City Research Online

City St George's, University of London

Citation: Černý, A. (2004). Dynamic Programming and Mean-Variance Hedging in Discrete Time. Applied Mathematical Finance, 11(1), pp. 1-25. doi: 10.1080/1350486042000196164

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/16293/

Link to published version: https://doi.org/10.1080/1350486042000196164

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to City Research Online policy.

City Research Online: http://openaccess.city.ac.uk/ publications@citystgeorges.ac.uk

Dynamic Programming and Mean-Variance Hedging in Discrete Time *

Aleš Černý (a.cerny@imperial.ac.uk)
The Business School, Imperial College London

Forthcoming in Applied Mathematical Finance
First draft: April 1999, this version October 2003
typo in Theorem 3 (time subscripts in expression for ε_t) corrected 15/10/2004

Abstract. In this paper we solve the general discrete time mean-variance hedging problem by dynamic programming. Thanks to its simple recursive structure our solution is well suited for computer implementation. On the theoretical side, we show how the variance-optimal measure arises in our dynamic programming solution and how one can define conditional expectations under this (generally non-equivalent) measure. We are then able to relate our result to the results of previous studies in continuous time, namely Rheinländer and Schweizer (1997), Gourieroux et al. (1998), and Laurent and Pham (1999).

Keywords: mean-variance hedging, discrete time, dynamic programming, incomplete market, arbitrage

JEL classification code: G11, C61
Mathematics subject classification: 90A09, 90C39

1. Introduction

This paper gives a dynamic programming solution to the general discrete time mean-variance hedging problem, a solution which from the practical point of view is well suited for computer implementation thanks to its recursive structure. We show how the optimal strategy hedging is implemented on a spreadsheet for the case with leptokurtic IID stock returns. On the theoretical side, we show how the variance-optimal measure arises in the dynamic programming solution and how one defines conditional expectations under this (generally non-equivalent) measure. We are then able to relate our result to the results of previous studies in continuous time.

The mean-variance hedging in continuous time has been tackled via Galtchouk-Kunita-Watanabe decomposition under a suitable, so-called variance-optimal, martingale measure. The problem was first formulated in Duffie and Richardson (1991), Schweizer (1992) obtained

* I wish to thank Martin Schweizer and three anonymous referees for valuable comments that led to improvements in the paper. I am solely responsible for any errors in the manuscript.

the first ground breaking result under the assumption of a so-called *constant investment opportunity set.* This special case has the property that the variance-optimal measure coincides with the so-called minimal martingale measure of Föllmer and Schweizer (1991). A fully general solution was finally obtained by Rheinländer and Schweizer (1997) and Gourieroux et al. (1998), the latter using an elegant numeraire method. Laurent and Pham (1999) used the framework of Gourieroux et al. (1998) coupled with duality theory and dynamic programming to calculate explicit characterization of the variance-optimal measure in stochastic volatility models.

Studies of mean-variance hedging in discrete time are relatively few. Schäl (1994) applies dynamic programming in the case of constant investment opportunity set to examine various intertemporal mean-variance criteria. Schweizer (1995) solves the general problem with one asset and non-stochastic interest rate. This solution, however, does not have fully recursive structure. Namely, it requires calculation of two processes $(\beta_t)_{t=0,1,\ldots,T-1}$ and $(\rho_t)_{t=0,1,\ldots,T-1}$ as conditional expectations of T-measurable variables at every node of the state space, which is computationally inefficient – in a recombining trinomial tree it requires in the order of 3^{T-t} operations at every node at time t and for all t. In the same situation, a fully recursive dynamic programming solution requires only 3 operations at every node and at all times. Such solution has been derived, independently of our work, by Bertsimas et al. (2001) for one basis asset and non-stochastic interest rate[1].

The present paper can be seen as an extension of Schäl (1994) to the case of non-constant investment opportunity set and several risky assets. The contribution of our paper is threefold. Firstly, unlike Schweizer (1995) and Bertsimas et al. (2001) we solve the hedging problem with stochastic interest rate (and an arbitrary number of basis assets). Secondly, we give a simple recursive solution which in Markov setting improves greatly on the computational efficiency compared to the result of Schweizer (1995). Last but not least, by suitably defining the conditional expectation under the non-equivalent variance-optimal measure we are able to link our discrete time results to the continuous-time results of Rheinländer and Schweizer (1997), Gourieroux, Laurent, and Pham (1998), and Laurent and Pham (1999).

[1] Bertsimas et al. (2001) solve the hedging problem with $r_t = 0$, claiming that this entails no loss of generality. What they mean is that with stochastic interest rate they are able to minimize $\mathrm{E}_0\left[\left(V_T^{x,\theta} - H_T\right)^2 / \left(S_T^0\right)^2\right]$. Of course, the economically interesting problem is that of minimizing $\mathrm{E}_0\left[\left(V_T^{x,\theta} - H_T\right)^2\right]$ to which Bertsimas et al. (2001) cannot provide an answer within their setup when interest rates are stochastic.

The paper is organized as follows: In the first section we present the main result. The second section gives an explicit example with fat-tailed return distribution and weekly rebalancing period. In the third section we show how the variance-optimal measure arises in the dynamic programming solution, how one defines conditional expectations under this measure and how the existence of the variance-optimal measure is related to the no-arbitrage assumption. The final section relates our result to the results of previous studies in discrete and continuous time.

1.1. Notation

Let us have a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_t\}_{t\in\mathcal{T}})$ with E_t^P denoting the expectation conditional on the information at time t

$$\begin{aligned} \mathrm{E}_t^P[X] &\equiv \mathrm{E}^P[X|\mathcal{F}_t], \\ \mathcal{T} &= \{0,1,\ldots,T\}. \end{aligned}$$

To keep the technicalities at minimum we will assume that Ω is finite[2], that is the information structure can be represented by a tree. The conditional expectation $\mathrm{E}_t^P[X]$ assigns one value to each node at time t in the information tree.

All processes defined in the next section, except for the cumulative discount S^0, are adapted. For a process $\{X_t\}_{t\in\mathcal{T}}$ being adapted means that the value of X_t is known at time t but that generally the value of X_t is uncertain as of $t-1$. The process S^0 is predictable, that is S_t^0 is known already at time $t-1$.

Measurability is important for manipulation of conditional expectations, for example we will frequently use the fact that

$$\mathrm{E}_t^P[X_t Y] = X_t \mathrm{E}_t^P[Y],$$

and vice versa. The only other tool we need is the law of iterated expectations

$$\mathrm{E}_t^P\left[\mathrm{E}_s^P[X]\right] = \mathrm{E}_t^P[X],$$

whenever $s \geq t$. Matrix transpose is denoted by asterisk*.

2. Mean-square hedging in discrete time

Suppose that there are n basis assets with $\mathbb{R}^n$-valued price process $S = (S_t)_{t=0,\ldots,T}$ and dividend process $\delta = (\delta_t)_{t=1,\ldots,T}$. Assume further

[2] Interested reader should consult Schweizer (1996) for questions of integrability.

that there is a risk-free short term borrowing and lending with return $R_{\mathrm{f}\,t}$ and define the cumulative discount

$$\begin{aligned} S^0_t &= R_{\mathrm{f}\,0} \times R_{\mathrm{f}\,1} \times \ldots \times R_{\mathrm{f}\,t-1}, \\ S^0_0 &= 1. \end{aligned}$$

Let us define the discounted gain process of the basis assets X

$$\begin{aligned} X_t &\equiv \frac{S_t}{S^0_t} + \sum_{i=1}^{t} \frac{\delta_i}{S^0_i}, \\ \Delta X_t &\equiv X_t - X_{t-1}. \end{aligned}$$

Process $V^{x,\theta}$ denotes the wealth obtained by self-financing strategy with initial investment x and with shares of risky investment given by the portfolio process $\theta = (\theta_t)_{t=0,\ldots,T-1}$.

$$V^{x,\theta}_t = R_{\mathrm{f}\,t-1} V^{x,\theta}_{t-1} + S^0_t \theta^*_{t-1} \Delta X_t \tag{1}$$

$$\Delta\left(\frac{V^{x,\theta}}{S^0}\right)_{t-1} = \theta^*_{t-1} \Delta X_t \tag{2}$$

$$\frac{V^{x,\theta}_t}{S^0_t} = x + \sum_{i=0}^{t-1} \theta^*_i \Delta X_{i+1}.$$

In what follows it is important to capture the fact that some processes are unaffected by the choice of the trading strategy.

DEFINITION 1. *A random variable $X : \Omega \to \mathbb{R}$ is called exogenous if for every fixed $\omega \in \Omega$ the value $X(\omega)$ does not depend on the choice of $V^{x,\theta}_0(\omega), \theta_0(\omega), \ldots, \theta_{T-1}(\omega)$.*

We assume that processes X and S^0 are exogenous; this is a standard assumption in finance literature, its relaxation has been studied prominently by Frey and Stremme (1997); see also Černý (1999). Given an exogenous and $\mathcal{F}_T$-measurable payoff H_T the best mean-square hedge for H_T is given by the initial wealth x and portfolio weights θ which are found by minimizing the expected square replication error $\mathrm{E}^P_0 \left[V^{x,\theta}_T - H_T\right]^2$

$$\min_{x,\theta_0,\ldots,\theta_{T-1}} \mathrm{E}^P_0 \left[\left(S^0_T \left(x + \sum_{i=0}^{T-1} \theta^*_i \Delta X_{i+1} \right) - H_T \right)^2 \right]$$

$$x \text{ is } \mathcal{F}_0\text{-measurable}, \ \theta_i \text{ is } \mathcal{F}_i\text{-measurable}, \ i = 0, 1, \ldots, T-1.$$

THEOREM 2 (Optimal controls). *Let k_t and H_t be $\mathcal{F}_t$-measurable and exogenous, with $k_t > 0$ a.s. Assume further that there is no arbitrage among the basis assets, which in particular implies $R_{f\,t-1} > 0$, and that the excess returns of basis assets are linearly independent.*

Then the matrix $\mathrm{E}^P_{t-1}\left[k_t \Delta X_t \Delta X_t^*\right]$ is invertible at each node of the information set $\mathcal{F}_{t-1}$ and the problem

$$\min_{x,\theta_0,\ldots,\theta_{t-1}} \mathrm{E}^P_0\left[k_t\left(V_t^{x,\theta} - H_t\right)^2\right] \tag{3}$$

has the same optimal controls $x, \theta_0, \ldots, \theta_{t-2}$ as the problem

$$\min_{x,\theta_0,\ldots,\theta_{t-2}} \mathrm{E}^P_0\left[k_{t-1}\left(V_{t-1}^{x,\theta} - H_{t-1}\right)^2\right] \tag{4}$$

with $k_{t-1} > 0$ and H_{t-1} being $\mathcal{F}_{t-1}$-measurable and exogenous

$$\frac{k_{t-1}}{R^2_{f\,t-1}} = \mathrm{E}^P_{t-1}\left[k_t\right] - \mathrm{E}^P_{t-1}\left[k_t \Delta X_t^*\right]\left(\mathrm{E}^P_{t-1}\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1} \mathrm{E}^P_{t-1}\left[k_t \Delta X_t\right], \tag{5}$$

$$H_{t-1} = \frac{\mathrm{E}^P_{t-1}\left[\left(k_t - \mathrm{E}^P_{t-1}\left[k_t \Delta X_t^*\right]\left(\mathrm{E}^P_{t-1}\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1} k_t \Delta X_t\right)\frac{H_t}{R_{f\,t-1}}\right]}{\frac{k_{t-1}}{R^2_{f\,t-1}}}, \tag{6}$$

and the dynamically optimal value of θ_{t-1} is given as

$$\theta^D_{t-1} = -\left(\mathrm{E}^P_{t-1}\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1} \mathrm{E}^P_{t-1}\left[k_t \Delta X_t\left(\frac{V^{x,\hat{\theta}}_{t-1}}{S^0_{t-1}} - \frac{H_t}{S^0_t}\right)\right]. \tag{7}$$

Proof. See Appendix A. ■

THEOREM 3 (Value function). *Let $k_T > 0$ and H_T be $\mathcal{F}_T$-measurable random variables. Then*

$$\min_{\theta_t,\ldots,\theta_{T-1}} \mathrm{E}^P_t\left[k_T\left(V_T^{x,\theta} - H_T\right)^2\right] = k_t\left(V_t^{x,\theta} - H_t\right)^2 + \varepsilon_t^2$$

where $\{\varepsilon\}$ is a well-defined $\mathcal{F}$-adapted exogenous process satisfying

$$\begin{aligned}
\varepsilon_t^2 &= \mathrm{E}^P_t\left[\varepsilon^2_{t+1}\right] + \mathrm{E}^P_t\left[k_{t+1}\left(R_{f\,t}H_t + S^0_{t+1}\left(\bar{\theta}^D_t\right)^* \Delta X_{t+1} - H_{t+1}\right)^2\right], \\
\varepsilon_T^2 &= 0,
\end{aligned}$$

where $\bar{\theta}^D_{t-1}$ is obtained from (7) by substituting H_{t-1} for V^{x,θ^D}_{t-1},that is,

$$\bar{\theta}^D_{t-1} = -\left(\mathrm{E}^P_{t-1}\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1} \mathrm{E}^P_{t-1}\left[k_t \Delta X_t\left(\frac{H_{t-1}}{S^0_{t-1}} - \frac{H_t}{S^0_t}\right)\right]. \tag{8}$$

Proof. See Appendix A. ■

COROLLARY 4. *Theorems 2 and 3 provide complete characterization of the solution to the mean-variance hedging problem. A repeated application of the theorems starting from T with $k_T = 1$ gives us all values of k_t and H_t for $0 \leq t \leq T$. Further, at the end of the backward run we learn that the optimal value of initial wealth is $\hat{x} = H_0$ and that the expected squared hedging error is ε_0^2. In a forward run from time 0 we can then recover the optimal portfolio and optimal hedging wealth from (7) and (1). If X_t is Markov, the interest rate is non-stochastic and H_T is a European contingent claim, $H_T = H(X_T)$, then it follows from (5) and (6) that the processes k_t, H_t, $\bar{\theta}_{t-1}^D$ and ε_t^2 depend only on one state variable, X_t. The optimal wealth $V_t^{H_0,\theta^D}$, however, is - save for very special cases - path dependent, see Schweizer (1998).*

3. Discrete Time Hedging with IID Returns

In this section we illustrate how the dynamic programming solution can be usefully applied in practice. We will consider trading once a week for 6 weeks. We assume that weekly log-returns are spaced regularly with 2% gap, see Figure 1.

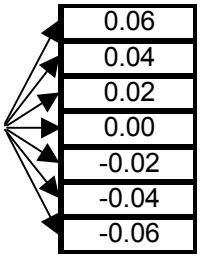

Figure 1. Logarithm of weekly stock returns.

The conditional objective probabilities of movement in the lattice are calibrated from historical data, based on the FTSE 100 weekly returns in the period 1984-2001, see Figure 2.

Obviously, in a more practically-minded application one will trade more frequently, the spacing of log-returns will be finer and the returns will be more leptokurtic. All these concerns are easily accommodated in the present framework. Crucially, one may want to abandon the assumption of IID returns to allow for periods of high and low volatility, bull/bear markets etc. These features can be easily handled within the recursive solution of Theorems 2 and 3, if one is willing to accept a

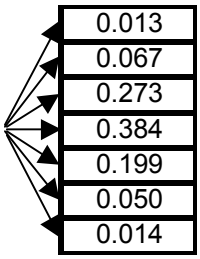

Figure 2. Conditional objective probabilities of stock price movement.

larger number of state variables. Such extensions, however, are beyond the scope of this illustrative section.

Our aim is to hedge a European call option with 6 weeks to expiry, rehedging once a week. We will assume that the initial value of the index is $S_0 = 100$ and that the option is struck at the money $K = 100$. The resulting stock price lattice is depicted in Figure 3.

3.1. Mean value process H

This process could be computed directly from equation (6), but in an IID model it is most conveniently constructed with the help of special risk-neutral probabilities called *variance-optimal* probabilities, see equation (18). The variance-optimal probabilities in turn are computed from the distribution of excess return. The variance-optimal measure will be denoted $\tilde{P}$ to distinguish it from the objective probability measure P. The corresponding change of measure is given by the formula

$$\frac{\mathrm{d}\tilde{P}}{\mathrm{d}P} = m_{1|0} m_{2|1} \dots m_{T|T-1} \tag{9}$$

$$m_{t+1|t} \triangleq \frac{q_{t+1|t}}{p_{t+1|t}} = \frac{1 - a\left(R_{t+1} - R_{\mathrm{f}}\right)}{b} \tag{10}$$

$$a = \frac{\mathrm{E}_t^P\left[R_{t+1} - R_{\mathrm{f}}\right]}{\mathrm{E}_t^P\left[\left(R_{t+1} - R_{\mathrm{f}}\right)^2\right]} \tag{11}$$

$$b = 1 - \frac{\left(\mathrm{E}_t^P\left[R_{t+1} - R_{\mathrm{f}}\right]\right)^2}{\mathrm{E}_t^P\left[\left(R_{t+1} - R_{\mathrm{f}}\right)^2\right]}. \tag{12}$$

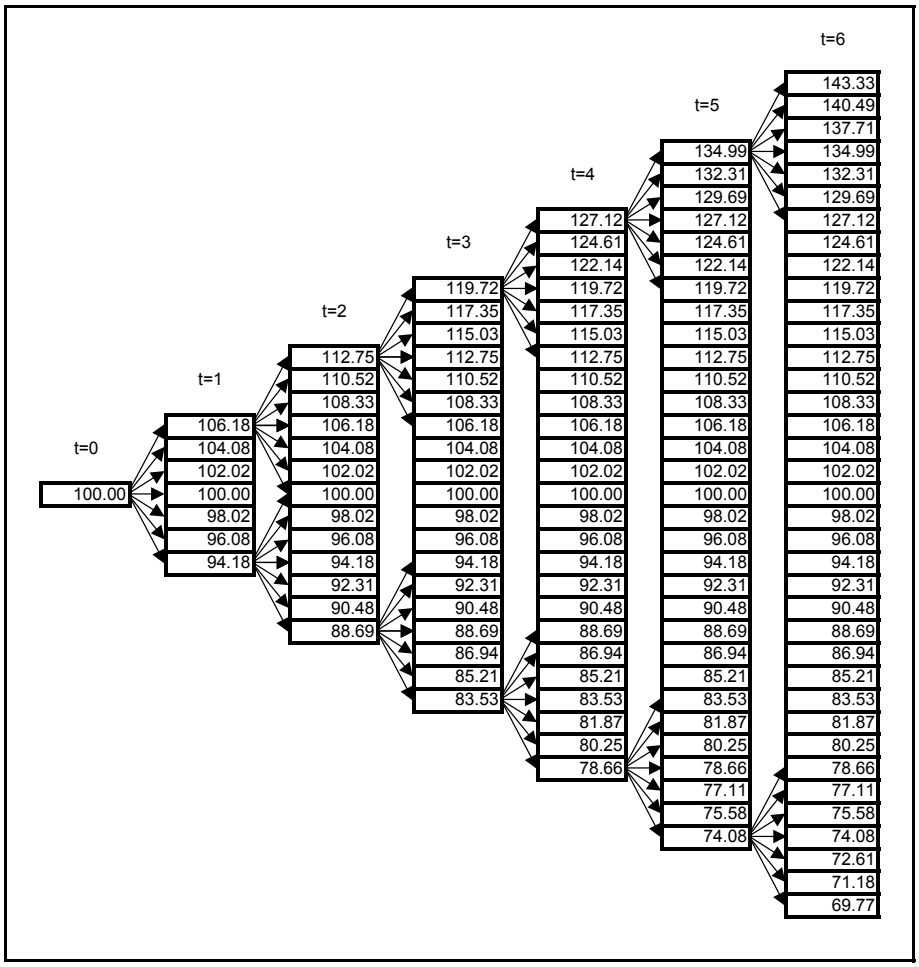

Figure 3. Stock price lattice.

Numerically,

$$R_{t+1} = \left[\, e^{0.06} \;\; e^{0.04} \;\; e^{0.02} \;\; e^{0.00} \;\; e^{-0.02} \;\; e^{-0.04} \;\; e^{-0.06} \,\right],$$

$$R_{\mathrm{f}} = 1.00075,$$

$$R_{t+1} - R_{\mathrm{f}}$$

$$= \left[\, 6.108 \;\; 4.006 \;\; 1.945 \;\; -0.075 \;\; -2.056 \;\; -3.997 \;\; -5.899 \,\right] \times 10^{-2},$$

$$\mathrm{E}_t^P \left[R_{t+1} - R_{\mathrm{f}}\right] = 1.58 \times 10^{-3}, \quad \mathrm{E}_t^P \left[(R_{t+1} - R_{\mathrm{f}})^2\right] = 4.72 \times 10^{-4},$$

$$a = \frac{1.58 \times 10^{-3}}{4.72 \times 10^{-4}} = 3.35$$

$$b = 1 - \frac{1.58^2 \times 10^{-6}}{4.72 \times 10^{-4}} = 0.9947$$

$$m_{t+1|t} = [\, 0.7995 \;\; 0.8704 \;\; 0.9398 \;\; 1.0079 \;\; 1.0746 \;\; 1.1400 \;\; 1.2041 \,]$$

$$\begin{aligned} q_{t+1|t} &= m_{t+1|t} p_{t+1|t} \\ &= [\, 0.010 \;\; 0.058 \;\; 0.257 \;\; 0.387 \;\; 0.214 \;\; 0.057 \;\; 0.017 \,] \end{aligned}$$

The risk-neutral probabilities q and the option payoff H_T define the mean value process $\{H_t\}_{t=0,1,\ldots,T}$ as follows

$$H_t = \mathrm{E}_t^{\tilde{P}} \left[\frac{H_T}{R_{\mathrm{f}}^{T-t}} \right]$$

In our special case with IID returns and deterministic interest rate the conditional variance-optimal probabilities $q_{t+1|t}$ coincide with the risk-neutral probabilities of one-period Markowitz CAPM model. Thus H_{T-1} is the CAPM price of the option at time $T-1$, H_{T-2} is the CAPM price of H_{T-1} at time $T-2$ and so on.

The value of H_t is computed recursively using the risk-neutral probabilities and starting from the last period as in the complete market case

$$\begin{aligned} H_t &= \mathrm{E}_t^{\tilde{P}} \left[\frac{H_{t+1}}{R_{\mathrm{f}}} \right], \\ t &= T-1, \ldots, 0. \end{aligned} \tag{13}$$

However, formula (13) differs from its complete market counterpart in one important aspect. While in a complete market there is a self-financing portfolio with value H_t that perfectly replicates H_{t+1}, in an incomplete market such a portfolio generally does not exist.

The mean value process H_t is depicted in Figure 4, together with corresponding continuous Black–Scholes value for comparison. Consider, for example the middle node at $t = 1$. The Black–Scholes formula dictates

$$C(S,K,r,\sigma,\tau) = S\Phi\left(\frac{\ln \frac{S}{K} + (r + \frac{\sigma^2}{2})\tau}{\sigma\sqrt{\tau}} \right) - e^{-r\tau} K \Phi\left(\frac{\ln \frac{S}{K} + (r - \frac{\sigma^2}{2})\tau}{\sigma\sqrt{\tau}} \right)$$

with

$$\begin{aligned} S &= 100, \qquad K = 100, \\ r &= \ln R_{\mathrm{f}} = 7.5 \times 10^{-4}, \\ \sigma &= \sqrt{\mathrm{Var}_t^P(\ln R_{t+1})} = 2.165 \times 10^{-2}, \\ \tau &= 6 - 1 = 5, \end{aligned}$$

resulting in $C = 2.12$ as compared to $H = 2.09$.

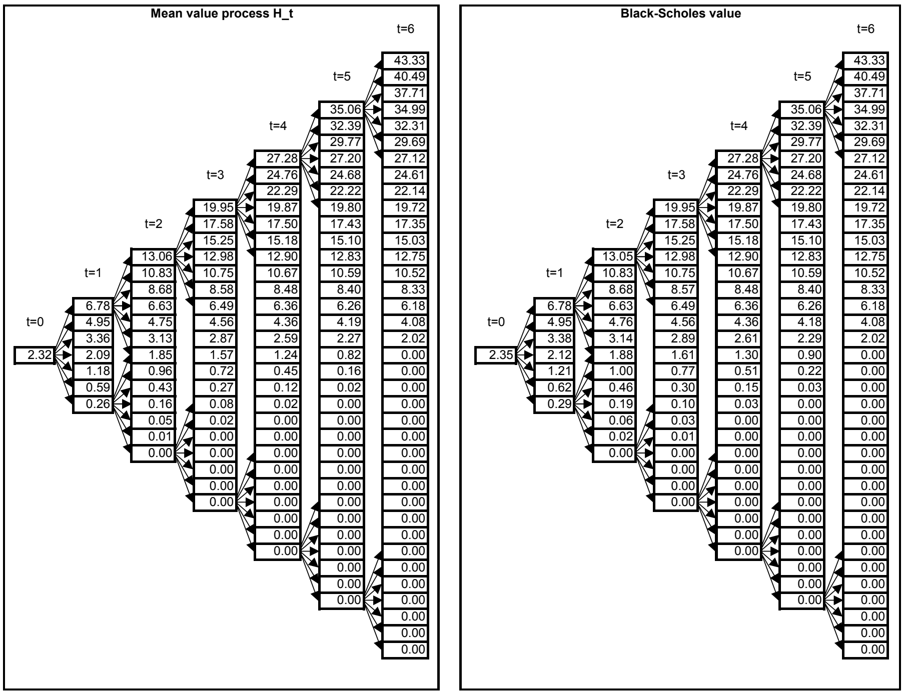

Figure 4. Comparison of mean value process H with continuous-time Black–Scholes option prices.

3.2. Black–Scholes delta and optimal hedging strategy

It transpires from the proof of Theorems 2 and 3 that the dynamically optimal hedging strategy θ_t^D is obtained from the minimization of the one-step-ahead hedging error $\mathrm{E}_t^P\left[(V_{t+1} - H_{t+1})^2\right]$. Using the self-financing condition $V_{t+1} = R_\mathrm{f} V_t + \theta_t^D S_t \left(R_{t+1} - R_\mathrm{f}\right)$ the squared error can be written as

$$\mathrm{E}_t^P\left[\left(R_\mathrm{f} V_t + \theta_t^D S_t \left(R_{t+1} - R_\mathrm{f}\right) - H_{t+1}\right)^2\right] \tag{14}$$

and it is clear that the optimal value of θ_t depends not only on H_{t+1} but also on V_t.

The nature of the self-financing portfolio means that once we arrive at time t we cannot chose V_t, it is given by our past trading strategy and realizations of stock prices. But it makes sense to inquire what value of V_t we *would* prefer *if* we had the choice. It turns out that the

optimal pair V_t, θ_t minimizing (14) is $V_t = H_t$, $\theta_t = \bar{\theta}_t^D$

$$\bar{\theta}_t^D = \frac{\mathrm{E}_t^P \left[(H_{t+1} - R_\mathrm{f} H_t) (R_{t+1} - R_\mathrm{f}) \right]}{S_t \mathrm{E}_t^P \left[(R_{t+1} - R_\mathrm{f})^2 \right]}, \tag{15}$$

where $\bar{\theta}_t^D$ is the discrete time analogy of the continuous time Black–Scholes delta. It turns out that in the IID model $\bar{\theta}_t^D$ represents so-called locally optimal hedge. The evaluation of locally optimal and dynamically optimal hedging errors appears in Černý (2002). $\bar{\theta}_t^D$ can be obtained mechanically by simplifying equation (8).

Effectively, the Black–Scholes hedging strategy assumes that the value of the hedging portfolio is always at its optimum H_t. In an incomplete market this is obviously not always the case, therefore the dynamically optimal strategy makes an adjustment for the difference between V_t and H_t

$$\theta_t^D = \bar{\theta}_t^D + R_\mathrm{f} a \frac{H_t - V_t}{S_t}.$$

The coefficient a is computed from (11), numerically $a = 3.35$. When the self-financing portfolio is above the target value H_t the delta is adjusted downwards and vice versa. Chapter 3 in Černý (2004) shows that a represents the optimal proportion of investment in the stock per unit of investor's risk tolerance.

The locally optimal delta is easily computed from formula (15), because we already know the values of H in all nodes. Figure 5 compares $\bar{\theta}_t^D$ with its continuous-time counterpart,

$$\theta_t^{\mathrm{BS}} = N \left(\frac{\ln S/K + \left(r + \frac{\sigma^2}{2} \right) (T - t)}{\sigma \sqrt{T - t}} \right). \tag{16}$$

3.3. Squared error process

The mean value process H_t represents the target value the hedging portfolio V_t is trying to achieve. It follows from Theorem 3 that $V_t = H_t$ minimizes the expected squared replication error as seen at time t, $\mathrm{E}_t^P \left[(V_T - H_T)^2 \right]$. The size of the error in the ideal case $V_t = H_t$ is measured by the *squared error process* ε_t^2 and is computed recursively as follows

$$\begin{aligned} \varepsilon_t^2 &= \mathrm{E}_t^P \left[\varepsilon_{t+1}^2 \right] + k_{t+1} \mathrm{ESRE}_t^P (H_{t+1}) \\ \varepsilon_T &= 0 \end{aligned} \tag{17}$$

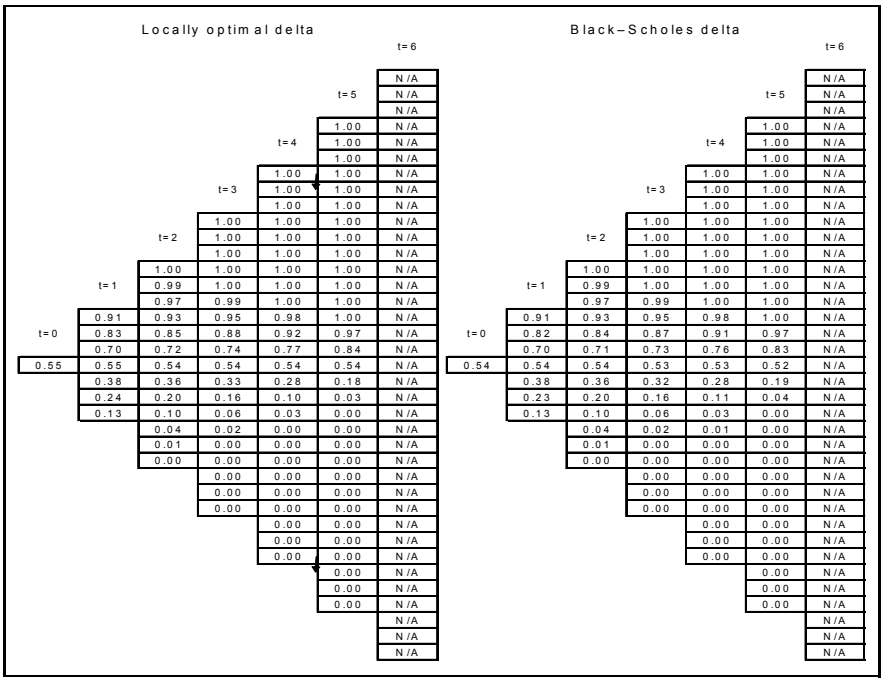

Figure 5. Comparison between the locally optimal delta $\bar{\theta}^D$ from equation (15) and the continuous-time Black–Scholes delta in (16).

The term $\mathrm{ESRE}_t^P(H_{t+1})$ is the *one-period* Expected Squared Replication Error from hedging payoff H_{t+1} using the risk-free bank account and the stock as basis assets

$$\mathrm{ESRE}_t^P(H_{t+1}) = \mathrm{E}_t^P\left[\left(R_\mathrm{f}H_t + \bar{\theta}_t^D S_t\left(R_{t+1} - R_\mathrm{f}\right) - H_{t+1}\right)^2\right].$$

In our model with IID returns and non-stochastic interest rate process k becomes deterministic

$$k_t = R_\mathrm{f}^{2(T-t)} b^{T-t} = 0.9962^{T-t}.$$

Recall from (12) that $b = 1 - \frac{\left(\mathrm{E}_t^P[R_{t+1}-R_\mathrm{f}]\right)^2}{\mathrm{E}_t^P\left[(R_{t+1}-R_\mathrm{f})^2\right]}$. An econometrician would interpret $\frac{\left(\mathrm{E}_t^P[R_{t+1}-R_\mathrm{f}]\right)^2}{\mathrm{E}_t^P\left[(R_{t+1}-R_\mathrm{f})^2\right]}$ as the non-central R^2 from the regression of the risk-free rate onto the excess return. Naturally, the excess return performs very poorly in explaining the variation in the risk-free rate, consequently the R^2 will be small and b will be very close to 1, which is one reason why the expected squared error of the dynamically optimal hedge is only marginally smaller than the expected squared error of the locally optimal hedge, see Černý (2002). The coefficient b can also be

interpreted in terms of the market Sharpe Ratio SR,

$$b = (1 + \mathrm{SR}^2)^{-1};$$

the quantity $1 - \sqrt{b} \approx \frac{1}{2}\mathrm{SR}^2$ measures the percentage increase in investor's certainty equivalent wealth per unit of risk tolerance, see Chapter 3 in Černý (2004).

Figure 6 depicts the one-period expected squared hedging errors. Intuitively we know that options far in the money and far out of the money can be replicated perfectly. Indeed, we observe that the replication error is the largest at the money and that it goes down to zero for very high and very low stock prices. There is a simple relationship between the one-step ahead expected squared hedging error and the option *gamma.* This link is best seen by comparing $\mathrm{ESRE}_t^P(H_{t+1})$ with formula (18) in Toft (1996) where transaction costs are ignored.

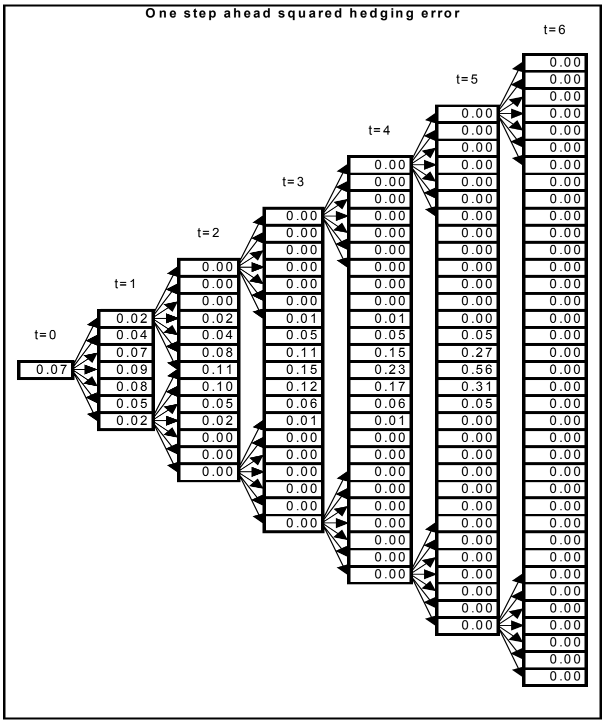

Figure 6. One-period expected squared hedging errors.

The total expected squared replication error ε_t^2 can be computed recursively from (14). Numerical values are shown in Figure 7.

To summarize, if one sells one option at $H_0 = 2.32$ and hedges it optimally to maturity then one has entered a risky position with expected payoff zero and standard deviation of $\varepsilon_0 = \sqrt{0.6} = 0.77$.

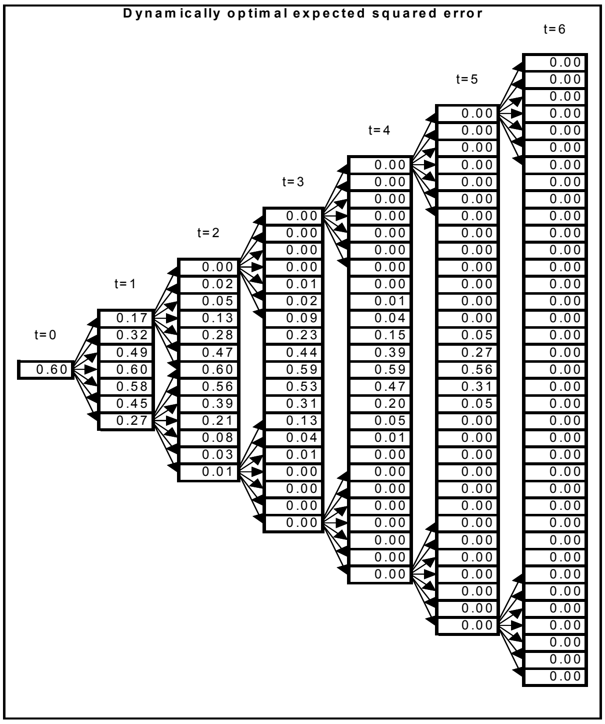

Figure 7. The smallest possible expected squared hedging error to maturity, ε_t^2, corresponding to a perfectly balanced hedging portfolio, $V_t = H_t$.

4. The variance-optimal measure

Theorem 2 does not use any duality theory and the only technical difficulty is to find the self-preserving recursive structure (3), (4). Nevertheless, a martingale measure, known as the *variance-optimal measure*, emerges from the solution in equation (6).

Let us define

$$\tilde{m}_{t|t-1} \equiv R_{\mathrm{f}\,t-1}^2 \frac{k_t - \mathrm{E}_{t-1}^P\left[k_t \Delta X_t^*\right]\left(\mathrm{E}_{t-1}^P\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1} k_t \Delta X_t}{k_{t-1}} \tag{18}$$

We notice that $\mathrm{E}_{t-1}^P\left[\tilde{m}_{t|t-1}\right] = 1$ by virtue of equation (5), and moreover $\tilde{m}_{t|t-1}$ is $\mathcal{F}_t$-measurable. Therefore $\tilde{m}_{t|t-1}$ can be thought of as a one-step conditional change of measure. The problem is that $\tilde{m}_{t|t-1}$ is *not* necessarily *strictly positive*, and thus one has to be careful when defining conditional expectations.

We can define the density process

$$\tilde{m}_t = \tilde{m}_{t|t-1} \ldots \tilde{m}_{1|0} \tag{19}$$

and write

$$\frac{\mathrm{d}\tilde{P}}{\mathrm{d}P} = \tilde{m}_T.$$

Since $\tilde{m}_{t|t-1}$can be equal to 0 with positive probability, the ratio $\frac{\tilde{m}_T}{\tilde{m}_t}$ is not well defined. But if we formally agree that

$$\frac{\tilde{m}_T}{\tilde{m}_t} \equiv \tilde{m}_{T|T-1} \dots \tilde{m}_{t+1|t}$$

then we can write

$$\mathrm{E}_t^{\tilde{P}}[X] = \frac{\mathrm{E}_t^P[\tilde{m}_T X]}{\tilde{m}_t} \equiv \mathrm{E}_t^P\left[\tilde{m}_{T|T-1} \dots \tilde{m}_{t+1|t} X\right] \tag{20}$$

as if $\tilde{P}$ were equivalent to P. Note that by Theorem 1 each of the variables $\tilde{m}_{t+1|t}$ is well defined. Thus although $\tilde{P}$ is a signed measure and generally not equivalent to P, equation (20) grants a well defined conditional expectation under $\tilde{P}$.

Another interesting property of $\tilde{P}$ is that it turns the discounted gain process into a martingale

$$\begin{aligned}\mathrm{E}_{t-1}^{\tilde{P}}[\Delta X_t^*] &\equiv \mathrm{E}_{t-1}^P\left[m_{t|t-1}\Delta X_t^*\right] = R^2_{\mathrm{f}\,t-1} \\ \times \frac{\mathrm{E}_{t-1}^P[k_t \Delta X_t^*] - \mathrm{E}_{t-1}^P[k_t \Delta X_t^*]\left(\mathrm{E}_{t-1}^P[k_t \Delta X_t \Delta X_t^*]\right)^{-1}\mathrm{E}_{t-1}^P[k_t \Delta X_t \Delta X_t^*]}{k_{t-1}} \\ &= 0.\end{aligned}$$

Finally, recursive application of equation (6) together with the law of iterated expectations imply

$$\begin{aligned}\frac{H_t}{S_t^0} &= \mathrm{E}_t^P\left[\tilde{m}_{T|T-1} \dots \tilde{m}_{t+1|t}\frac{H_T}{S_T^0}\right] = \mathrm{E}_t^{\tilde{P}}\left[\frac{H_T}{S_T^0}\right] \\ \hat{x} &= H_0 = \mathrm{E}_0^{\tilde{P}}\left[\frac{H_T}{S_T^0}\right].\end{aligned} \tag{21}$$

4.1. Special case: minimal martingale measure

When $k_t = \mathrm{E}_t^P\left[R^2_{\mathrm{f}\,t} \dots R^2_{\mathrm{f}\,T-1}\right]$ we obtain one-step conditional change of measure of the so-called minimal martingale measure. When the risk-free rate is non-stochastic the minimal change of measure becomes

$$\hat{m}_{t|t-1} = \frac{1 - \mathrm{E}_{t-1}^P[\Delta X_t^*]\left(\mathrm{E}_{t-1}^P[\Delta X_t \Delta X_t^*]\right)^{-1}\Delta X_t}{1 - \mathrm{E}_{t-1}^P[\Delta X_t^*]\left(\mathrm{E}_{t-1}^P[\Delta X_t \Delta X_t^*]\right)^{-1}\mathrm{E}_{t-1}^P[\Delta X_t]}. \tag{22}$$

4.2. Variance-optimal measure and arbitrage

With $\tilde{P}$ being a martingale measure the equation (21) looks like a no-arbitrage pricing formula, and therefore H_0 is referred to by Schäl (1994) as the *fair hedging price* of the contingent claim H_T, while Schweizer (1995) calls H_0 *approximation price*. Schweizer correctly acknowledges, however, that the term *price* is misleading since H_0 is not necessarily a no-arbitrage price of H_T.

A related point is that the *variance-optimal measure exists* even *when there is arbitrage* among the basis assets. Take a simple binomial tree example with $R_\mathrm{f} = 1$, no dividends, where $\frac{S_{t+1}}{S_t}$ can only take two values, +4 and +2 with equal probability $\frac{1}{2}$. Since in both states the risky return is greater than the risk-free rate there is arbitrage. The model is actually complete, so there is only one candidate for the variance-optimal 'probabilities'

$$\begin{aligned} 4q_\mathrm{up} + 2q_\mathrm{down} &= 1 \\ q_\mathrm{up} + q_\mathrm{down} &= 1 \\ q_\mathrm{up} &= -\frac{1}{2} \\ q_\mathrm{down} &= \frac{3}{2}. \end{aligned}$$

One can show that the hedging procedure described in Theorem 2 will go through fine, with $k_t = 0.1^{T-t}$. This means that *mean-variance hedging is not arbitrage-proof*, in the sense that the algorithm described in Corollary 4 will not automatically come to a halt if there is arbitrage among the basis assets.

As an aside we highlight the result of Schachermayer (1993) who showed that the minimal measure $\hat{P}$ (which in the above example is equal to the variance-optimal measure $\tilde{P}$) may not exist even when there is no arbitrage. This would seem to contradict our Theorem 1, however, Schachermayer's result has to do with integrability and we assumed away all integrability problems by taking Ω finite.

5. Comparison with previous studies

5.1. Definitions of variance-optimal measure

It follows from Section 4 that the natural definition of the variance-optimal measure $\tilde{P}$ is

$$\frac{\mathrm{d}\tilde{P}}{\mathrm{d}P} = \tilde{m}_T = \tilde{m}_{T|T-1} \dots \tilde{m}_{1|0}. \tag{23}$$

where $\tilde{m}_{t|t-1}$ was given in equation (18). Let us see how this definition relates to that of other authors, namely Schweizer (1995), Gourieroux et al. (1998), and Laurent and Pham (1999).

To simplify expression (23) let us define

$$a_{t-1} \equiv \left(\mathrm{E}_{t-1}^{P}\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1} \mathrm{E}_{t-1}^{P}\left[k_t \Delta X_t\right], \tag{24}$$

whereby the expression for the conditional change of measure becomes

$$\tilde{m}_{t|t-1} = \frac{R_{\mathrm{f}\,t-1}^2 k_t (1 - a_{t-1}^* \Delta X_t)}{k_{t-1}}. \tag{25}$$

For the unconditional change of measure we then have

$$\tilde{m}_T = \prod_{t=1}^{T} \frac{R_{\mathrm{f}\,t-1}^2 k_t (1 - a_{t-1}^* \Delta X_t)}{k_{t-1}} = \frac{\left(S_T^0\right)^2}{k_0} \prod_{t=1}^{T} (1 - a_{t-1}^* \Delta X_t). \tag{26}$$

Since by construction $\mathrm{E}_0^P\left[\tilde{m}_T\right] = 1$, an immediate consequence of the above is

$$k_0 = \mathrm{E}_0^P\left[\left(S_T^0\right)^2 \prod_{t=1}^{T} (1 - a_{t-1}^* \Delta X_t)\right]. \tag{27}$$

Similarly, the fact that

$$\mathrm{E}_t^P\left[\tilde{m}_{t+1|t}\tilde{m}_{t+2|t+1}\ldots\tilde{m}_{T|T-1}\right] =$$

$$= \mathrm{E}_t^P\left[\tilde{m}_{t+1|t}\underbrace{\mathrm{E}_{t+1}^P\left[\tilde{m}_{t+2|t+1}\ldots\underbrace{\mathrm{E}_{T-1}^P\left[\tilde{m}_{T|T-1}\right]}_{=1}\right]}_{=1}\right] = 1$$

implies

$$k_t = \mathrm{E}_t^P\left[\left(\frac{S_T^0}{S_t^0}\right)^2 \prod_{j=t+1}^{T} (1 - a_{j-1}^* \Delta X_j)\right]. \tag{28}$$

Plugging this result back into equation (24) we have

$$
\begin{aligned}
a_{t-1} = & \left(\mathrm{E}_{t-1}^{P} \left[\Delta X_t \Delta X_t^* \mathrm{E}_t^{P} \left[\left(\frac{S_T^0}{S_t^0} \right)^2 \prod_{j=t+1}^{T} (1 - a_{j-1}^* \Delta X_j) \right] \right] \right)^{-1} \\
& \times \mathrm{E}_{t-1}^{P} \left[\Delta X_t \mathrm{E}_t^{P} \left[\left(\frac{S_T^0}{S_t^0} \right)^2 \prod_{j=t+1}^{T} (1 - a_{j-1}^* \Delta X_j) \right] \right] \\
= & \left(\mathrm{E}_{t-1}^{P} \left[\Delta X_t \Delta X_t^* \left(\frac{S_T^0}{S_t^0} \right)^2 \prod_{j=t+1}^{T} (1 - a_{j-1}^* \Delta X_j) \right] \right)^{-1} \\
& \times \mathrm{E}_{t-1}^{P} \left[\Delta X_t \left(\frac{S_T^0}{S_t^0} \right)^2 \prod_{j=t+1}^{T} (1 - a_{j-1}^* \Delta X_j) \right]. \qquad (29)
\end{aligned}
$$

The above expression implies that for zero interest rate the process a coincides with the adjustment process β of Schweizer (1995). Furthermore, Schweizer defines the variance-optimal measure $\tilde{P}$ as

$$
\frac{\mathrm{d}\tilde{P}}{\mathrm{d}P} = \frac{\prod_{t=1}^{T} (1 - a_{t-1}^* \Delta X_t)}{\mathrm{E}_0^{P} \left[\prod_{t=1}^{T} (1 - a_{t-1}^* \Delta X_t) \right]},
$$

and this definition by virtue of (27) coincides with our definition of the variance-optimal measure (26) when the interest rate is zero.

An alternative interpretation of formula (26) comes from the feedback form (7) when we set $H_T = 0$ and keep $x = 1$:

$$
\tilde{\theta}_{t-1} = -\frac{V_{t-1}^{1,\tilde{\theta}}}{S_{t-1}^0} a_{t-1}, \qquad (30)
$$

whereby it becomes clear that $a_t = \tilde{a}_t$ in Laurent and Pham (1999). Here $V_{t-1}^{1,\tilde{\theta}}$ is interpreted as the optimal wealth from the self-financing strategy that minimizes $\mathrm{E}_0^{P} \left[V_T^{x,\theta} \right]^2$ with initial wealth $x = 1$, $\tilde{\theta}$ being the corresponding optimal investment strategy. From (1) we can write

$$
\begin{aligned}
V_t^{1,\tilde{\theta}} &= R_{\mathrm{f}\,t-1} V_{t-1}^{1,\tilde{\theta}} + S_t^0 \tilde{\theta}_{t-1} \Delta X_t \\
&= R_{\mathrm{f}\,t-1} V_{t-1}^{1,\tilde{\theta}} \left(1 - a_{t-1}^* \Delta X_t \right)
\end{aligned}
$$

and hence

$$
\begin{aligned}
V_T^{1,\tilde{\theta}} &= \prod_{t=1}^{T} R_{\mathrm{f}\,t-1} (1 - a_{t-1}^* \Delta X_t) = \qquad (31) \\
&= V_t^{1,\tilde{\theta}} \frac{S_T^0}{S_t^0} \prod_{j=t+1}^{T} (1 - a_{j-1}^* \Delta X_j). \qquad (32)
\end{aligned}
$$

Therefore, from (31) and (26) the variance-optimal change of measure can be expressed equivalently as

$$\tilde{m}_T = \frac{S_T^0 V_T^{1,\tilde{\theta}}}{k_0} = \frac{S_T^0 V_T^{1,\tilde{\theta}}}{\mathrm{E}_0^P\left[S_T^0 V_T^{1,\tilde{\theta}}\right]}$$

as in Gourieroux et al. (1998), equation (4.4).

One can use (32) to write (28) as

$$k_t = \frac{\mathrm{E}_t^P\left[V_T^{1,\tilde{\theta}} S_T^0\right]}{V_t^{1,\tilde{\theta}} S_t^0}.$$

On the other hand equation (4.7) in Laurent and Pham (1999) implies

$$j_t = \frac{V_t^{1,\tilde{\theta}} S_t^0}{\mathrm{E}_t^P\left[V_T^{1,\tilde{\theta}} S_T^0\right]}$$

and hence our process k is the same as their process $i \equiv \frac{1}{j}$.

5.2. Feedback form

Schweizer (1995) finds a feedback solution of the following form

$$\theta_t^D = \rho_t - \beta_t V_t^{H_0,\theta^D} \tag{33}$$

$$\rho_t = \frac{\mathrm{E}_{t-1}^P\left[\Delta X_t \prod_{j=t+1}^T \left(1-\beta_j \Delta X_j\right) H_T\right]}{\mathrm{E}_{t-1}^P\left[\left(\Delta X_t\right)^2 \prod_{j=t+1}^T \left(1-\beta_j \Delta X_j\right)^2\right]}. \tag{34}$$

assuming $R_{\mathrm{f}\,t} = 1$. To verify that (33) is equivalent to (7) we have to show that $\rho_t = \left(\mathrm{E}_{t-1}^P\left[\left(\Delta X_t\right)^2 k_t\right]\right)^{-1} \mathrm{E}_{t-1}^P\left[k_t \Delta X_t H_t\right]$. First we employ the law iterated expectations

$$\begin{aligned}
\rho_t &= \frac{\mathrm{E}_{t-1}^P\left[\Delta X_t \prod_{j=t+1}^T \left(1-\beta_j \Delta X_j\right) H_T\right]}{\mathrm{E}_{t-1}^P\left[\left(\Delta X_t\right)^2 \prod_{j=t+1}^T \left(1-\beta_j \Delta X_j\right)^2\right]} \\
&= \frac{\mathrm{E}_{t-1}^P\left[\Delta X_t \left[\mathrm{E}_t^P \prod_{j=t+1}^T \left(1-\beta_j \Delta X_j\right) H_T\right]\right]}{\mathrm{E}_{t-1}^P\left[\left(\Delta X_t\right)^2 \mathrm{E}_t^P\left[\prod_{j=t+1}^T \left(1-\beta_j \Delta X_j\right)^2\right]\right]}
\end{aligned}$$

using the fact that ΔX_t is $\mathcal{F}_t$-measurable. From Lemma 3 in Schweizer (1996) one has

$$\mathrm{E}_t^P\left[\prod_{j=t+1}^{T}\left(1-\beta_j\Delta X_j\right)^2\right]=\mathrm{E}_t^P\left[\prod_{j=t+1}^{T}\left(1-\beta_j\Delta X_j\right)\right]. \tag{35}$$

On the other hand, equation (28) in this paper claims

$$k_t=\mathrm{E}_t^P\left[\prod_{j=t+1}^{T}\left(1-\beta_j\Delta X_j\right)\right]. \tag{36}$$

Thus we have

$$\begin{aligned}\rho_t &= \frac{\mathrm{E}_{t-1}^P\left[k_t\Delta X_t\mathrm{E}_t^P\left[H_T\prod_{j=t+1}^{T}\frac{k_j}{k_{j-1}}\left(1-\beta_j\Delta X_j\right)\right]\right]}{\mathrm{E}_{t-1}^P\left[\left(\Delta X_t\right)^2k_t\right]}\\ &= \frac{\mathrm{E}_{t-1}^P\left[k_t\Delta X_t\mathrm{E}_t^P\left[\tilde{m}_{t|t-1}\ldots\tilde{m}_{T|T-1}H_T\right]\right]}{\mathrm{E}_{t-1}^P\left[\left(\Delta X_t\right)^2k_t\right]}\\ &= \frac{\mathrm{E}_{t-1}^P\left[k_t\Delta X_tH_t\right]}{\mathrm{E}_{t-1}^P\left[\left(\Delta X_t\right)^2k_t\right]}\end{aligned}$$

as required by virtue of (6), (18) and (25).

The continuous time solution in feedback form was obtained for the first time in Schweizer (1992) under the assumption of constant opportunity set, which implies $\tilde{P}=\hat{P}$. A general proof for continuous semimartingales was given in Gourieroux, Laurent, and Pham (1998) and Rheinländer and Schweizer (1997). Let us rewrite (7) in a more familiar form, similar to equation (0.1) in Rheinländer and Schweizer (1997).

$$\begin{aligned}\hat{\theta}_{t-1} &= -\left[\mathrm{E}_{t-1}^Pk_t\Delta X_t\Delta X_t^*\right]^{-1}\mathrm{E}_{t-1}^Pk_t\Delta X_t\left[\frac{V_{t-1}^{H_0,\hat{\theta}}}{S_{t-1}^0}-\frac{H_t}{S_t^0}\right]=\\ &= R_{\mathrm{f}\,t-1}a_{t-1}\frac{V_{t-1}^{H_0,\hat{\theta}}-H_{t-1}}{S_{t-1}^0}+\\ &\quad+\left(\mathrm{E}_{t-1}^P\left[k_t\Delta X_t\Delta X_t^*\right]\right)^{-1}\mathrm{E}_{t-1}^P\left[k_t\Delta X_t\Delta\frac{H_t}{S_t^0}\right]\end{aligned}$$

It follows from the results of Rheinländer and Schweizer (1997) that for continuous gain process the second term can be interpreted as the integrand of $\mathrm{d}X_t$ in the Galtchouk-Kunita-Watanabe decomposition

of $\frac{H_T}{S_T^0}$ under $\tilde{P}$. This is essentially made possible by the fact that for continuous processes the randomness of k in the above formula vanishes and that both processes X and $\frac{H}{S^0}$ are martingales under $\tilde{P}$. In jump-diffusion limit, however, such interpretation is no longer possible, even if k_t is deterministic.

It seems more difficult to relate directly our formula (7) to the continuous time result of Gourieroux et al. (1998), Theorem 5.1. We can only refer to Rheinländer and Schweizer (1997) who show the equivalence between their equation (0.1) and the Theorem 5.1 of Gourieroux et al. (1998).

5.3. Continuous limit

It is interesting to see what happens to the solution in continuous time limit if we assume that all processes become Itô processes. When X is a Markov process this procedure can be formalized using the locally consistent Markov chain approximations championed in Kushner and Dupuis (2001).

Assume that the stochastic differential equation for the discounted gain process satisfies

$$X_{t+\mathrm{d}t} - X_t \equiv \mathrm{d}X_t = \mu \mathrm{d}t + \sigma \mathrm{d}B_t,$$

and the process k follows

$$\mathrm{d}k_t = \mu_k \mathrm{d}t + \sigma_k \mathrm{d}B_t.$$

The continuous-time counterpart of equation (5) reads

$$k_t = (1 + r_t \mathrm{d}t)^2 \times \left(\mathrm{E}_t^P \left[k_{t+\mathrm{d}t} \right] - \mathrm{E}_t^P \left[k_{t+\mathrm{d}t} \mathrm{d}X_t^* \right] \left(\mathrm{E}_t^P \left[k_{t+\mathrm{d}t} \mathrm{d}X_t \mathrm{d}X_t^* \right] \right)^{-1} \mathrm{E}_t^P \left[k_{t+\mathrm{d}t} \mathrm{d}X_t \right] \right). \tag{37}$$

From the Itô's lemma we have

$$\begin{aligned} \mathrm{E}_t^P \left[k_{t+\mathrm{d}t} \right] &= k_t + \mu_k \mathrm{d}t + o(\mathrm{d}t) \\ \mathrm{E}_t^P \left[k_{t+\mathrm{d}t} \mathrm{d}X_t \right] &= k_t \mu \mathrm{d}t + \sigma \sigma_k^* \mathrm{d}t + o(\mathrm{d}t) \\ \mathrm{E}_t^P \left[k_{t+\mathrm{d}t} \mathrm{d}X_t \mathrm{d}X_t^* \right] &= k_t \sigma \sigma^* \mathrm{d}t + o(\mathrm{d}t), \end{aligned}$$

and collecting the $\mathrm{d}t$ terms in (37) we obtain

$$\begin{aligned} 0 &= 2kr + \mu_k - \left(k\mu + \sigma\sigma_k^* \right)^* \left(k\sigma\sigma^* \right)^{-1} \left(k\mu + \sigma\sigma_k^* \right) \\ 0 &= 2r + \frac{\mu_k}{k} - \left(\mu + \sigma \frac{\sigma_k^*}{k} \right)^* \left(\sigma\sigma^* \right)^{-1} \left(\mu + \sigma \frac{\sigma_k^*}{k} \right). \end{aligned} \tag{38}$$

Similarly, from (30) we have

$$a = (\sigma\sigma^*)^{-1}\left(\mu + \sigma\frac{\sigma_k^*}{k}\right).$$

To see the market price of risk defining the variance-optimal measure we have to look at the continuous time limit of equation (18),

$$\tilde{m}_{t+\mathrm{d}t|t} = (1 + r_t\mathrm{d}t)^2$$
$$\times\frac{k_{t+\mathrm{d}t} - \mathrm{E}_t^P\left[k_{t+\mathrm{d}t}\mathrm{d}X_t^*\right]\left(\mathrm{E}_t^P\left[k_{t+\mathrm{d}t}\mathrm{d}X_t\mathrm{d}X_t^*\right]\right)^{-1}k_{t+\mathrm{d}t}\mathrm{d}X_t}{k_t} =$$
$$= 1 + \left(\frac{\sigma_k}{k}\left(I - \sigma^*\left(\sigma\sigma^*\right)^{-1}\sigma\right) - \mu^*\left(\sigma\sigma^*\right)^{-1}\sigma\right)\mathrm{d}B.$$

Taking into account that $\tilde{m}_{t+\mathrm{d}t|t} = \frac{\tilde{m}_{t+\mathrm{d}t}}{\tilde{m}_t}$ we can rearrange the above equation to obtain

$$\mathrm{d}\tilde{m}_t = \tilde{m}_t\left\{\frac{\sigma_k}{k}\left(I - \sigma^*\left(\sigma\sigma^*\right)^{-1}\sigma\right) - \mu^*\left(\sigma\sigma^*\right)^{-1}\sigma\right\}\mathrm{d}B$$

whereby we identify

$$\tilde{\lambda} \equiv -\left[I - \sigma^*\left(\sigma\sigma^*\right)^{-1}\sigma\right]\frac{\sigma_k^*}{k} + \sigma^*\left(\sigma\sigma^*\right)^{-1}\mu$$

as the *variance-optimal* market price of risk. Recall that the variable

$$\lambda \equiv \sigma^*\left(\sigma\sigma^*\right)^{-1}\mu$$

is known as the minimal market price of risk.

There are three interesting special cases:

1. k is *non-stochastic.* In this case $\sigma_k = 0$ and consequently

$$\begin{aligned}\tilde{\lambda} &= \lambda\\ a &= \left(\sigma\sigma^*\right)^{-1}\mu\end{aligned}$$

2. $\mathrm{d}k$ is *perfectly correlated* with the gain process $\mathrm{d}X$. In this case

$$\left[I - \sigma^*\left(\sigma\sigma^*\right)^{-1}\sigma\right]\sigma_k^* = 0$$

and therefore

$$\tilde{\lambda} = \lambda,$$

that is the variance-optimal measure coincides with the minimal martingale measure, but we are unable to say more about the process k.

3. dk is *uncorrelated* with the gain process dX ($\sigma\sigma_k^* = 0$). Then

$$\tilde{\lambda} = \lambda - \frac{\sigma_k^*}{k}$$

and from (38) we have

$$0 = 2r + \frac{\mu_k}{k} - \mu^* \left(\sigma\sigma^*\right)^{-1} \mu = 2r + \frac{\mu_k}{k} - \lambda^2.$$

Note that the Markov chain approximation results coincide with the rigorously derived results of Laurent and Pham (1999), who in the cases 2. and 3. are able to provide more explicit characterization of the process k (called i in their paper).

6. Conclusion

This paper presents a dynamic programming solution to the general mean-variance hedging problem in discrete time. Our analysis shows that in discrete time a natural solution exists which does not require the use of either variance-optimal martingale measure or duality theory.

Comparing the continuous time limit of our results with the results obtained for continuous semimartingales by martingale and duality methods we observe that our solution gives good explicit characterization of the variance-optimal market price of risk, however, it is not as detailed as the results of Laurent and Pham (1999). One merit of our solution is that it provides the kind of minimalistic recursive structure suitable for computer implementation, and that it offers a simple non-technical context in which the results of previous studies are easier to analyze and understand. The results presented here are a natural stepping stone to the analysis of hedging with discontinuous price processes.

Appendix

A. Proof of Theorems 2 and 3

1) Linear independence of basis assets means $\theta_{t-1}^* \Delta X_t = 0$ a.s.$\Rightarrow \theta_{t-1}^* = 0$ where θ_{t-1}^* is $\mathcal{F}_{t-1}$-measurable. By contradiction assume that the matrix $\mathrm{E}_{t-1}^P \left[k_t \Delta X_t \Delta X_t^*\right]$ is singular at a particular information node at $t-1$. Then there is $\theta_{t-1}^* \neq 0$ such that

$$\begin{aligned} 0 &= \theta_{t-1}^* \mathrm{E}_{t-1}^P \left[k_t \Delta X_t \Delta X_t^*\right] \theta_{t-1} \\ &= \mathrm{E}_{t-1}^P \left[\left(\theta_{t-1}^* \sqrt{k_t} \Delta X_t\right)^2\right] \end{aligned}$$

at that node, which is only possible if $\theta^*_{t-1}\sqrt{k_t}\Delta X_t = 0$ at the node in question. Since by assumption $\sqrt{k_t} > 0$ a.s., we have $\theta^*_{t-1}\Delta X_t = 0$ and from the linear independence of random variables ΔX_t it follows that $\theta^*_{t-1} = 0$ which contradicts $\theta^*_{t-1} \neq 0$. Hence $\mathrm{E}^P_{t-1}\left[k_t \Delta X_t \Delta X^*_t\right]$ is invertible at every node of the information set $\mathcal{F}_{t-1}$.

2) We wish to find the optimal control $x \in \mathcal{F}_0, \theta_i \in \mathcal{F}_i$, $i = 0, 1, \ldots, T-1$, to minimize

$$\min_{x,\theta_0,\ldots,\theta_{t-1}} \mathrm{E}^P_0 \left[k_t \left(S^0_t \left(V^{x,\theta}_0 + \sum_{i=0}^{t-1} \theta_i \Delta X_{i+1} \right) - H_t \right)^2 \right].$$

given that $k_t > 0$ and H_t are $\mathcal{F}_t$-measurable and exogenous in the sense of Definition 1. Bellman's principle of optimality dictates

$$\begin{aligned} &\min_{x,\theta_0,\ldots,\theta_{t-1}} \mathrm{E}^P_0 \left[k_t \left(V^{x,\theta}_t - H_t \right)^2 \right] = \\ &= \min_{x,\theta_0,\ldots,\theta_{t-2}} \mathrm{E}^P_0 \left[\min_{\theta_{t-1}} \mathrm{E}^P_{t-1} \left[k_t \left(V^{x,\theta}_t - H_t \right)^2 \right] \right] = \\ &= \min_{x,\theta_0,\ldots,\theta_{t-2}} \mathrm{E}^P_0 \left[\min_{\theta_{t-1}} \mathrm{E}^P_{t-1} \left[k_t \left(R_{\mathrm{f}\,t-1} V^{x,\theta}_{t-1} + S^0_t \theta_{t-1} \Delta X_t - H_t \right)^2 \right] \right], \end{aligned}$$

where the last equality follows from the definition of self-financing strategy (1).

The partial problem

$$J_{t-1} \triangleq \min_{\theta_{t-1}} \mathrm{E}^P_{t-1} \left[k_t \left(R_{\mathrm{f}\,t-1} V^{x,\theta}_{t-1} + S^0_t \theta_{t-1} \Delta X_t - H_t \right)^2 \right] \tag{39}$$

is simply a least squares regression. Let us recall that the abstract problem

$$\hat{U} = \min_{\beta} \mathrm{E}\left[Y + X^* \beta \right]^2$$

has the following solution

$$\begin{aligned} \hat{\beta} &= -\left(\mathrm{E}\left[XX^*\right]\right)^{-1} \mathrm{E}\left[XY\right] \\ \hat{U} &= \mathrm{E}\left[Y^2\right] - \left(\mathrm{E}\left[XY\right]\right)^* \left(\mathrm{E}\left[XX^*\right]\right)^{-1} \mathrm{E}\left[XY\right], \end{aligned}$$

provided that the random variables X are linearly independent. Thus the problem (39) is solved by setting

$$\begin{aligned} \hat{U} &= J_{t-1} \\ Y &\equiv \sqrt{k_t} \left(R_{\mathrm{f}\,t-1} V^{x,\theta}_{t-1} - H_t \right) \\ X &\equiv \sqrt{k_t} \Delta X_t \\ \beta &\equiv S^0_t \theta_{t-1}, \end{aligned}$$

whereby, after collecting all the powers of V, we obtain

$$J_{t-1} = k_{t-1}\left(V_{t-1}^{x,\theta} - H_{t-1}\right)^2 + \mathrm{E}_{t-1}^P\left[k_t H_t^2\right] - k_{t-1}H_{t-1}^2 - \\ - \left(\mathrm{E}_{t-1}^P\left[k_t \Delta X_t H_t\right]\right)^* \left(\mathrm{E}_{t-1}^P\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1} \mathrm{E}_{t-1}^P\left[k_t \Delta X_t H_t\right] \quad (40)$$

$$\theta_{t-1}^D = -\frac{1}{S_t^0}\left(\mathrm{E}_{t-1}^P\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1} \mathrm{E}_{t-1}^P\left[k_t \Delta X_t \left(R_{\mathrm{f}\,t-1} V_{t-1}^{x,\theta^D} - H_t\right)\right]$$

with k_{t-1}, H_{t-1} defined in (5), (6).

3) It remains to be shown that the process H_{t-1} is well defined, that is $k_{t-1} > 0$ a.s. By assumption $k_t > 0$ a.s. Note from (5) that k_{t-1} can be written equivalently as

$$k_{t-1} = R_{\mathrm{f}\,t-1}^2 \mathrm{E}_{t-1}^P\left[\left(\sqrt{k_t} - \theta^* \sqrt{k_t}\Delta X_t\right)^2\right], \quad (41)$$

with $\theta^* = \mathrm{E}_{t-1}^P\left[k_t \Delta X_t^*\right]\left(\mathrm{E}_{t-1}^P\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1}$. We continue the proof by contradiction: suppose that $k_{t-1} = 0$ at a particular node. From (41) this is only possible if $\sqrt{k_t} - \theta^*\sqrt{k_t}\Delta X_t = 0$ a.s. for a conditional distribution at this node. Since by assumption $\sqrt{k_t} > 0$ it must be that $1 = \theta^* \Delta X_t$ for the conditional distribution at this node, which contradicts our assumption of no arbitrage at that node. Hence it must be that $k_{t-1} > 0$ a.s.

4) By induction assumption k_t and H_t are exogenous. Clearly, the conditional expectation of an exogenous variable is still exogenous, and any function of exogenous variables is again exogenous. Since the formulae (5), (6) only involve functions and conditional expectations of exogenous variables it follows that k_{t-1} and H_{t-1} are exogenous. The same reasoning implies that the last three terms in equation (40) are exogenous, that is the problem

$$\min_{x,\theta_0,\ldots,\theta_{t-1}} \mathrm{E}_0^P\left[k_t\left(V_t^{x,\theta} - H_t\right)^2\right] = \\ = \mathrm{E}_0^P\left[k_t H_t^2 - k_{t-1}H_{t-1}^2 - \right. \\ \left. - \left(\mathrm{E}_{t-1}^P\left[k_t \Delta X_t H_t\right]\right)^* \left(\mathrm{E}_{t-1}^P\left[k_t \Delta X_t \Delta X_t^*\right]\right)^{-1} \mathrm{E}_{t-1}^P\left[k_t \Delta X_t H_t\right]\right] + \\ + \min_{x,\theta_0,\ldots,\theta_{t-2}} \mathrm{E}_0^P\left[k_{t-1}\left(V_{t-1}^{x,\theta} - H_{t-1}\right)^2\right]$$

has the same optimal controls as the problem

$$\min_{\theta_0,\ldots,\theta_{t-2}} \mathrm{E}_0^P\left[k_{t-1}\left(V_{t-1}^{x,\theta} - H_{t-1}\right)^2\right],$$

which completes the proof of Theorem 2.

4) As a by-product of the above calculation we have learnt that

$$\min_{\theta_{t-1}} \mathrm{E}_{t-1}^{P}\left[k_t\left(V_t^{x,\theta}-H_t\right)^2\right]=k_{t-1}\left(V_{t-1}^{x,\theta}-H_{t-1}\right)^2+\mathrm{ESRE}_{t-1}^{P}\left(H_t\right), \tag{42}$$

where the one-step ahead expected squared replication error is given by

$$\mathrm{ESRE}_{t-1}^{P}\left(H_t\right)\equiv \mathrm{E}_{t-1}^{P}\left[k_t H_t^2\right]-k_{t-1}H_{t-1}^2$$
$$-\left(\mathrm{E}_{t-1}^{P}\left[k_t\Delta X_t H_t\right]\right)^{*}\left(\mathrm{E}_{t-1}^{P}\left[k_t\Delta X_t\Delta X_t^{*}\right]\right)^{-1}\mathrm{E}_{t-1}^{P}\left[k_t\Delta X_t H_t\right].$$

A simple manipulation demonstrates the identity

$$\mathrm{ESRE}_{t-1}^{P}\left[H_t\right]=\mathrm{E}_{t-1}^{P}\left[k_t\left(R_{\mathrm{f}\,t-1}H_{t-1}+S_t^0\left(\bar{\theta}_{t-1}^{D}\right)^{*}\Delta X_t-H_t\right)^2\right],$$

where $\bar{\theta}_t^D$ is given by (8).

5) A recursive application of (42) starting from $t=T$ gives

$$\min_{\theta_{t-1}} \mathrm{E}_{t-1}^{P}\left[k_T\left(V_T^{x,\theta}-H_T\right)^2\right]=k_{t-1}\left(V_{t-1}^{x,\theta}-H_{t-1}\right)^2+\varepsilon_{t-1}^2$$

where ε is given recursively

$$\begin{aligned}\varepsilon_{t-1}^2 &= \mathrm{E}_{t-1}^{P}\left[\varepsilon_t^2\right]+\mathrm{ESRE}_{t-1}^{P}\left[H_t\right],\\ \varepsilon_T &= 0.\end{aligned}$$

which completes the proof of Theorem 3.

References

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

Černý, A. (1999). Currency crises: Introduction of spot speculators. *International Journal of Finance and Economics 4*(1), 75–89.

Černý, A. (2002). Derivatives without differentiation: Optimal hedging on Markov chains. Technical report, Imperial College Management School.

Černý, A. (2004). *Mathematical Techniques in Finance: Tools for Incomplete Markets*. Princeton University Press.

Duffie, D. and H. Richardson (1991). Mean–variance hedging in continuous time. *Annals of Applied Probability 1*, 1–15.

Föllmer, H. and M. Schweizer (1991). Hedging of contingent claims under incomplete information. In M. H. A. Davis and R. J. Elliott (Eds.), *Applied Stochastic Analysis*, pp. 389–414. New York: Gordon and Breach.

Frey, R. and A. Stremme (1997). Market volatility and feedback effects from dynamic hedging. *Mathematical Finance 7*(4), 351–374.

Gourieroux, C., J.-P. Laurent, and H. Pham (1998). Mean-variance hedging and numéraire. *Mathematical Finance 8*(3), 179–200.

Kushner, H. J. and P. Dupuis (2001). *Numerical Methods for Stochastic Control Problems in Continuous Time* (2nd ed.). Springer Verlag.

Laurent, J. P. and H. Pham (1999). Dynamic programming and mean-variance hedging. *Finance and Stochastics 3*(1), 83–110.

Rheinländer, T. and M. Schweizer (1997). On L^2 projections on a space of stochastic integrals. *Annals of Probability 25*(4), 1810–1831.

Schachermayer, W. (1993). A counterexample to several problems in the theory of asset pricing. *Mathematical Finance 3*, 217–229.

Schäl, M. (1994). On quadratic cost criteria for option hedging. *Mathematics of Operations Research 19*(1), 121–131.

Schweizer, M. (1992). Mean-variance hedging for general claims. *The Annals of Applied Probability 2*(1), 171–179.

Schweizer, M. (1995). Variance-optimal hedging in discrete time. *Mathematics of Operations Research 20*(1), 1–32.

Schweizer, M. (1996). Approximation pricing and the variance-optimal martingale measure. *The Annals of Probability 24*(1), 206–236.

Schweizer, M. (1998). Risky options simplified. *International Journal of Theoretical and Applied Finance 2*(1), 59–82.

Toft, K. B. (1996). On the mean-variance tradeoff in option replication with transactions costs. *The Journal of Financial and Quantitative Analysis 31*(2), 233–263.