiv 0$, $\forall t = 0, \ldots, T$, and in particular the final hedging error $w(T) - p(T)$ is identically zero.*

Proof: The proposition is proved by induction on the trading period t backwards in time, showing that $e(t-1) = (1+r)^{-1}e(t)$, with $e(t) = w(t) - p(t)$, is a deterministic variable. Consider at time $T-1$ the portfolio

$$e(T-1) = w(T-1) - p(T-1)$$

obtained by gathering the wealth $w(T-1)$ and short selling the option for its price $p(T-1)$. As by assumption the wealth $e(T)$ of the new portfolio at time T has zero variance, the trading move $u(T-1)$ is riskless during the time period $[(T-1)\Delta_T, T\Delta_T]$. As such, in the absence of arbitrage conditions, it must earn the interest rate r

$$e(T-1) = (1+r)^{-1}e(T),$$

which makes $e(T-1)$ also a deterministic variable. Assume now that $e(k)$ is deterministic for all $k = t$, $t+1, \ldots, T$ and consider at time $t-1$ the new portfolio

$$e(t-1) = w(t-1) - p(t-1)$$

obtained by gathering the wealth $w(t-1)$ and short selling the option for its price $p(t-1)$. Following the same reasoning, in the absence of arbitrage conditions we prove that

$$e(t-1) = (1+r)^{-1}e(t). \tag{16}$$

By induction, $\{e(t)\}_{t=0}^{T}$ is a finite sequence of deterministic variables satisfying the recursive relation (16). In particular, $e(t) = (1+r)^t(w(0) - p(0)) = 0$, $\forall t = 0, \ldots, T$. □

Proposition 3.1 implies that if the minimum variance criterion (15) is used and perfect hedging is achieved, then $w(t)$ should track $p(t)$ perfectly, at each time t. The opposite is also trivially true: if a hedging strategy exists such that $w(t)$ tracks $p(t)$ perfectly, then in particular $w(T) - p(T) = 0$. Hence, from a system theoretical perspective, the dynamic option hedging problem is equivalent to designing a control strategy that lets the output $w(t)$ of the stochastic linear system (13), affected by stochastic multiplicative noise on the input matrix $b(t) = [b_1(t) \ \ldots \ b_n(t)]$, track the reference signal $p(t)$ generated by the linear reference model (8), as illustrated in figure 1. A particular control strategy based on *receding horizon control* (also referred to as *model predictive control*) ideas is presented in the next section.

4. Stochastic model predictive control

At a generic trading time $t = 0, 1, \ldots, T-1$ let the portfolio composition $u(t)$ be chosen by solving the following finite-time stochastic dynamic optimization problem

$$\min_{\{u(k, m_k)\}} \mathrm{Var}_{m_T}[w(T, m_T) - p(T, m_T)], \tag{17a}$$

$$\begin{aligned} \text{s.t. } & w(k+1, m_{k+1}) \\ & \quad = (1+r)w(k, m_k) + \sum_{i=0}^{n} b_i(k, m_{k+1})u_i(k, m_k), \\ & \quad k = t, \ldots, T-1, \end{aligned} \tag{17b}$$

$$w(t, m_t) = w(t), \tag{17c}$$

where m_k represents a generic realization of the state of the considered asset market at time k (determined by a finite stochastic noise sequence $\{z(t+1), \ldots, z(t+k)\}$, corresponding to the realization of future asset prices $x(t+1), \ldots, x(t+k)$). Also, expectations (i.e. the variance of the final hedging error being minimized) are taken in the real-world measure with respect to m_T, conditioned on the values $z(0), \ldots, z(t)$ already realized. In particular, the deterministic variable $m_t = m(t)$ represents the current state of the market.

Typically, in *stochastic optimization* approaches to option pricing and hedging, the set $\mathcal{M}_t$ of possible future market states is discrete and arranged on a binary or ternary tree. Note that, in (17), the number of optimization vectors $u(k, m_k) \in \mathbb{R}^n$ depends linearly on the number M of elements of $\mathcal{M}_t$, the number of remaining trading periods $T - t$, and the number of assets n. Note that, for $t = 0$, if the initial portfolio value $w(0)$ is also treated as an optimization variable, (17) also provides the optimal option price $p(0)$.

Stochastic programming (SP) is a popular technique for solving optimization problems under uncertainty (Birge and Louveaux 1997, Sahinidis 2004). In particular, scenario enumeration methods have been proposed to solve SP. The idea is to enumerate a certain number M of scenarios (or, alternatively, of tree nodes), where each scenario corresponds to the realization of a certain sequence of stochastic variables and has a probability π_j of occurring, $j = 1, \ldots, M$, $\pi_j > 0$, $\pi_j \leq 1$, $\sum_{k=1}^{M} \pi_j = 1$. Problem (17) can be solved through multi-stage stochastic programming when the possible realizations of $m(k)$ can be aggregated according to the scenario tree $\mathcal{M}_t$. The drawback of the procedure is that the problem becomes computationally infeasible when the horizon T is long (because of long expiration dates and/or frequent trading every small Δ_T time units), since the number of nodes in the tree is exponential in T. Moreover, the procedure must be repeated before each trading instant t over the remaining horizon $[t, T]$ to take into account unmodeled price dynamics and disturbance realizations, and each time only the first optimal trading move $u_i^*(t)$ is actually used.

In very simple cases, problem (17) may be solved through stochastic dynamic programming techniques Bertsimas *et al.* (2001) to obtain off-line an explicit state feedback policy $u_i^*(t) = K_i(w(t), x(t), y(t))$, $t = 0, \ldots, T-1$. Assuming that an *option pricing engine* is available, for instance based on Monte Carlo simulation and empirical computation of expectations of the payoff function, in this paper we assume that $w(0) = p(0) = \alpha(1+r)^T \tilde{E}[p(T)]$ is already determined and propose a stochastic model predictive strategy for hedging, based on the pricing engine, that solves (17) approximately and recursively over shorter time horizons.

4.1. SMPC algorithm for option hedging

Besides discretizing more roughly in the probability space (=smaller number M of scenarios), another way of alleviating the complexity of (17) is to decrease the optimization horizon from $[t, T]$ to $[t, \min\{t+N, T\}]$, $N \geq 1$. Such a practice is used typically in *receding horizon* control (also called *model predictive control*, MPC). Several approaches to *stochastic* MPC have been proposed in the literature for stochastic linear systems (Schwarm and Nikolaou 1999, Li *et al.* 2000, Batina *et al.* 2002, van Hessem and Bosgra 2002, Muñoz de la Peña *et al.* 2005, Couchman *et al.* 2006, Primbs 2007a, b, Oldewurtel *et al.* 2008), Markov jump linear systems (Blackmore *et al.* 2007, Bernardini and Bemporad 2009), and stochastic hybrid systems (Bemporad and Di Cairano 2005). An alternative way based on 'fans' rather than 'trees' was proposed by Meindl (2006) and Meindl and Primbs (2008) based on clustering scenarios to reduce the number of optimization variables. Although numerically appealing, the main drawback of such an approach is that optimal results may reflect artificial arbitrage opportunities induced by the clustering algorithm, as the *causality constraint* that makes control moves $u(k)$ that only depends on m_k may not be satisfied at all $k = t, t+1, \ldots, t+\bar{N}$, $\bar{N} = \min\{N, T-t\}$.

In this paper we propose to remap the stochastic dynamic optimization problem (17) from the optimization horizon $[t, T]$ into $[t, t+\bar{N}]$, by assuming that 'perfect hedging' occurs between time $t+\bar{N}+1$ and T,

$$\min_{\{u(k,m_k)\}} \mathrm{Var}_{m_{t+\bar{N}}}[w(t+\bar{N}, m_{t+\bar{N}}) - p(t+\bar{N}, m_{t+\bar{N}})], \quad (18a)$$

$$\text{s.t. } w(k+1, m_{k+1}) = (1+r)w(k, m_k) + \sum_{i=0}^{n} b_i(k, m_{k+1})u_i(k, m_k), \quad k = t, \ldots, t+\bar{N}, \quad (18b)$$

$$w(t, m_t) = w(t). \quad (18c)$$

As highlighted in figure 2, the idea is to find the optimal trading moves $u(k, m_k)$ up to time $k = t+\bar{N}$, and assume that optimal trading moves $u(k, m_k)$ exist for $k = t+\bar{N}+1, \ldots, T$ that provide perfect hedging $w(k) = p(k) = (1+r)^{k-T}\tilde{E}[p(T)]$. We emphasize that the cost function in (18) is computed using the real-world probability measure; the risk-neutral measure is used only to express option prices.

For a given complexity (that is, the number M of future market states), there is obviously a trade-off between the length of the horizon $\bar{N}$ and the 'arity' (number of branches at each node) of the tree describing the set $\mathcal{M}_t$ of future market states.

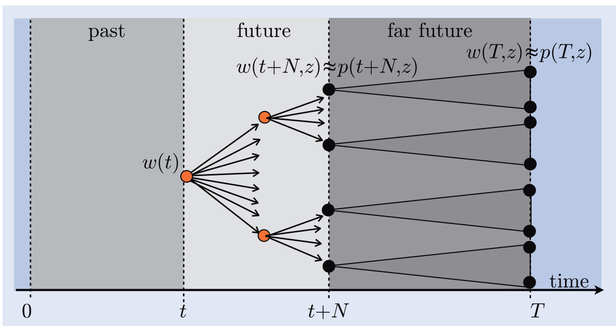

Figure 2. Remapping the stochastic dynamic optimization onto a shorter horizon under the 'perfect hedging' assumption.

In the special case $N = \bar{N} = 1$, problem (18) becomes the one-step-ahead minimum variance problem

$$\min_{\{u(t)\}} \mathrm{Var}_{m_{t+1}}[w(t+1, m_{t+1}) - p(t+1, m_{t+1})], \tag{19a}$$

$$\text{s.t. } w(t+1, m_{t+1}) = (1+r)w(t) + \sum_{i=0}^{n} b_i(t, m_{t+1})u_i(t). \tag{19b}$$

The reason for focusing on the formulation (19) is that only one vector $u(t)$ is optimized, which drastically limits the number of optimization variables to the number of trading assets n. Hence, the number M of scenarios can be quite large, as no further branching takes place after time $t+1$. The drawback of (19) with respect to (17) is that the option price $p(t+1, m_{t+1})$ must be evaluated for each future market state m_{t+1}, i.e. for all M considered scenarios. However, by optimizing the sample variance of $w(t+1) - p(t+1)$, problem (19) can be rewritten as the following very simple *least squares* problem

$$\min_{u(t)} \sum_{j=1}^{M} \left(w^j(t+1) - p^j(t+1) - \frac{1}{M}\sum_{i=1}^{M} w^i(t+1) - p^i(t+1) \right) \Bigg)^2, \tag{20}$$

where $w^j(t+1) = (1+r)w(t) + \sum_{i=0}^{n} b_i^j(t)u_i(t)$ are the future wealths of the portfolio for each scenario $j = 1, \ldots, M$. The values $b_i^j(t) = x_i^j(t+1) - x_i(t)$ are obtained through *Monte Carlo simulation* of the dynamical model (1), each one corresponding to a different realization of the disturbance $[z^x(\tau),\ z^y(\tau)]$ in the time interval $[t\Delta_T,\ (t+1)\Delta_T)$ given the current market state $m(t)$. The option pricing engine is used to generate the corresponding future option prices $p^j(t+1)$, $j = 1, \ldots, M$.

The proposed SMPC algorithm is summarized by algorithm 4.1, which is solved at each trading instant $t = 0, \ldots, T-1$.

Algorithm 4.1: SMPC algorithm for dynamic option hedging:

---

(1) Let t = current hedging date, $w(t)$ = current wealth of portfolio, $m(t)$ = current market state;
(2) Use Monte Carlo simulation to generate M scenarios of future market states $m^1(t+1), \ldots, m^M(t+1)$;
(3) Use a pricing engine to generate the corresponding future option prices $p^1(t+1), \ldots, p^M(t+1)$;
(4) Solve the least square problem (20) to minimize the sample variance of $w(t+1) - p(t+1)$;
(5) Rebalance the portfolio according of the optimal solution $u^*(t)$ of problem (20);
(6) End.

4.1.1. Pricing future option values. An option pricing engine is needed at step 2 of algorithm 4.1 to compute the future option prices $p^1(t+1), \ldots, p^M(t+1)$ over the generated scenarios, which may be a bottleneck of the proposed SMPC approach for exotic options. Several approaches exist to option pricing, such as those based on Monte Carlo simulation. If each option evaluation requires the simulation of L scenarios, then one has to simulate ML paths on-line at each trading period t to build the optimization problem (20), which may be a time-consuming task. Although advanced techniques exist for parallel computation of Monte Carlo simulations (see, for example, Tian *et al.* (2008)), alternative *off-line* function approximation techniques can be used to obtain option prices for each future scenario. The idea is to construct a function that returns the option price as a function of $m(t)$ (that is, of the current asset parameters and of other option-related quantities such as $\{x(t_k)\}$ in 'Napoleon cliquet' options). In this paper we use a function approximation inspired by the Monte Carlo method of Longstaff and Schwartz (2001) for pricing American derivates, in which the continuation value (i.e. the option value at a future date) is estimated by a regression of the discounted payoff on a base of functions of some state variables. This methodology proved to have superior performance with respect to other classical general-purpose function approximation methods.

4.1.2. Alternative cost functions. In the absence of transaction costs, the main advantage of the formulation in (20) is that the optimization problem is a simple least square problem in n variables that can be solved very quickly. Alternatively, one can modify the cost function in (20) to obtain a *linear programming* problem formulation in several ways:

- *Max deviation* of wealth $w(t+1)$ from option price $p(t+1)$:

$$\min \max_{j=1,\ldots,M} |w^j(t+1) - p^j(t+1)|. \tag{21}$$

- *Average absolute deviation* of wealth $w(t+1)$ from option price $p(t+1)$:

$$\min \frac{1}{M}\sum_{j=1}^{M} |w^j(t+1) - p^j(t+1)|. \tag{22}$$

- *Max shortfall* of wealth $w(t+1)$ from option price $p(t+1)$:

$$\min \max_{j=1,\ldots,M} \max\{p^j(t+1) - w^j(t+1), 0\}. \tag{23}$$

- *Average shortfall* of wealth $w(t+1)$ from option price $p(t+1)$:

$$\min \frac{1}{M}\sum_{j=1}^{M} \max\{p^j(t+1) - w^j(t+1), 0\}. \tag{24}$$

Moreover, the cost function can be further modified to penalize large long or short positions by adding penalties on squares (or absolute values) of $x^j(t)u^j(t)$ or $x^j(t+1)u^j(t)$, and indirectly on transaction costs by adding penalties on squares (or absolute values) of $x^j(t)(u^j(t) - u^j(t-1))$.

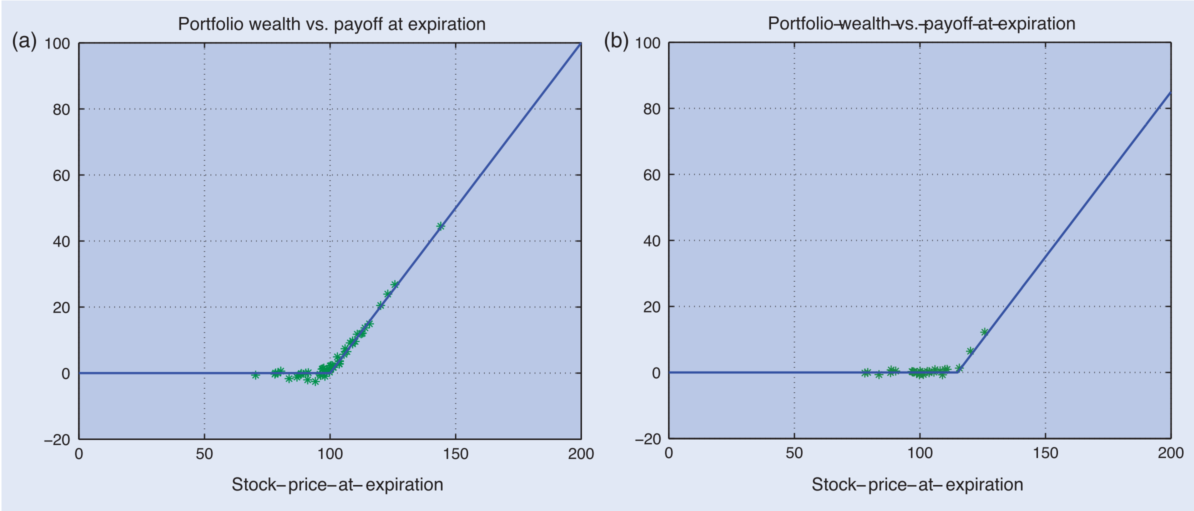

Figure 3. Hedging a European call using SMPC based on the BS model: payoff function $p(T)$ and final wealth $w(T)$ (€) as a function of the stock price $x_1(T)$ at expiration (€). (a) ATM option, $k=€100$ and (b) OTM option, $k=€115$.

5. Simulation results

In this section we test the SMPC algorithm 4.1 on different options and asset price models.

All simulation were performed on a MacBook Air with a 1.86 GHz Intel Core 2 Duo processor and 2 Gb RAM running MatlabR2009 under MS Windows, using the following parameters: $M=100$ is the number of scenarios, $N=1$ is the prediction horizon, $\Delta_T=1$ week is the time interval between consecutive reallocations of the portfolio, $T=24$ is the maturity in number of weeks of the option, and $r_a=4\%$ is the annualized continuously compounded interest rate so that $r=e^{0.04(1/54)}-1=0.00074102$ is the return of the risk-free investment over Δ_T. In every example shown below, we test the hedging strategy over 50 simulations of randomly generated market evolutions.

We will consider a single stock $x_1(t)$ with initial spot price $x_1(0)=€100$. For European call options (9), we will consider the strike price $K=€100$, unless indicated otherwise. The number of traded assets is $n=1$ when only the underlying stock is traded, or $n=2$ when the European call option with expiration at time $t+T$ and strike price $x_1(t)(1+r)^{T-t}$ is also traded in the portfolio. For 'Napoleon cliquet' options (10), we consider $N_{\text{fix}}=3$ fixing dates, with $t_0=0$, $t_1=8$, $t_2=16$, $t_3=24$ weeks, and coupon $C=0.1$. When Monte Carlo simulation is used to price 'Napoleon cliquet' options, $L=1000$ scenarios are evaluated to compute the expected payoff.

We will consider the log-normal stock price model (2) with† $\mu=r_a$, $dz_1^x\sim\mathcal{N}(0,1)$ and volatility $\sigma=0.2$, which will also be referred to as the Black–Scholes (BS) model, and Heston's (H) model (3), with initial variance $y_1(0)=0.04$, and parameters $\theta_1=0.04$, $\kappa_1=1$, $\omega_1=0.3$ and $\rho_1=-0.5$. In all simulations we assume that the value of market volatility is estimated exactly.

5.1. European call option

We first test the SMPC strategy (20) to replicate a European call option, only trading the risk-free asset and the underlying stock ($n=1$). The BS model (2) is used both in the MPC formulation and to generate actual market prices in the simulation. An analytical pricing formula is used to compute future asset values $p^j(t+1)$, $j=1,\dots,M$. The results are depicted in figure 3(a), where each point denotes the final wealth value $w(T)$ at the end of a simulation. Note that we are considering here an at-the-money (ATM) option ($K=x_1(0)=€100$). The average CPU time to execute algorithm 4.1 is 7.52 ms. The actual paths $w(t)$, $p(t)$, $t=0,1,\dots,T$, are shown in figure 4(a), where paths corresponding to the same simulation are depicted in the same color. The numerical values obtained in one of the simulations are reported in table 1, where $u_0(t)$ denotes the amount of risk-free asset in the portfolio. The last column reports the quantity $x_1(t)\Delta(t)$, where $\Delta(t)=\partial p(t)/\partial x_1(t)$ is computed by numerical differentiation of the Black–Scholes analytical pricing formula. It is apparent that the MPC approach provides a trading move very close to the classical 'delta hedging' $u_1(t)=\Delta(t)x_1(t)$, as must be expected from Black–Scholes theory for exact hedging. A comparison between MPC hedging and delta hedging, which takes an average CPU time of 0.2 ms per time step, in a simulation of the market under the same stock price realizations is reported in figure 5. In both cases the hedging errors $e(T)=w(T)-p(T)$ are within $\pm€2.5$, and the difference between hedging errors achieved by SMPC and delta hedging is within $\pm€0.13$.

In order to test the proposed SMPC hedging algorithm on out-of-the-money (OTM) options, we consider the same European call with strike price changed to $K=€115$, which roughly corresponds to one standard

†In this particular case, the probability measure used for the asset price and portfolio dynamics coincides with the risk-neutral one. However, the reader should note that the approach of this paper relies on the real-world probability measure for the asset price and portfolio dynamics.

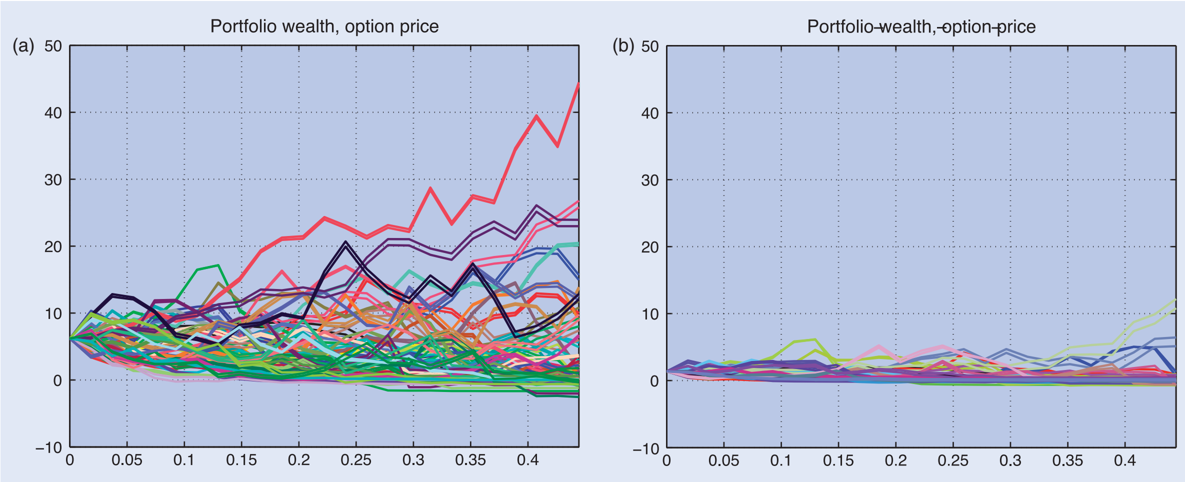

Figure 4. Trajectories of the wealth $w(t)$ (€) and the option price $p(t)$ (€) corresponding to the hedging results of figure 3. The same colors correspond to the same simulation. (a) ATM option, $k = €100$ and (b) OTM option, $k = €115$.

Table 1. Hedging of a European call based on the Black–Scholes model: results of a sample simulation ($u_0(t)$ is the amount of risk-free asset; $\Delta(t) = \partial p(t)/\partial x_1(t)$ is computed by numerical differentiation of the analytical pricing formula). Note the similarity between the last two columns.

t	$x_1(t)$	$w(t)$	$p(t)$	$u_0(t)$	$x_1(t)u_1(t)$	$x_1(t)\Delta(t)$
0.0000	100.000	6.196	6.196	−52.152	58.348	57.926
0.0185	101.367	6.955	6.865	−56.091	63.046	62.628
0.0370	96.897	4.134	4.261	−42.629	46.762	46.307
0.0556	94.582	2.985	3.108	−35.080	38.065	37.607
0.0741	93.057	2.345	2.415	−29.877	32.222	31.771
0.0926	93.371	2.431	2.395	−30.200	32.632	32.165
0.1111	94.295	2.732	2.591	−32.518	35.250	34.760
0.1296	88.192	0.426	0.859	−14.985	15.411	15.053
0.1481	90.411	0.803	1.199	−19.776	20.579	20.147
0.1667	88.586	0.373	0.754	−14.236	14.609	14.234
0.1852	87.683	0.214	0.544	−11.312	11.526	11.186
0.2037	90.998	0.641	1.000	−18.744	19.385	18.910
0.2222	94.742	1.425	1.867	−30.734	32.158	31.555
0.2407	99.890	3.149	3.945	−52.320	55.469	54.841
0.2593	102.720	4.682	5.466	−64.736	69.418	68.857
0.2778	99.723	2.609	3.439	−51.468	54.077	53.379
0.2963	99.591	2.499	3.147	−50.513	53.012	52.268
0.3148	98.178	1.709	2.233	−42.460	44.169	43.336
0.3333	100.471	2.709	3.142	−55.135	57.845	57.034
0.3519	102.804	4.012	4.363	−69.359	73.371	72.719
0.3704	97.457	0.144	1.202	−34.892	35.037	33.884
0.3889	97.789	0.238	1.030	−34.692	34.930	33.564
0.4074	98.881	0.602	1.089	−41.289	41.891	40.275
0.4259	97.699	0.071	0.308	−22.850	22.921	20.300
0.4444	96.002	−0.344	0.000	−0.344	0.000	0.000

deviation from ATM in the Black–Scholes world given the volatility σ of the underlying. In this case the ratio $x_1(0)/K \simeq 0.87$, which is considered OTM according to the classification of Rubinstein (1985), and $\Delta(0) = 0.1982$, which is also considered OTM by Bollen and Whaley (2004).

Next, we adopt Heston's model (3) for stock prices to replicate a European call option with strike price $K = €100$, only trading the risk-free asset and the

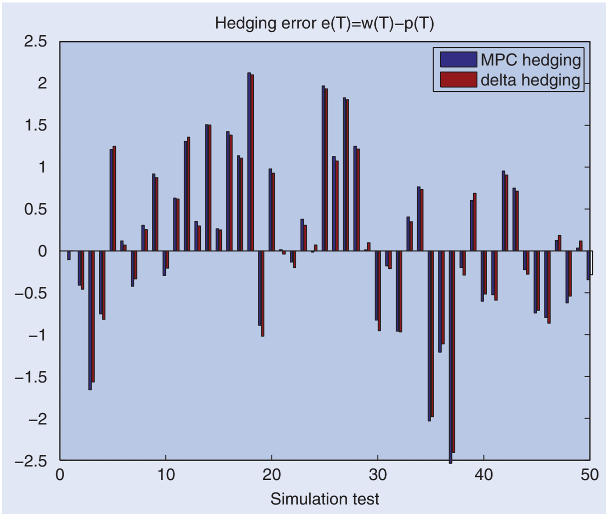

Figure 5. Comparison of the hedging errors $e(T) = w(T) - p(T)$ (€) corresponding to the hedging results of figure 3: SMPC vs. delta hedging of a European call based on the Black–Scholes model.

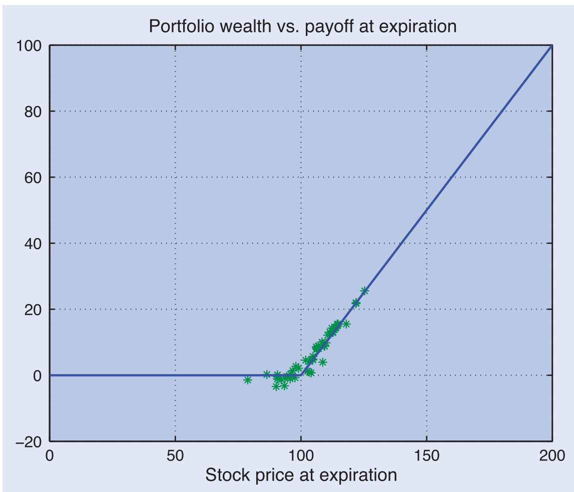

Figure 6. Hedging a European call using SMPC based on the Heston model: payoff function $p(T)$ and final wealth $w(T)$ (€) as a function of the stock price $x_1(T)$ at expiration (€).

underlying stock ($n = 1$). The analytical pricing formula (Heston 1993) is used to compute future asset values $p^j(t+1)$, $j = 1, \ldots, M$. The results are depicted in figure 6. The average CPU time to execute algorithm 4.1 is 85.5 ms. A comparison between MPC hedging and delta hedging, which takes an average CPU time of 1.85 ms per time step, in a simulation of the market under the same stock price realizations is reported in figure 7. In both cases the hedging errors $e(T) = w(T) - p(T)$ are within ±€5.5, and the difference between hedging errors achieved by SMPC and delta hedging is within ±€1.42.

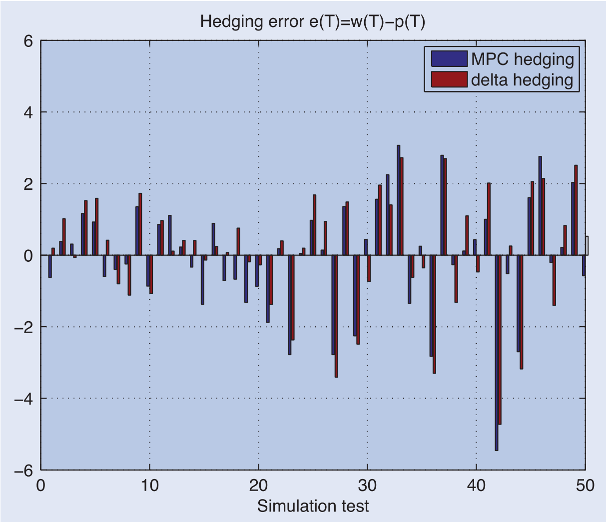

Figure 7. Comparison of the hedging errors $e(T) = w(T) - p(T)$ (€) corresponding to the hedging results of figure 6: SMPC vs. delta hedging of a European call based on the Heston model.

5.2. 'Napoleon cliquet' option

As described in the previous section, for plain vanilla options the SMPC approach produces hedging results comparable to standard (and simpler) delta-hedging approaches. The advantages of the SMPC strategy become apparent when replicating path-dependent exotic options. Consider first a 'Napoleon cliquet' option with payoff (10) based on the Black–Scholes stock price model (2). Figure 8 compares the results obtained by (a) adopting the SMPC strategy by only trading the risk-free asset and the underlying stock ($n = 1$) and Monte Carlo simulation to price future asset values $p^j(t+1)$, $j = 1, \ldots, M$, (b) adopting the SMPC strategy trading also the European call $x_2(t)$ associated with the same stock ($n = 2$), and (c) using delta hedging. The results are depicted in figure 8. The average CPU time to execute algorithm 4.1 is (a) 1.40 s, (b) 1.62 s and (c) 2.41 ms.

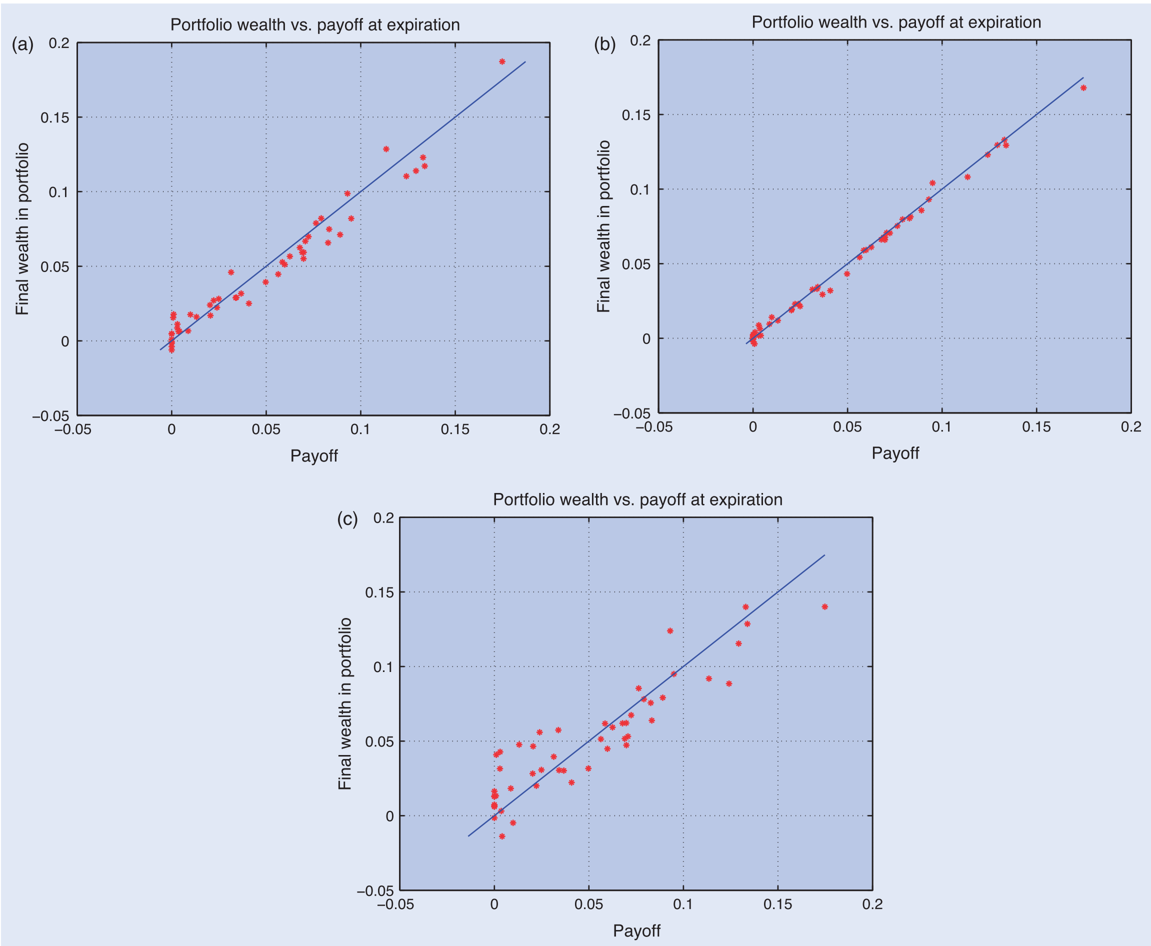

Figure 8. Hedging a 'Napoleon cliquet' option using SMPC based on the Black–Scholes model: final wealth $w(T)$ (€) vs. payoff $p(T)$ (€) at expiration. (a) SMPC $n = 1$; (b) SMPC $n = 2$ and (c) Delta hedging.

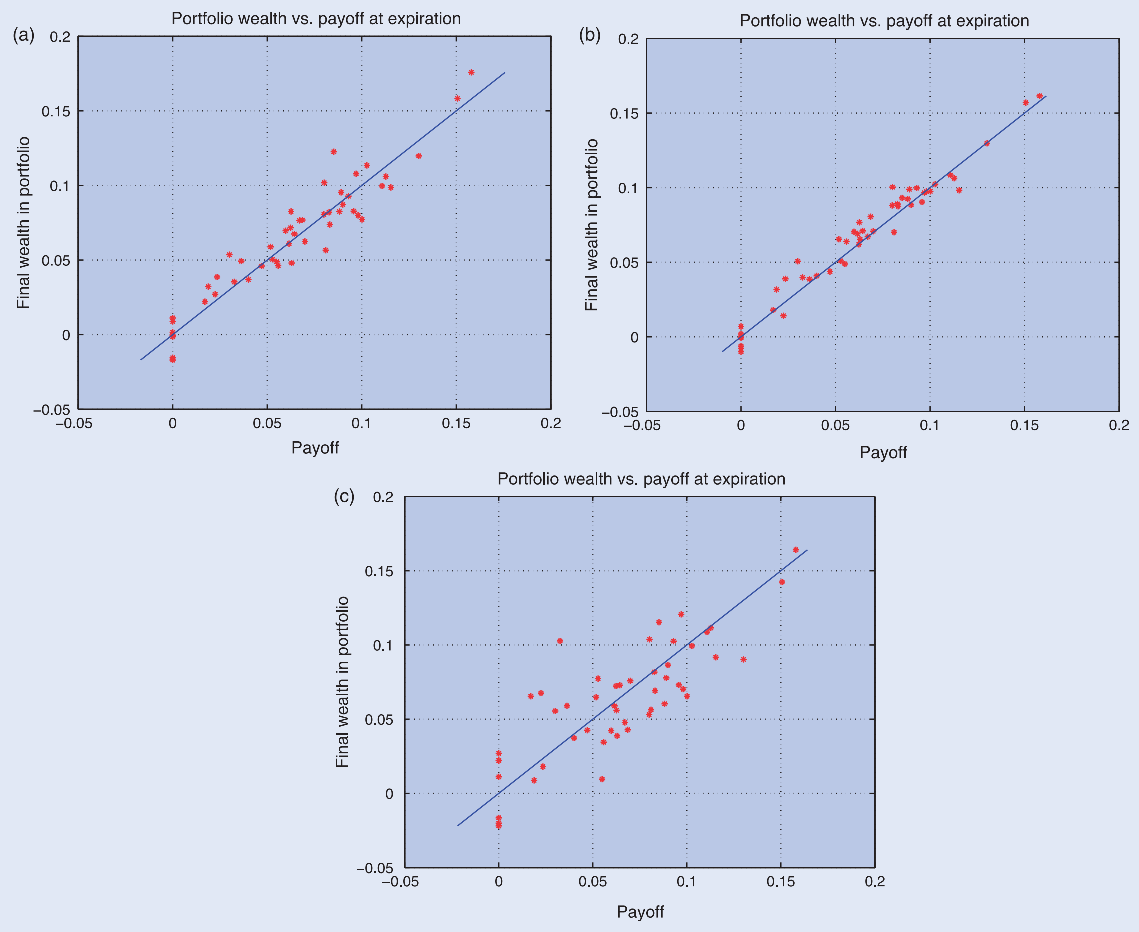

Figure 9. Hedging a 'Napoleon cliquet' option using SMPC based on the Heston model: final wealth $w(T)$ (€) vs. payoff $p(T)$ (€) at expiration. (a) SMPC $n=1$; (b) SMPC $n=2$ and (c) Delta hedging.

We estimate the option price $p(t)$ as a function of the spot price $x_1(t)$ and of the spot prices at past fixing dates $x_1(t_0), \ldots, x_1(t_k)$, with $t_k \leq t < t_{k+1}$, using the approximation method of Longstaff and Schwartz (2001). This takes about 76.6 s of *off-line* computation. Hedging results are almost indistinguishable from those obtained using on-line Monte Carlo simulation to price future options, with a drastic reduction of *on-line* CPU time to 50.5 ms ($n=1$) and 59.2 ms ($n=2$).

When Heston's model (3) is used, the Longstaff–Schwartz approximation method takes about 156 s of off-line CPU time to estimate the option price $p(t)$ as a function of the spot price $x_1(t)$, its variance $y_1(t)$, and of the spot prices at past fixing dates $x_1(t_0), \ldots, x_1(t_k)$, with $t_k \leq t < t_{k+1}$. On-line CPU time is 220 ms ($n=1$) and 277 ms ($n=2$), and the corresponding hedging results are depicted in figure 9. For comparison, figure 9(c) reports the results obtained using delta hedging, which is computed by numerical differentiation and whose on-line CPU time is 156 ms.

6. Conclusions

After highlighting how the dynamic hedging problem of financial options can be recast as a stochastic control problem, in this paper we have proposed a stochastic model predictive control approach based on a minimum variance criterion to rebalance periodically the portfolio underlying the option. We have shown that the tool is very versatile for dynamic option hedging. In fact, being based on Monte Carlo simulation, it can handle multiple assets, very general exotic options and payoff functions, and rather general stock price models. The computational demand of the SMPC approach is mostly due to pricing future option values, a task that can be alleviated by approximating the pricing function off-line.

The potential use of SMPC by financial institutions is twofold. It can be used on-line to suggest trading moves to traders, or off-line to run extensive simulations and quantify the average hedging error for a given market model and option type.

The results of this paper have recently been extended to test the robustness of the approach with respect to market modeling errors and to decrease the load of option pricing by sampling scenarios from a given probability distribution (Bemporad *et al.* 2010), and to handle proportional transaction costs and different risk measures in a numerically efficient way (Bemporad *et al.* 2011).

Acknowledgements

The authors sincerely thank Alexandre Radicchi and Eric Remo Robert Reynolds for stimulating discussions.

We also acknowledge MPS Capital Services for providing financial support for this research.

References

Batina, I., Stoorvogel, A.A. and Weiland, S., Optimal control of linear, stochastic systems with state and input constraints. Paper presented at the 41st IEEE Conference on Decision and Control, 2002.

Bemporad, A., Model-based predictive control design: New trends and tools. Paper presented at the 45th IEEE Conference on Decision and Control, San Diego, CA, 2006, pp. 6678–6683.

Bemporad, A. and Di Cairano, S., Optimal control of discrete hybrid stochastic automata. In *Hybrid Systems: Computation and Control (Lecture Notes in Computer Science 3414)*, edited by M. Morari and L. Thiele, pp. 151–167, 2005 (Springer: Berlin).

Bemporad, A. and Morari, M., Robust model predictive control: A survey. In *Robustness in Identification and Control*, edited by A. Garulli, A. Tesi, and A. Vicino, pp. 207–226, 1999 (Springer: Berlin).

Bemporad, A., Morari, M., Dua, V. and Pistikopoulos, E.N., The explicit linear quadratic regulator for constrained systems. *Automatica*, 2002, **38**(1), 3–20.

Bemporad, A., Morari, M. and Ricker, N.L., *Model Predictive Control Toolbox for Matlab – User's Guide*, 2004 (The Mathworks, Inc.). Available online at: http://www.mathworks.com/access/helpdesk/help/toolbox/mpc/

Bemporad, A., Gabbriellini, T., Puglia, L. and Bellucci, L., Scenario-based stochastic model predictive control for dynamic option hedging. Paper presented at the 49th IEEE Conference on Decision and Control, Atlanta, GA, 2010, pp. 6089–6094.

Bemporad, A., Puglia, L. and Gabbriellini, T., A stochastic model predictive control approach to dynamic option hedging with transaction costs. Paper presented at the American Control Conference, San Francisco, CA, 2011, pp. 3862–3867.

Bernardini, D. and Bemporad, A., Scenario-based model predictive control of stochastic constrained linear systems. Paper presented at the 48th IEEE Conference on Decision and Control, 2009, pp. 6333–6338.

Bertsimas, D., Kogan, L. and Lo, A.W., Hedging derivative securities and incomplete markets: An ε-arbitrage approach. *Oper. Res.*, 2001, **49**(3), 372–397.

Birge, J.R. and Louveaux, F., *Introduction to Stochastic Programming*, 1997 (Springer: New York).

Black, F. and Scholes, M., Pricing of options and corporate liabilities. *J. Polit. Econ.*, 1973, **81**(3), 637–654.

Blackmore, L., Bektassov, A., Ono, M. and Williams, B.C., Robust optimal predictive control of jump Markov linear systems using particles. In *Hybrid Systems: Computation and Control (Lecture Notes in Computer Science 4416)*, edited by A. Bemporad, A. Bicchi, and G. Buttazzo, pp. 104–117, 2007 (Springer: Berlin).

Bollen, N.P. and Whaley, R.E., Does net buying pressure affect the shape of implied volatility functions. *J. Finance*, 2004, **59**, 711–753.

Boyle, P., Options: A Monte Carlo approach. *J. Financ. Econ.*, 1977, **4**(3), 323–338.

Couchman, P.D., Cannon, M. and Kouvaritakis, B., Stochastic MPC with inequality stability constraints. *Automatica*, 2006, **42**(12), 2169–2174.

Cox, L.C., Ingersoll, J.E. and Ross, S.A., A theory of the term structure of interest rates. *Econometrica*, 1985, 385–407.

Duffie, D., *Dynamic Asset Pricing Theory*, 1996 (Princeton University Press: Princeton, NJ).

Edirisinghe, C., Naik, V. and Uppal, R., Optimal replication of options with transactions costs and trading restrictions. *J. Financ. Quantit. Anal.*, 1993, **28**(1), 117–138.

Fedotov, S., Stochastic optimization approach to options pricing. Paper presented at the American Control Conference, 1999, pp. 1450–1453.

Gondzio, J., Kouwenbergb, R. and Vorst, T., Hedging options under transaction costs and stochastic volatility. *J. Econ. Dynam. Control*, 2003, **27**, 1045–1068.

Herzog, F., Keel, S., Dondi, G., Schumann, L.M. and Geering, H.P., Model predictive control for portfolio selection. Paper presented at the American Control Conference, Minneapolis, MN, 2006, pp. 1252–1259.

Heston, S.L., A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.*, 1993, **6**, 327–343.

Heston, S.L. and Nandi, S., A closed-form GARCH option valuation model. *Rev. Financ. Stud.*, 2000, **11**(3), 585–625.

Høyland, K. and Wallace, S.W., Generating scenario trees for multistage decision problems. *Mgmt Sci.*, 2001, **47**(2), 295–307.

Hull, J.C., *Options, Futures, and Other Derivatives*, 6th ed., 2006 (Prentice Hall: Upper Saddle River, NJ).

Hull, J. and White, A., The pricing of options on assets with stochastic volatilities. *J. Finance*, 1987, **42**(2), 281–300.

Klaassen, P., Financial asset-pricing theory and stochastic programming models for asset/liability management: A synthesis. *Mgmt Sci.*, 1998, **44**(1), 31–48.

Klaassen, P., Comment on "Generating scenario trees for multistage decision problems". *Mgmt Sci.*, 2002, **48**(11), 1512–1516.

Kolmanovsky, I., Siverguina, I. and Lygoe, B., Optimization of powertrain operating policy for feasibility assessment and calibration: Stochastic dynamic programming approach. Paper presented at the American Control Conference, 2002, pp. 1425–1430.

Kouwenberg, R. and Vorst, T., Dynamic portfolio insurance: A stochastic programming approach. *Econometric Institute and Department of Finance, 9909*, 1998 (Erasmus University: Rotterdam).

Li, P., Wendt, M. and Wozny, G., Robust model predictive control under chance constraints. *Comput. Chem. Engng*, 2000, **24**, 829–834.

Longstaff, F.A. and Schwartz, E.S., Valuing American options by simulation: A simple least-squares approach. *Rev. Financ. Stud.*, 2001, **114**(1), 113–147.

Mayne, D.Q. and Rawlings, J.B., *Model Predictive Control: Theory and Design*, 2009 (Nob Hill Publishing: Madison, WI).

Meindl, P.J., Portfolio optimization and dynamic hedging with receding horizon control, stochastic programming, and Monte Carlo simulation. PhD thesis, Department of Management Science & Engineering, Stanford University, 2006.

Meindl, P. and Primbs, J., Dynamic hedging of single and multi-dimensional options with transaction costs: A general utility maximization approach. *Quantit. Finance*, 2008, **8**(3), 299–312.

Merton, R.C., The theory of rational option pricing. *Bell J. Econ. Mgmt Sci.*, 1973, **4**, 141–183.

Muñoz de la Peña, D., Bemporad, A. and Alamo, T., Stochastic programming applied to model predictive control. Paper presented at the 44th IEEE Conference on Decision and Control and European Control Conference, Sevilla, Spain, 2005, pp. 1361–1366.

Oldewurtel, F., Jones, C.N. and Morari, M., A tractable approximation of chance constrained stochastic MPC based on affine disturbance feedback. Paper presented at the 47th IEEE Conference on Decision and Control, Cancun, Mexico, 2008, pp. 4731–4736.

Petrelli, A., Zhang, J., Siu, O., Chatterjee, R. and Kapoor, V., *Optimal dynamic hedging of cliquets*, DefaultRisk.com 2008.

Pochart, B. and Bouchaud, J.P., Option pricing and hedging with minimum local expected shortfall. *Quantit. Finance*, 2004, **4**(5), 607–618.

Potters, M., Bouchaud, J.P. and Sestovic, D., Hedged Monte-Carlo: Low variance derivative pricing with objective probabilities. *Physica, A*, 2001, **289**(3/4), 517–525.

Primbs, J.A., Stochastic Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise. Paper presented at the American Control Conference, New York, NY, 2007a, pp. 4470–4475.

Primbs, J.A., A soft constraint approach to stochastic receding horizon control. Paper presented at the 46th IEEE Conference on Decision and Control, 2007b, pp. 4797–4802.

Primbs, J.A., Dynamic hedging of basket options under proportional transaction costs using receding horizon control. *Int. J. Control*, 2009, **82**(10), 1841–1855.

Primbs, J.A. and Sung, C.H., Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise. *IEEE Trans. Automat. Control*, 2009, **54**(2), 221–230.

Rawlings, J.B., Tutorial overview of model predictive control. *IEEE Control Syst. Mag.*, 2000, 38–52.

Rubinstein, M., Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August 31, 1978. *J. Finance*, 1985, **40**, 455–479.

Sahinidis, N.V., Optimization under uncertainty: State-of-the-art and opportunities. *Comput. Chem. Engng*, 2004, **28**(6/7), 971–983.

Schwarm, A.T. and Nikolaou, M., Chance-constrained model predictive control. *AIChE J.*, 1999, **45**(8), 1743–1752.

Tian, X., Benkrid, K. and Gu, X., High performance Monte-Carlo based option pricing on FPGAs. *Engng Lett.*, 2008, **16**(3), 434–442.

van Hessem, D.H. and Bosgra, O.H., A conic reformulation of Model Predictive Control including bounded and stochastic disturbances under state and input constraints. Paper presented at the 41st IEEE Conference on Decision and Control, Las Vegas, NV, 2002, pp. 4643–4648.

Wang, W. and Boyd, S., Fast model predictive control using online optimization. *Paper presented at the 17th IFAC World Congress*, 2008, pp. 6974–6979.

Zhao, Y. and Ziemba, W.T., A stochastic programming model using an endogenously determined worst case risk measure for dynamic asset allocation. *Math. Program.*, 2001, **89**, 293–309.