Optimal Investment Problems and Volatility Homogenization Approximations

Mattias Jonsson

Department of Mathematics
University of Michigan
Ann Arbor, MI 48109-1109.

Ronnie Sircar

Operations Research & Financial Engineering Department
Princeton University
Princeton, NJ 08544.

December 17, 2001

Contents

Abstract

We describe some stochastic control problems in financial engineering arising from the need to find investment strategies to optimize some goal. Typically, these problems are characterized by nonlinear Hamilton-Jacobi-Bellman partial differential equations, and often they can be reduced to linear PDEs with the Legendre transform of convex duality. One situation where this cannot be achieved is in a market with stochastic volatility. In this case, we discuss an approximation using asymptotic analysis in the limit of fast mean-reversion of the process driving volatility. Simulations illustrate that marginal improvement can be achieved with this approach even when volatility is not fluctuating that rapidly.

1 Introduction

The purpose of this article to introduce to students and researchers schooled in differential equations and scientific computing some problems involving optimal investment strategies that arise in financial engineering. These can often be characterized by nonlinear second-order Hamilton-Jacobi-Bellman (HJB) equations in three or more space dimensions, so that it is difficult to solve them numerically within the time-frame that the solution is required. We discuss an approximation method based on asymptotic analysis that is effective for models in which asset prices have randomly varying volatility. This method supposes that volatility is fluctuating rapidly with respect to the timescale of the optimization problem. The upshot of the singular perturbation analysis is to approximate the changing volatility by two averages when forecasting future expected performance. One homogenized volatility is the root-mean-square average that is common in financial calculations; the second is a harmonic average that arises from the particular structure of the HJB equations in this context. We show from

simulation that even when the volatility is not fluctuating so rapidly, as the analysis assumes, the homogenization approximations serve to improve average performance, albeit marginally, over the competing constant volatility theory.

1.1 Background

In the financial services industry, stochastic modeling of prices has long been recognized as crucial to everyday management of risk, particularly that associated with derivative instruments. By derivatives, we refer to contracts such as options whose payoffs depend in some way on the behaviour of the price of some underlying asset (stock, commodity, exchange rate, for example) over some period of time. The canonical example, a call option on, say, a stock, pays nothing if the stock price X_T on a specified *expiration date T* in the future is less than the specified *strike price K*, and the difference $X_T - K$ otherwise. It is completely characterized by its payoff function

$$h(X_T) = (X_T - K)^+,$$

where $x^+ = \max\{x, 0\}$, the positive-part function.

Such securities with potentially unbounded payouts (for example if X_T becomes very large in the case of a call) can bring huge and rapid returns on relatively small investments if the stock moves the right way for the holder. But they are often associated with huge losses (for example for the writer of the call option). Large trading organizations have been using (to varying degrees) mathematical models for hedging their derivatives risk, at least since the Black-Scholes-Merton (BSM) methodology introduced in 1973 [2, 15].

In this context, let us remark that in some mathematical market models, it is possible to replicate the payoff of any derivative by continuous trading in the underlying. Such markets, which include the BSM model, are called *complete* and are much easier to analyze than the incomplete ones. In this article we will be largely interested in stochastic volatility, a feature that renders the market incomplete.

A recurrent problem in managing risks associated with derivatives positions is to maximize over suitably defined trading strategies

$$I\!\!E\{U(X_T, V_T)\},$$

the expectation of a *utility function U* of the future value of a controlled random *wealth* (or portfolio value) process $(V_t)_{0 \le t \le T}$ where the utility also depends on the (possibly vector-valued) state variable (X_t) at the future time T. In the situations of interest here, (X_t) is an asset price, possibly driven by other stochastic variables like volatility, U is a concave function of V_T, and the state-dependence of the utility function arises from risk associated with a derivative security $h(X_T)$.

The control variable, denoted (π_t), represents the amount the investor has in the stock. The rest $V_t - \pi_t$ is in the bank earning interest at rate r. The wealth process (V_t) therefore evolves according to

$$dV_t = \frac{\pi_t}{X_t}\, dX_t + r(V_t - \pi_t)\, dt, \tag{1}$$

which reads that the change in the portfolio value is the number of stocks held times the change in the stock price, plus interest earned on the rest in the bank. We assume the initial

wealth v is given. Recent work on existence and uniqueness of optimal solutions under various models and constraints is summarized in [13].

   We first give some practical examples that can be described in this framework, and then describe one approach to this problem using dynamic programming.

1.2   Examples

In a complete market, there exists an amount $v^\star$ of initial capital and a strategy $(\pi_t^\star)$ such that

$$V_T = h(X_T) \qquad \text{with probability 1.}$$

In this case, $(\pi_t^\star)$ is called the replicating strategy for the derivative. Alternatively, it is the *hedging* strategy for a short position in the derivative since it is often used to cancel *exactly* (in theory) the risk of a short derivative position. The quantity $v^\star$ is called the *Black-Scholes price* of the derivative.

   The problems we describe here are concerned with the gap

$$h(X_T) - V_T$$

which may not be zero with probability 1 either because the initial capital $v < v^\star$ or because of market incompleteness.

1.2.1   Partial Hedging

Suppose the investor has written (sold) a derivative, but wants to insure or hedge against his future (uncertain) liability with a trading strategy in the stock. The problem is to find a strategy $(\pi_t)_{0 \le t \le T}$ such that, starting with initial capital v, the terminal value of the *hedging portfolio* V_T comes as close as possible to $h(X_T)$. We want the performance of the strategy to be penalized for falling short, with the actual size of the shortfall taken into account, but when the strategy overshoots the target, the size of the overshoot has no bearing on the measure of *risk*. This problem has been studied under various assumptions recently by Föllmer-Leukert [7] and Cvitanić & Karatzas [3], where existence and uniqueness results are established. In Section 7, we discuss *computation* of optimal strategies when volatility is changing randomly.

   If we are given enough initial capital v, we can make sure that the terminal wealth V_T exceeds $h(X_T)$ with probability 1. The smallest such v is called the *superhedging price* of the derivative $h(X_T)$. Clearly, the partial hedging problem is only interesting if v is strictly smaller than the superhedging price.

   For the explicit computations and examples here, the penalty function we use is

$$I\!E\{\frac{1}{p}((h(X_T) - V_T)^+)^p\},$$

the *one-sided* pth moment. The parameter $p > 1$ allows control of one's risk aversion. We will present simulations corresponding to the case $p = 1.1$ in Section 7. The limiting case $p = 1$ has a particular economic significance since the penalty function then corresponds to a coherent measure of risk (see [3]), but also leads to certain implementation difficulties. This is why we pick $p = 1.1$.

We also insist that $V_t \geq 0$ for all $0 \leq t \leq T$ almost surely for the strategy to be admissible. That is, the hedging portfolio is bounded below by zero, a portfolio constraint. The problem can be reformulated as a (state-dependent) utility maximization with

$$U(x, v) = \frac{1}{p} \left[h(x)^p - ((h(x) - v)^+)^p \right].$$

For fixed x, this is concave in v on $(0, \infty)$, strictly concave on $(0, h(x))$ and satisfies $U(x, 0) = 0$, $U(x, v) = \frac{1}{p} h(x)^p$ for $v \geq h(x)$, see Figure 3.

We call this problem *partial hedging* of the derivative risk. It is the main example that we pick to illustrate the asymptotic analysis in Section 6; it is also treated in [12].

1.2.2 Mean-Square Pricing

The problem is to find the initial wealth level such that the minimum variance

$$\inf_{(\pi_t)} I\!E\{(h(X_T) - V_T)^2\}$$

is minimized as a function of v. This is then defined to be the option price. In our framework, this corresponds to the utility function

$$U(x, v) = h(x)^2 - (h(x) - v)^2.$$

Such a quadratic state-dependent utility is considered in [5, 1], for example. In complete markets the expectation above is zero for the optimal v because the option can be replicated exactly with initial capital $v^\star$, the Black-Scholes price. The technique of mean-square pricing is most interesting in incomplete markets, such as markets with stochastic volatility.

1.2.3 Options Based Portfolio Insurance

Fund managers are interested in selling trading opportunities in which the clients can take advantage of a rising market but are insured against losses in the sense that the value of the fund is bounded below by a number that is either constant or depends on a future asset price. More precisely, the client is guaranteed that the portfolio value V_T at time $T > 0$ is at least $h(X_T)$, i.e. the value of an option with payoff $h(X_T) \geq 0$. This is studied in complete markets in [6], for example, and corresponds to the utility function

$$U(x, v) = \begin{cases} \tilde{U}(v) & \text{if } v \geq h(x) \\ -\infty & \text{otherwise} \end{cases}$$

where $\tilde{U}$ is an increasing, concave (state-independent) utility function.

1.2.4 Utility-Indifference Pricing

In this mechanism for pricing, described in [4] for example, the (selling) price of a derivative with payoff $h(X_T)$ is defined to be the extra initial compensation the seller would have to receive so that he/she is *indifferent with respect to maximum expected utility* between having the liability of the short derivative position or just trading in the stock.

Let
$$H(v) = \sup_{\pi} I\!E\{\tilde{U}(V_T - h(X_T)) \mid \text{initial capital is } v\},$$

where $\tilde{U}$ is an increasing concave utility function and we show only the v dependence of H. Here the seller can trade in the stock, but at time T has liability $-h(X_T)$. Thus the relevant state-dependent utility function is

$$U(x, v) = \tilde{U}(v - h(x)),$$

where $\tilde{U}$ is an increasing, concave (state-independent) utility function. Let

$$M(v) = \sup_{\pi} I\!E\{\tilde{U}(V_T) \mid \text{initial capital is } v\},$$

the solution to the same problem without the derivative position. Then the (seller's) price P of the derivative security is defined by

$$H(v + P) = M(v).$$

In a complete market, this mechanism recovers the Black-Scholes price $v^\star$ which does not depend on the wealth level v of the seller. The same desirable wealth-independent property can be attained in an incomplete market by assuming an exponential utility $\tilde{U}(v) = -e^{-\gamma v}$, where γ is a risk-aversion coefficient.

2 The Merton Problem

To begin our exposition of the methods for dealing with utility maximization problems, we begin with the simplest and original continuous-time stochastic control problem in finance, the Merton optimal asset allocation problem. The original reference is [14] and this is reprinted in the book [16]. Here we concentrate on maximizing expected utility at a final time horizon T, mainly because we are later interested in problems associated with derivative contracts that have a natural terminal time; other versions of the asset allocation problem consider optimizing utility of consumption over all times, possibly with an infinite horizon.

An excellent mathematical introduction to this type of stochastic control problem, as well as stochastic modelling and connections to differential equations, is the book by Øksendal [17].

2.1 The BSM Model

The Black-Scholes-Merton (BSM) model takes the price of the underlying asset (for example a stock) to have a deterministic growth component (measured by an average rate of return μ) and uncertainty or risk generated (for mathematical convenience) by a Brownian motion and quantified by a volatility parameter σ. The price X_t at time t satisfies the stochastic differential equation
$$\frac{dX_t}{X_t} = \mu \, dt + \sigma \, dW_t, \tag{2}$$

where (W_t) is a standard Brownian motion. This model is simple and extremely tractable for a number of important problems (particularly the problem of *pricing* and hedging derivative

securities), and as such, it has had enormous impact over the past thirty years. It is the primary example of a complete, continuous-time market.

An investor starts with capital v at time $t = 0$ and has the choice to invest in the stock or put his/her money in the bank earning interest at the (assumed constant) interest rate r. He continuously balances his portfolio, adjusting the weights between these two choices so as to maximize his expected utility of wealth at time T. That is, he has a utility function U that is increasing (he prefers more wealth to less) and concave (representing risk-aversion) through which he values one investment strategy over another.

If π_t is the dollar amount of stock that the investor holds at time t, we define his wealth process (V_t) by (1). The investor looks for a strategy $\pi_t^\star$ that maximizes

$$\mathbb{E}\{U(V_T)\}$$

starting with initial wealth v and under the constraint of nonbankruptcy: $V_t \geq 0$ for all $0 \leq t \leq T$ (almost surely).

Without loss of generality, we shall henceforth take the interest rate r to be zero. This can be justified simply by measuring all capital values *in units of the money market account*, i.e. replacing (X_t), (V_t) and (π_t) by $(X_t e^{-rt})$, $(V_t e^{-rt})$ and $(\pi_t e^{-rt})$, respectively. We recover the same equations as long as we relabel $\mu - r$ as μ. With this convention, the equations for (X_t, V_t) in terms of the control π_t are

$$\frac{dX_t}{X_t} = \mu\,dt + \sigma\,dW_t, \tag{3}$$

$$dV_t = \mu\pi_t\,dt + \sigma\pi_t\,dW_t, \tag{4}$$

where μ is now the excess growth over the risk free rate. Notice that the process (V_t) is autonomous, i.e. it does not directly involve (X_t).

2.2 Dynamic Programming and the Bellman Equation

One approach to this problem is to observe that if π_t were chosen to be a (nice) function of V_t, the controlled process (V_t) defined by (4) would be a Markov process. Therefore, *if* the optimal strategy were of this so-called *feedback form*, we could take advantage of the structure and study the evolution of the optimal expected utility as a function of the starting wealth and starting time. To this end, we define the *value function*

$$H(t, v) = \sup_{(\pi_s)_{s \in [t, T]}} \mathbb{E}\{U(V_T) \mid V_t = v\},$$

where by hypothesis, H is a function of only the starting time t and wealth v (and not, for instance, details of the path). We shall also use the shorthand

$$\mathbb{E}_{t,v}\{\cdot\} = \mathbb{E}\{\cdot \mid V_t = v\}.$$

Notice that (X_t) has vanished from the problem because everything can be stated in terms of the process (V_t). This is a direct consequence of the geometric Brownian motion model (3) assumed here. If μ or σ were general functions of X_t, the reduction in dimension would not be possible. In the optimization problems involving derivative contracts, X_T appears in the utility function and, again, we cannot eliminate the x variable.

The following is a loose derivation of the Bellman equation for the value function H. The idea is to divide the control process (π_s) over the the time interval $[t, T)$ into the (constant) control π_t over $[t, t+dt)$ and the rest (π_s) over $s \in [t+dt, T)$. We then optimize over these two parts separately. Formally, conditioning on the wealth at time $t + dt$ and using the iterated expectations formula, this looks like

$$
\begin{aligned}
H(t, v) &= \sup_{(\pi_s)_{s \in [t,T)}} \mathbb{E}_{t,v} \left\{ \mathbb{E}\{U(V_T) \mid V_{t+dt}\} \right\} \\
&= \sup_{\pi_t} \mathbb{E}_{t,v} \left\{ \sup_{(\pi_s)_{s>t}} \mathbb{E}\{U(V_T) \mid V_{t+dt}\} \right\} \\
&= \sup_{\pi_t} \mathbb{E}_{t,v} \left\{ H(t + dt, V_{t+dt}) \right\} \\
&= \sup_{\pi_t} \mathbb{E}_{t,v} \left\{ H(t, v) + H_t(t, v)\, dt + \mathcal{L}_v H(t, v)\, dt + \sigma \pi_t H_v(t, v)\, dW_t \right\},
\end{aligned}
$$

where $\mathcal{L}_v$ denotes the infinitesimal generator of (V_t), namely

$$
\mathcal{L}_v = \frac{1}{2} \sigma^2 \pi_t^2 \frac{\partial^2}{\partial v^2} + \mu \pi_t \frac{\partial}{\partial v},
$$

and we have used Itô's formula in the last step. Note that $\mathcal{L}_v$ depends on the control π, but we do not denote the dependence in this notation. The last term is zero since $\mathbb{E}\{dW_t\} = 0$. We therefore obtain

$$
H_t + \sup_{\pi_t} \mathcal{L}_v H = 0. \tag{5}
$$

Notice that the computation exploits the Markovian structure to reduce the optimization over the whole time period to successive optimizations over infinitesimal time intervals. This description of the optimal strategy for the long run, that is to do as well as one can over the short run, is called the Bellman principle.

The Bellman partial differential equation (5) applies in the domain $t < T$ and $v > 0$ with the terminal condition

$$
H(T, v) = U(v)
$$

and the boundary condition $H(t, 0) = 0$ enforcing the bankruptcy constraint.

If we can find a smooth solution $H(t, v)$ to which Itô's formula can be applied, it follows from a verification theorem that H gives the maximum expected utility and the optimal strategy is given by a Markov control $\pi_t = \pi(V_t)$. For details, see [8].

In our case, the Bellman equation is

$$
H_t + \sup_{\pi} \left(\frac{1}{2} \sigma^2 \pi^2 H_{vv} + \mu \pi H_v \right) = 0.
$$

The internal optimization is simply to find the extreme point of a quadratic. We shall assume (and it can be shown rigorously) that the value function inherits the convexity of the utility function. Moreover, under reasonable conditions, it is *strictly* convex for $t < T$ even if U is not. This follows from the diffusive part of the equation. Therefore $H_{vv} < 0$. The optimization over π is not constrained because we have not assumed any constraints

on the trading strategies themselves, as long as the wealth stays positive, and therefore the quadratic is maximized by

$$\pi^\star = -\frac{\mu H_v}{\sigma^2 H_{vv}}. \tag{6}$$

Substituting with this π, we can rewrite the Bellman equation as

$$H_t - \frac{\mu^2}{2\sigma^2}\frac{H_v^2}{H_{vv}} = 0. \tag{7}$$

2.3 Solution by Legendre Transform

We further take advantage of the assumed convexity of the value function to define the Legendre transform

$$\hat{H}(t,z) = \sup_{v>0}\{H(t,v) - zv\},$$

where $z > 0$ denotes the dual variable to v. The value of v where this optimum is attained is denoted by $g(t,z)$, so that

$$g(t,z) = \inf\{v > 0 \mid H(t,v) \geq zv + \hat{H}(t,z)\}.$$

The two functions $g(t,z)$ and $\hat{H}(t,z)$ are closely related and we shall refer to either one of them as the (convex) dual of H. In this article, we will work mainly with the function g, as it is easier to compute numerically and suffices for the purposes of computing optimal trading strategies. The function $\hat{H}$ is related to g by $g = -\hat{H}_z$.

At the terminal time, we denote

$$\hat{U}(x,z) = \sup\{U(v) - zv \mid 0 < v < \infty\}$$
$$G(x,z) = \inf\{v > 0 \mid U(v) \geq zv + \hat{U}(z)\}.$$
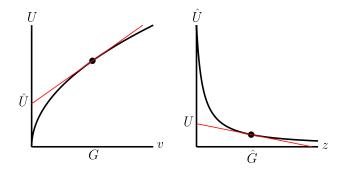
This is illustrated in Figure 1.

Figure 1:   *The Legendre transform. The left curve is the graph of $U(v)$ the utility function and the tangent line has slope z for $v = G(z)$ and has vertical intercept $\hat{U}(z)$. The right curve is the graph of $\hat{U}(z)$ and the tangent line has slope $-v$ at the point $z = \hat{G}(v)$ and has vertical intercept $U(v)$.*

Assuming that H is strictly concave and smooth as a function of v, we have that

$$H_v(t, g(t,z)) = z,$$

or $g = H_v^{-1}$, whereby, in economics parlance, g is sometimes referred to as the inverse of marginal utility.

By differentiating this with respect to t and z we recover the following rules relating the derivatives of the value function to the dual function g:

$$H_{tv} = -\frac{g_t}{g_z} \qquad H_{vv} = \frac{1}{g_z} \qquad H_{vvv} = -\frac{g_{zz}}{g_z^3},$$

where the left is evaluated at $(t, g(t, z))$. Differentiating equation (7) with respect to v and substituting gives an autonomous equation for g:

$$g_t + \frac{\mu^2}{2\sigma^2} z^2 g_{zz} + \frac{\mu^2}{\sigma^2} z g_z = 0,$$
$$g(T, z) = G(z),$$

in $t < T$ and $z > 0$. The key observation is that this is now a *linear* PDE.

2.4 Optimal Strategy

We are not usually interested in the value function, but rather in the optimal investment strategy. From (6), we can compute the optimal stock holding as a feedback formula in terms of derivatives of the value function. In terms of the dual function g, it is given by

$$\pi_t^\star = -\frac{\mu}{\sigma^2} z g_z(t, z).$$

Thus for a given Merton problem, we solve the linear PDE for g and recover the investment amount $\pi^\star$ from its first derivative. All that remains is to use the value of the dual variable z that corresponds to the current time and wealth level (t, v). This is obtained from the relation

$$g(t, z) = v.$$

2.5 Example: Power Utility

The success of the Merton approach is due primarily to an appealing explicit formula for certain simple utility functions. For an isoelastic (or power) utility function of the form

$$U(v) = v^\gamma / \gamma, \qquad 0 < \gamma < 1,$$

the Legendre duals G and $\hat{U}$ are given by

$$G(z) = z^{\frac{1}{\gamma-1}} \quad \text{and} \quad \hat{U}(z) = \frac{1-\gamma}{\gamma} z^{\frac{\gamma}{\gamma-1}}.$$

These are illustrated in Figure 2. In this case, the linear PDE for g admits a separable solution of the form

$$g(t, z) = z^{\frac{1}{\gamma-1}} u(t),$$

for some function $u(t)$ we can compute. It follows that for a given (t, v),
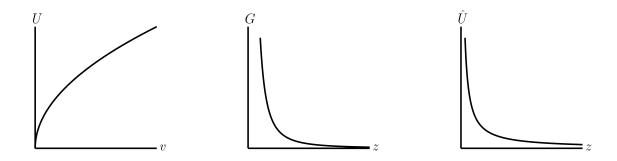
$$g(t, z) = v$$

Figure 2: *Terminal conditions for the Merton problem with power utility.*

and the optimal strategy is given by

$$\pi_t^\star = -\frac{\mu}{\sigma^2} z g_z = -\frac{\mu}{\sigma^2} \frac{1}{\gamma - 1} g = \frac{\mu}{\sigma^2(1-\gamma)} v.$$

That is, the optimal strategy is to hold the fraction

$$M = \frac{\mu}{\sigma^2(1-\gamma)}$$

of current wealth in the risky asset (the stock) and to put the rest in the bank. As the stock price rises, this strategy says to sell some stock so that the fraction of the portfolio comprised of the risky asset remains the same.

The fraction M is known as the Merton ratio. More importantly, as we see next, this *fixed-mix* result generalizes to multiple securities as long as they are also assumed to be geometric Brownian motions.

2.6 The Multi-Dimensional Merton Problem

We have a market with $n \geq 1$ stocks where the ith stock price $(X_t^{(i)})$ (in units of the money market account) is modeled by

$$\frac{dX_t^{(i)}}{X_t^{(i)}} = \mu_i \, dt + \sum_{j=1}^n \sigma_{ij} dW_t^{(j)},$$

where the μ_i are the growth rates less the risk free rate and $(W_t^{(i)})$ are independent Brownian motions. The price processes are correlated through the diffusion terms. We denote by (X_t) the vector price process, μ is the vector of return rates, and σ is the $n \times n$ volatility matrix.

An investor chooses amounts $\pi_t = (\pi_t^{(1)}, \cdots, \pi_t^{(n)})$ to invest in each stock, and puts the rest in the bank. Again (π_t) denotes the vector control process. His wealth process is

$$dV_t = \pi_t^T \mu \, dt + \pi_t^T \sigma \, dW_t,$$

where ^{T} denotes transpose. Recall that in the one-dimensional Merton's problem, the stock price process disappeared from the problem. This happens in the multidimensional case, too, and the Bellman equation for the value function $H(t, v)$ becomes

$$H_t + \sup_\pi \left(\frac{1}{2} \pi^T \sigma \sigma^T \pi H_{vv} + \pi^T \mu H_v \right) = 0.$$

This can be transformed to a linear PDE using the Legendre transform. In the case of power utility, $U(v) = v^\gamma/\gamma$, the optimal strategy is given by

$$\pi_t^\star = \frac{1}{1-\gamma}(\sigma\sigma^T)^{-1}\mu v.$$

Thus, the optimal strategy is again to hold fixed fractions of current wealth in the risky assets (the stocks) and to put the rest in the bank. The fractions are given by the vector of Merton ratios $\frac{1}{1-\gamma}(\sigma\sigma^T)^{-1}\mu$.

3 State-Dependent Utility Maximization: Constant Volatility

In this section, we describe briefly the generalization of the utility maximization problem to the case when the process modeling the price of the stock cannot be eliminated. We assume still that it is a geometric Brownian motion, but that the utility depends on its final value, as in the examples outlined in Section 1.2.

The value function is

$$H(t,x,v) = \sup_\pi I\!\!E\{U(X_T, V_T) \mid X_t = x, V_t = v\}$$

and it is conjectured to satisfy the Bellman equation

$$\begin{aligned}
0 &= H_t + \mu x H_x + \frac{1}{2}\sigma^2 x^2 H_{xx} + \sup_\pi\left(\pi\mu H_v + \pi\sigma^2 x H_{xv} + \frac{1}{2}\pi^2\sigma^2 H_{vv}\right) \\
&= H_t + \mu x H_x + \frac{1}{2}\sigma^2 x^2 H_{xx} - \frac{(\mu H_v + \sigma^2 H_{xv})^2}{2\sigma^2 H_{vv}}.
\end{aligned} \tag{8}$$

3.1 Convex Duals

Proceeding as in the Merton problem, we introduce the duals

$$\hat{H}(t,x,z) = \sup\{H(t,x,v) - zv \mid 0 < v < \infty\}$$
$$g(t,x,z) = \inf\{v > 0 \mid H(t,x,v) \geq zv + \hat{H}(t,x,z)\}.$$

Notice that the we take the convex duals with respect to the variable v only. What is of interest is how the extra state variable x affects the transformation of the Bellman PDE. The transformation rules for the derivatives are given by

$$H_v = z \qquad \begin{aligned} H_t &= \hat{H}_t \\ H_x &= \hat{H}_x \end{aligned} \qquad \begin{aligned} H_{vv} &= -\frac{1}{\hat{H}_{zz}} \\ H_{xv} &= -\frac{\hat{H}_{xz}}{\hat{H}_{zz}} \end{aligned} \qquad H_{xx} = \hat{H}_{xx} - \frac{\hat{H}_{xz}^2}{\hat{H}_{zz}}.$$

This implies that

$$\hat{H}_{xx} = H_{xx} - \frac{H_{xv}^2}{H_{vv}}. \tag{9}$$

If we rewrite (8) as

$$H_t + \mu x H_x + \frac{1}{2}\sigma^2 x^2 \left(H_{xx} - \frac{H_{xv}^2}{H_{vv}} \right) - \mu x \frac{H_v H_{xv}}{H_{vv}} - \frac{\mu^2}{2\sigma^2}\frac{H_v^2}{H_{vv}} = 0,$$

we see that the first nonlinear term is transformed into a linear term because of (9), and the third is as in the Merton problem described in Section 2.3 where it became linear in the dual variables. Moreover the same happens to the second nonlinear term and we get

$$\hat{H}_t + \frac{1}{2}\sigma^2 x^2 \hat{H}_{xx} + \frac{1}{2}\frac{\mu^2}{\sigma^2}z^2 \hat{H}_{zz} + \mu x H_x - \mu x z \hat{H}_{xz} = 0. \tag{10}$$

Note that the first nonlinear term would not have been transformed to a linear one if the (V_t) process was not perfectly correlated with the (X_t) process. This is the situation in the stochastic volatility models we discuss later because volatility is not a traded asset.

We can recover the PDE for the other dual function $g(t, x, z)$ using

$$g = -\hat{H}_z$$

and differentiating the PDE for $\hat{H}$ with respect to z to give

$$g_t + \frac{1}{2}\sigma^2 x^2 g_{xx} + \frac{\mu^2}{\sigma^2}z g_z + \frac{1}{2}\frac{\mu^2}{\sigma^2}z^2 g_{zz} - \mu x z g_{xz} = 0. \tag{11}$$

3.2 Optimal Strategy

Having solved the *linear* PDE for either g or $\hat{H}$, we recover the optimal strategy from

$$\pi^* = -\frac{(\mu H_v + \sigma^2 x H_{xv})}{\sigma^2 H_{vv}}$$

$$= \frac{\mu}{\sigma^2}z\hat{H}_{zz} - x\hat{H}_{xz}$$

$$= xg_x - \frac{\mu}{\sigma^2}z g_z.$$

The value of z to be used when the stock and wealth are (x, v) at time t are found from

$$g(t, x, z) = v.$$

3.3 Example: Partial Hedging

In this problem, we are partially hedging a derivative with payoff $h(X_T) \geq 0$ at time T, as explained in Section 1.2.1. We assume that v is strictly less than the superhedging price of the derivative. In the present case of constant volatility, the market is complete and the superhedging price is the Black-Scholes price $v^\star$. It is the (uniquely determined) initial capital for a replicating strategy for the derivative.

Recall that the state-dependent utility function is

$$U(x, v) = \frac{1}{p}\left(h^p - \left((h - v)^+ \right)^p \right),$$

where $h = h(x)$ and $p > 1$. This has convex duals

$$G(x, z) = \left(h - z^{\frac{1}{p-1}}\right)^+$$

$$\hat{U}(x, z) = \left(\frac{1}{p}h^p + \frac{p-1}{p}z^{\frac{p}{p-1}} - zh\right)\mathcal{H}\left(h - z^{\frac{1}{p-1}}\right),$$

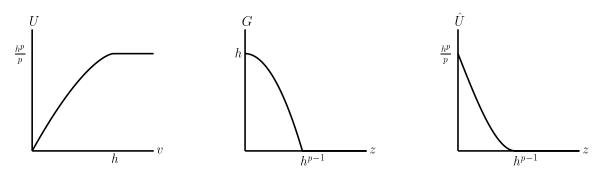where $\mathcal{H}$ is the Heaviside (step) function. See Figure 3.

Figure 3: *Terminal conditions for partial hedging.*

The linear PDEs (10) and (11) have a probabilistic interpretation as the expectation of a function of a (two-dimensional) Markov process, but this can be reduced to a function of just the terminal stock price (because of the degeneracy of (11)). It can be shown that

$$g(t, x, z) = \mathbb{E}^\star_{t,x}\{\left(h(X_T) - \tilde{c}X_T^{-\kappa}\right)^+\}, \tag{12}$$

where $\kappa = \frac{\mu}{\sigma^2(p-1)}$, and $\mathbb{E}^\star$ denotes expectation with respect to the (so-called risk-neutral) probability measure $\mathbb{P}^\star$ under which

$$dX_t = \sigma X_t dW_t^\star,$$

with $(W_t^\star)$ a $\mathbb{P}^\star$-Brownian motion, and

$$\tilde{c} = \tilde{c}(t, x, z) = zx^{-\mu/\sigma^2}\exp\left(\frac{1}{2}(\frac{\mu^2}{\sigma^2} - \mu)(T - t)\right).$$

This representation is useful in finance where linear diffusion PDEs are associated with pricing equations for derivative securities. We do not stress this interpretation here, but merely comment that the optimal strategy is to trade the stock in such a way so as to replicate the *target wealth*, which here is a European derivative contract with the modified payoff function

$$\left(h(X_T) - \tilde{c}X_T^{-\kappa}\right)^+.$$

This is illustrated in Figure 4. The number z is determined by

$$g(t, x, z) = \text{current wealth}.$$

For computational purposes, the only requirement is to solve the PDE for g or $\hat{H}$. It turns out that g is more amenable to explicit computation.
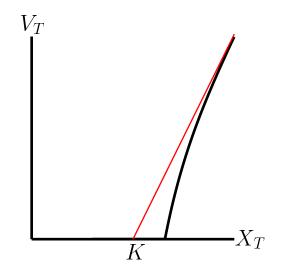
Figure 4: *Target wealth for partial hedging of a European call option. The thin line is the payoff function of the original call. When partial hedging, they payoff is replaced by the thick line.*

3.4 Explicit Computation for a Call Option

Sometimes we can get closed formulas for the function $g(t, x, z)$ and its derivatives, and hence the optimal strategy in the utility maximization. This greatly increases computational speed and is the case, for instance, when we are partially hedging a call option $h(x) = (x - K)^+$. Indeed, let us define $\tilde{d}_2 = \tilde{d}_2(t, x, z)$ as the (unique) solution to

$$xe^{-\frac{1}{2}\sigma^2\tau - \sigma\sqrt{\tau}\tilde{d}_2} - K - z^{\frac{1}{p-1}} e^{\frac{\mu^2\tau}{2(p-1)\sigma^2} + \frac{\mu\sqrt{\tau}\tilde{d}_2}{(p-1)\sigma}} = 0,$$

Given K and values of t, x, z, this can be solved numerically, since the left hand side is a strictly decreasing function of $\tilde{d}_2$. We then get from (12):

$$g(t, x, z) = \int_{-\infty}^{\tilde{d}_2} \left(xe^{-\frac{1}{2}\sigma^2\tau - \sigma\sqrt{\tau}\xi} - K - z^{\frac{1}{p-1}} e^{\frac{\mu^2\tau}{2(p-1)\sigma^2} + \frac{\mu\sqrt{\tau}\xi}{(p-1)\sigma}} \right) \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}\xi^2} \, d\xi$$

$$= xN(\tilde{d}_1) - KN(\tilde{d}_2) - z^{\frac{1}{p-1}} e^{\frac{p\mu^2\tau}{2(p-1)^2\sigma^2}} N(\tilde{d}_3), \tag{13}$$

where

$$\tilde{d}_1 = \tilde{d}_2 + \sigma\sqrt{\tau} \qquad \tilde{d}_3 = \tilde{d}_2 - \frac{\mu\sqrt{\tau}}{(p-1)\sigma}.$$

This is similar to the celebrated Black-Scholes formula for the price of a call option, and in fact reduces to this as $z \to 0$.

In special cases such as $p = 2$ one can also get a closed formula for $\hat{H}(x, z)$, but we only need g in order to find the optimal strategy:

$$\frac{\pi^*(t, x, z)}{x} = g_x - \frac{\mu z}{\sigma^2 x} g_z$$

$$= N(\tilde{d}_1) + \frac{\mu z^{\frac{1}{p-1}}}{(p-1)\sigma^2 x} \exp\left(\frac{p\mu^2\tau}{2(p-1)^2\sigma^2} \right) N(\tilde{d}_3).$$

This strategy was derived in [7] and we shall refer to it as the Föllmer-Leukert strategy. It is implemented in the simulations of Section 7.

4   Stochastic Volatility Models

The BSM model supposes that asset price volatility is constant. This is contradicted by empirical evidence of fluctuating historical volatility. More crucially, volatility as a measure of risk is seen by traders as the most important variable (after the price of the underlying asset itself) in driving probabilities of profit or loss.

Stochastic volatility models which replace σ in (2) by a random process (σ_t) arise not just because of empirical evidence of historical volatility's "random characteristics", but also from considerations of market hysteria, uncertainty in estimation, or they could be used to simulate non-Gaussian (heavy-tailed) returns distributions. They describe a much more complex market than the Black-Scholes model, and this is reflected in the increase in difficulty of the derivative hedging problems we shall describe. They were introduced in the academic literature in the late 1980's [11] and are popular in the industry today.

A model for stock prices in which volatility (σ_t) is a random process starts with the stochastic differential equation

$$\frac{dX_t}{X_t} = \mu \, dt + \sigma_t \, dW_t, \tag{14}$$

the analogue of (2).

A key aim in the modeling is to say as little about volatility as possible so that we are not tied to a specific model. Since volatility is not observed directly, there is a paucity of consistent econometric information about its behaviour. One feature that most empirical studies point out, and which squares with common experience, is that volatility is *mean-reverting*: it is not wandering into far-flung excursions, but seems to be pulled upwards when it is low and downwards when it is high. It is convenient to model volatility as a function of a simple mean-reverting (ergodic) Markov diffusion process (Y_t), for example an Ornstein-Uhlenbeck (OU) process. We discuss the class of models

$$\sigma_t = f(Y_t); \qquad dY_t = \alpha(m - Y_t) \, dt + \beta \, d\hat{Z}_t, \tag{15}$$

where f is a positive bounded function through which the generality of possible volatility models is obtained. In fact the asymptotic results for derivative pricing in [10] are insensitive to all but a few general features of f, and the way the method there is calibrated means that this function never has to be chosen.

In (15), $(\hat{Z}_t)$ is a Brownian motion modeling the fine-scale volatility fluctuations that is correlated with the other Brownian motion (W_t). Itô diffusions provide a simple way to model the much observed "leverage effect" that volatility and stock price shocks are negatively correlated: when volatility goes up, stock prices tend to fall. The instantaneous correlation coefficient ρ, where

$$I\!E\{dW_t \, d\hat{Z}_t\} = \rho \, dt$$

measures this asymmetry in the probability distribution of future stock prices: $\rho < 0$ generates a fatter left tail.

4.1 Fast Mean-Reverting Stochastic Volatility

The effects of *fast mean-reversion* in volatility were studied in [10]. Mean-reversion is mathematically described by ergodicity and refers to the characteristic *time* it takes for an ergodic process to get back to the mean-level of its long-run distribution. The separation of scales that we shall exploit is that while stock prices change on a tick-by-tick time-scale (minutes), volatility changes more slowly, on the scale of a few days, but still fluctuates rapidly over the time-scale (months) of a derivative contract. This phenomenon of bursty or clustering volatility is characterized as *fast mean-reversion* in the models we look at. That is, the volatility process is <u>fast</u> mean-reverting with respect to the long time-scale (months) of reference. (It is slow mean-reverting with respect to the tick time-scale, by which it is sometimes described).

The important parameter in (15) is α, the rate of mean-reversion of the volatility-driving process (Y_t). Fast mean-reversion describes the limit α tending to infinity with $\beta^2/2\alpha$, the variance of the long-run distribution of the OU process, fixed. A detailed study of high-frequency $S\&P500$ data, where the large rate of volatility mean-reversion was established by a variety of methods, appears in [9]. As an illustration, Figure 5 shows simulated volatility
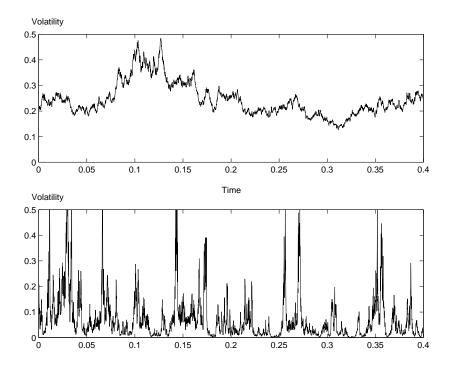

Figure 5: *Simulated volatility for small ($\alpha = 1$, top) and large ($\alpha = 200$, bottom) rates of mean-reversion for the OU model, with the choice $f(y) = e^y$. Note how high volatility appears in short bursts in the latter case.*

paths for one of the models above in which $\alpha = 1$ in the top graph and $\alpha = 200$ in the bottom graph. Of course volatility is not directly observable as a time-series, which complicates estimation issues and makes it desirable to have a theory that is robust to volatility modeling.

In Section 6, we exploit this separation of scales to construct an approximation for the state-dependent utility optimization problem under stochastic volatility, which is discussed in the next section.

5 Utility Maximization under Stochastic Volatility

In this section, we look at the state-dependent utility maximization problem and how it is modified for a market with stochastic volatility. Under the class of mean-reverting stochastic volatility models (14)-(15), we study as before the Bellman equation of dynamic programming: the value function

$$H(t, x, y, v) = \sup_{\pi} I\!\!E\{U(X_T, V_T) \mid X_t = x, V_t = v, Y_t = y\}$$

is conjectured to satisfy:

$$H_t + \mathcal{L}_{x,y}H - \frac{(\mu H_v + f(y)^2 x H_{xv} + \rho\beta f(y)H_{yv})^2}{2f(y)^2 H_{vv}} = 0, \tag{16}$$

where

$$\mathcal{L}_{x,y}H = \mu x H_x + \alpha(m - y)H_y + \frac{1}{2}f(y)^2 x^2 H_{xx} + \rho\beta f(y)x H_{xy} + \frac{1}{2}\beta^2 H_{yy},$$

the infinitesimal generator of the process (X_t, Y_t). The domain is $x > 0$, $-\infty < y < \infty$, $v > 0$ and $t < T$ and the terminal condition is

$$H(T, x, y, v) = U(x, v).$$

If a smooth solution (to which Itô's formula can be applied) can be found, a verification theorem shows that the optimal strategy is given in feedback form by

$$\pi_t^\star = -\frac{\mu H_v + f(y)^2 x H_{xv} + \rho\beta f(y)H_{yv}}{f(y)^2 H_{vv}}. \tag{17}$$

We proceed to study the dual optimization problem for the Legendre transform (with respect to v) of the value function. Defining

$$\hat{H}(t, x, y, z) = \sup\{H(t, x, y, v) - zv \mid 0 < v < \infty\}$$

$$g(t, x, y, z) = \inf\left\{v > 0 \mid H(t, x, y, v) \geq zv + \hat{H}(t, x, y, z)\right\},$$

the dual functions g and $\hat{H}$ satisfy simpler-looking equation. In the case of g it reads

$$g_t + \mathcal{L}_{x,y}g + \frac{1}{2}\frac{\mu^2}{f(y)^2}z^2 g_{zz} - \mu x z g_{xz} - \frac{\rho\beta\mu}{f(y)}z g_{yz} + \frac{\mu^2}{f(y)^2}z g_z$$

$$-\mu x g_x - \frac{\rho\beta\mu}{f(y)}g_y - \frac{1}{2}\beta^2(1 - \rho^2)\left[\frac{2g_y g_{yz}}{g_z} - \frac{g_y^2}{g_z^2}g_{zz}\right] = 0, \tag{18}$$

which is still nonlinear, but only because of the last bracketed term. In the case of a complete market, meaning nonrandom ($\beta = 0$) or perfectly correlated ($|\rho| = 1$) volatility, notice that the nonlinear term disappears and the work done by the Legendre transform is to reveal that the optimization problem is simply a linear pricing problem in disguise. However this complete reduction is not possible with stochastic volatility. Nonetheless, the transform has done some work in isolating the nonlinearity due to the fact that volatility is not traded.

6   Asymptotics for Utility Maximization

In this section, we study the effect of uncertain volatility on the optimal strategies for state-dependent utility maximization. We take advantage of *fast mean-reversion* and use a singular perturbation analysis to find a relatively simple trading strategy that approximates the optimal one. The analysis in Section 5 still has not yielded a way to compute the optimal strategy short of solving one of the nonlinear PDEs (16) or (18) which have three spatial dimensions. One of the benefits of the approach described here is easing of this dimensional burden. Another one is robustness—as we will see we do not need to know all the parameters in the model for the purposes of the approximate strategy.

In the zero-order approximation derived here, two kinds of average (or homogenized) volatilities emerge: $\bar{\sigma} := \langle f^2 \rangle^{1/2}$ and $\sigma_\star := \langle f^{-2} \rangle^{-1/2}$, where $\langle \cdot \rangle$ denotes a particular averaging procedure described below.

6.1   Singular Perturbation Analysis

We introduce the scaling

$$
\begin{aligned}
\alpha &= 1/\varepsilon \\
\beta &= \sqrt{2}\,\nu/\sqrt{\varepsilon}
\end{aligned}
$$

where $0 < \varepsilon << 1$ and ν is fixed, to model fast mean-reversion (clustering) in market volatility. Recall that α measures the characteristic speed of mean-reversion of (Y_t) and ν^2 is the variance of the long-run distribution, measuring the typical size of the fluctuations of the volatility-driving process.

Then $g = g^\varepsilon$ satisfies the PDE (18), which we re-write with the new notation as:

$$
\left(\frac{1}{\varepsilon}\mathcal{L}_0 + \frac{1}{\sqrt{\varepsilon}}\mathcal{L}_1 + \mathcal{L}_2 \right) g^\varepsilon + \frac{\nu^2}{\varepsilon}(1 - \rho^2)\mathrm{NL}^\varepsilon = 0, \tag{19}
$$

where we define

$$
\mathcal{L}_0 = \nu^2 \frac{\partial^2}{\partial y^2} + (m - y)\frac{\partial}{\partial y} \tag{20}
$$

$$
\mathcal{L}_1 = \sqrt{2}\,\rho\nu \left(f(y)x\frac{\partial^2}{\partial x \partial y} - \frac{\mu}{f(y)}z\frac{\partial^2}{\partial y \partial z} - \frac{\mu}{f(y)}\frac{\partial}{\partial y} \right) \tag{21}
$$

$$
\mathcal{L}_2 = \frac{\partial}{\partial t} + \frac{1}{2}f(y)^2 x^2 \frac{\partial^2}{\partial x^2} + \frac{\mu^2}{f(y)^2}\left(\frac{1}{2}z^2\frac{\partial^2}{\partial z^2} + z\frac{\partial}{\partial z} \right) - \mu xz\frac{\partial^2}{\partial x \partial z}, \tag{22}
$$

and the nonlinear part is

$$
\mathrm{NL}^\varepsilon = -\frac{\partial}{\partial z}\left(\frac{(g_y^\varepsilon)^2}{g_z^\varepsilon} \right) = -\left[\frac{2g_y^\varepsilon g_{yz}^\varepsilon}{g_z^\varepsilon} - \left(\frac{g_y^\varepsilon}{g_z^\varepsilon} \right)^2 g_{zz}^\varepsilon \right].
$$

Notice that $\mathcal{L}_0$ is the usual (scaled) OU generator and $\mathcal{L}_1$ takes derivatives in y and kills functions that do not depend on y.

The approach is now to think of the actual market (and our optimization problem) as embedded in a family of similar problems parametrized by (small) values of ε. For $\varepsilon = 0$,

volatility is mean-reverting "infinitely fast" and can be replaced by some average as far as expectations are concerned. However two different averages are needed for different facets of the optimal strategy. This principal approximation may be sufficient for many purposes. It can be improved by perturbation or expansion around $\varepsilon = 0$, but the cost is greater reliance on model specification. We refer to [12] for details.

6.1.1   Expansion

We look for an expansion

$$g^\varepsilon(t, x, y, z) = g^{(0)}(t, x, y, z) + \sqrt{\varepsilon} g^{(1)}(t, x, y, z) + \varepsilon g^{(2)}(t, x, y, z) + \cdots$$

for small ε. There is a similar expansion for the other dual function $\hat{H}$ and the two are related by $\hat{H}_z = -g$.

6.1.2   Term of Order $1/\varepsilon$

Inserting the expansion for g and comparing terms of order $1/\varepsilon$ gives

$$\nu^2 g_{yy}^{(0)} + (m - y)g_y^{(0)} + \nu^2(1 - \rho^2)\frac{\partial}{\partial z}\left(\frac{(g_y^{(0)})^2}{g_z^{(0)}}\right) = 0.$$

An obvious solution to this is any smooth function $g^{(0)}$ that does not depend on y. It is an interesting and crucial fact that the solution $g^{(0)}$ that we are looking for is of this kind, i.e. does not depend on y. To see this, it is easier to work with the PDE for the zero-order term $\hat{H}^{(0)}$ of $\hat{H}$ (recall that $g = -\hat{H}_z$ so $g^{(0)} = -\hat{H}_z^{(0)}$). This reads

$$\nu^2 \hat{H}_{yy}^{(0)} + (m - y)\hat{H}_y^{(0)} - \nu^2(1 - \rho^2)\frac{(\hat{H}_{yz}^{(0)})^2}{\hat{H}_{zz}^{(0)}} = 0. \tag{23}$$

Denoting $u = \hat{H}_y$ and using the concavity property $\hat{H}_{zz} > 0$ we may write (23) as an ordinary differential inequality:

$$\nu^2 u_y + (m - y)u \geq 0. \tag{24}$$

Let us define

$$\Phi(y) = \frac{1}{\sqrt{2\pi\nu^2}} e^{-(y-m)^2/2\nu^2}, \tag{25}$$

the density of the $\mathcal{N}(m, \nu^2)$ distribution. Using this, we can re-write (24) as

$$\frac{1}{\Phi(y)}\frac{\partial}{\partial y}\left(\Phi(y)u(y, z)\right) \geq 0.$$

By integrating we get

$$u(y, z) \geq u(m, z)e^{\frac{(y-m)^2}{2\nu^2}} \quad \text{for } y \geq m$$

$$u(y, z) \leq u(m, z)e^{\frac{(y-m)^2}{2\nu^2}} \quad \text{for } y \leq m.$$

We conclude that $u(m, z) = 0$, because otherwise $u(m, y)$ would grow too fast as $y \to \pm\infty$. This implies that any solution to (23) is independent of y. This is because if we specify the value q of any solution at the point $y = m$,

$$\hat{H}^{(0)}(m, z) = q,$$

then the constant function

$$\tilde{H}(y, z) \equiv q \qquad \text{for all } y$$

is also a solution to (23) with $\tilde{H}_y(m, z) = 0$. From uniqueness of solutions it follows that $\hat{H}^{(0)} \equiv \tilde{H}$ and so $\hat{H}^{(0)}$ and $g^{(0)} = -\hat{H}_z^{(0)}$ do not depend on y, so

$$g^{(0)} = g^{(0)}(t, x, z).$$

6.1.3 Term of Order $1/\sqrt{\varepsilon}$

At the order $1/\sqrt{\varepsilon}$,

$$\mathcal{L}_1 g^{(0)} + \mathcal{L}_0 g^{(1)} = 0,$$

which implies $g^{(1)}$ also does not depend on y because $\mathcal{L}_1 g^{(0)} = 0$ and $\mathcal{L}_0$ has null space spanned by constants. This is a general property of generators of "nice" ergodic processes like the OU.

Since both $g^{(0)}$ and $g^{(1)}$ do not depend on y, the nonlinear term is effectively

$$\mathrm{NL}^{\varepsilon} = \mathcal{O}(\varepsilon^2)$$

and only contributes to the asymptotics when we compare order ε and higher. This fact is crucial to the further analysis. We will go as far as order one (ε^0) here, so we are dealing essentially with linear asymptotics (except for the very first equation).

6.1.4 Zero-order Term

At order 1, we have

$$\mathcal{L}_0 g^{(2)} + \mathcal{L}_1 g^{(1)} + \mathcal{L}_2 g^{(0)} = 0.$$

The middle term is zero because $g^{(1)}$ does not depend on y. We have a Poisson equation (in y) for $g^{(2)}$. The solvability condition is that $\mathcal{L}_2 g^{(0)}$ must be centered with respect to the invariant distribution of the OU process (Y_t) (equivalently, orthogonal to the null space of the adjoint of $\mathcal{L}_0$, the Fredholm alternative). Therefore

$$\langle \mathcal{L}_2 g^{(0)} \rangle = \langle \mathcal{L}_2 \rangle g^{(0)} = 0, \tag{26}$$

where $\langle \cdot \rangle$ denotes the averaging

$$\langle \Psi \rangle = \int \Psi \Phi = \frac{1}{\sqrt{2\pi\nu^2}} \int_{-\infty}^{\infty} \Psi(y) \, e^{-\frac{(m-y)^2}{2\nu^2}} \, dy,$$

that is, the average with respect to the $\mathcal{N}(m, \nu^2)$ distribution, the invariant or long-run distribution of the OU process (Y_t).

The averaged operator is

$$\langle \mathcal{L}_2 \rangle = \frac{\partial}{\partial t} + \frac{1}{2}\bar{\sigma}^2 x^2 \frac{\partial^2}{\partial x^2} + \frac{\mu^2}{\sigma_\star^2}\left(\frac{1}{2}z^2\frac{\partial^2}{\partial z^2} + z\frac{\partial}{\partial z}\right) - \mu x z \frac{\partial^2}{\partial x \partial z},$$

where we define

$$\bar{\sigma}^2 = \langle f^2 \rangle \tag{27}$$

$$\frac{1}{\sigma_\star^2} = \langle \frac{1}{f^2} \rangle. \tag{28}$$

The terminal condition is

$$g^{(0)}(T, x, z) = G(x, z).$$

The problem for $g^{(0)}(t, x, z)$ is similar to the constant volatility problem (11), with two important differences:

1. The zero-order approximation $g^{(0)}$ depends not just on the usual long-run average historical volatility $\bar{\sigma}$, but also on the *harmonically-averaged volatility* $\sigma_\star$ defined by (28). Thus the asymptotic approximation of the optimal strategy will depend on estimating this unusual volatility too. By Jensen's inequality, $\sigma_\star \leq \bar{\sigma}$ and equality holds if and only if volatility is constant a.s.

2. The "homogenized" operator $\langle \mathcal{L}_2 \rangle$ is nondegenerate even though $\mathcal{L}_2$ is degenerate. As a result, $g^{(0)}(t, x, z)$ is the expectation of a functional of a *two-dimensional* Brownian motion, unlike the expectation in (12). In other words, the zero-order asymptotic approximation is not simply the complete market problem with constant averaged volatility.

The consequences of this are discussed in Section 6.2.

6.1.5   Zero-order Strategy

The optimal zero-order strategy is given by

$$\pi^\star = \left(x\frac{\partial}{\partial x} - \frac{\mu}{f(y)^2}z\frac{\partial}{\partial z}\right)g^{(0)}. \tag{29}$$

Notice that this *does* depend on tracking volatility $f(y)$ even though the corrected minimum expected loss does not (to zero-order). However, it does not depend on the difficult to estimate parameters of a volatility model, namely α, β, ρ, m, nor on a specific choice of model f.

6.1.6   Interpretation and Estimation of $\sigma_\star$

One possible way to estimate $\sigma_\star$ is to use the Taylor expansion

$$\frac{1}{\sigma_t^2} = \frac{1}{\bar{\sigma}^2 + (\sigma_t^2 - \bar{\sigma}^2)}$$

$$\approx \frac{1}{\bar{\sigma}^2}\left[1 - \left(\frac{\sigma_t^2 - \bar{\sigma}^2}{\bar{\sigma}^2}\right) + \left(\frac{\sigma_t^2 - \bar{\sigma}^2}{\bar{\sigma}^2}\right)^2\right],$$

so that

$$\langle \frac{1}{\sigma_t^2} \rangle \approx \frac{\langle \sigma_t^4 \rangle}{\bar{\sigma}^6}.$$

The long-run volatility $\bar{\sigma}$ and the fourth-moment $\langle \sigma_t^4 \rangle$ can be estimated stably from the second and fourth-moments of high-frequency historical returns. There is no need to specify a volatility model $f(Y_t)$.

This rough estimator also shows that

$$\frac{\bar{\sigma}^2}{\sigma_\star^2} \approx \frac{\langle \sigma_t^4 \rangle}{\bar{\sigma}^4}, \tag{30}$$

and so $\bar{\sigma}/\sigma_\star$ is a measure of excess kurtosis.

As an indication of the type of volatility observed in the market, we plot in Figure 6 the time-series of the VIX index produced by the CBOE, which is a measure of the daily volatility of US equity indices. It is produced from option implied volatilities and is only used here to gauge the order of volatility fluctuations. Taking this as a realization of volatility, we
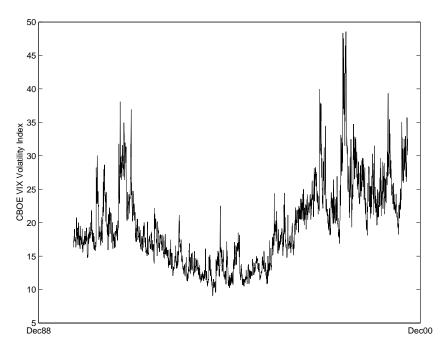

Figure 6: *CBOE VIX volatility index Dec 1988-Dec 2000 (in percentage units).*

estimate the two average volatilities:

$$\bar{\sigma} \approx 0.220 \qquad \sigma_\star \approx 0.177,$$

which gap describes wide fluctuations in this measure of volatility, as is seen from the picture.

6.2 Explicit Computations

In order to implement the zero order strategy we need to know the function $g^{(0)}(t, x, z)$ and its derivatives with respect to x and z. While these three quantities could be computed

numerically from a discretization of the linear PDE (26), it is quicker to use the Feynman-Kac representation (see e.g. [17]) of the solution. We refer to [12] for further details on the following argument (in the case of partial hedging). It is possible to write

$$g^{(0)}(t, x, z) = I\!E \left\{ g\left(t, x, za; \bar{\sigma}\right) \right\}, \tag{31}$$

where $g(t, x, z; \bar{\sigma})$ is the solution to the constant volatility ($= \bar{\sigma}$) problem (13) and the expectation is over η drawn from a standard $\mathcal{N}(0, 1)$ distribution. Here

$$a = a(\eta) = e^{\frac{\mu}{\sigma_\star} \sqrt{1 - \hat{\rho}^2} \sqrt{\tau} \eta + \frac{1}{2} \frac{\mu^2}{\sigma_\star^2} (1 - \hat{\rho}^2) \tau},$$

where the parameter $\hat{\rho}$ is given by

$$\hat{\rho} := \frac{\sigma_\star}{\bar{\sigma}}, \tag{32}$$

and satisfies $0 < \hat{\rho} \leq 1$ by Jensen's inequality. This parameter can be thought of as a correlation coefficient and is a measure of how much volatility is fluctuating. In particular, $\hat{\rho} = 1$ if and only if volatility is constant a.s. In the case $f(y) = e^y$ in (15), known as the expOU model, for example, the approximation (30) turns out to be exact and $\hat{\rho} = e^{-2\nu^2}$, where $\nu^2 = \beta^2/2\alpha$ is a measure of the *size* of volatility fluctuations.

We can also use (31) to compute derivatives of $g^{(0)}$. For instance, we get

$$g_x^{(0)}(t, x, z) = I\!E \left\{ g_x\left(t, x, za; \bar{\sigma}\right) \right\}$$
$$g_z^{(0)}(t, x, z) = I\!E \left\{ a g_z\left(t, x, za; \bar{\sigma}\right) \right\},$$

which further yields a formula for the optimal zero-order strategy using (29). The main point of the formulas above is that if we have an efficient way of computing the constant-volatility solution (such as in the case of partial hedging of a call option), then we can fairly easily compute the zero-order approximation through (31). Specifically, we may compute $g^{(0)}(t, x, z)$ (and its derivatives) on a grid in z as weighted averages of the constant volatility solutions $g(t, x, z; \bar{\sigma})$ on a corresponding grid.

7 Simulations

In this section we present numerical results that illustrate the approximation described above. We focus on partial hedging of a call option. In particular, we demonstrate the performance of the zero-order strategy suggested by the asymptotics and the Föllmer-Leukert strategy derived in [7] for a constant volatility market (see Section 3.4). However, in the latter at every rebalancing of the portfolio, we update the level of the volatility as would be done in market practice. We assume throughout the level of volatility is perfectly observable. A reasonable approximation of current volatility can usually be obtained each day (for example from near-to-expiration implied volatilities or by averaging intraday returns), but it is harder to assemble a reasonable time-series including intraday volatility levels on which to perform detailed model calibration.

7.1   Simulation Details

For the simulation, we discretized the explicit model

$$\frac{dX_t}{X_t} = \mu\,dt + e^{Y_t}\,dW_t$$

$$dY_t = \alpha(m - Y_t)\,dt + \nu\sqrt{2\alpha}\left(\rho\,dW_t + \sqrt{1 - \rho^2}\,dB_t\right)$$

$$dV_t = \pi_t \frac{dX_t}{X_t},$$

where $r = 0.03, \mu = 0.15 - r$, $K = 100$, $\rho = -0.2$ and (W_t) and (B_t) are independent Brownian motions. In the (Y_t) process, we fixed $\bar{\sigma} = 0.1$ and $\nu = 0.25$ and ran hedging tests for various α. Changing α does not change $\bar{\sigma}$ or $\sigma_\star$ because the invariant density of the OU process does not change if we hold $\bar{\sigma}$ and ν fixed. The characteristic speed of the process varies with α, but not the characteristic size of volatility fluctuations. In this case, with f the exponential function,

$$\sigma_\star = \bar{\sigma}e^{-2\nu^2} = 0.0882.$$

We used initial stock price $X_0 = 100$ (at the money) and initial capital equal to 25% of the Black-Scholes price with constant volatility $\bar{\sigma}$.

To compute the two strategies along the path we did as follows.

- For the Föllmer-Leukert strategy, given t, X_t and V_t we solved numerically the equation $g(t, X_t, z; \bar{\sigma}) = V_t$ for z, where g satisfies (13). The hedging ratio $\Delta = \pi/X$ is given by

$$\Delta_t = g_x(t, X_t, z; \bar{\sigma}) - \frac{\mu}{f(Y_t)^2}\frac{z}{X_t}g_z(t, X_t, z; \bar{\sigma})$$

  and the value of the portfolio was updated using $dV_t = \Delta_t\,dX_t$. Notice that the level of volatility is updated in this hedge as would be done in practice.

- For the zero-order strategy we did as follows. Given t, X_t and V_t we solved the equation $g^{(0)}(t, X_t, z) = V_t$ for z, with $g^{(0)}$ from (31). The hedging ratio was then chosen as

$$\Delta_t = g_x^{(0)}(t, X_t, z) - \frac{\mu}{(\exp(Y_t)^2)}\frac{z}{X_t}g_z^{(0)}(t, X_t, z)$$

  and the value of the portfolio updated using $dV_t = \Delta_t\,dX_t$.

Throughout, we used $p = 1.1$; see Section 1.2.1 for a comment on this choice. The time horizon of the problem is $T = 0.5$ years, so we are hedging a six-month option. We used a time-discretization of 200 even intervals over this period, corresponding roughly to rehedging twice a day. This might be more frequent than in practice, but we wanted to compare the two strategies rather than deal with significant discretization issues.

7.2   Large α Simulation

We first compared the strategies when the volatility *is* fast mean-reverting as assumed by the asymptotics. We set $\alpha = 200$ in annualized units which corresponds to a typical mean-reversion time of about one day. We computed estimates for the expected losses

$$L = \mathbb{E}\{(\frac{1}{p}(X_T - K - V_T)^+)^p\}$$

for the two strategies.

With 60000 paths, we got $L_1 = 2.979$ (Föllmer-Leukert) and $L_2 = 2.903$ (zero-order), so the zero-order strategy outperformed the Föllmer-Leukert strategy by approximately 2.6%. This is illustrated in Figure 7. One natural question is whether the improvement in the
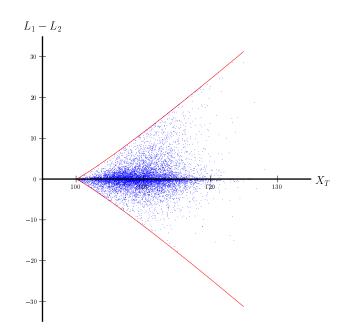

Figure 7: *Difference in losses between Föllmer-Leukert strategy and zero-order strategy as a function of the terminal stock price for* 60,000 *paths with* $\alpha = 200$.

expected loss comes at significant cost to the variance. From these simulations, the answer is no (for the Föllmer-Leukert strategy, the standard deviation of the loss was 4.46, and it was 4.21 for the zero-order strategy.)

7.3 Small α Simulation

In some markets, the rate of mean-reversion may not be so rapid. Nonetheless, the zero-order strategy may still have some effect as a crude approximation to the full stochastic volatility strategy. We compared the strategies when the volatility is mean-reverting with $\alpha = 1$ in annualized units which corresponds to a typical mean-reversion time of about one year.

With 60000 paths, we got $L_1 = 3.156$ (Föllmer-Leukert) and $L_2 = 3.111$ (zero-order), so the zero-order strategy outperformed the Föllmer-Leukert strategy by approximately 1.4%. This is illustrated in Figure 8.

The standard deviations were 3.88 for the Föllmer-Leukert strategy and again it was lower (3.68) for the zero-order strategy.
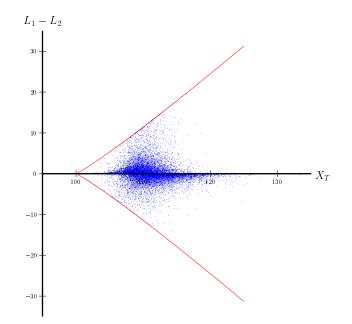
Figure 8: *Difference in losses between Föllmer-Leukert strategy and zero-order strategy as a function of the terminal stock price for* 60,000 *paths with* $\alpha = 1$.

8 Conclusions

In this article we have reviewed some stochastic optimization problems that are the subject of much current research in financial mathematics. These present significant computational challenges, particularly within realistic market models such as those with stochastic volatility that we have focused on here.

Our approach is to use a homogenization approximation exploiting rapid mean-reversion in the stochastic process driving volatility. This leads to a strategy that, as demonstrated by simulation, performs better than the strategy based on assuming volatility is constant (the usual root-mean-square average $\bar{\sigma}$), but with minimal implementation overhead. The key ingredient is the harmonically averaged volatility $\sigma_\star$ which encapsulates the effect of mean-reverting stochastic volatility for these problems. Obviously, the degree of suboptimality of this approximation will be extremely model dependent, but since volatility models and parameters are not specified with much confidence in practice, this robust method is preferable to solving a high-dimensional Bellman equation whose parameters are anyway uncertain.

Additionally, as is often found with homogenization approximations, the zero-order strategy is an improvement *even outside the regime of fast mean-reversion.* This is demonstrated in the small α simulations above.

References

[1] D. Bertsimas, L. Kogan, and A. Lo. Pricing and hedging derivative securities in incomplete markets: an epsilon-arbitrage approach. Technical report, MIT Sloan School,

1999.

[2] F. Black and M. Scholes. The Pricing of Options and Corporate Liabilities. *J. Political Econ.*, 81:637–659, 1973.

[3] J. Cvitanić and I. Karatzas. On dynamic measures of risk. *Finance and Stochastics*, 3(4), 1999.

[4] M. Davis, V. Panas, and T. Zariphopoulou. European pricing with transaction costs. *SIAM J. Control and Optim.*, 31:470–93, 1993.

[5] D. Duffie and H. Richardson. Mean-variance hedging in continuous time. *Annals of Applied Probability*, 1:1–15, 1991.

[6] N. ElKaroui, M. Jeanblanc-Picque, and V. Lacoste. Optimal portfolio management with American capital guarantee. Technical report, CMAP, Ecole Polytechnique, November 2000. Preliminary version.

[7] H. Főllmer and P. Leukert. Efficient hedging: Cost versus shortfall risk. *Finance and Stochastics*, 4(2), 2000.

[8] W. H. Fleming and H. M. Soner. *Controlled Markov Processes and Viscosity Solutions.* Springer-Verlag, 1993.

[9] J.-P. Fouque, G. Papanicolaou, K. R. Sircar, and K. Solna. Mean-Reversion of S&P 500 Volatility. *Submitted*, 2000.

[10] J.-P. Fouque, G. Papanicolaou, and K.R. Sircar. *Derivatives in Financial Markets with Stochastic Volatility.* Cambridge University Press, 2000.

[11] J. Hull and A. White. The Pricing of Options on Assets with Stochastic Volatilities. *J. Finance*, XLII(2):281–300, June 1987.

[12] M. Jonsson and K. R. Sircar. Partial Hedging in a Stochastic Volatility Environment. *Submitted*, 2000.

[13] I. Karatzas and S. Shreve. *Methods of Mathematical Finance.* Springer-Verlag, 1998.

[14] R. C. Merton. Lifetime portfolio selection under uncertainty: the continous-time case. *Rev. Econom. Statist.*, 51:247–257, 1969.

[15] R. C. Merton. Theory of rational option pricing. *Bell Journal of Economics*, 4(1):141–183, Spring 1973.

[16] R. C. Merton. *Continuous-Time Finance.* Blackwell, 1992.

[17] B. Øksendal. *Stochastic Differential Equations, 5th ed.* Springer-Verlag, 1998.