Maximising the Probability of a Perfect Hedge using an Imperfectly Correlated Instrument

David Hobson* Jeremy Penn†

First Version: 1st April 2004
This Version: 21st January 2005

Abstract

Let X^ϕ denote the trading wealth generated using a strategy ϕ, and let C_T be a contingent claim which is not spanned by the traded assets. Consider the problem of finding the strategy which maximises the probability of terminal wealth meeting or exceeding the claim value at some fixed time horizon, i.e. of finding $\sup_\phi \mathbb{P}^x(X_T^\phi \geq C_T)$. This problem is sometimes referred to as the quantile hedging problem.

We consider the quantile hedging problem when the traded asset and the contingent claim are correlated geometric Brownian motions. This fits with several important examples. One of the benefits of working with such a concrete model is that although it is incomplete we can still do calculations. In particular, we can consider some detailed issues such as the impact of the timing at which information about C_T is revealed.

Keywords: Hedging strategies; stochastic control; Brownian motion; policy improvement; timely information
Mathematics Subject Classification: 60H30; 60J65; 91B28; 93E20
Journal of Economic Literature Classification: G11; G13; D81; C61

*masdgh@bath.ac.uk, Department of Mathematical Sciences, University of Bath, Claverton Down, Bath, BA2 7AY, UK.
†Credit Suisse First Boston. Supported, in part, by EPSRC grant number 9980098X.

1 Introduction

This article is concerned with optimal trading behaviour in complete and incomplete markets, when replication and super-replication are not possible.

Suppose we have a trading wealth X^{ϕ}, which evolves according to the strategy ϕ, and a contingent claim C that we shall be obliged to meet at some time horizon T. We shall be concerned with maximising the chance of meeting the claim, that is with finding $\sup_{\phi} \mathbb{P}(X^{\phi}_T \geq C_T)$. We are interested in several different cases. The claim may be correlated to the trading wealth or be independent of it; it may be dependent on some asset or not; it may be revealed at the outset, at the horizon time or continuously up to the horizon time.

There are many situations of genuine practical interest where questions of this kind might arise. Trading activity in banks and other financial institutions is often determined by traders trying to reach targets of one kind or another, often linked to their bonus packages. They are not concerned with how much they might fall short of or exceed the target, only whether they reach it or not. Similarly portfolio managers and pension fund managers are often concerned with exceeding a certain benchmark (for example a market index) by a fixed percentage.

In particular, we shall consider cases where the trading wealth is that from trading on one stock and the contingent claim depends on another stock, which may be only partially correlated with, or even independent of, the tradeable stock. This is a good way of looking at a situation where a financial institution has entered into an agreement to pay a claim based on some stock, but cannot hedge using that stock because of a lack of liquidity or because of regulatory restrictions.

A number of authors have looked at the problem of maximising the probability of meeting a claim. Kulldorff (1993) controls the volatility of a geometric-Brownian trading wealth to reach a constant claim. He considers both discrete and continuous time. The former is essentially a linear programming problem that will not concern us. In the latter, he obtains a partial differential equation for the value function, that is the probability of meeting the claim given a particular wealth and time to the horizon. Heath (1993) considers the same problem but uses an approach based on the Neyman-Pearson Lemma.

Browne (1999*b*) also has a constant claim. He has n geometric-Brownian assets and also obtains a partial differential equation for his value function. He analyses the region in which borrowing takes place, that is wealth and time-to-go pairs for which the holding in cash is negative (though total portfolio value will always be positive).

Spivak & Cvitanić (1999) also consider traded assets which are driven by exponential Brownian motions. In the complete market setting they consider the problem of quantile hedging a general contingent claim. They consider the dual problem, and find that they should meet the claim when it is small and ensure that they are left with nothing when the claim is large. Spivak & Cvitanić (1999) also consider an incomplete market version of the model in which the return rate on the risky asset is unknown.

The quantile hedging problem for a general semi-martingale traded asset is considered in Föllmer & Leukert (1999). They reduce the problem to that of finding the maximal success set and then apply the Neyman-Pearson Lemma. They also consider incomplete markets and have an interesting example in an exponential Brownian motion setting in which volatility suffers a shock at some time T, and the agent wishes to quantile hedge a claim at some later time T'. By solving the problem explicitly over the period $(T, T']$ the problem can be reduced to a quantile hedging problem over $[0, T]$ in which the claim to be replicated depends upon a random quantity, the value of which is revealed at the horizon time T.

Föllmer & Leukert (2000) and Schulmerich & Trautmann (2003) consider the related problem of minimising the expected shortfall, $\mathbb{E}[(C_T - X_T)^+]$.

Our analysis is different to each of the above papers, because we consider incomplete markets, and the incompleteness is introduced in a completely different way. We introduce incompleteness by defining the claim value relative to a second Brownian motion which is independent of the traded asset. This model is in the spirit of Browne (1999a), and the literature on basis risk. However, Browne (1999a) considers a perpetual, infinite horizon problem and various investment criteria other than quantile hedging.

The approach we shall take will be structured as follows. In Section 2 we analyse the situation in which the value of the claim to be replicated is made known either at the outset ($t = 0$) or at the horizon time ($t = T$). We do this both in the complete market and in the case of an independent claim. This analysis provides a benchmark for interesting comparisons with the case where information about the claim value is revealed over time.

The main novelty of our study is in the fact that there is partial correlation between our tradeable asset and our contingent claim. We have a claim with value C and a traded asset with price P given by

$$\frac{dC_t}{C_t} = \sigma_t dW_t + \nu_t dt, \qquad C_0 = c_0 \tag{1}$$

$$\frac{dP_t}{P_t} = \eta_t(\rho_t dW_t + \sqrt{1-\rho_t^2} dB_t) + \mu_t dt, \tag{2}$$

(where B and W are independent Brownian motions). The agent chooses a strategy ϕ, depending on information available at time t, which determines how the trading wealth X evolves:

$$\frac{dX_t^\phi}{X_t^\phi} = \phi_t \frac{dP_t}{P_t}, \qquad X_0 = x_0.$$

This model provides a natural way of specifying the evolution of the claim value and a good way of introducing correlation.

We will first derive a general partial differential equation for this category of problems. We show the equivalence of optimising the success probability and optimising the expected ratio of wealth to claim, i.e. we show that $\sup_\phi \mathbb{P}(X_T^\phi \geq C_T) = \sup_\phi \mathbb{E}[U_T \wedge 1]$. Such wealth or success ratios feature in Föllmer & Leukert (1999).

In Section 4, we show how a more general multi-dimensional formulation can be reduced to a one-dimensional problem, and how a further transformation can be used to reduce the problem with correlation to the uncorrelated case. In Section 5 we consider a couple of related problems and show that they too can be reduced to the form we are considering. In particular, in our context, the expected shortfall problem can be reduced to the quantile hedging problem.

From Section 6 onwards we focus on the one-dimensional (i.e. single traded asset) quantile hedging problem in which the claim is uncorrelated with the traded asset. By the above remarks this is with no loss of generality.

In the complete market case we obtain solutions by reducing the case with non-zero claim volatility to that with constant claim (a special case of the independent claim case). When the claim is constant we have an algebraic formula for the value function or success probability.

We are also able to obtain, in Section 7, explicit solutions for the independent claim problem in the restricted case of a zero-drift tradeable asset. However it seems to be impossible to give a closed form solution to the quantile hedging problem when the asset has non-zero drift. Hence, in Section 8 we turn to a numerical approach for solving the independent claim problem. We first solve the zero-drift case numerically as the comparison between exact and numerical solutions is useful later when we cannot obtain an exact solution. When we consider the case of an independent asset with drift, we use a policy improvement algorithm, as the partial differential equation to be solved is non-linear. The dual problem can also be solved numerically, and this gives us a further check on the accuracy of our results.

In Section 9 we consider the impact that the time at which information about the claim C is revealed has upon the probability of super-replicating the claim. We compare the odds of success when a similar claim is revealed at time 0, just before the end of trading (time $T-$), just after trading finishes T, or continuously through time. As expected we find that the sooner information about C is revealed, the greater the probability of super-replicating, but we are also able to quantify the relative advantage of early information. In Section 10 we return to the complete markete case and consider the impact that the timing of volatility has upon the probability of meeting the claim.

The key feature of our model is that we allow information about the claim to be hedged to be *revealed continuously* over time and that we have *partial* correlation between tradeable asset and contingent claim. This is a realistic model for many problems which arise in real options.

2 Claims Revealed Instantaneously in Time

In this section we begin with an original market that is complete. By assumption, the wealth X^ϕ from trading in this market is restricted to be non-negative. However in this section, we do not need to assume anything about the dynamics of any underlying assets.

Now we introduce a non-negative contingent claim C. This together with the tradeable asset may or may not form a complete market. We consider the two cases separately.

Consider first the case where the tradeable assets and the contingent claim form a complete market. This case can be motivated by thinking of a trader who is concerned with trading so as to outperform the market, possibly by some given proportion. This is considered in both Föllmer & Leukert (1999) and Spivak & Cvitanić (1999).

Denote the sample space by Ω and take the real-world probability measure to be $\mathbb{P}$. We assume there is an equivalent measure $\tilde{\mathbb{P}}$ such that $\tilde{\mathbb{E}}[X^\phi] \leq x_0$ for any attainable wealth. We shall write $Z = (d\tilde{\mathbb{P}}/d\mathbb{P})|_{\mathcal{F}}$, where $\mathcal{F}$ is the information available at time T.

Applying a direct Lagrangian approach (which is essentially a simpler version of the argument used below for the incomplete market) we find that the optimal wealth is

$$X^* = C\mathbb{1}_{\{C \leq c^*\}} = \begin{cases} C & \text{if} \quad C \leq c^* \\ 0 & \text{otherwise} \end{cases} ,$$

where c^* is such that

$$\int_{\{C \leq c^*\}} ZC\, \mathbb{P}(d\omega) = x_0.$$

In particular the optimal strategy has an 'all-or-nothing' flavour.

Now suppose that the claim is independent of the the traded asset and that the strategies available to the agent are now allowed to depend on the claim value. The fact that the agent is not allowed to use information about the claim value is representative of the idea that the claim value is only revealed at the maturity time T. Föllmer & Leukert (1999) consider this case. Let the sample space be partitioned into a sample space pertaining to the wealth process X and one pertaining to the claim C, that is let $\Omega = \left(\Omega' \times \Omega^C\right)$. We can identify Ω^C with $\mathbb{R}_+$. The information available at time T is similarly partitioned, $\mathcal{F} = \left(\mathcal{F}' \times \mathcal{F}^C\right)$. We assume there is an equivalent measure $\tilde{\mathbb{P}}$ such that $\tilde{\mathbb{E}}[X^\phi] \leq x_0$ for any ϕ and the law of C is unchanged. Write $Z = (d\tilde{\mathbb{P}}/d\mathbb{P})|_{\mathcal{F}'}$.

We first consider non-negative claims whose distribution functions F_C are concave on $\mathbb{R}_+$. Conditional on $X(\omega')$, the probability of meeting the claim is $F_C(X(\omega'))$. The Lagrangian for the problem is

$$L(X,\lambda) = \int_{\Omega'} \left\{ F_C(X(\omega')) - \lambda(Z(\omega')X(\omega') - x_0) \right\} \mathbb{P}(d\omega').$$

For given ω' we choose X to maximise

$$F_C(X(\omega')) - \lambda(Z(\omega')X(\omega') - x_0),$$

subject to our requirement that wealth be non-negative.

Differentiating and using the Lagrangian Sufficiency Theorem we find that for non-negative λ the finite maximum is obtained by taking the terminal wealth to be

$$X^* = I(\lambda Z),$$

where I is the left-continuous version of $(F_C')^{-1}$. (I is well defined by the assumed concavity of F_C.) The value of λ is now chosen such that

$$\int_{\Omega'} Z I(\lambda Z)\, \mathbb{P}(d\omega') = x_0.$$

Turn now to a general claim and denote its distribution function by F_C. For a function $f : \mathbb{R}_+ \to \mathbb{R}$ refer to $\tilde{f} = \inf\{g : g \geq f,\, g \text{ concave}\}$ as the *concave relaxation* of f. We can take $\tilde{F}_C$ to be the concave relaxation of F_C and solve the $\tilde{F}_C$-problem. Denote by $\tilde{X}^*$ the solution to this problem. Now in the $\tilde{F}_C$-problem for given $\omega' \in \Omega'$ we wish to maximise

$$\tilde{F}_C(X(\omega')) - \lambda(Z(\omega')X(\omega') - x_0). \tag{3}$$

Where F_C and $\tilde{F}_C$ do not coincide (3) is linear in X. Over a linear portion we can move X either up or down so as to strictly not decrease (3). Hence we can modify our solution $\tilde{X}^*$ to give a solution X^* which is still optimal for $\tilde{F}_C$, and which only takes values for which F_C and $\tilde{F}_C$ coincide. Then for any X,

$$\mathbb{E}[F_C(X^*)] = \mathbb{E}[\tilde{F}_C(X^*)] \geq \mathbb{E}[\tilde{F}_C(X)] \geq \mathbb{E}[F_C(X)]$$

by the $\tilde{F}_C$-optimality of X^* and the definition of the concave relaxation. So in fact X^* is F_C-optimal. This analysis has all been fairly abstract, but we calculate the relevant quantities for our main example (the lognormal case) in Section 9 below.

The solutions we have obtained in this Section will give us useful intuition for problems in later sections for claims whose value evolves over time. In particular, in the case of the complete market, the theme of meeting the claim or having nothing will recur.

3 The Bellman Equation

In this section we consider the situation where the claim is the terminal value of a stochastic process which evolves over time. This is more realistic for practical applications, like our examples of a trader trying to meet his targets or a company hedging against a claim on a restricted or illiquid stock. It also allows us to specify some correlation between the claim and the trading wealth. However, the cost of specifying how the claim is revealed is that we must make our model much more specific.

Henceforth we will assume that the claim and the traded asset price process evolve according to geometric Brownian motions in the following manner:

$$\begin{aligned} \frac{dC_t}{C_t} &= \sigma_t dW_t, \quad C_0 = c_0, && (4) \\ \frac{dP_t}{P_t} &= \rho_t dW_t + \sqrt{1-\rho_t^2} dB_t + \mu_t dt, && (5) \end{aligned}$$

where W_t and B_t are independent Brownian motions. To keep notation simple we assume that interest rates are zero or equivalently that we are working with forward prices. The parameter ρ gives the correlation between the contingent claim and the tradeable asset. Here σ_t, μ_t and ρ_t are deterministic. A strategy ϕ_t, chosen by the agent and adapted to the natural filtration of B and W, controls how the wealth evolves:

$$\frac{dX_t^\phi}{X_t^\phi} = \phi_t \frac{dP_t}{P_t}, \tag{6}$$

with $X_0 = x_0$. Two of the key reasons for choosing a geometric-Brownian model are that it is both tractable and a reasonable model for the price of a financial security.

Note that in comparison with (2) we have taken the volatility η of the traded asset to be the unit constant. We show in the next section that this is with no loss of generality, since η can always be absorbed into the strategy ϕ. From now on we will also assume that σ, μ and ρ are constants which enables us to simplify the notation.

It will be convenient to reduce the number of variables in our problem by considering the ratio $U_t = X_t/C_t$. Then by Itô's formula

$$\frac{dU_t}{U_t} = dW_t(\phi_t\rho - \sigma) + dB_t\phi_t\sqrt{1-\rho^2} + dt(\phi_t\mu + \sigma^2 - \phi_t\sigma\rho), \tag{7}$$

so that, given ϕ_t, U is an autonomous diffusion. Since the terminal condition can also be rewritten in terms of U we find that we can define the value function, in terms of U (and t) alone,

$$V(u,t;\sigma,\mu,\rho) \equiv V(u,t) = \sup_{(\phi_s):t\le s\le T} \mathbb{E}_t\left[\mathbb{1}_{\{U_T\ge 1\}} | U_t = u\right].$$

The value function is a super-martingale under any policy ϕ and a martingale under the optimal policy ϕ^*. Standard arguments give that the Hamilton-Jacobi-Bellman (HJB) equation for this problem is

$$0 = \sup_{\phi} \left\{ \frac{1}{2} V_{uu} u^2 \phi^2 + \phi (V_u u(\mu - \sigma\rho) - V_{uu} u^2 \sigma\rho) \right\} + \sigma^2 (V_u u + \frac{1}{2} V_{uu} u^2) + \dot{V}. \tag{8}$$

We have boundary conditions $V(0,t) = 0$, meaning that if we have no money we will definitely lose, and $V(u,t) \to 1$ as $u \to \infty$, meaning that as we have more and more money, or a smaller and smaller claim, we become increasingly certain of winning.

If $V_{uu} < 0$ then we have a finite maximum at

$$\phi^* = \frac{V_u(\mu - \sigma\rho)}{V_{uu} u} - \sigma\rho,$$

and the HJB equation reduces to

$$\sigma^2 (V_u u + \frac{1}{2} V_{uu} u^2) - \frac{\{V_u(\mu - \sigma\rho) - V_{uu} u \sigma\rho\}^2}{2 V_{uu}} + \dot{V} = 0. \tag{9}$$

We are interested in finding the supremum over strategies ϕ of $\mathbb{P}\left(X_T^\phi \geq C_T \right)$ which corresponds to having a payoff function of $V(u,T) = \mathbb{1}_{\{u \geq 1\}}$. However, we will now show that, in our setting, for $t < T$, we obtain the same value function if we use the (continuous) boundary condition $u \wedge 1$.

The payoff $V(u,T) = u \wedge 1$ is continuous and concave whereas our original payoff is not. Föllmer & Leukert (1999) justified the use of this payoff by a heuristic comparison with the Neyman-Pearson theory. In our setting, where all the fundamental processes are continuous (almost surely), we can prove that this change of boundary condition is appropriate.

Theorem 1 *We have the following equivalence of payoffs for our model with dynamics given by (4)–(5), for $t < T$:*

$$\sup_{\phi_s : t \leq s \leq T} \mathbb{P}\left(U_T^\phi \geq 1 \middle| \mathcal{F}_t \right) = \sup_{\phi_s : t \leq s \leq T} \mathbb{E}\left[U_T^\phi \wedge 1 \middle| \mathcal{F}_t \right].$$

Proof. First suppose $x_0 < c_0$ and suppose temporarily that the time horizon is very short. Choose, for $0 \leq t < \varepsilon$, the policy

$$\theta_t = \frac{\mathrm{sgn}(\mu)}{X_t \sqrt{\varepsilon - t}} \mathbb{1}_{\{0 \leq X_t \leq c_0\}},$$

giving

$$dX_t^\theta = \left\{ \frac{\mathrm{sgn}(\mu)}{\sqrt{\varepsilon - t}} dB_t + \frac{|\mu|}{\varepsilon - t} dt \right\} \mathbb{1}_{\{0 \leq X_t \leq c_0\}}.$$

Define

$$dY_t = \frac{\text{sgn}(\mu)}{\sqrt{\varepsilon - t}} dB_t \mathbb{1}_{\{0 \le Y_t \le c_0\}}, \quad Y_0 = x_0.$$

Then $X^\theta \ge Y$ and it follows that

$$\sup_\phi \mathbb{P}\left(X_\varepsilon^\phi \ge c_0\right) \ge \mathbb{P}\left(X_\varepsilon^\theta \ge c_0\right) \ge \mathbb{P}\left(Y_\varepsilon \ge c_0\right).$$

But Y is a martingale and by construction Y_t is almost sure to exit $(0, c_0)$ by time ε. Thus $\mathbb{P}(Y_\varepsilon \ge c_0) = x_0/c_0$.

On the other hand, if $x_0 \ge c_0$ then taking $\theta = 0$, we have $\mathbb{P}\left(X_\varepsilon^\theta \ge c_0\right) = 1$ and, combining these two cases

$$\sup_\phi \mathbb{P}(X_\varepsilon^\phi \ge c_0) \ge \frac{x_0}{c_0} \wedge 1.$$

Our intuition for this bound is that as the time horizon is short, we are unable to take advantage of the drift. Instead we gamble on reaching c_0, with martingale odds.

We are in fact not concerned with with $\mathbb{P}\left(X_\varepsilon^\phi \ge c_0\right)$ but rather with $\mathbb{P}\left(X_\varepsilon^\phi \ge C_\varepsilon\right)$. Consider the chance of the claim moving a relatively long way in a short time,

$$p(\varepsilon) = \mathbb{P}\left(C_\varepsilon > c_0(1 + \varepsilon^{\frac{1}{4}})\right) = 1 - \Phi\left(\frac{\log(1 + \varepsilon^{\frac{1}{4}}) + \frac{1}{2}\sigma_t^2 \varepsilon}{\sigma \varepsilon^{\frac{1}{2}}}\right),$$

where Φ is the normal distribution function. Note that $p(\varepsilon) \to 0$ as $\varepsilon \downarrow 0$. It follows that for general $x_0 \ge 0$,

$$\begin{aligned} \sup_\phi \mathbb{P}(X_\varepsilon \ge C_\varepsilon) &\ge \sup_\phi \mathbb{P}\left(X_\varepsilon \ge c_0(1 + \varepsilon^{\frac{1}{4}}),\, C_\varepsilon \le c_0(1 + \varepsilon^{\frac{1}{4}})\right) \\ &\ge \frac{1 - p(\varepsilon)}{1 + \varepsilon^{\frac{1}{4}}} \left(\frac{x_0}{c_0} \wedge 1\right). \end{aligned} \tag{10}$$

Now we wish to extend this result to arbitrary time horizons. Suppose the strategy π is such that $(\pi_t)_{0 \le t \le T-\varepsilon}$ achieves within ε of the maximal value for $\mathbb{E}[U_{T-\varepsilon} \wedge 1]$ and $(\pi_t)_{T-\varepsilon < t \le T}$ achieves within ε of the maximal value for $\mathbb{P}(X_T^\theta \ge C_T | \mathcal{F}_{T-\varepsilon})$, then

$$\mathbb{P}(X_T^\pi \ge C_T) = \mathbb{E}[\mathbb{P}(X_T^\pi \ge C_T | \mathcal{F}_{T-\varepsilon})] \ge \left(\frac{1 - p(\varepsilon)}{1 + \varepsilon^{\frac{1}{4}}}\right)\left(\mathbb{E}[U_{T-\varepsilon}^\pi \wedge 1] - \varepsilon\right) - \varepsilon,$$

using (10). So

$$\begin{aligned}
\sup_\phi \mathbb{P}(X_T^\phi \geq C_T) &\geq \liminf_{\varepsilon\downarrow 0} \sup_\phi \left(\frac{1-p(\varepsilon)}{1+\varepsilon^{\frac{1}{4}}}\right)\left(\mathbb{E}[U^\phi_{T-\varepsilon}\wedge 1]-\varepsilon\right)-\varepsilon \\
&\geq \sup_\phi \liminf_{\varepsilon\downarrow 0} \left(\frac{1-p(\varepsilon)}{1+\varepsilon^{\frac{1}{4}}}\right)\left(\mathbb{E}[U^\phi_{T-\varepsilon}\wedge 1]-\varepsilon\right)-\varepsilon \\
&\geq \sup_\phi \mathbb{E}[\liminf_{\varepsilon\downarrow 0}(U^\phi_{T-\varepsilon}\wedge 1)] \\
&= \sup_\phi \mathbb{E}[U^\phi_T \wedge 1],
\end{aligned}$$

using Fatou's Lemma and dominated convergence.

However, we also have $\mathbb{1}_{\{U_T\geq 1\}} \leq U_T \wedge 1$ so the reverse inequality is trivial and

$$\sup_\phi \mathbb{P}(U^\phi_T \geq 1) = \sup_\phi \mathbb{E}[U^\phi_T \wedge 1].$$

□

4 Model Generalisations

Suppose that instead of specifying that the claim and price processes evolved using (4)–(5) we used the following model. The claim is as before except that it has a constant drift κ (the results extend easily to a deterministic drift)

$$\frac{dC_t}{C_t} = \sigma dW_t + \kappa dt, \quad C_0 = c_0. \tag{11}$$

The single asset P is replaced by n assets indexed by i and solving the equations

$$\frac{dP^i_t}{P^i_t} = \eta^i \sum_{j=1}^n \Pi^{ij} dZ^j_t + \mu^i dt, \tag{12}$$

where the Brownian motions Z^j are independent. We suppose that the dependence between W and Z^j is given by the relationship

$$dW = \sum_{j=1}^n \rho^j dZ^j + \overline{\rho} d\overline{Z}, \tag{13}$$

where $\overline{Z}$ is a further Brownian motion independent of the Z^j, and again all parameters are constant with $\sum_j (\rho^j)^2 + \overline{\rho}^2 = 1$. By assumption the market without the claim C is complete, and the number of traded assets equals the number of Brownian motions Z^j. Further, to exclude degeneracy and arbitrage we assume that Π is non-singular. If W is linearly dependent on these Z^j then the market remains complete with the introduction of C

and $\overline{\rho} = 0$. Otherwise the introduction of the claim C makes the market incomplete.

In order to emphasise the similarities with previous sections we can rewrite (12) in the form

$$\frac{dP_t^i}{P_t^i} = \gamma^i dW_t + \sum_{j=1}^{n} \Gamma^{ij} dB_t^j + \mu^i dt, \tag{14}$$

where the Brownian motions B^j are a reparameterisation of the Z^j, but with the additional property that they are each independent of W.

We have

$$\begin{pmatrix} \Gamma & \gamma \\ 0 & \sigma \end{pmatrix} \begin{pmatrix} dB \\ dW \end{pmatrix} = \begin{pmatrix} D\Pi & 0 \\ \sigma\rho & \sigma\overline{\rho} \end{pmatrix} \begin{pmatrix} dZ \\ d\overline{Z} \end{pmatrix}, \tag{15}$$

where $D = D(\eta^1; \eta^2; \dots; \eta^n)$, the diagonal matrix with entries $\eta^1, \eta^2, \dots, \eta^n$. From the quadratic variation of

$$D\Pi dZ = \Gamma dB + \gamma dW,$$

we obtain

$$\eta\Pi\Pi^T\eta^T = \Gamma\Gamma^T + \gamma\gamma^T,$$

which we shall use in Theorem 2.

Given the formulation (11) - (13) we can define B and W by inverting on the left in (15). Conversely, given B and W we can define

$$dZ = \Pi^{-1}D^{-1}\Gamma dB + \Pi^{-1}D^{-1}\gamma dW,$$

and, if $\overline{\rho} \neq 0$,

$$d\overline{Z} = -\overline{\rho}^{-1}\rho\Pi^{-1}D^{-1}\Gamma dB + \overline{\rho}^{-1}dW.$$

If $\overline{\rho} = 0$, then $\overline{Z}$ is not necessary for the fomulation in (13).

The trading strategy of the agent is represented by the vector χ_t^i so that the dynamics of the trading wealth are given by

$$\frac{dX_t}{X_t} = \sum_{i=1}^{n} \chi_t^i \frac{dP_t^i}{P_t^i}, \quad X_0 = x_0. \tag{16}$$

Theorem 2 *The value function $J_n(x/c, t_0; \sigma, \kappa, \mu, [\gamma, \Gamma]) = \sup \mathbb{P}(X_T \geq C_T | X_{t_0} = x, C_{t_0} = c)$ of the n-dimensional problem, with price process dynamics given by (12), can be expressed in terms of the value function of our original one-dimensional problem $V(\frac{x}{c}, t_0; \sigma, \mu, \rho)$, with price dynamics (5),*

$$J_n(u, t_0; \sigma, \kappa, \mu, [\gamma, \Gamma]) = V\left(ue^{(T-t)(\sigma\gamma^T G^{-1}\mu - \kappa)}, t_0; \tilde{\sigma}, \tilde{\mu}, 0\right),$$

where

$$\tilde{\sigma}^2 = \sigma^2\left(1 - \gamma^T G^{-1}\gamma\right), \quad \tilde{\mu} = \left(\nu^T G^{-1}\nu\right)^{\frac{1}{2}},$$

and

$$G = \eta\Pi\Pi^T\eta^T = \Gamma\Gamma^T + \gamma\gamma^T, \quad \nu = \mu - \sigma\gamma.$$

There is a similar expression for the optimal policy,

$$\chi_t^*(u, t; \sigma, \kappa, \mu, [\gamma, \Gamma]) \\ = (\nu G^{-1}\nu)^{-1/2}\left(\phi^*(ue^{(T-t)\left(\sigma\gamma^T G^{-1}\mu - \kappa\right)}, t; \tilde{\sigma}, \tilde{\mu}, 0)G^{-1}\nu\right) + \sigma G^{-1}\gamma.$$

By assumption the traded assets are linearly independent so that $G = \eta\Pi\Pi^T\eta^T$ is non-singular and G is invertible as the theorem requires.

To prove the theorem we use the following lemma which follows directly from the Cauchy-Schwartz Inequality.

Lemma 1 *For $x, y \in \mathbb{R}^d$ and Λ a positive definite symmetric $d \times d$ matrix we have*

$$x^T y \le \left(x^T\Lambda x\right)^{\frac{1}{2}}\left(y^T\Lambda^{-1}y\right)^{\frac{1}{2}},$$

with equality when $x = c\Lambda^{-1}y$ for some $c > 0$.

Proof of Theorem 2. We prove the theorem when the current time is zero. Applying Ito's Formula to $\tilde{U} = X/C$, then analogously to (7),

$$\frac{d\tilde{U}_t}{\tilde{U}_t} = \sqrt{(\chi_t^T\gamma - \sigma)^2 + \chi_t^T\Gamma\Gamma^T\chi_t}\,d\beta_t + (\chi_t^T(\mu - \sigma\gamma) + \sigma^2 - \kappa)dt, \qquad (17)$$

for β a Brownian motion.

Now if we take $\tilde{\chi}_t = \chi_t - \sigma G^{-1}\gamma$ we have

$$\tilde{\chi}_t^T G\tilde{\chi}_t + \sigma^2(1 - \gamma^T G^{-1}\gamma) = \chi_t^T\Gamma\Gamma^T\chi_t + (\chi_t^T\gamma - \sigma)^2,$$

and so (17) becomes

$$\frac{d\tilde{U}_t}{\tilde{U}_t} = \sqrt{\tilde{\sigma}^2 + \tilde{\chi}_t^T G\tilde{\chi}}\,d\beta_t + (\tilde{\chi}_t^T\nu + \tilde{h})dt,$$

where

$$\tilde{h} = \tilde{\sigma}^2 + \sigma\gamma^T G^{-1}\mu - \kappa.$$

Now define ϕ by $\phi_t^2 = \tilde{\chi}_t^T G\tilde{\chi}_t$ and $U = U^\phi$ via $U_0 = x/c$ and

$$\frac{dU_t^\phi}{U_t^\phi} = \sqrt{\tilde{\sigma}^2 + \phi_t^2}\,d\beta_t + (\phi_t\tilde{\mu} + \tilde{h})dt.$$

We have $\tilde{\chi}_t^T\nu \le \phi_t\tilde{\mu}$ using the first part of Lemma 1. Consequently $\tilde{U}^\chi \le U^\phi$ using a stochastic comparison theorem. Further if we take

$$\tilde{\chi}_t^* = (\nu^T G^{-1}\nu)^{-1/2}\left(\phi_t^* G^{-1}\nu\right),$$

then, using the second part of Lemma 1, $\tilde{\chi}^*$ is optimal. Hence

$$J_n(x/c, 0; \sigma, \kappa, \mu, [\gamma, \Gamma]) = \sup_\phi \mathbb{E}[U_T^\phi \geq 1 | X_0/C_0 = x/c]$$

If we define

$$\hat{U}_t = U_t^\phi e^{(T-t)\left(\sigma\gamma^T G^{-1}\mu - \kappa\right)},$$

then we find

$$\frac{d\hat{U}_t}{\hat{U}_t} = \sqrt{\tilde{\sigma}^2 + \phi^2} d\beta_t + \left(\phi_t \tilde{\mu} + \tilde{\sigma}^2\right) dt.$$

This is the same as (7), but with $\rho = 0$. Hence

$$\sup_\phi \mathbb{E}[U_T^\phi \geq 1 | X_0/C_0 = u] = V\left(ue^{T\left(\sigma\gamma^T G^{-1}\mu - \kappa\right)}, 0; \tilde{\sigma}, \tilde{\mu}, 0\right).$$

We can also read off the optimal strategy in the same fashion. □

This result is useful even in the one-dimensional case.

Corollary 1 *The value function and optimal policy in the imperfectly correlated case can be expressed in terms of those for the case where the claim is fully independent of the tradeable asset,*

$$V(u, t_0; \sigma, \mu, \rho) = V(ue^{(T-t_0)\sigma\mu\rho}, t_0; \sigma\sqrt{1-\rho^2}, \mu - \rho\sigma, 0),$$

$$\phi^*(u, t_0; \sigma, \mu, \rho) = \phi^*(ue^{(T-t_0)\sigma\mu\rho}, t_0; \sigma\sqrt{1-\rho^2}, \mu - \rho\sigma, 0) + \sigma\rho.$$

Proof. Take $n = 1$ and $\gamma = \rho$, $\Gamma = (\sqrt{1-\rho^2})$ in Theorem 2. □

We note that ρ near 1 corresponds to transformed claim volatility small, i.e. transformed claim nearly constant, which corresponds intuitively to a nearly complete market.

It should be noted that, at the expense of a more complicated notation, each of the results Theorem 2 and Corollary 1 can be extended to the case where the parameters are non-constant but deterministic. See Penn (2003), and Section 9 below.

5 Related Problems

In this section we describe two problems that are related to our perfect hedging problem.

The first problem we consider is that of minimising the expected shortfall in a hedge, $\mathbb{E}[(C_T - X_T)^+]$. This problem is considered in Föllmer & Leukert (2000) (within the framework of Föllmer & Leukert (1999)) and also by Schulmerich & Trautmann (2003) who use a linear programming approach. We shall show that in our context the expected shortfall problem reduces to our problem of maximising the probability of a perfect hedge under modified dynamics.

Suppose that, as before, the contingent claim and the traded-asset price process are given by geometric Brownian motions,

$$\frac{dC_t}{C_t} = \sigma dW_t, \quad C_0 = c_0 \tag{18}$$

$$\frac{dP_t}{P_t} = \rho dW_t + \sqrt{1-\rho^2}\, dB_t + \mu dt, \tag{19}$$

with a strategy ϕ_t chosen by the agent determining how the wealth evolves,

$$\frac{dX_t^\phi}{X_t^\phi} = \phi_t \frac{dP_t}{P_t}, \quad X_0 = x_0. \tag{20}$$

The minimal expected shortfall is given by $V_E(x,c,t) = V_E(x,c,t;\sigma,\mu,\rho)$.

Theorem 3 *For the model dynamics (18)-(19) we have*

$$V_E(x,c,t_0) = c\left(1 - V\left(\frac{x}{c}e^{-(T-t_0)\sigma^2}, t_0; \sigma, \mu+\rho\sigma, \rho\right)\right) \tag{21}$$

$$= c\left(1 - V\left(\frac{x}{c}e^{(T-t_0)(\sigma\rho\mu-(1-\rho^2)\sigma^2)}, t_0; \sigma\sqrt{1-\rho^2}, \mu, 0\right)\right). \tag{22}$$

Proof. Note that

$$\mathbb{E}[(C_T - X_T)^+] = \mathbb{E}[C_T(1-U_T)^+] = \mathbb{E}[C_T] - \mathbb{E}[C_T(U_T \wedge 1)].$$

Let $\mathbb{Q}$ be defined via $d\mathbb{Q}/d\mathbb{P} = C_T/\mathbb{E}[C_T]$. Then the problem of minimising $\mathbb{E}[(C_T - X_T)^+]$ over strategies ϕ is equivalent to maximising

$$\mathbb{E}^{\mathbb{Q}}[U_T \wedge 1].$$

Hence the expected shortfall problem reduces to the quantile hedging problem with modified dynamics.

For full details see Penn (2003). □

The second related problem is concerned with optimal portfolio choice. The aim in Karatzas (2002) is to find the portfolio which maximises the probability of reaching unit wealth, before bankruptcy and before the horizon time T. (We should note though that Karatzas (2002) imposed a no-short-sales constraint, $\phi \geq 0$, on his portfolio choices which we will ignore.) The problem becomes to find

$$V_K(x,t) = \sup_{\phi_s : t \leq s \leq T} \mathbb{P}\left(X_s \geq 1 \text{ for some } s \in [t,T] | X_t = x\right).$$

The boundary conditions,

$$V_K(0,t) = 0, \quad V_K(1,t) = 1, \quad V_K(x,T) = \mathbb{1}_{\{x \geq 1\}},$$

are modifications of the conditions in our perfect hedging problem. This problem could be solved using the ideas of this paper. Note however that the no-short-sales constraint makes a big difference to the solution of the problem since it prevents the construction of the policy required to prove the payoff equivalence of Theorem 1.

6 A Complete Market

We return to our original one-dimensional perfect-hedging problem (with price process dynamics given by (5)) and first consider the complete market case. By comparison with the earlier results we either know or are able to conjecture the optimal solution. We can then obtain the corresponding value function and show that it satisfies our partial differential equation.

Consider the case of a constant claim. This is the case studied by Kulldorff (1993), see Theorem 6, and Heath (1993). See also Browne (1999*b*), Remark 3.2.

Proposition 1 *For a constant claim but a tradeable asset with non-zero drift we have, for $t < T$, the value function*

$$V(u,t;0,\mu,0) = \Phi\left(\Phi^{-1}(u) + |\mu|\sqrt{T-t}\right). \tag{23}$$

Proof. Without loss of generality we may assume $t = 0$. Recall that $\tilde{\mathbb{P}}$ is the minimal martingale measure and that $Z = d\tilde{\mathbb{P}}/d\mathbb{P}$. Suppose we aim to obtain

$$X_T^* = c\mathbb{1}_{\{Z_T \le z^*\}}, \tag{24}$$

where z^* is such that the budget constraint,

$$\mathbb{E}[Z_T X_T | \mathcal{F}_t] = x,$$

is satisfied. Now

$$Z_s = \exp\left\{-\mu\beta_s - \frac{1}{2}\mu^2 s\right\},$$

where $d\beta = \rho dW + \sqrt{1-\rho^2}dB$ defines a standard Brownian motion. So,

$$\begin{aligned} \mathbb{P}(Z_T \le y) &= \mathbb{P}\left(\exp\left\{-\mu\beta_T - \frac{1}{2}\mu^2 T\right\} \le y\right) \\ &= \Phi\left(\frac{\log y + \frac{1}{2}\mu^2 T}{|\mu|\sqrt{T}}\right), \end{aligned} \tag{25}$$

and, the budget constraint, becomes

$$c\int_{-\infty}^{z^*} \frac{y}{|\mu|\sqrt{T}} \frac{1}{y\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\frac{\left(\log y + \frac{1}{2}\mu^2 T\right)^2}{\mu^2 T}\right\} dy = x,$$

With a few lines of algebra we can rearrange this to give

$$\Phi\left(\frac{\log z^* - \frac{1}{2}\mu^2 T}{|\mu|\sqrt{T}}\right) = \frac{x}{c}.$$

Hence, we find using (25) that for the target wealth given in (24)

$$\mathbb{P}(X_T \geq c) = \mathbb{P}(Z_T \leq z^*) = \Phi\left(|\mu|\sqrt{T} + \Phi^{-1}(u)\right). \tag{26}$$

Direct calculation now shows that this satisfies (9). To show that this gives the value function it remains to show that it satisfies the boundary condition at $t = T$. Substituting $t = T$ into (23) gives

$$\Phi\left(\Phi^{-1}(u)\right) = u \wedge 1.$$

However, Theorem 1 shows that for any time $t < T$ the value function for the problem with terminal payoff $V(u,T) = u \wedge 1$ is the same as that for the problem with terminal payoff $V(u,T) = \mathbb{1}_{\{u \leq 1\}}$. Consequently we do in fact have the correct value function. □

We can obtain the optimal policy from the value function,

$$\phi_t^* = \frac{V'\mu}{V''u} = \frac{\Phi'(\Phi^{-1}(u))}{\sigma\sqrt{T-t}}.$$

We note that, surprisingly, the optimal policy is independent of μ. This observation was also made by Browne (1999*b*).

In a complete market we can combine Proposition 1 with Corollary 1 to deduce a corresponding result.

Proposition 2 *In the complete market case with non-zero claim volatility and non-zero asset drift we have the value function*

$$V(u,t;\sigma,\mu,1) = \Phi\left(\Phi^{-1}(ue^{\mu\sigma(T-t)}) + |\mu - \sigma|\sqrt{T-t}\right). \tag{27}$$

Proposition 2 shows that Theorem 2 reduces the complete market problem to a case with independent assets.

A plot of the value function (27) is given in Figure 1. We note the flat region of the plot where success is almost guaranteed. This region is wider when there is more time to go.

7 An Independent Claim and a Zero-drift Asset

We have seen in Corollary 1 that we can focus on the case when the claim is independent of the tradeable asset. If also there is no asset drift then we can again solve for the value function. In this case there is no asset drift that a trading strategy might take advantage of, and it is best to put off investing in risky stock for as long as possible.

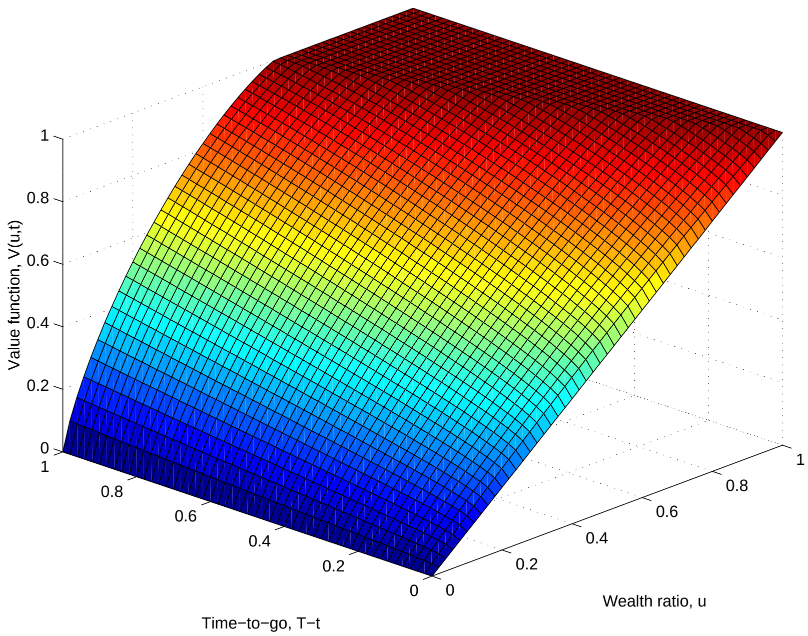

Figure 1: The Value Function in the Complete-Market Case when $\mu = 0.1$, $\sigma = 1$. Note that for all t, $V(u, t) = 1$ for u greater than some time-dependent threshold value.

Proposition 3 *When the tradeable asset is zero-drift and the claim is independent of it the value function is given by*

$$V(u,t) = \Phi\left(\frac{\log u + \frac{1}{2}\sigma^2(T-t)}{\sigma\sqrt{T-t}}\right) + \left(1 - \Phi\left(\frac{\log u + \frac{3}{2}\sigma^2(T-t)}{\sigma\sqrt{T-t}}\right)\right) u e^{\sigma^2(T-t)}. \tag{28}$$

Proof. Suppose, as we suggested above, we defer all trading until just before T then, if $x_0 < C_T$, trade so as to have wealth C_T with probability $\frac{x_0}{C_T}$ and have zero otherwise. For this strategy

$$\begin{aligned}
\mathbb{P}(X_T \geq C_T | C_t) &= \int_0^\infty \mathbb{P}(C_T \in dy)\left(1 \wedge \frac{x_0}{y}\right) \\
&= \Phi\left(\frac{\log u + \frac{1}{2}\sigma^2(T-t)}{\sigma\sqrt{T-t}}\right) \\
&\quad + \int_{x_0}^\infty \frac{x_0}{y^2\sigma\sqrt{T-t}}\Phi'\left(\frac{\log\frac{y}{c} + \frac{1}{2}\sigma^2(T-t)}{\sigma\sqrt{T-t}}\right) dy,
\end{aligned}$$

which gives (28) on making the substitution $y = ce^r$.

This can be confirmed to satisfy (9). Substituting $t = T$, gives $u \wedge 1$ and so, using Theorem 1, the result follows. □

A plot of (28) is given in Figure 2. We note that although there is a relatively flat region for large wealth ratio it is not as clearly demarcated as in the complete market case, (27). For the last small time interval before the horizon, there is a sharp increase in the success probability for wealth ratio near 1, as it becomes clear that the chance of the claim escaping from reach in the time left is small.

8 A Numerical Approach

Our main purpose in this section is to develop numerical methods to solve the problem with zero correlation between the traded asset and the claim in the case where the traded asset has drift.

First we consider the zero-drift case for which we already have explicit formulae. The HJB equation in this case is

$$V_u u + \frac{1}{2}V_{uu}u^2 - \frac{1}{\sigma^2}\dot{V} = 0.$$

In order to remove the dependence on the wealth ratio, u, in the coefficients of this equation we make the transformation $Y(w,t) = V(e^w, t)$ which gives

$$Y_w + Y_{ww} - \frac{2}{\sigma^2}\dot{Y} = 0. \tag{29}$$

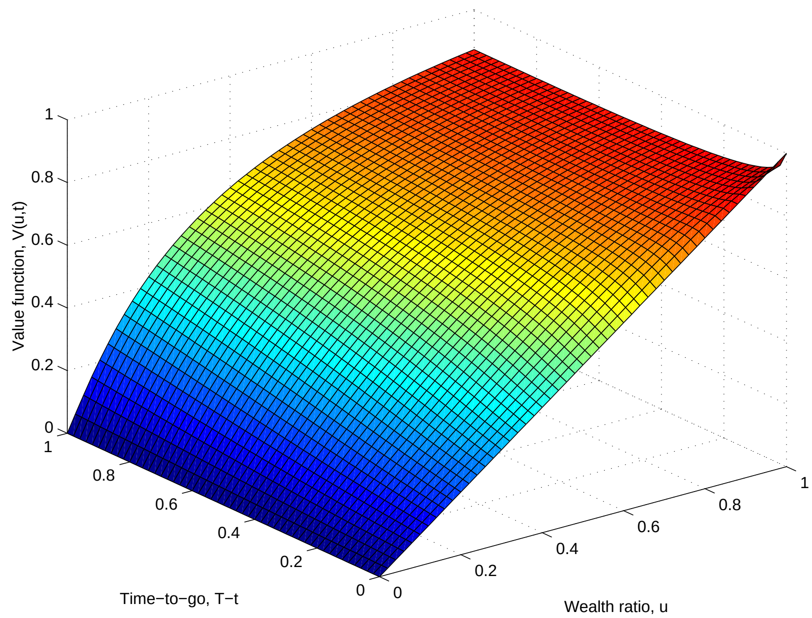

Figure 2: The Value Function in the Continuous-time Independent Asset Zero-drift Case. Note that for all u and t, $V(u,t) < 1$.

The boundary conditions become $Y(w,T) = e^w \wedge 1$, $Y(w,t) \to 1$ as $w \to \infty$ and $Y(w,t) \to 0$ as $w \to -\infty$.

We solve this problem numerically using a Crank-Nicholson finite difference scheme. In this simple case we can use simple boundary conditions: on the small-wealth-ratio boundary we use the approximation $Y(w_{\min}, t) = 0$, and on the large-wealth-ratio boundary we use the success probability from the strategy of doing nothing,

$$Y(w_{\max}, t) = \Phi\left(\frac{w + \frac{1}{2}\sigma^2(T-t)}{\sigma\sqrt{T-t}}\right).$$

For suitable parameter values the error in the numerical calculation of the super-replication probability is found to be of the order of 3×10^{-5}. Analysis of this simple problem for which we know the exact answer allows us to choose appropriate mesh sizes and grid boundaries.

Turn now to the case of an asset with non-zero drift. The HJB equation for the value function here is

$$V_u u + \frac{1}{2}V_{uu}u^2 - \frac{\mu^2 V_u^2}{2\sigma^2 V_{uu}} - \frac{1}{\sigma^2}\dot{V} = 0, \tag{30}$$

and under the transformation $Y(w,t) = V(e^w,t)$ this becomes

$$Y_w + Y_{ww} - \frac{\mu^2}{\sigma^2}\frac{Y_w^2}{(Y_{ww} - Y_w)} - \frac{2}{\sigma^2}\dot{Y} = 0. \qquad (31)$$

This time the discretisation of the pde gives a non-linear system of simultaneous equations.

Recall that if the optimal policy is ϕ^* then the transformed value function satisfies

$$\frac{1}{2}Y_{ww}(\phi^{*2} + \sigma^2) + Y_w\left(\mu\phi^* + \frac{1}{2}\sigma^2 - \frac{1}{2}\phi^{*2}\right) - \dot{Y} = 0, \qquad (32)$$

(substituting

$$\phi^* = \frac{Y_w \mu}{Y_w - Y_{ww}}, \qquad (33)$$

recovers (31)). Now (32) is linear and so given any policy ϕ we can, numerically, obtain the success probability corresponding to using that policy. This allows us to use a policy improvement scheme. Given an initial policy, ϕ^0, we can obtain its success probability, Y^0, then using a discretised form of (33), we can use this to obtain an improved policy, ϕ^1 and so on. At each stage of the policy improvement scheme the success probability is increased. Consequently the approximations to the optimal value function function converge pointwise. This is reflected in the numerical calculations.

For the problem with drift it turns out that great care is needed to calculate the appropriate boundary conditions, particularly when the wealth ratio is small. Although the pde in (32) is fairly standard, finding the solution is a delicate matter. Fortunately, the primal policy improvement method gives a lower bound on the value function, and if we solve the dual problem (by a similar method) then we get an upper bound (see Penn (2003)). In this way we can be confident about the accuracy of our solutions. (The difference between our numerical solutions of the primal and dual problems is about 5×10^{-5}. In contrast, both solutions give value improvements of the order 5×10^{-3}. Hence the relative error of the numerical solution is at most 1%.)

Figure 3 gives a plot of this value improvement in the case where $\sigma = 1$ and $\mu = 0.1$. The value improvement is small, it decreases to zero for both large and small wealth ratio, and it decreases when there is little time to go. Furthermore, for fixed $w = \log u$ the value improvement is monotonic in time.

The optimal policy, ϕ^* is shown in Figure 4. There is some holding in stock throughout time, though, as we would expect, more near the horizon time. Note that in the plot the optimal policy we calculate is not decreasing in w for small time-to-go. It seems likely that this is a boundary effect caused by the use of a finite grid, rather than a true feature of the problem.

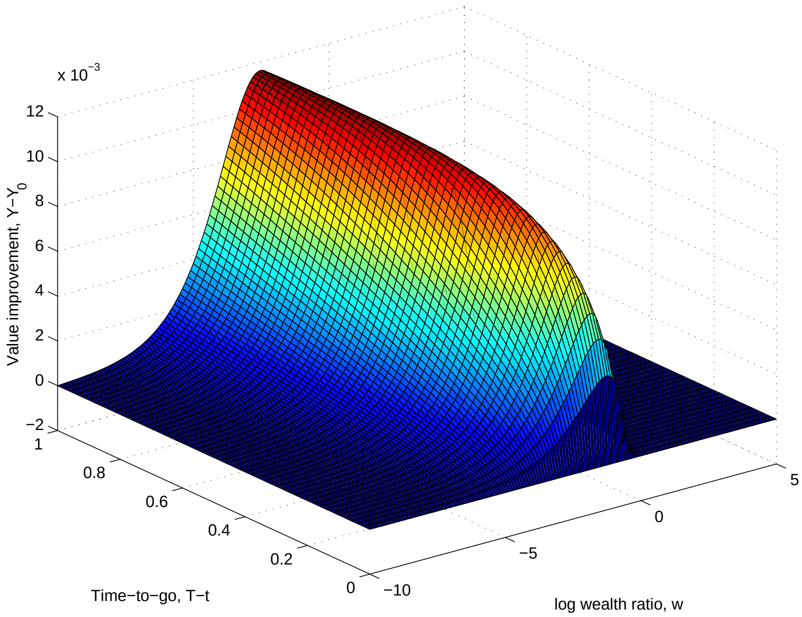

Figure 3: The Optimal Value Improvement when $\sigma = 1$ and $\mu = 0.1$.

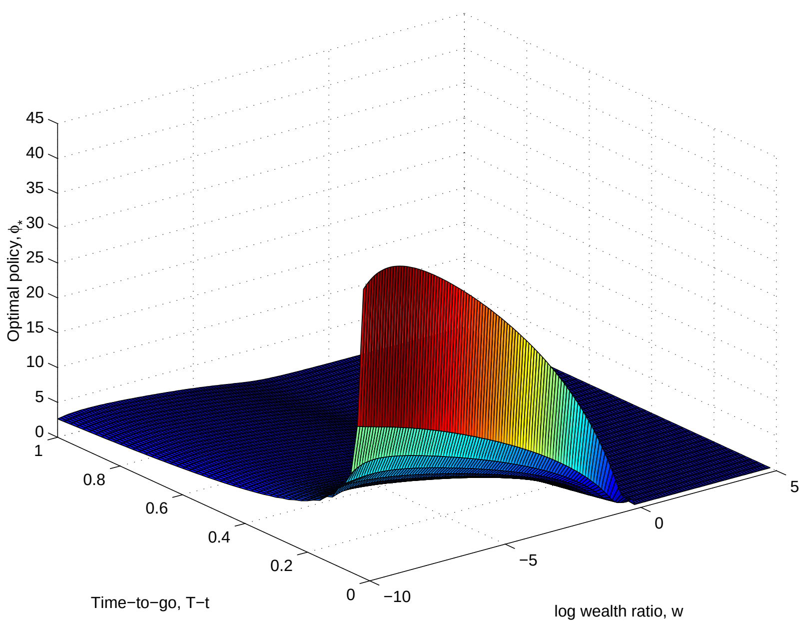

Figure 4: The Optimal Policy when $\sigma = 1$ and $\mu = 0.1$.

9 The Value of Timely Information

In Section 2 we considered the quantile hedging problem in the case where the contingent claim was independent of the traded asset, and its value was revealed at time T. In this section we investigate the impact of the time at which information about the claim is revealed on the probability of successfully super-replicating the claim.

In order to ensure a fair comparison, we will always take the claim to have the same distribution. In particular let C have a lognormal distribution given by

$$C = c_0 e^{\sigma\sqrt{T}G - \frac{1}{2}\sigma^2 T},$$

where G is a $N(0,1)$ random variable which is independent of the stochastic process driving the traded asset. This distribution is equal to that used in the Sections 3 and 5-8, at least when $\rho = 0$. The cumulative distribution function of the claim is

$$F_C(y) = \Phi\left(\frac{\log\frac{y}{c_0} + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}\right).$$

Note that this function is not concave. Denote the sigma-algebra generated by the random variable G by $\sigma(G)$. We will only use this notation in the following paragraph, so that the confusion with the volatility σ should be minimal.

Recall that (in the zero correlation case) B is the Brownian motion generating the price process of the traded asset, and let $\{\mathcal{F}_t^B\}_{0\le t\le T}$ be the filtration generated by this asset. The problem is to maximise the probability of a perfect hedge when the trading strategy is taken to lie in one of several different classes. If ϕ_t is adapted to $\mathcal{G}_t^0 = \mathcal{F}_t^B \vee \sigma(G)$ then we say the claim is revealed at time 0, since for all $t \ge 0$ the strategy can depend on the value of G. More generally define $\mathcal{G}_t^\gamma = \mathcal{F}_t^B$ for $t < \gamma$ and $\mathcal{G}_t^\gamma = \mathcal{F}_t^B \vee \sigma(G)$ for $t \ge \gamma$. In saying that the claim is revealed at time γ we mean that ϕ must be chosen adapted to $\{\mathcal{G}_t^\gamma\}_{0\le t\le T}$. Now we describe what we mean by saying that the claim is revealed continuously. Let W be a second Brownian motion independent of B, and identify W_T with $\sqrt{T}G$. Let $\mathcal{F}^W$ denote the filtration generated by W. If $\mathcal{G}_t^{cts} = \mathcal{F}_t^B \vee \mathcal{F}_t^W$ and ϕ must be chosen to be adapted to $\{\mathcal{G}_t^{cts}\}_{0\le t\le T}$ then we say we are in the continuously revealed claim case.

Define $V^{(\gamma)}(u)$ to be the probability of succesfully superreplicating the claim given initial wealth $x_0 = uc_0$, and given that the claim is revealed at time γ. (A superscript in this section will denote dependence on the time at which the claim is revealed.) First we consider the case when the claim is revealed at time T, after the finish of the period when trading is possible.

Set

$$k(\xi, T) = \sqrt{2\sigma^2 T - 2\log\left(\xi c_0 \sigma\sqrt{2\pi T}\right)}.$$

Obtaining the concave relaxation of F_C and inverting gives, after some simplification,

$$I(\xi) = \mathbb{1}_{\{\xi<\tilde{\xi}\}} c_0 e^{-\frac{3}{2}\sigma^2 T} \exp\left\{\sigma\sqrt{T}\, k(\xi, T)\right\},$$

where $\tilde{\xi}$ is the solution to

$$\Phi\left(k(\tilde{\xi}, T) - \sigma\sqrt{T}\right) = \frac{1}{\sigma\sqrt{T}}\Phi'\left(k(\tilde{\xi}, T) - \sigma\sqrt{T}\right).$$

Proceeding as in Section 2, we find that the Lagrange multiplier $\lambda = \lambda^{(T)}$ satisfies

$$\mathbb{E}[e^{-\mu B_T - \frac{1}{2}\mu^2 T} I(\lambda^{(T)} e^{-\mu B_T - \frac{1}{2}\mu^2 T})] = x_0.$$

We can evaluate such expectations by numerical quadrature and solve for λ by binary search. Finally,

$$V^{(T)}(u_0) = \mathbb{P}\left(I\left(\lambda^{(T)} e^{-\mu B_T - \frac{1}{2}\mu^2 T}\right) \geq c_0 e^{\sigma\sqrt{T}G - \frac{1}{2}\sigma^2 T}\right),$$

(recall that the function I depends on x_0, c_0 and T) which again can be evaluated by quadrature.

Now consider the case where the contingent claim is still random, but is revealed just after the outset, at time 0+. Proposition 1 tells us that given a claim value of $C = c$, the optimal success probability is

$$\sup_{\phi_t : 0 \leq t \leq T} \mathbb{P}\left(X_T^\phi \geq C_T \middle| X_0 = x_0,\, C = c\right) = \Phi\left(\Phi^{-1}\left(\frac{x_0}{c}\right) + |\mu|\sqrt{T}\right).$$

and that averaging over the possible values of c, the success probability is

$$V^{(0+)}\left(\frac{x_0}{c_0}\right) = \mathbb{E}\left[\Phi\left(\Phi^{-1}\left(\frac{x_0}{c_0} e^{-\sigma\sqrt{T}G + \frac{1}{2}\sigma^2 T}\right) + |\mu|\sqrt{T}\right)\right].$$

It follows that

$$V^{(0+)}\left(\frac{x_0}{c_0}\right) = \int_0^1 \left(1 - \Phi\left(\frac{\log\left(\Phi\left(\Phi^{-1}(y) - |\mu|T\right)\right) - \log\frac{x_0}{c_0} - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}\right)\right) dy,$$

which we can evaluate by numerical quadrature.

Next suppose that the contingent claim is revealed just before maturity, that is at time T^-. Unlike the case where the claim is revealed at maturity, time T, we are able to trade in response to the claim, if only for an instant. Consequently the result of Theorem 1 applies. So we wish to maximise

$$\sup_{X_T^\phi} \mathbb{E}\left[\frac{X_T^\phi}{C_T} \wedge 1\right],$$

subject to the budget constraint $\tilde{\mathbb{E}}\left[X_T^\phi\right] = x_0$.

Time at which claim is revealed	Success probability	
	$\ln x_0 = -0.5$	$\ln x_0 = -0.1$
0+	0.7772	0.8647
$T-$	0.7624	0.8546
T	0.5001	0.6555
continuously	0.7657	0.8566

Table 1: The dependence of success probability on the timing of information.

Now we can take a similar approach to that in Section 2. We find that

$$V^{(T-)}\left(\frac{x_0}{c_0}\right) = \mathbb{E}\left[I\left(\lambda^{(T-)}e^{-\mu B_T - \frac{1}{2}\mu^2 T}\right)e^{-\sigma W_T + \frac{1}{2}\sigma^2 T} \wedge 1\right]$$

where

$$I(\xi) = c_0 \exp\left\{\sigma\sqrt{T}\Phi^{-1}(1 - c_0\xi) - \frac{3}{2}\sigma^2 T\right\}\mathbb{1}_{\left\{\xi \le \frac{1}{c_0}\right\}}.$$

We can now find $\lambda = \lambda^{(T-)}$ by a binary search and then evaluate $V^{(T-)}$ by numerical quadrature.

Table 1 shows how the success probability depends on the timing of the claim. The parameter values are $c_0 = 1$, $\sigma = 1$, $T = 1$, $\mu = 0.1$ and initial wealths $x_0 = e^{-0.5}$ and $x_0 = e^{-0.1}$. (These values have been chosen to give success probabilities which are not too close to 0 or 1.) The value in the continuous case is as constructed from the numerical solution of the HJB equation in the previous section.

The first striking observation is that the success probability when the claim is revealed after the last opportunity for trading is an order of magnitude smaller than the success probability when there is still time to trade after the claim is revealed. Thus there is a crucial advantage in knowing the value of the claim, and being able to gamble on an all or nothing basis at the close of the trading period.

The second observation is that, as is to be expected, the sooner the claim is revealed the higher the success probability. In other words probabilities are increasing in the available information. Note also that the improvement in success probability when comparing the continuously revealed claim with the claim revealed at $T-$ is approximately 20% of the improvement in success probability when comparing claims revealed at 0+ wih $T-$. This figure of 1 in 5 varies only slightly with the value of initial wealth.

Note that it is also possible to calculate the success probability when the claim is revealed at an intermediate time γ. The results for this calculation are given in Figure 5. We see that the sooner the claim is revealed the greater the success probability, but that the advantage of seeing the claim sooner is strongest when the claim is revealed at a time near $T-$.

Figure 6 shows how the success probability depends on μ, in each of the three cases where the claim is revealed at the outset, continuously and just

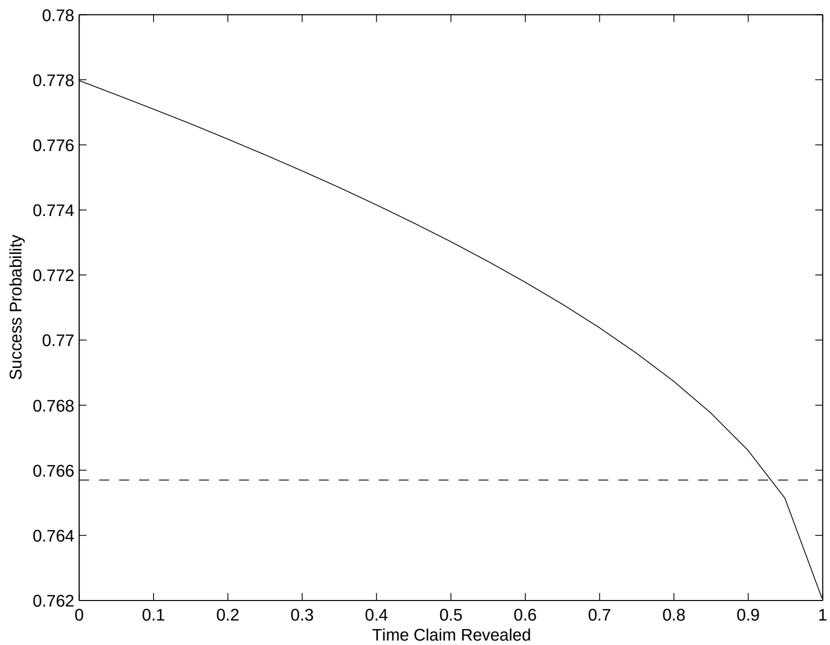

Figure 5: Success Probabilities as a function of the time at which the claim is revealed, when $x_0 = e^{-0.1}$ and $c_0 = 1$. The dashed horizontal line shows the success probability in the case when information is revealed continuously through time.

before the maturity. As expected, as the information is delayed, the success probability decreases, and as the drift μ increases the success probability increases. In each case, as μ decreases to zero, the three success probabilities converge to the same value (modulo a small numerical error). Rather suprisingly, when the claim is revealed at time 0, the success probability is almost exactly linear in μ, but the other two curves display convexity in μ. This near linearity results form the fact that the curve represents the expected value of a function which is part convex and part concave in μ, and for the parameter values we have chosen the effects from these regions cancel.

10 The timing of volatility

In the above analysis we have considered the case where either the final value of the claim is revealed continuously through time at a uniform rate, or the claim value is revealed at a single instant. The market was incomplete, and the claim was uncorrelated with the traded asset. Now we briefly consider the complete market case.

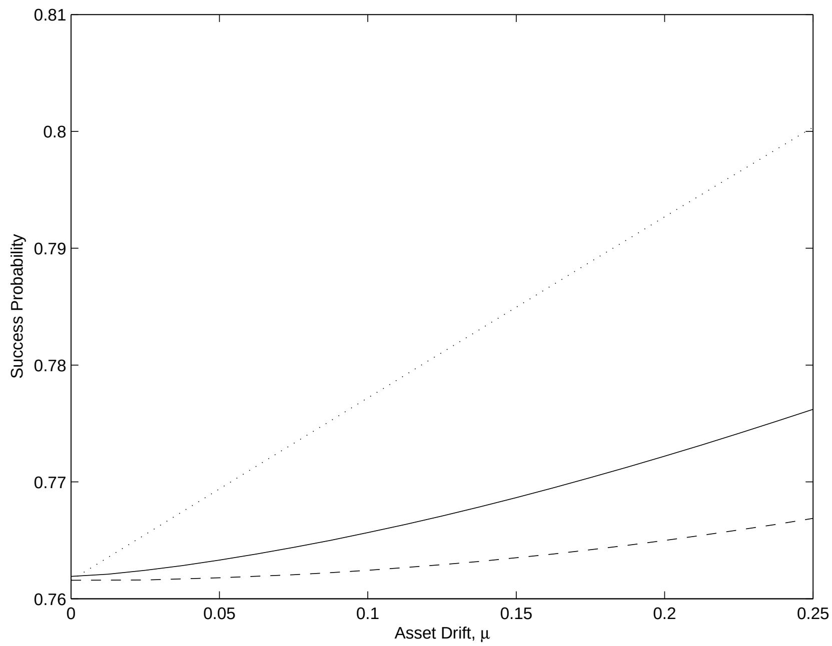

Figure 6: Success Probabilities as a function of the drift of the traded asset. Here $x_0 = e^{-0.5}$; other parameters as in the text. The three lines correspond to (starting with the highest) a claim revealed at time $0+$, a claim revealed continuously through time, and a claim revealed at $T-$.

For a fixed total claim variance we consider how the value function in the continuously-revealed case depends on the timing of volatility. If we have time-dependent volatility (σ_t) and constant asset drift μ the success probability is

$$V(u,0) = \Phi\left(\Phi^{-1}\left(ue^{\mu\int_0^T \sigma_s\, ds}\right) + \sqrt{\int_0^T (\sigma_s - \mu)^2\ ds}\right).$$

(This is an easy extension of Proposition 2.)

Fix the total variance $\int_0^T \sigma_s^2\, ds = (\sigma^*)^2 T$, or equivalently the terminal distribution of the claim, and consider the success probability as a function of the integrated volatility $\Sigma = \int_0^T \sigma_s\, ds$. We wish to consider how the success probability

$$\tilde{V}(\Sigma) = \Phi\left(\Phi^{-1}\left(ue^{\mu\Sigma}\right) + \sqrt{(\sigma^*)^2 T + \mu^2 T - 2\mu\Sigma}\right).$$

depends on Σ, and in particular, to find the value of Σ which maximises this expression. The fact that we can write the success probability as a function of Σ shows that we are not concerned about the timing of volatility only about the total volatility, (recall that the integral of the squared volatility is fixed). To simplify the analysis we take $\sigma^* = T = 1$.

For given asset drift μ, we can find the value $\Sigma = \Sigma^*$ which maximises $\tilde{V}$. The optimal Σ^* is a function of the wealth ratio u. Figure 7 shows Σ^* for varying u and μ. The white regions correspond to where $\Sigma^* = 0$, the black regions correspond to $\Sigma^* = 1$. The grey regions correspond to intermediate values, $0 < \Sigma^* < 1$. Note that we obtain $\Sigma^* = 1$ when $\sigma \equiv 1$, whereas $\Sigma^* = 0$ arises in the limit when the volatility is concentrated at a single instant. Intermediate values correspond to volatilities which are neither fully concentrated, nor constant.

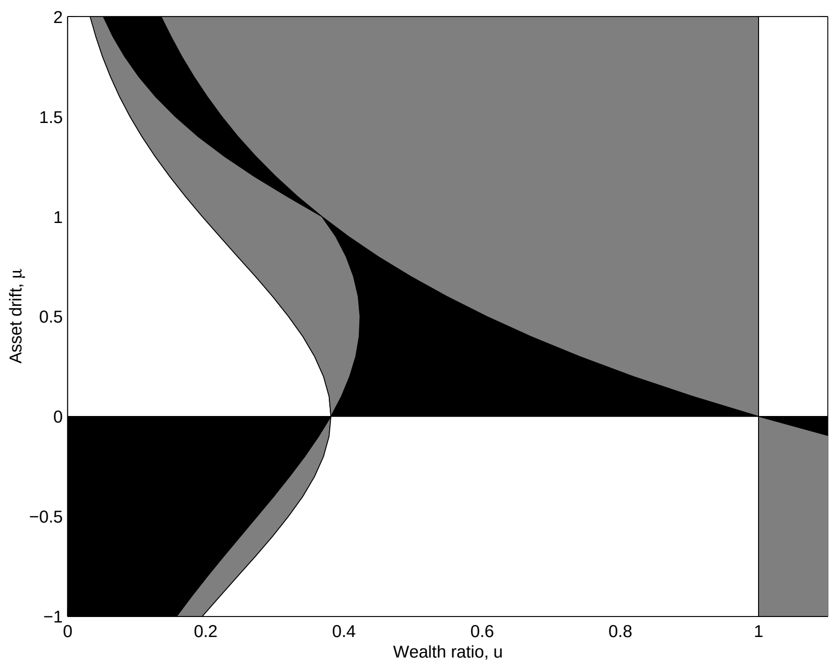

Figure 7: The Optimal Distribution of Volatility as a Function of Wealth Ratio and Asset Drift μ

Let us consider the economic interpretation of this picture. When $\mu = 1$ the contingent claim is a martingale under the pricing measure and so the expression for the success probability simplifies greatly. This is why we find that we have distinct behaviour for $\mu = 1$, in the sense that in that case it is never most beneficial to have constant volatility.

We notice that crossing the line $\mu = 0$ black regions change to white and vice versa. This makes sense as we change from having a long position in the stock to a short position or vice versa. Consequently we might expect our preference for constant ($\Sigma^* = 1$) or spikey ($\Sigma^* = 0$) volatility to change. On the other hand, as u varies the behaviour is smooth in that white changes to black via grey and vice versa.

As we cross the line $u = 1$, that is as we change from having less money than we need to meet the claim to having more, the behaviour changes. For positive μ, we change from intermediate distribution of volatility to taking volatility spikey. If we have more money than we need we want sudden jumps in the claim value that we can then respond to rather than constant change in the claim value. For negative μ the opposite is true.

We also see a change in behaviour as we cross the line $u = e^{-\mu}$. We saw in Proposition 2 that the factor $e^{\mu \int_t^T \sigma_s \, ds}$ reduces the general complete market case to the constant claim case so we would expect on crossing this line to change from taking our claim volatility constant to taking a more varied distribution of volatility.

It remains to consider the region where $\mu > 0$ and $u < e^{-\mu}$. For relatively small wealth we find that, as was the case when we had very large wealth, we prefer to have sudden jumps in the claim value which we respond to over time rather than having constant change. As u increases this effect diminishes, with the desired Σ^* changing smoothly to 1.

11 Conclusions and Extensions

In this paper we investigated the problem facing an agent who seeks to meet a contingent claim, but who has insufficient funds to finance a replicating portfolio (in a complete market) or super-hedging portfolio (in an incomplete market). In particular we suppose the agent aims to maximise the probability of meeting the claim — the so-called quantile hedging problem.

Our attention has focused on markets where both the traded asset and the claim are exponential Brownian motions. (The extension to jump processes is considered in Xu (2003).) There were several reasons for our choice including tractability, but also including realism and applicability. Moreover such a framework allows us to consider the impact of correlation, and of the timing of information.

The general problem with several traded assets can be reduced to a single traded asset. Further, in our context, other problems such as minimising expected shortfall, can also be recast into a quantile hedging problem with modified parameters. Hence the simple structure which is the focus of the research is a paradigm for a wider set of problems.

In a complete market, and in a market with zero correlation and traded assets which are martingales, we can find explicit analytic solutions for the value function. We can also show that the general problem in which there is correlation between the asset and the claim can be reduced to the uncorrelated case.

For the incomplete problem with a traded asset with a non-zero drift, we use numerical methods to obtain a solution for the value function. We use a method of policy improvement to solve a sequence of linear partial differential equations.

Of particular interest is a comparison of the optimal strategies. The optimal strategy in the case with (positive) drift is to hold a positive, state and time dependent, proportion of wealth in the risky asset. This proportion increases as the time to go decreases, or as the ratio of wealth to claim decreases. The case of zero drift can be thought of as a special case in which all the trading takes place at the last instant, and then only if the claim value exceeds current wealth.

We can also compare receiving information on the claim continuously through time to the case where the value of the claim is revealed instantaneously. Most strikingly, there is a significant increase in the success probability if the agent still has trading opportunities after the final value of the claim is revealed. Otherwise the dependence of the success probability on the timing of information is small, but the sooner the information about the claim is revealed, the higher the chance of super-replicating the claim.

References

Browne, S. (1999*a*), 'Beating a moving target: optimal portfolio strategies for outperforming a stochastic benchmark', *Finance and Stochastics* **3**, 275–294.

Browne, S. (1999*b*), 'Reaching goals by a deadline: Digital options and continuous-time active portfolio management', *Advances in Applied Probability* **31**, 551–577.

Föllmer, H. & Leukert, P. (1999), 'Quantile hedging', *Finance and Stochastics* **3**, 251–273.

Föllmer, H. & Leukert, P. (2000), 'Efficient hedging: Cost versus shortfall risk', *Finance and Stochastics* **4**, 117–146.

Heath, D. (1993), A continuous time version of Kulldorff's result. Unpublished Manuscript.

Karatzas, I. (2002), 'Probabilistic aspects of portfolio analysis and optimization', Second World Congress of Bachelier Finance Society.

Kulldorff, M. (1993), 'Optimal control of favourable games with a time limit', *SIAM Journal of Control and Optimization* **31**, 52–69.

Penn, J. W. V. (2003), 'Hedging: imperfect correlation, timely information and stochastic volatility', PhD Thesis, University of Bath.

Schulmerich, M. & Trautmann, S. (2003), 'Local expected shortfall-hedging in discrete time', *European Finance Review* **7**, 75–102.

Spivak, G. & Cvitanić, J. (1999), ‘Maximising the probability of a perfect hedge’, *Annals of Applied Probability* **9**, 1303–1328.

Xu, M. (2003), ‘Minimising shortfall risk using duality approach - an application to partial hedging in incomplete markets’, Preprint, Carnegie-Mellon University.