49th IEEE Conference on Decision and Control
December 15-17, 2010
Hilton Atlanta Hotel, Atlanta, GA, USA

Scenario-based Stochastic Model Predictive Control for Dynamic Option Hedging

Alberto Bemporad†, Tommaso Gabbriellini‡, Laura Puglia†, Leonardo Bellucci‡

Abstract— **For a rather broad class of financial options, this paper proposes a stochastic model predictive control (SMPC) approach for dynamically hedging a portfolio of underlying assets. By employing an option pricing engine to estimate future realizations of option prices on a finite set of one-step-ahead scenarios, the resulting stochastic optimization problem is easily solved as a least-squares problem at each trading date with as many variables as the number of traded assets and as many constraints as the number of predicted scenarios. After formulating the dynamic hedging problem as a stochastic control problem, we test its ability to replicate the payoff at expiration date for plain vanilla and exotic options. We show not only that relatively small hedging errors are obtained in spite of price realizations, but also that the approach is robust with respect to market modeling errors.**

I. Introduction

Issuing derivative contracts requires dynamically trading a self-financing portfolio of more liquid and simpler securities so as to match the option payoff for every possible state of the market. A popular and well studied class of derivatives 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*. 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 get the right to exercise the option, and (ii) the dynamic strategy for managing this money by creating a portfolio and for 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.

The seminal works [1], [2] and their extensions to models with stochastic volatility [3] 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. *Simulation* is another method often used by investment firms to price options [4], [5]. A (large) set of scenarios for the future prices of the underlying assets is generated by Monte Carlo simulation; 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, in this paper 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 expiration date. The problem can be cast as a stochastic optimal control problem and rely on the Hamilton-Jacobi-Bellman partial differential equation. This category includes *multi-stage stochastic programming* approaches, in which the pricing and hedging problem is solved as a stochastic linear programming problem [6]–[8]. The approach is often limited by numerical reasons. In fact, the number of nodes in the tree is exponential in the number of trading periods, which typically limits the number of branches at each node to two or three. *Stochastic dynamic programming* approaches [9], [10] 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 numerical explosion when the number of trading periods is large and several assets are traded.

This paper attacks the hedging problem from a feedback control viewpoint and proposes stochastic model predictive control (SMPC) ideas [11]–[13] to design a dynamic hedging strategy. 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 in financial applications only very recently, such as in [14] for portfolio optimization, and in [15], [16] for option pricing and hedging.

In this paper we propose a novel SMPC approach to dynamic option hedging based on a minimum variance criterion that requires a simple *least-squares optimization* to evaluate the optimal trading moves, by extending results proposed in [17]. To be able to handle very general stock price models and exotic payoffs, for which no analytic hedging policy exist, a pricing engine is used on-line to generate a finite number of future scenarios of option prices, rather than analytically deriving expected values from pricing models as in [14], [15]. To evaluate each option price, the pricing engine employs either Monte Carlo simulation (*on-line* computations), or *off-line* function approximation 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.

Corresponding author. †Department of Mechanical and Structural Engineering, University of Trento, Via Mesiano 77, 38100 Trento, Italy. E-mail: bemporad@ing.unitn.it, laura.puglia@ing.unitn.it

‡MPS Capital Services, Modelling & Quantitative Analysis, Viale Mazzini 23, Siena, Italy. E-mail: tommaso.gabbriellini@mpscapitalservices.it, leonardo.bellucci@mpscapitalservices.it

978-1-4244-7744-9/10/$26.00 ©2010 IEEE

The paper is organized as follows. In Section II 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 III we state the hedging problem as a stochastic model predictive control problem, and formulate a solution algorithm. The performance of the approach and its robustness with respect to modeling errors of the underlying asset market are tested in Section IV on a European call and on two exotic options.

II. Model formulation

A. Asset price dynamics

Consider the problem of dynamically hedging a European option[1] $\mathcal{O}$ 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

$$dx_i(\tau) = \mu_i^x(x_i(\tau), y_i(\tau))d\tau + \sigma_i^x(x(\tau), y(\tau))dz_i^x \quad (1a)$$
$$dy_i(\tau) = \mu_i^y(y_i(\tau))d\tau + \sigma_i^y(y(\tau))dz_i^y \quad (1b)$$

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

$$dx_i(\tau) = (\mu d\tau + \sigma dz_i^x)x_i(\tau) \quad (2)$$

and Heston's model [3].

In this paper we are interested in evaluating $x_i(\tau)$, $y_i(\tau)$ at certain trading dates $\tau = t\Delta_T$, where $t \in \mathbb{Z}$, $t \geq 0$, denotes a discrete-time index[2]. To this end, we need to discretize (1) into difference equations either through the exact integration of (1), for example

$$x_i(t+1) = e^{(\mu - \frac{1}{2}\sigma^2)\Delta t + \sigma\sqrt{\Delta t} z_i^x(t)} x_i(t) \quad (3)$$

or by numerical integration. 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.

B. 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 number of sampling steps. The payoff $p(T)$ of $\mathcal{O}$ is described by a function $\mathcal{P}$:

$$p(T) = \mathcal{P}(m(T)) \quad (4)$$

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

$$p(t) = (1+r)^{t-N} \tilde{E}\left[\mathcal{P}(m(T)) | m(t)\right] \quad (5)$$

where $m(T)$ is the state of the market at time t and $\mathcal{P}(m(T))$ is the expected value of the payoff in the risk-neutral measure, given the market at time t. In (5) $r = e^{r_a \Delta T} - 1$ 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 (5) can be also restated recursively as $p(t) = (1+r)^{-1}\tilde{E}[p(t+1)|m(t)]$). For instance, for a European call option on a single stock x with strike price K, we have

$$p(T) = \max\{x(T) - K, 0\} \quad (6)$$

$m(t) = \{x(t), y(t)\}$, and $p(t) = e^{-r(N-t)}\tilde{E}\left[\max\{x(T) - K, 0\} | x(t), y(t)\right]$. In particular, for log-normal price models, $m(t) = x(t)$. For "Napoleon cliquet" path-dependent exotic options

$$p(T) = \max\left\{0, C + \min_{i \in \{1, \dots, N_{\text{fix}}\}} \frac{x(t_i) - x(t_{i-1})}{x(t_{i-1})}\right\} \quad (7)$$

where t_i, $i = 1, \dots, N_{\text{fix}}$ are the fixing dates, and C is a fixed value. In this case $m(t) = \{x(t_0), \dots, x(t_k), x(t), y(t)\}$, where k is the fixing index such that $t_k \leq t < t_{k+1}$. For *weak* path-dependent "Barrier" exotic options

$$\begin{aligned} p(T) &= \begin{cases} \max(x(T) - K, 0) & \text{if } x(t) < x_u, \ \forall t \leq T \\ 0 & \text{otherwise} \end{cases} \\ &= \begin{cases} \max(x(T) - K, 0) & \text{if } x_\ell(t) = 0 \\ 0 & \text{if } x_\ell(t) = 1 \end{cases} \end{aligned} \quad (8)$$

where x_u define the upper barrier level, and $x_\ell(t) \in \{0, 1\}$ is a logic state with dynamics $x_\ell(t+1) = x_\ell(t)$ OR $[x(t) \geq x_u]$, $x_\ell(0) = 0$. In this case $m(t) = \{x(t), x_\ell(t), y(t)\}$.

C. Portfolio dynamics

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

$$w(t) = u_0(t) + \sum_{i=1}^n x_i(t)u_i(t) \quad (9)$$

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

$$w(t+1) = (1+r)w(t) + \sum_{i=1}^n b_i(t)u_i(t) \quad (10)$$

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 $\mathcal{O}$, $w(0) = (1+r)^{-N}\tilde{E}[p(T)|x(0), y(0)]$.

[1] The approach can be easily extended to other options as described in [9], such as to Asian options.

[2] The results presented in this paper can be immediately extended to non-uniform trading intervals Δ_T.

We remark 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 (10) is a first-order linear stochastic discrete-time system.

III. Stochastic MPC formulation

Based on the models developed in Section II, the dynamic option problem can be reformulated in system theoretical terms as a stochastic reference tracking and disturbance rejection problem, in which the wealth $w(t) \in \mathbb{R}$ is 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 $p(t)$ 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 $x(t)$.

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 variance of the hedging error, that we address next in a *stochastic model predictive control* (SMPC) setting.

We employ a stochastic programming approach to SMPC as in [18] to formulate the minimum variance problem by enumerating a certain number M of scenarios (or, alternatively, of tree nodes) of future asset price realizations. Each scenario corresponds to the realization of a certain sequence of stochastic variables and has a probability π_j of occurring, $j = 1, \ldots, M$, $\pi_j > 0$, $\pi_j \leq 1$, $\sum_{k=1}^{M} \pi_j = 1$. To limit the number of scenarios, one could discretize more roughly in the probability space. Another way of reducing complexity 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). In this paper we choose the special case $N = 1$, and adopt the *terminal condition* of "perfect hedging" between time $t + 1$ and T. 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

$$\min_{\{u(t)\}} \ \mathrm{Var}_{m_{t+1}} \left[w(t+1, m_{t+1}) - p(t+1, m_{t+1}) \right] \tag{11a}$$

$$\text{s.t.} \ w(t+1, m_{t+1}) = (1+r)w(t) + \sum_{i=0}^{n} b_i(t, m_{t+1}) u_i(t) \tag{11b}$$

In (11) m_{t+1} represents a generic realization of the state of the considered asset market at time $t+1$. This is determined by the stochastic noise realization $z(t+1)$, corresponding to the realization of future asset prices $x(t + 1)$. Expectations (i.e., the variance of the final hedging error being minimized) are taken in the real-world measure with respect to the current state $m(t)$ of the market, conditioned on the values $z(0), \ldots, z(t)$ already realized.

Problem (11) is a one-step-ahead minimum variance problem. The reason for focusing on the formulation (11) 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 $t + 1$. By optimizing the sample variance of $w(t + 1) - p(t + 1)$, problem (11) can be rewritten as the following very simple *least squares* problem

$$\min_{u(t)} \sum_{j=1}^{M} \pi^j \left(e^j(t+1) - \left(\frac{1}{M} \sum_{i=1}^{M} e^i(t+1) \right) \right)^2 \tag{12}$$

where $e^j(t + 1) = w^j(t + 1) - p^j(t + 1)$, $w^j(t + 1) = (1 + r)w(t) + \sum_{i=0}^{n} b_i^j(t) u_i(t)$ are the future wealths of the portfolio for each scenario $j = 1, \ldots, M$, and π^j is the corresponding probability, $\pi^j \geq 0$, $\sum_{i=1}^{M} \pi^j = 1$.

Each scenario corresponds 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 III.1, which is solved at each trading instant $t = 0, \ldots, T - 1$.

1. Let t=current hedging date, $w(t)$= current wealth of portfolio, $m(t)$=current market state;
2. Generate M scenarios of future market states $m^1(t + 1)$, …, $m^M(t + 1)$, with corresponding probabilities π^1, …, π^M;
3. Use a pricing engine to generate the corresponding future option prices $p^1(t + 1)$, …, $p^M(t + 1)$;
4. Solve the least square problem (12) 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 (12);
6. End.

Algorithm III.1: SMPC algorithm for dynamic option hedging

A. Scenario generation

The values $b_i^j(t) = x_i^j(t + 1) - x_i(t)$ can be obtained through *Monte Carlo simulation* of the dynamical model (1) and $\pi^j = \frac{1}{M}$. Alternatively, future asset scenarios can be generated based on the discretization (computed off-line) of the normal distribution of $[z^x(\tau), z^y(\tau)]$. Consider for simplicity the special case of a log-normal prediction model ($z^y(t) = 0$) with a single asset. We use model (3) to compute discrete probabilities associated with scenarios, under the assumption that $z^x(t)$ has a normal Gaussian distribution $\pi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}$ (zero mean, unit variance). For a fixed grid $\bar{z}_j$, $j = 1, \ldots, M$, $z_1 = -\infty$, $z_j = -3\sigma \left(1 - \frac{2(j-2)}{M-3}\right)$, $z_M = +\infty$, we obtain

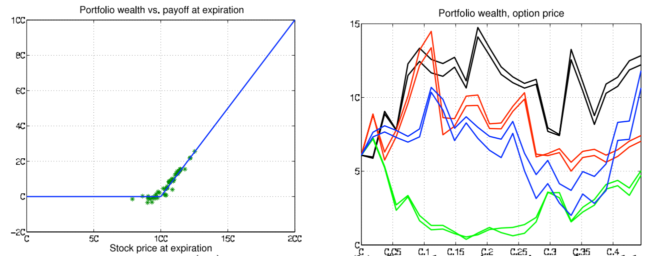

(a) payoff function $p(T)$ and final wealth $w(T)$ (€) as a function of the stock price $x_1(T)$ at expiration (€)

(b) Sample trajectories of wealth $w(t)$ and option price $p(t)$

Fig. 1. Hedging a European call using SMPC based on Heston's model (values in €)

$$
\begin{aligned}
\pi^j &= \int_{\bar{z}_j}^{\bar{z}_{j+1}} \pi(z)dz = \frac{1}{2}\left(\operatorname{erf}\left(\frac{\bar{z}_{j+1}}{\sqrt{2}}\right) - \operatorname{erf}\left(\frac{\bar{z}_j}{\sqrt{2}}\right)\right) \\
x^j(t+1) &= \frac{1}{\pi^j}\int_{\bar{z}_j}^{\bar{z}_{j+1}} x(z)\pi(z)dz \\
&= x(t)\frac{k_1}{2\pi^j}e^{\frac{k_2}{2}}\left(\operatorname{erf}\left(\frac{\bar{z}_{j+1}-k_2}{\sqrt{2}}\right) - \operatorname{erf}\left(\frac{\bar{z}_j - k_2}{\sqrt{2}}\right)\right) \\
p^j(t+1) &= p(t+1, m(z_j))
\end{aligned} \tag{13}
$$

where $k_1 = e^{\mu - \frac{1}{2}\sigma^2\Delta t}$, $k_2 = \sigma\sqrt{\Delta t}$, $x(z) = x(t)k_1 e^{\frac{k}{2}} z$, and $m(z)$ is the state of the asset market.

B. Pricing future option values

An option pricing engine is needed at step 3. of Algorithm III.1, 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. In alternative, *off-line* function approximation techniques can used to construct a function that returns the option price as a function of $m(t)$. In this paper we use a function approximation inspired by the Monte Carlo method of Longstaff and Schwartz [19] for pricing American derivates, in which the continuation value (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.

C. Simulation results

In this section we test the SMPC algorithm III.1 on different options and asset price models. All simulation were performed on an Asus 1.70 GHz Intel Pentium R processor and 2 Gb RAM running MATLAB™ R2007b under Windows XP, using the following parameters: $M = 100$ scenarios (unless specified differently), prediction horizon $N = 1$, $\Delta_T = 1$ week is the time interval between consecutive reallocations of the portfolio, $T = 24$ weeks is the maturity of the option, $r_a = 4\%$ is the annualized continuously compounded interest rate so that $r = e^{0.04\frac{1}{54}} - 1$ is the return of the risk free investment over Δ_T. In each example the hedging strategy is tested over $N_s = 1000$ simulations of randomly generated market evolutions.

We consider a single stock $x_1(t)$ with initial spot price $x_1(0) = 100$ €. For European call options (6), we consider the strike price $K = 100$ €. The number of traded assets is $n = 1$ when only the underlying stock is traded, or $n = 2$ when also the European call option with expiration at time $t\Delta T$ and strike price $x_1(t)(1+r)^{T-t}$ is also traded in the portfolio. For "Napoleon cliquet" options (7), we consider $N_{\text{fix}} = 3$ fixing dates, with $t_0 = 0$, $t_1 = 8$, $t_2 = 16$, $t_3 = 24$ weeks, and coupon $C = 0.1$. For barrier options, we consider an UP-AND-OUT option with barrier $x_u = 120$ €, where the barrier level is checked only at trading instants. When Monte Carlo simulation is used to price "Napoleon cliquet" and Barrier options, $L = 1000$ simulations are evaluated to compute each expected payoff value.

We consider the log-normal stock price model (2) with[3] $\mu = r_a$, $dz_1^x \sim \mathcal{N}(0,1)$ and volatility $\sigma = 0.5$, which will be also referred to as Black-Scholes (BS) model, and Heston (H) model [3], with initial variance $y_1(0) = 0.25$, and parameters $\theta_1 = 0.25$, $\kappa_1 = 1$, $\omega_1 = 0.3$, $\rho_1 = -0.5$. In all simulations we assume that the value of market volatility is estimated exactly.

1) European call option: We first test the SMPC strategy (12) to replicate a European call option, only trading the risk-free asset and the underlying stock ($n = 1$). Heston's model [3] is used both in the MPC formulation and to generate actual market prices in simulation. Only the risk-free asset and the underlying stock are traded ($n = 1$). The analytical pricing formula [3] is used to compute future asset values $p^j(t+1)$, $j = 1, \ldots, M$. The results are depicted in Figure 1, where only the first 50 simulations are reported in Figure III-B and 4 simulations in Figure III-B. The empirical distribution of the hedging error[4] computed on all N_s simulations is depicted in Figure 2 (purple line). The average CPU time to execute Algorithm III.1 is 81.2 ms. The average hedging error $E[e(T)] = -0.0511$ €, $E[|e(T)] = 1.9907$, $\max|e(T)| = 14.5699$. For comparison, Figure 2 also shows the error distribution when delta hedging[5] is applied (green line), which takes an average CPU time of 2.5 ms per time step. In each simulation, the difference between the hedging error $e(T)$ achieved by SMPC and the one obtained by delta hedging is within ± 3.75 €.

2) Exotic options: The advantage of using the SMPC strategy becomes more evident when replicating exotic options. We use again Heston's model [3] both as a market model and a prediction model for stock prices. For the "Napoleon Cliquet" option, we only consider the case $n = 2$ and we use Longstaff-Schwartz's off-line approximation (calibrated in 251.5310 s) 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 $x_1(t_0)$, ..., $x_1(t_k)$,

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

[4]Hedging errors are sampled with the Freedman-Diaconis rule [20].

[5]By letting $\Delta = \frac{\partial p}{\partial x}$, in *Delta hedging* at each time step the portfolio contains a quantity $-\Delta$ of asset x. In our simulations Δ is computed by differentiating the pricing formula [3] numerically.

with $t_k \leq t < t_{k+1}$. On-line CPU time is 0.4391 s per time step (for comparison, when using on-line Monte Carlo simulation to compute future options CPU time is 2.49 s). Hedging results are reported in the third and fourth rows of Table I, where for comparison in the fifth row we also show the results obtained through delta hedging.

For the barrier option, off-line pricing approximation takes 114.016 s to estimate $p(t)$ as a function of $x_1(t)$ and its variance $y_1(t)$. On-line CPU time is 428.8 ms ($n = 2$). Hedging results are reported in the last two rows of Table II.

SMPC model	$E[e(T)]$	$E[\|e(T)\|]$	$\max \|e(T)\|$	CPU (ms)
Fixed Black	0.0031	0.0080	0.0561	1256.28
Implied Black	0.0031	0.0079	0.0560	1293.7
Heston (MC)	0.0032	0.0075	0.0516	6717.48
Heston (LS)	0.0025	0.0110	0.4159	439.1
Δ hedging	-0.0032	0.0176	0.1344	33.7

TABLE I

NAPOLEON CLIQUET OPTION (FINAL HEDGING ERROR $e(T)$ IN €, MC=MONTE CARLO ONLINE PRICING, LS=LONGSTAFF&SCHWARTZ OFFLINE OPTION PRICE APPROXIMATION)

SMPC model	$E[e(T)]$	$E[\|e(T)\|]$	$\max \|e(T)\|$	CPU (ms)
Fixed Black	-0.0324	0.6965	14.5866	113.12
Implied Black	-0.0320	0.6956	14.5829	155.64
Heston	0.1936	0.7961	16.0870	363.76
Δ hedging	-1.3060	2.4335	18.9145	103.0

TABLE II

BARRIER OPTION (FINAL HEDGING ERROR $e(T)$ IN €)

IV. ROBUSTNESS WITH RESPECT TO MARKET MODELING ERRORS

Generating future scenarios of asset prices requires a model of their stochastic and dynamic evolution. Getting such a model is often a complex task and unavoidably affected by inaccuracy. This is due to the fact that we are trying to enclose a huge net of complicated relationships, in addition to a large source of randomness, in a small box. As complicated as the model can be, one will never be able to catch the exact dynamics of the assets, and in any case a very complicated model would lose the advantages of modelization. Therefore, in general, the asset price model will always be different from the way the real world behaves, and one must find a compromise, by using a simple enough model which allows one to keep computational complexity as low as possible.

In the previous sections we have assumed that the actual prices behave according to the same model we use to predict their evolution (nominal conditions). The hedging error was exclusively due to randomness. In this section we test numerically the robustness of the SMPC algorithm not only with respect to price stochasticity, but also when real and prediction model mismatch. In particular, we assume that real assets evolve following Heston's model [3], while the simpler Black and Scholes model (2) is used to generate future scenarios in SMPC.

The tool that will be used to concile the two models is the calibration of the lognormal model (2) using the so-called *implied volatility*, which is the market's view of future actual volatility and is updated at each trading period from observed

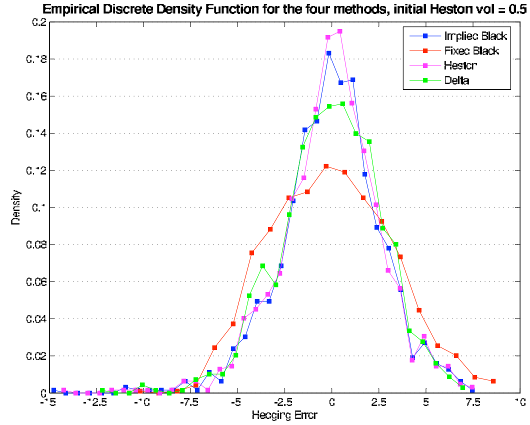

Fig. 2. Comparing the empirical distribution of hedging errors among the four methods: *Heston*, *Implied BS*, *Fixed BS* and *Delta Hedging*.

market prices of plain vanilla options (generated by the Heston's model [3] in our setting) by inverting (numerically) the Black-Scholes pricing formula. Such a value of implied volatility will be used in our simple prediction model (3).

A. Simulation results

Assume the real market evolves according to Heston's model with initial volatility $\sigma = 0.5$ ($y_1(0) = 0.25$), and that, to avoid bias in hedging errors due to wrong initial pricing, the initial wealth of the portfolio is computed correctly using Heston's model and exact $y_1(0)$. For SMPC we consider instead three different models:

1) *Fixed Black-Scholes*: The log-normal model (2) is used to generate future scenarios in SMPC, setting the volatility to a fixed arbitrary value, different from the actual;
2) *Implied Black-Scholes*: at each prediction step the estimated implied volatility is used in (2);
3) *Heston*: nominal case, both the SMPC model and the real market model coincide, and the actual volatility is observed exactly.

1) European Call: We first test the robustness of the SMPC algorithm on a European call option, only trading the risk-free asset and the underlying stock ($n = 1$). The analytical pricing formula [3] is used to compute future asset values $p^j(t+1)$, $j = 1, \ldots, M$. Figure 2 shows the empirical discrete density function of the hedging error $e(T) = w(T) - p(T)$ in the presence of modeling errors. Note that all four distributions are bell-shaped. We can easily see that the density of *Fixed BS* (red line) has fatter tails than the others and that *Implied BS* (blue line) better follows the distribution of *Heston* (=the exact model, purple line). While *Fixed BS* and *Implied BS* take approximately the same CPU time (9.6 ms and 10.2 ms per time step, respectively), *Heston* (nominal conditions) takes 81.2 ms per time step. Standard *delta hedging* is faster: only takes 2.5 ms per time step, because it simply uses finite differences.

The benefits of resorting to the discretization of the normal distribution as in (13) with respect to Monte Carlo simulation $\pi^j = \frac{1}{M}$ are highlighted in Table III, where the *Implied BS* method is used to hedge in the SMPC algorithm. Note that the average final hedging error and the average absolute hedging error obtained when only 3 scenarios, weighted with the corresponding probabilities as in (13), are used in SMPC are very similar to the case with $M = 100$ scenarios generated by Monte Carlo, but with evident savings of CPU time.

2) Exotic options: The robustness with respect to modeling errors in the case of path-dependent "Napoleon cliquet" options with payoff (7) is highlighted in (the first and second rows of) Table I, where we use $M = 100$ equally probably scenarios generated by using Longstaff-Schwartz's off-line approximation. For exotic options we only consider the case $n = 2$, that is, trading both the asset and its associated call option. While all methods perform similarly, it is apparent the computational benefits of hedging using the log-normal model, in spite of the modeling error. Note that, although delta hedging is the fastest algorithm, its performance in terms of $E[|e(T)|]$ deteriorates by almost 50% with respect to Implied Black and almost 60% with respect to SMPC based on Heston's model; partly this is because delta hedging does not include options in the portfolio ($n = 1$). Similar results are obtained on the UP-AND-OUT Barrier option, as shown in Table II.

V. CONCLUSIONS

After recasting the dynamic hedging problem of financial options 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 showed that the tool is very versatile for dynamic option hedging, as it can handle multiple assets, very general exotic options and payoff functions, and rather general stock price models, and is also robust with respect to market modeling assumptions. The computational demand of the SMPC approach is mostly due to pricing future option values, a task which can be alleviated in three ways: (i) by approximating the pricing function off-line, (ii) by using a simplified log-normal model (with implied volatility), and (iii) by sampling the uniform distribution instead of generating random and equally probably samples using Monte Carlo simulation. In this paper we assumed that transaction costs are negligible. Current research is extending the results of this paper to cope with transaction costs, based on either quadratic or linear programming formulations of the SMPC problem.

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.

M	π^j	$E[e(T)]$	$E[\|e(T)\|]$	$\max\|e(T)\|$	CPU (ms)
100	$\frac{1}{M}$	-0.0587	2.0296	15.0929	10.2
8	$\frac{1}{M}$	-0.1177	3.7149	18.2113	3.4
8	Eq. (13)	-0.0914	2.1517	13.7776	4.4
5	$\frac{1}{M}$	-0.1763	5.3472	20.4058	3.2
5	Eq. (13)	-0.0962	2.1697	13.5410	3.8
3	$\frac{1}{M}$	-0.1717	5.2603	20.5207	3.1
3	Eq. (13)	-0.0501	2.0153	15.2368	3.4

TABLE III

MONTECARLO VS. DISCRETIZATION OF PROBABILITY DENSITY FUNCTION IN GENERATING SCENARIOS (EUROPEAN CALL, ERRORS EXPRESSED IN €)

ACKNOWLEDGEMENTS

The authors thank Alexandre Radicchi and Eric Remo Robert Reynolds for stimulating discussions. We also ackowledge MPS Capital Services for providing financial support for this research.

REFERENCES

[1] F. Black and M. Scholes, "Pricing of options and corporate liabilities," *Journal of Political Economy*, vol. 81, no. 3, pp. 637–654, 1973.

[2] R. Merton, "The theory of rational option pricing," *Bell Journal of Economics and Management Science*, vol. 4, pp. 141–183, 1973.

[3] S. Heston, "A closed-form solution for options with stochastic volatility with applications to bond and currency options," *Review of Financial Studies*, vol. 6, pp. 327–343, 1993.

[4] P. Boyle, "Options: A Monte Carlo approach," *Journal of Financial Economics*, vol. 4, no. 3, pp. 323–338, 1977.

[5] J. Hull and A. White, "The pricing of options on assets with stochastic volatilities," *Journal of Finance*, vol. 42, no. 2, pp. 281–300, 1987.

[6] J. Gondzio, R. Kouwenbergb, and T. Vorst, "Hedging options under transaction costs and stochastic volatility," *Journal of Economic Dynamics & Control*, vol. 27, pp. 1045–1068, 2003.

[7] P. Klaassen, "Financial asset-pricing theory and stochastic programming models for asset/liability management: A synthesis," *Management Science*, vol. 44, no. 1, pp. 31–48, 1998.

[8] Y. Zhao and W. Ziemba, "A stochastic programming model using an endogenously determined worst case risk measure for dynamic asset allocation," *Mathematical Programming*, vol. 89, pp. 293–309, 2001.

[9] D. Bertsimas, L. Kogan, and A. Lo, "Hedging derivative securities and incomplete markets: an ϵ-arbitrage approach," *Operations Research*, vol. 49, no. 3, pp. 372–397, 2001.

[10] S. Fedotov, "Stochastic optimization approach to options pricing," in *Proc. American Contr. Conf.*, Chicago, IL, June 1999, pp. 1450–1453.

[11] P. Couchman, M. Cannon, and B. Kouvaritakis, "Stochastic MPC with inequality stability constraints," *Automatica*, vol. 42, no. 12, pp. 2169–2174, 2006.

[12] D. Muñoz de la Peña, A. Bemporad, and T. Alamo, "Stochastic programming applied to model predictive control," in *Proc. 44th IEEE Conf. on Decision and Control and European Control Conf.*, Sevilla, Spain, 2005, pp. 1361–1366.

[13] J. Primbs and C. Sung, "Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise," *IEEE Trans. Automatic Control*, vol. 54, no. 2, pp. 221–230, 2009.

[14] F. Herzog, S. Keel, G. Dondi, L. Schumann, and H. Geering, "Model predictive control for portfolio selection," in *Proc. American Contr. Conf.*, Minneapolis, MN, 2006, pp. 1252–1259.

[15] J. Primbs, "Dynamic hedging of basket options under proportional transaction costs using receding horizon control," *Int. J. Control*, vol. 82, no. 10, pp. 1841–1855, 2009.

[16] P. Meindl and J. Primbs, "Dynamic hedging of single and multi-dimensional options with transaction costs: A general utility maximization approach," *Quantitative Finance*, vol. 8, no. 3, pp. 299–312, Apr. 2008.

[17] A. Bemporad, L. Bellucci, and T. Gabbriellini, "Dynamic option hedging via stochastic model predictive control based on scenario simulation," 2009, submitted for journal publication.

[18] D. Bernardini and A. Bemporad, "Scenario-based model predictive control of stochastic constrained linear systems," in *Proc. 48th IEEE Conf. on Decision and Control*, 2009, to appear.

[19] F. Longstaff and E. Schwartz, "Valuing american options by simulation: A simple least-squares approach," *Review of Financial Studies*, vol. 114, no. 1, pp. 113–147, 2001.

[20] D. Freedman and P. Diaconis, "On the histogram as a density estimator: l_2 theory," *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 57, no. 4, pp. 453–476, 1981.