Modelling liquidity and its effects on price

L. C. G. Rogers* and Surbjeet Singh
University of Cambridge and University of Bath

January 26, 2006

1 Introduction

After credit risk, liquidity risk is probably the next most important risk faced by the finance industry; and yet the study of liquidity is far less advanced. This may be in part due to the fact that there is no agreed definition of what liquidity is, even in qualitative terms; everyone would agree that the effect of illiquidity is to make it difficult or costly to trade large volumes of the underlying asset in small times, but there are different approaches to modelling this.

One modelling philosophy is that trading large amounts moves the price, and the papers of [7], [8], [16], [15], [17] are examples of this viewpoint, where the stock price responds instantaneously to the amount of the stock held by a single large trader. Though such models have a flavour of liquidity, we regard them rather as models of feedback effects. There is evidence that the actions of a large trader can influence the price of the underlying (see, for example, [12] and [11]). A large trader can try to 'corner the market'. One way to do this in a commodity market is to take a huge long futures position and at the same time buy up the underlying commodity. As expiry approaches the investors who are short the futures contract may find that there is not enough supply to meet their demand and hence the price is pushed up. One such alleged case of this was the activities of the Hunt brothers in the silver market in 1979-1980. Their trading caused the price to increase from $9 per ounce to $50 per ounce. Another way to 'corner the market', this time involving shares, is to buy up a large supply and then lend some to investors who want to go short. When these shares are sold on the market the agent buys them up and then calls in the short shares. Since the agent has limited the supply of shares by buying a large amount this pushes the price up. These types of manipulation involve taking huge positions in the underlying and documented examples of this type of activity have shown large price increases followed by a crash.

Though interesting in their own right, such feedback models have drawbacks which make them unsuitable as models of liquidity. One of these is the 'free round trip' phenomenon, discussed in [17]; if the large agent rapidly sells and then buys back a large amount of stock, he can force the price instantaneously to drop, and if this round trip is not costly (as is the case in some studies), then the large agent could make profits by selling down-and-out calls, and

*Corresponding author: Statistical Laboratory, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK; L.C.G.Rogers@statslab.cam.ac.uk

subsequently knocking them out by a round trip. Another problem with such feedback models is that they typically present the solution to a hedging problem in feedback form, exhibiting the hedge as a function of time, and current stock price - but if the initial portfolio is not at the exactly correct value, it is not clear how it is to be moved to that value. A more serious problem with such models is that if the actions of *one* agent may affect the price, then logically the actions of *all* agents may affect the price, and the resulting analysis of the inter-related behaviour of the many agents in the market becomes impossibly cumbersome; see, for example, [9] for a partial analysis of such a setup.

Our viewpoint here is to start from a simple discrete-time microeconomic story, and try to derive the dynamics of an agent's wealth process from that. Dividing the time axis into equal intervals of length Δt, we suppose that the number of shares offered for sale or purchase in that interval is of the order of Δt. An agent who wishes to buy a large number of shares in any one time interval will have to pay an inflated price, because he is trading in a shallow market. On the other hand, the amount he actually trades, itself of order Δt, is small relative to the total number of shares in the firm, so will have negligible impact on the price of the shares in the next period. What transpires is that the impact of liquidity modelled in this fashion is like a transaction cost, but not one which is proportional to the amounts traded. In the conventional proportional transaction cost model (see [14], [5]), the bid-offer spread (the difference between the lowest offer price and highest bid price) is proportional to the current asset price, but the depth of the limit order book is not taken into account. Our approach is to model illiquidity as a transaction cost taking into account the depth of the book and the trade size; [4] and [18] explore further aspects of the model introduced here.

The approach of [13] is closest in spirit to ours; the assumption there is that the holding of stock must be a finite-variation process with bounded derivative, which may be thought of as a special limiting case of the model we propose. Other models which feature transaction costs include [2] and [3]. However, in both their models transaction costs can be completely avoided by following a continuous trading strategy of bounded variation. This seems rather unrealistic. The paper of Bakstein & Howison [1] shares a number of features with ours; one main difference is that it leads to feedback effects, which we are trying to eliminate.

The layout of the remainder of the paper is as follows. In Section 2, we present the discrete-time model underlying our analysis, deriving a continuous-time model by formally passing to the limit as the time step decreases to zero. In Section 3, we consider how a European put option would be priced by CARA-utility-indifference; we find that the utility-indifference price satisfies a PDE which cannot be solved in closed form. Accordingly, in Section 4, we seek an asymptotic for the price and the hedging strategy that should be used, presenting in Section 5 some numerical results to assess the accuracy of the asymptotic.

2 Modelling liquidity.

We shall consider a single asset, whose price at time t is denoted by S_t; for simplicity of exposition, we shall assume a zero interest rate for the time being. Fixing now some $\Delta t > 0$, we let p_n denote $\log S(n\Delta t)$, which we suppose evolves as a random walk,

$$p_n = p_{n-1} + \xi_n, \tag{2.1}$$

where the ξ_n are independent and identically distributed. We imagine that in each period $\lambda_a \Delta t$ shareholders consider whether to sell their shares, and $\lambda_b \Delta t$ agents consider whether to buy shares. A hedger also comes to the market with the intention of buying ΔH_n shares during the n^{th} period $((n-1)\Delta t, n\Delta t] \equiv (t_{n-1}, t_n]$. The log-price $\tilde{p}_n$ at which he will trade is determined by the equalisation of supply and demand:

$$\lambda_a \Delta t\, f_s(\tilde{p}_n - p_n) = \lambda_b \Delta t\, f_d(\tilde{p}_n - p_n) + \Delta H_n, \tag{2.2}$$

where the supply function f_s is continuous and strictly increasing, the demand function f_d is continuous and strictly decreasing. We therefore find that the log-price at which the hedger trades is determined by

$$\tilde{p}_n - p_n = \psi\left(\frac{\Delta H_n}{\Delta t}\right), \tag{2.3}$$

where ψ is the inverse function to $x \mapsto (\lambda_a f_s(x) - \lambda_b f_d(x))$. It is natural to suppose that $\psi(0) = 0$.

If we now let H_t denote the number of shares held by the hedger at time t, and K_t denote the amount of cash held by the hedger at time t, then the wealth[1] of the hedger at time t is $w_t = H_t S_t + K_t$. The change in wealth over the n^{th} period is therefore

$$\begin{aligned} w_{t_n} - w_{t_{n-1}} &= H_{t_{n-1}}(S_{t_n} - S_{t_{n-1}}) + S_{t_n}\Delta H_n - e^{\tilde{p}_n}\Delta H_n \\ &= H_{t_{n-1}}(S_{t_n} - S_{t_{n-1}}) - S_{t_n} h_n(\exp(\psi(h_n)) - 1)\Delta t, \end{aligned}$$

where we use the notation $h_n = \Delta H_n / \Delta t$. If we let $\Delta t \downarrow 0$, and suppose that H is differentiable, with derivative $h_t = dH_t/dt$, then we derive the (continuous-time) dynamics for wealth in the form

$$dH_t = h_t dt, \tag{2.4}$$

$$dw_t = H_t dS_t - h_t S_t f(h_t) dt, \tag{2.5}$$

where $f(x) \equiv \exp(\psi(x)) - 1$ is continuous and increasing, equal to 0 at 0.

We shall therefore now concentrate on the continuous-time dynamics (2.4-2.5), and let the discrete-time construction by which we arrived at it slip into the background. Various forms for f can be considered; since we intend to treat the illiquidity costs as small, we shall introduce the small parameter ε into the argument of f, leading to wealth dynamics

$$dw_t = H_t dS_t - h_t f(\varepsilon h_t) S_t dt. \tag{2.6}$$

It is natural to assume that f is continuous, increasing, bounded below by -1, and vanishes at 0.

In what follows, we shall suppose that S is a standard log-Browian motion, $dS_t = S_t(\sigma dW_t + \mu dt)$, and that the riskless rate r is a constant. The dynamics of the hedger's wealth then satisfy

$$dw_t = r w_t dt + H_t S_t(\sigma dW_t + (\mu - r)dt) - h_t f(\varepsilon h_t) S_t dt, \tag{2.7}$$

which we see reduces to the usual equation when there is no liquidity cost (that is, $f \equiv 0$).

[1] We are talking here of the *paper* or *nominal* wealth at time t, which is not the same as the *liquidation value* of the hedger's portfolio.

3 Utility-indifference pricing and the HJB equation

The model presented in Section 2 is incomplete, and so perfect replication is impossible. This means of course that there is no unique replication price as in the classical Black-Scholes paradigm, and we have to consider what is meant by a price. Alternative notions for price in incomplete markets have been advanced; the notion of super-replication, for example, is sometimes used, but turns out to be far too conservative. What we shall do here is to take the utility-indifference price, as introduced by Hodges & Neuberger [10], and studied by Davis, Panas & Zariphopoulou [6].

Informally, the utility-indifference (buy) price $\Pi_0(\eta)$ at time 0 of some bounded $\mathcal{F}_T$-measurable contingent claim η is obtained by evaluating[2]

$$\sup\big\{EU(X+\eta-p) : X \in \mathcal{X}_T\big\} \tag{3.1}$$

for p a constant parameter, and then varying p until the supremum matches $\sup\{EU(X) : X \in \mathcal{X}_T\}$, where $\mathcal{X}_T$ is the set of all wealths attainable at time T. In the standard Black-Scholes set-up, $\mathcal{X}_T$ depends only on w_0, the initial wealth, but with liquidity costs it will actually depend also on H_0, the initial holding of the stock. With liquidity costs, there is also an issue about what to do at time T; we shall assume that at time T the hedger may convert his holding of stock immediately into cash at no cost. Some convention must be made about time T, and this seems as good as any other; ignoring liquidity at time T may introduce some error, but it is of smaller order than the loss due to liquidity effects of hedging.

Writing $Y_t = w_t - H_t S_t$ for the cash held at time t, we shall now address the optimisation problem (3.1) by setting

$$V(t,H,Y,S,p) \equiv \sup E\big[\, U(w_T+\eta-p) \;\big|\; H_t=H, S_t=S, Y_t=Y \,\big], \tag{3.2}$$

and finding the equation to be satisfied by V. We shall simplify by assuming that the utility is CARA ($U(x) = -\exp(-\gamma x)$), and that $\eta = q(T,S_T)$ is a bounded contingent claim which depends only on the final value of the stock (such as a long or short put.) We shall let $q(t,S_t)$ be the Black-Scholes price at time t of the claim η; as is well known, q satisfies the PDE

$$\frac{1}{2}\sigma^2 S^2 q_{SS} + rSq_S + q_t - rq = 0. \tag{3.3}$$

The CARA assumption allows us to factor out the terms in wealth from (3.2), leaving

$$\begin{aligned} V(t,H,Y,S,p) &= \exp\{-\gamma(Ye^{r\tau}-p)\}\sup E\big[\, U(w_T+\eta) \;\big|H_t=H, S_t=S, Y_t=0\,\big] \\ &\equiv \exp\{-\gamma(Ye^{r\tau}-p)\}F^*(t,H,S). \end{aligned} \tag{3.4}$$

Here, we have used the abbreviation $\tau = T-t$.

Now the dynamics of the various processes are

$$\begin{aligned} dS &= S(\sigma dW + \mu dt), \\ dH &= hdt, \\ dY &= dw - d(HS) \\ &= dw - HdS - SdH \\ &= rYdt - hSf(\varepsilon h)dt - hSdt, \end{aligned}$$

[2] We are here thinking of the payment p being received at time T; if the payment were to be received at time 0, the we should replace p by pe^{rT} in (3.1).

so the HJB equation to be satisfied by V is

$$
\begin{aligned}
0 &= \sup_h \{ V_t + hV_H + (rY - hS(1+f(\varepsilon H)))V_Y + \frac{1}{2}\sigma^2 S^2 V_{SS} + \mu S V_S \} \\
&= \exp(-\gamma e^{r\tau} Y) \sup_h \{ F^*_t + hF^*_H + \gamma e^{r\tau} F^* Sh(1+f(\varepsilon h)) + \frac{1}{2}\sigma^2 S^2 F^*_{SS} + \mu S F^*_S \}, \qquad (3.5)
\end{aligned}
$$

taking $p = 0$ with no loss of generality. We intend to express F^* as some perturbation away from the case $\varepsilon = 0$, so we have to know what its form would be for the (Black-Scholes) case $\varepsilon = 0$. But in this case, the holding H_t of the stock can be freely changed to (or from) $H_t S_t$ in cash, and contingent claim η delivered at T is freely exchangeable for the Black-Scholes price $q(t, S_t)$ at time t. Thus (writing F^0 for the case $\varepsilon = 0$)

$$
\begin{aligned}
F^0(t,H,S) &= \sup E\big[U(w_T + \eta) \bigm| H_t = H, S_t = S, Y_t = 0 \big] \\
&= \exp\{-\gamma e^{r\tau}(HS + q(t,S))\} \sup E\big[U(w_T) \bigm| w_t = 0, S_t = S \big] \\
&= -\exp\big\{ -\gamma e^{r\tau}(HS + q(t,S)) - \frac{(\mu - r)^2 \tau}{2\sigma^2} \big\}, \qquad (3.6)
\end{aligned}
$$

as may be checked by straightforward calculations. The optimal wealth to invest in the stock is

$$
\frac{\mu - r}{\gamma \sigma^2 e^{r\tau}} - S q_S. \qquad (3.7)
$$

Thus we write the solution for general ε as

$$
F^*(t,H,S) = -\exp\{ -\gamma e^{r\tau}(HS + q(t,S)) - \frac{(\mu - r)^2 \tau}{2\sigma^2} + \gamma \varphi(t,H,S) \}, \qquad (3.8)
$$

where φ is going to be small in some sense, and non-negative, in view of the obvious inequality $F^* \le F^0$. Assuming that[3] the function $h \mapsto \psi(h) \equiv hf(h)$ is strictly convex, with conjugate function

$$
\tilde{\psi}(\lambda) = \inf_t \{ \psi(t) + \lambda t \}, \qquad (3.9)
$$

a few lines of calculation reduce the HJB equation (3.5) to

$$
\begin{aligned}
0 &= \frac{1}{2}\sigma^2 S^2 F^*_{SS} + \mu S F^*_S + F^*_t + \frac{\gamma e^{r\tau} S F^*}{\varepsilon} \tilde{\psi}\left(\frac{F^*_H + \gamma S e^{r\tau} F^*}{\gamma S e^{r\tau} F^*} \right) \qquad (3.10) \\
&= \frac{1}{2}\sigma^2 S^2 F^*_{SS} + \mu S F^*_S + F^*_t + \frac{\gamma e^{r\tau} S F^*}{\varepsilon} \tilde{\psi}\left(\frac{\varphi_H}{S e^{r\tau}} \right), \qquad (3.11)
\end{aligned}
$$

with the additional information that the optimal choice of h satisfies

$$
\psi'(\varepsilon h) = -\frac{F^*_H + \gamma S e^{r\tau} F^*}{\gamma S e^{r\tau} F^*} = -\frac{\varphi_H}{S e^{r\tau}}. \qquad (3.12)
$$

[3]We also assume that f is continuous, strictly increasing, and vanishing at zero, so that $\psi \ge 0$, $\tilde{\psi} \le 0$.

After some lengthy but straightforward calculations, using (3.3), we arrive at the relation

$$\frac{\frac{1}{2}\sigma^2 S^2 F^*_{SS} + \mu S F^*_S + F^*_t}{\gamma F^*} = \frac{1}{2}\sigma^2\gamma\left(Se^{r\tau}\tilde{H} - \frac{\mu - r}{\gamma\sigma^2} - S\varphi_S \right)^2 + \{\frac{1}{2}\sigma^2 S^2 \varphi_{SS} + rS\varphi_S + \varphi_t\},$$

where we write $\tilde{H} \equiv H + q_S$. Returning this to the HJB equation, we get that φ must satisfy

$$\frac{1}{2}\sigma^2 S^2 \varphi_{SS} + rS\varphi_S + \varphi_t + \frac{1}{2}\sigma^2\gamma\left(Se^{r\tau}\tilde{H} - \frac{\mu - r}{\gamma\sigma^2} - S\varphi_S \right)^2 + \frac{e^{r\tau}S}{\varepsilon}\,\tilde{\psi}\left(\frac{\varphi_H}{Se^{r\tau}}\right) = 0. \tag{3.13}$$

In the absence of liquidity, we saw that the optimal wealth to invest in the stock was $-Sq_S + (\mu - r)e^{-r\tau}/\gamma\sigma^2$, so if we set

$$H^0 \equiv H^0(t, S) = -q_S(t, S) + \frac{\mu - r}{\gamma\sigma^2 e^{r\tau} S}, \tag{3.14}$$

the PDE for φ becomes

$$\frac{1}{2}\sigma^2 S^2 \varphi_{SS} + rS\varphi_S + \varphi_t + \frac{1}{2}\sigma^2\gamma S^2 e^{2r\tau}(H - H^0 - e^{-r\tau}\varphi_S)^2 + \frac{e^{r\tau}S}{\varepsilon}\,\tilde{\psi}\left(\frac{\varphi_H}{Se^{r\tau}}\right) = 0. \tag{3.15}$$

4 Asymptotic analysis

Our aim in this section is to study the approximate form of the solution to the problem when ε is small. We shall for simplicity assume that

$$f(x) = x/2,$$

so that $\psi(x) = -\tilde{\psi}(x) = x^2/2$. Although this is not consistent with our previous assumptions on f, we shall see that it is only the behaviour of $\tilde{\psi}$ near zero that really matters, so the approximation may be expected to be reasonable.

We shall also assume that $\mu = r$; this assumption means that the target hedge ratio H^0 depends only on the derivative, and not on the rate of growth μ, which is notoriously difficult to estimate. Moreover, in the context of hedging the derivative, it would be hard to persuade a sceptical trader that his portfolio should not be zero even if he had no derivative to hedge.

Again for simplicity, we shall assume that $r = 0$; this entails no real loss of generality, because we could work instead with discounted stock prices, and discounted wealth, which would only change γ to γe^{rT}.

With these simplifying assumptions, (3.15) becomes

$$\frac{1}{2}\sigma^2 S^2 \varphi_{SS} + \varphi_t + \frac{1}{2}\sigma^2\gamma S^2 (H - H^0 - \varphi_S)^2 - \frac{\varphi_H^2}{2\varepsilon S} = 0, \tag{4.1}$$

where $H^0 = -q_S$, and q solves the Black-Scholes PDE with $r = 0$:

$$\frac{1}{2}\sigma^2 S^2 q_{SS} + q_t = 0. \tag{4.2}$$

Now for ε small, we expect that φ will be small, and we also expect that the value will change little with H, since any initial position can quickly and cheaply be turned into another quite different position. Looking at (4.1), the third term is quadratic in H, so is of much larger order than the first two; the only way the expression on the left can be zero is therefore if the final term $-\varphi_H^2/2\varepsilon S$ counterbalances the third term. This leads us to guess that we must have approximately

$$\varphi_H \simeq \delta\sqrt{\gamma}\,\sigma S^{3/2}(H - H^0),$$

where $\delta \equiv \sqrt{\varepsilon}$. From (3.12), we conclude that we ought to have

$$h \simeq h_0 \equiv -\sigma\sqrt{\frac{\gamma S}{\varepsilon}}(H - H^0). \tag{4.3}$$

Notice the clear intuitive content of the rule implied by (4.3): in the absence of liquidity costs, we would be holding the Black-Scholes hedge H^0, and when we are compelled to follow a finite-variation strategy, we simply mean-revert to the Black-Scholes hedge, with a strength proportional to $\sqrt{S/\varepsilon}$. This is the kind of rule that one might guess without the benefit of a model of liquidity (or even a definition!) but such an intuition would not help much, as it would give no guidance on how the strength of mean-reversion should be chosen.

What we intend to do now is to find the equation satisfied by the value when the policy h_0 is applied, and then to study how close it is to the value we would get if there were no liquidity costs. From the reasoning that led to (3.5), we find that the value for using policy h_0 is $\exp\{-\gamma Y\}F(t, H, S)$, where F solves

$$\mathcal{L}F \equiv F_t + h_0(F_H + \gamma SF) + \frac{1}{2}\gamma\varepsilon h_0^2 FS + \frac{1}{2}\sigma^2 S^2 F_{SS} = 0. \tag{4.4}$$

Recalling (3.6), that when there are no liquidity costs $F = F^0$, where

$$F^0(t, H, S) = -\exp(-\gamma(HS + q(t, S)),$$

we shall write $F = F^0 - g$, where g will be non-negative, and small. A few calculations turn (4.4) into

$$(-F^0)\gamma^2\sigma^2 S^2(H - H^0)^2 + g_t + h_0(g_H + \gamma Sg) + \frac{1}{2}\gamma\varepsilon h_0^2 Sg + \frac{1}{2}\sigma^2 S^2 g_{SS} = 0. \tag{4.5}$$

Substituting the form (4.3) of h_0 gives us more explicitly

$$g_t - \sigma\sqrt{\frac{\gamma S}{\varepsilon}}(H - H^0)(g_H + \gamma Sg) + \frac{1}{2}\gamma^2\sigma^2 S^2(H - H^0)^2 g + \frac{1}{2}\sigma^2 S^2 g_{SS} = F^0\gamma^2\sigma^2 S^2(H - H^0)^2. \tag{4.6}$$

In a more compact notation, this is

$$\mathcal{G}g - \frac{1}{\sqrt{\varepsilon}}\mathcal{A}g = \gamma^2\sigma^2 S^2(H - H^0)^2 F^0, \tag{4.7}$$

where

$$\begin{aligned}
\mathcal{G} &\equiv \frac{1}{2}\sigma^2 S^2\frac{\partial^2}{\partial S^2} + \frac{\partial}{\partial t} + \frac{1}{2}\gamma^2\sigma^2 S^2(H - H^0)^2, \\
\mathcal{A} &\equiv \sigma\sqrt{\gamma S}(H - H^0)\left(\frac{\partial}{\partial H} + \gamma S\right)
\end{aligned}$$

We plan now to perform a *formal* expansion of g as

$$g \equiv \sum_{m \geq 1} \varepsilon^{m/2} g^{(m)},$$

and see where this takes us. Substituting into (4.7) and comparing terms in the different powers of $\sqrt{\varepsilon}$ gives

$$-\mathcal{A} g^{(1)} = \gamma^2 \sigma^2 S^2 (H - H^0)^2 F^0, \tag{4.8}$$

$$\mathcal{G} g^{(m)} = \mathcal{A} g^{(m+1)} \qquad (m \geq 1). \tag{4.9}$$

We can therefore build the functions $g^{(m)}$ recursively, by solving a (first-order, linear, ordinary) differential equation at each step. To see how this works, we start with (4.8) which is

$$-\sigma \sqrt{\gamma S}(H - H^0)(g_H^{(1)} + \gamma S g^{(1)}) = F^0 \gamma^2 \sigma^2 S^2 (H - H^0)^2, \tag{4.10}$$

which is solved by

$$g^{(1)}(t, H, S) = e^{-\gamma H S} \{ f^{(1)}(t, S) + \frac{1}{2} (H - H^0)^2 (\gamma S)^{3/2} \sigma e^{-\gamma q(t,S)} \}. \tag{4.11}$$

Here, $f^{(1)}$ is some function of (t, S) that is to be determined. Quite how it is to be determined becomes clear when we look at (4.9) for $m = 1$; more fully, this equation is

$$\mathcal{G} g^{(1)} = \sigma \sqrt{\gamma S}(H - H^0)(g_H^{(2)} + \gamma S g^{(2)}). \tag{4.12}$$

Now it is clear that the right-hand side vanishes when $H = H^0$, and evaluating the left-hand side at $H = H^0$ gives the second-order linear PDE

$$\frac{1}{2} \sigma^2 S^2 f_{SS}^{(1)} + \gamma \sigma^2 S^2 q_S f_S^{(1)} + f_t^{(1)} + \frac{1}{2} \gamma^2 \sigma^2 S^2 q_S^2 f^{(1)} + 2\gamma^{3/2} e^{-\gamma q} q_t^2 / (\sigma \sqrt{S}) = 0 \tag{4.13}$$

for $f^{(1)}$. Under suitable regularity conditions on the contingent claim, together with the boundary condition $f^{(1)}(T, \cdot, \cdot) = 0$, the PDE (4.13) has a solution given by its Feynman-Kac representation, and moreover this solution is non-negative. We now suppose that this is the $f^{(1)}$ to be used in the definition of $g^{(0)}$. Thus the left-hand side of (4.12) is completely known, and solving (4.12) for $g^{(2)}$ is the same problem mathematically as solving (4.10). Once again, there is an undetermined function $f^{(2)}(t, S)$, which is fixed when we take (4.9) for $m = 2$. We may in principle continue in this way indefinitely, though there is of course no hope of explicit solutions.

There is a further issue with the formal expansion constructed in this way; the function g does not vanish at $t = T$. Indeed, from (4.11) we see that

$$g^{(1)}(T, H, S) = \frac{1}{2} e^{-\gamma H S} (H - H^0)^2 (\gamma S)^{3/2} \sigma e^{-\gamma q(T,S)}.$$

However, it is clear that this is non-negative, so the perturbation g that we calculate corresponds to the problem of optimally hedging if at time T we face this final non-negative cost. But what we were assuming was that at time T we face no cost at all, so the g constructed by the process described above should be an upper bound for the loss incurred in the problem as originally stated. What we learn then is that the loss due to illiquidity should be[4] $O(\sqrt{\varepsilon})$.

The issue of how we should account for liquidity at the expiry of the option is not clear-cut. Our assumption is that at that time liquidity costs are suspended, and the position can be liquidated immediately at no cost. Perhaps better would be to suppose that the agent must liquidate the holding of stock over a period. But how long will the agent be allowed to spend in liquidating the position? The first idea would be that he is compelled to borrow the money to pay the cash value of the option, and then accrues interest on the borrowing as he gradually liquidates the stock position. His optimal behaviour for this run-off period now constitutes a further optimisation problem, whose solution determines the boundary condition for the first optimisation problem, and the model begins to lose whatever charm it may have had in the first place. We shall not pursue these questions, but make just one observation; the costs of liquidation at or just after expiry must be $O(\varepsilon)$. Indeed, any policy that the agent might follow for $\varepsilon = \varepsilon_0$ could just as well be followed for $\varepsilon = \lambda\varepsilon_0$, at exactly λ times the cost.

Our view is that the analysis of this Section suggests an approximation to the cost of liquidity, and a natural approximation to the optimal hedging rule; in the next Section, we shall carry out a numerical analysis of an example to compare values from the numerical solution of the PDE (3.5) with the supposed asymptotics.

5 Numerical Results

The numerical results reported all refer to the pricing of a short European put option. Recalling that we take $\mu = r = 0$, (3.10) becomes

$$\frac{1}{2}S^2\sigma^2 F^*_{SS} + F^*_t = \frac{(F^*_H + \gamma S F^*)^2}{2\epsilon\gamma S F^*} \tag{5.1}$$

and (4.3) becomes

$$h_t = -\sqrt{\frac{S\gamma\sigma^2}{\varepsilon}}(H - H^0(t,S)). \tag{5.2}$$

In this section we solve the PDE (5.1) numerically, with the boundary condition $F^*(T,H,S) = \exp(-\gamma(HS + q(T,S))$ and compare the optimal control at time t=0 to (5.2). The equation (5.1) is a nonlinear PDE in two space variables and time, so it is difficult to solve; we use a mixed technique to solve it. For the left-hand side of (5.1) we use a Crank-Nicolson scheme, and for the right-hand side we use an explicit scheme. This has the advantage that at each timestep we are solving a series of 1-dimensional problems. We use natural boundary conditions for the S variable. In the H direction, when calculating the first derivative we use a 5-point central difference scheme when not including the boundary (H=-1 and H=1) and 6 points if we include the boundary. The S-grid has 101 points, ranging from 30/101 to 30; the H-grid has 51 points, ranging from -1 to 1. The parameter values we use are $\sigma = 0.2$, $\gamma = 0.01$, $\varepsilon = 0.002$, $T = 4$, $K = 15$. The number of timesteps we use is 80000.

Comparing the optimal control from the numerical procedure with the asymptotic control at the time 0, Figure 1, we have excellent agreement over a range of H-values.

[4] We do not claim that necessarily the losses $\sim \sqrt{\varepsilon}$; it seems likely that this is the case, but we have no evidence for this.

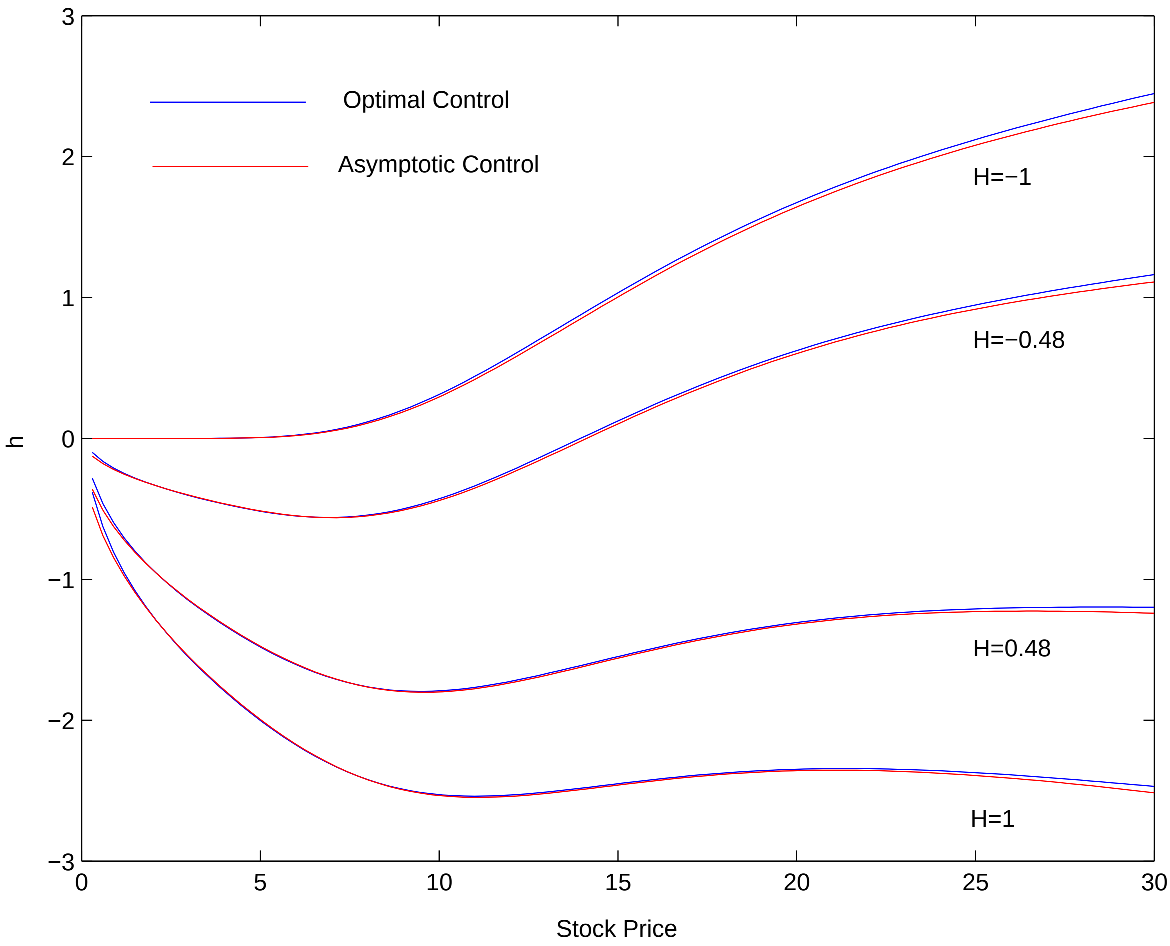

Figure 1: Optimal Control

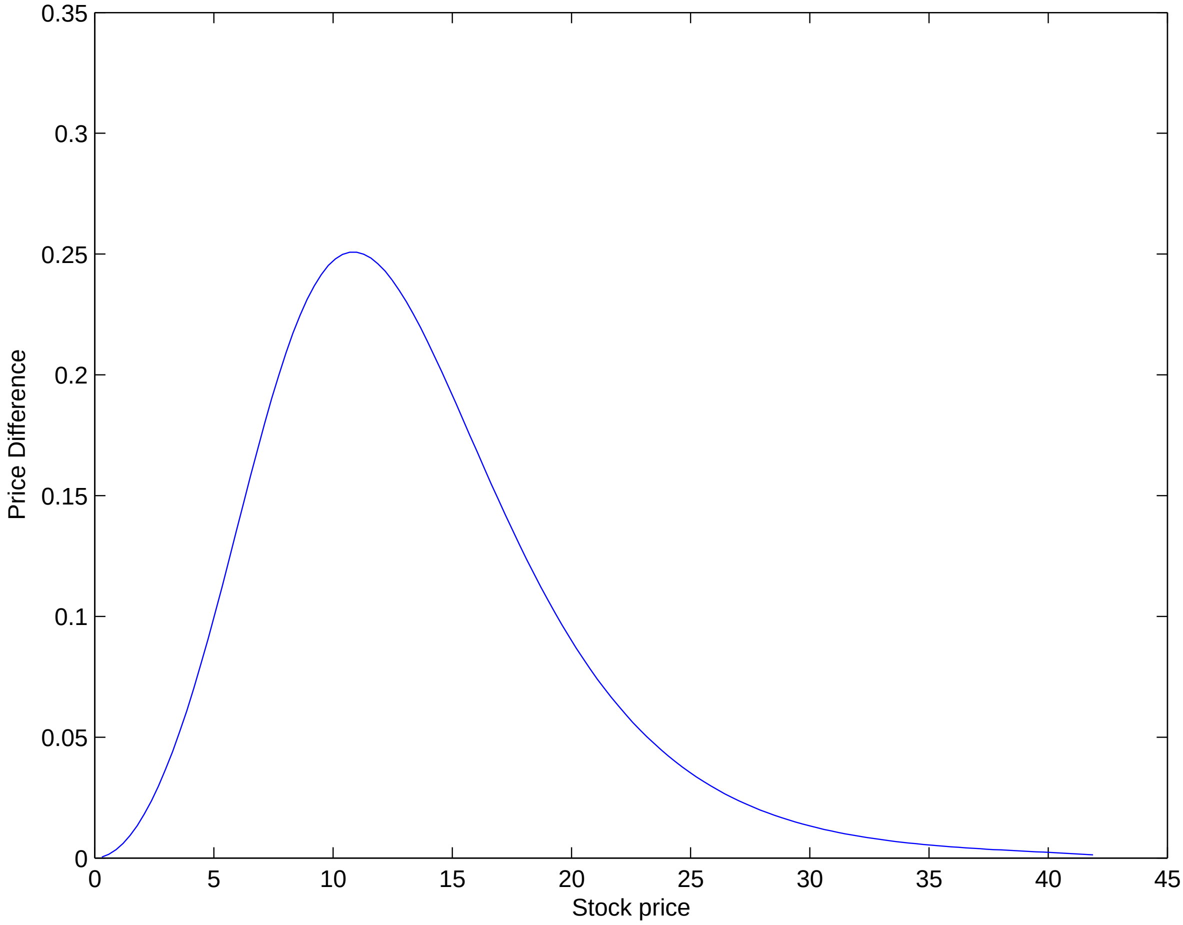

Figure 2: Illiquid Option Price minus Black Scholes Price

Changing two of the parameters, $\varepsilon = 2$ and $\gamma = 0.05$ so as to make the price difference a bit more substantial, we calculated the option price under liquidity costs and compared with the Black-Scholes price, this time using 160000 timesteps. Figure 2 shows the difference between the option price in the illiquid market and the Black Scholes price as a function of share price, evaluated at $H = 0$. The largest difference seems to be a little below the strike price, $K = 15$ with the difference tending to zero at the extremes of the share price. This is in accordance with what intuition would suggest it should be; near to the strike, the amount of hedging required is going to be greatest, and so the liquidity losses will also be greatest. Deep into (or out of) the money, the holding of stock will be almost constant, so there will be almost no hedging, and therefore almost no hedging losses.

The next figure, Figure 3, explores the effect of varying the volatility. It plots the price difference for two values of σ, $\sigma = 0.1, 0.05$, for comparison with Figure 2. All other paramter values are the same as for Figure 2.

The final figure, Figure 4, shows the option price difference as a function of time, plotted every 0.25 years out to 5 years, with $S = 50 * 30/101$, $H = 0$. Each step of 0.25 years was computed with 10000 timesteps, other parameters are as for the computation of Figure 2.

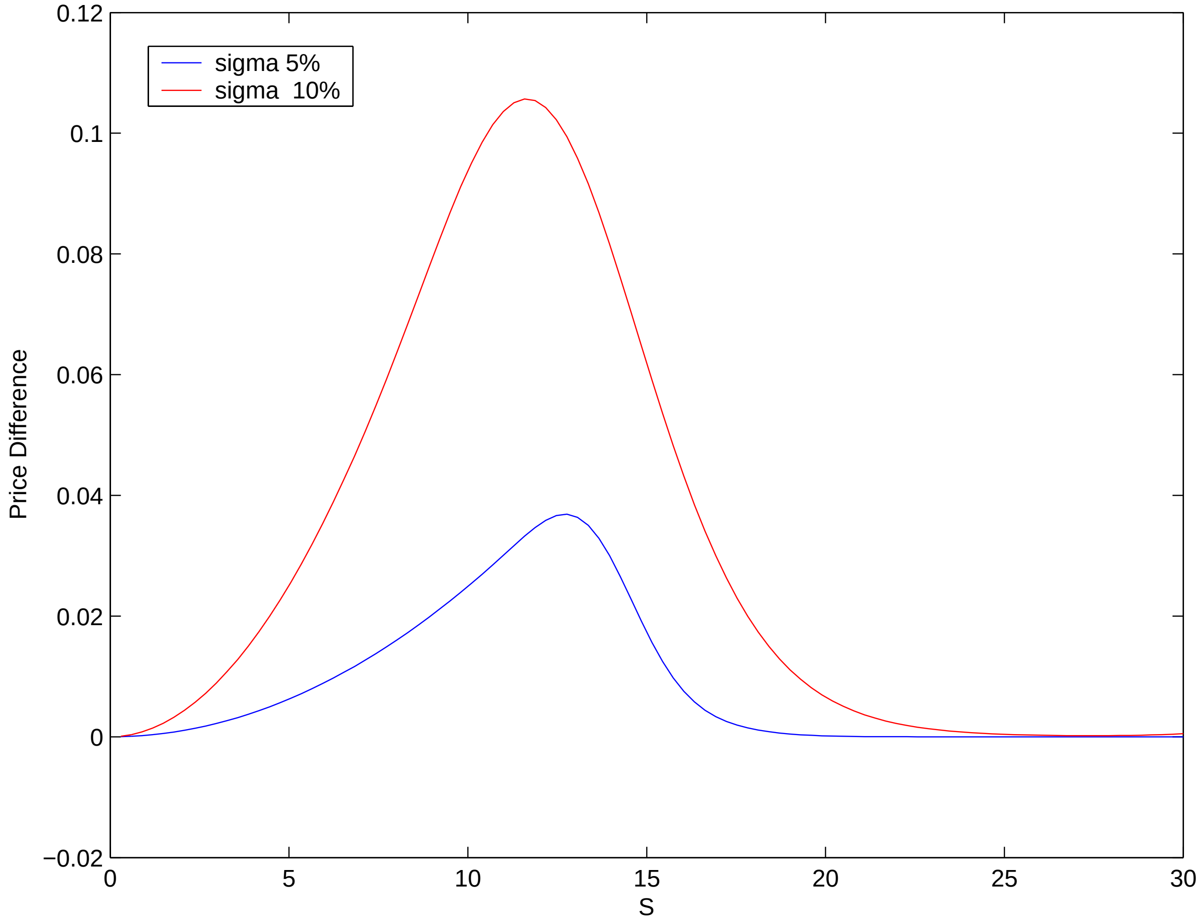

Figure 3: Illiquid Option Price minus Black Scholes Price, $\sigma = 0.1, 0.05$.

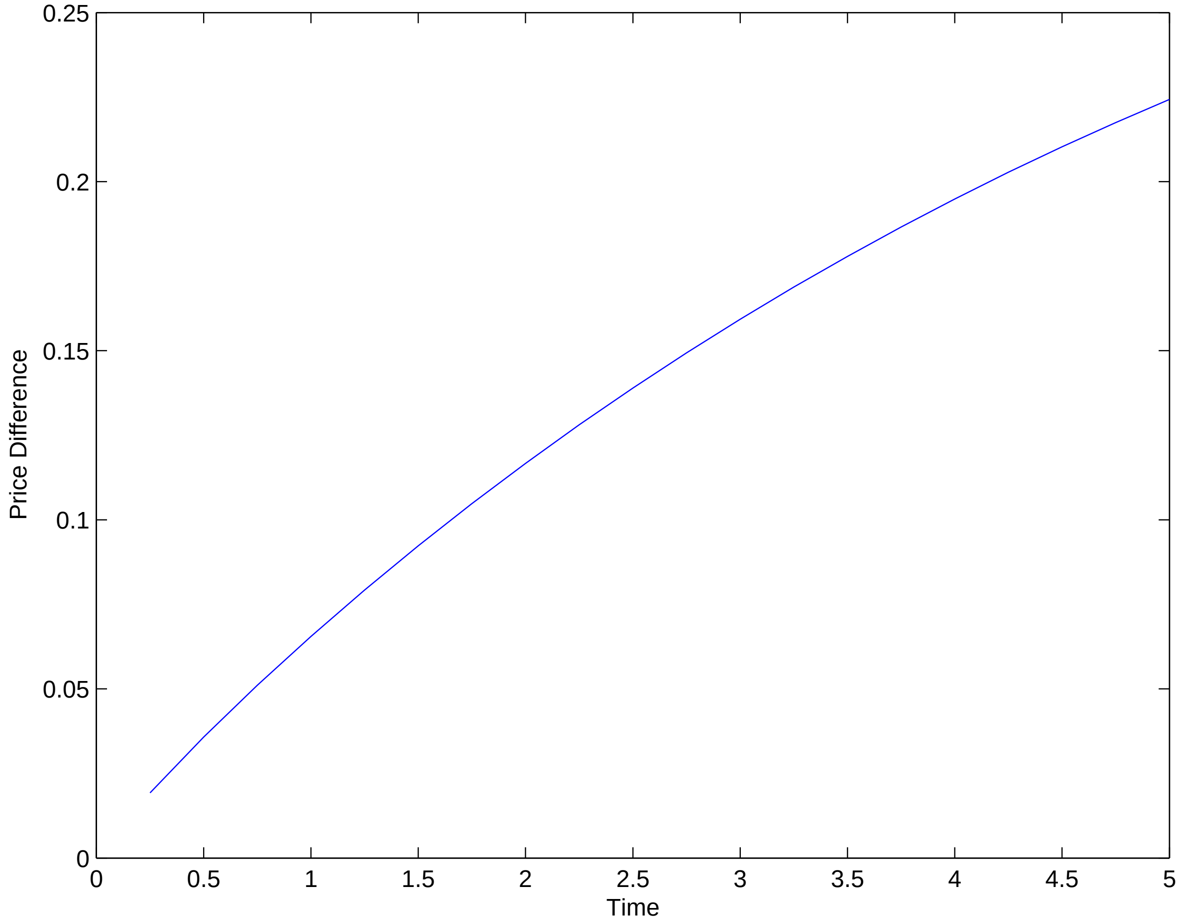

Figure 4: Illiquid Option Price minus Black Scholes Price, as a function of time

6 Conclusions.

Starting from a simple microeconomic story of equalisation of supply and demand, we have developed a model for the effect of liquidity, which we interpret as a cost to be paid if the hedger changes portfolio rapidly. In contrast to 'feedback' or 'large agent' models, the market price of the underlying is not affected by the actions of the hedger; the hedger's impatience to change his portfolio affects mainly himself, as he tries to move large quantities over the (quite thin) short-term market made up of the relatively few shares that are available in a short time period, but has little effect on the much larger number of shares in total. Think of punching a truck. This modelling standpoint now allows us to obtain simple dynamics, with their associated PDEs, but these are not soluble any way except numerically. Nevertheless, we present a formal asymptotic expansion whose conclusions are intuitively reasonable, and which agrees very closely with numerical values.

References

[1] D. Bakstein and S. Howison. A non-arbitrage liquidity model with observable parameters for derivatives. *Oxford University Preprint*, 2003.

[2] P. Bank and D. Baum. Hedging and portfolio optimization in financial markets with a large trader. *Mathematical Finance*, 14:1–18, 2004.

[3] U. Cetin, R. A. Jarrow, and P. Protter. Liquidity risk and arbitrage pricing theory. *Finance and Stochastics*, 8:311–341, 2004.

[4] U. Cetin and L. C. G. Rogers. Modelling liquidity effects in discrete time. *Mathematical Finance*, page ??, 2006.

[5] M. H. A. Davis and A. R. Norman. Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15:676–713, 1990.

[6] M. H. A. Davis, V. G. Panas, and T. Zariphopoulou. European option pricing with transaction costs. *SIAM Journal on Control and Optimization*, 31:470–493, 1993.

[7] R. Frey. Perfect option hedging for a large trader. *Finance & Stochastics*, 2:115–141, 1988.

[8] R. Frey and A. Stremme. Market volatility and feedback effects from dynamic hedging. *Mathematical Finance*, 7:351–374, 1997.

[9] J. P. Heritage and L. C. G. Rogers. Large investors, takeovers, and the rule of law. *Monte Carlo Methods and Applications*, 8:357–370, 2002.

[10] S. D. Hodges and A. Neuberger. Optimal replication of contingent claims under transaction costs. *Review of Futures Markets*, 8:222–239, 1989.

[11] J. C. Hull. *Options, Futures and Other Derivative Securities, 2nd Edition.* Prentice-Hall, Englewood Cliffs, NJ, 1993.

[12] R. A. Jarrow. Market manipulation, bubbles, corners, and short squeezes. *Journal of Financial and Quantitative Analysis*, 27:311–336, 1992.

[13] F. A. Longstaff. Optimal portfolio choice and the valuation of illiquid securities. *Review of Financial Studies*, 14:407–431, 2001.

[14] M. J. P. Magill and G. M. Constantinides. Portfolio selection with transaction costs. *Journal of Economic Theory*, 13:245–263, 1976.

[15] G. Papanicolaou and K. R. Sircar. General black-scholes models accounting for increased market volatility from hedging strategies. *Applied Mathematical Finance*, 5:45–82, 1998.

[16] E. Platen and M. Schweizer. On feedback effects from hedging derivatives. *Mathematical Finance*, 8:67–84, 1998.

[17] P. J. Schönbucher and P. Wilmott. The feedback effect of hedging in illiquid markets. *SIAM Journal on Applied Mathematics*, 61:232–272, 2000.

[18] S. Singh. *Modelling Liquidity and the Valuation of American Options using the Dual Method*. PhD thesis, University of Bath, 2005.