

Pricing Options in an Extended Black Scholes Economy

with Illiquidity: Theory and Empirical Evidence

U. Çetin

Department of Statistics,
London School of Economics and Political Science,
London, England, WC2A 2AE.

R. Jarrow*

Johnson Graduate School of Management,
Cornell University,
Ithaca, New York, 14853.

P. Protter[†]

School of Operations Research and Industrial Engineering,
Cornell University,
Ithaca, New York, 14853.

M. Warachka[‡]

School of Business,
Singapore Management University,
Singapore, 259756.

April, 2005

JEL Codes: G12, G13

Keywords: Option Pricing, Liquidity, Black Scholes

*Please address correspondence and requests for preprints to Robert Jarrow, 451 Sage Hall, Johnson Graduate School of Management, Cornell University, Ithaca, New York, 14853-6201. Email: raj15@cornell.edu
Phone: (607) 255-4729.

[†]Supported in part by NSF Grant DMS-0202958 and NSA Grant 904-00-1-0035.

[‡]Supported by Wharton-SMU Research Center, Singapore Management University. Assistance from the Center for Academic Computing at SMU is also gratefully acknowledged.

Abstract

This paper studies the pricing of options in an extended Black Scholes economy in which the underlying asset is not perfectly liquid. The resulting liquidity risk is modeled as a stochastic supply curve, with the transaction price being a function of the trade size. Consistent with the market microstructure literature, the supply curve is upward sloping with purchases executed at higher prices and sales at lower prices. Optimal discrete time hedging strategies are then derived. Empirical evidence reveals a significant liquidity cost intrinsic to every option.

Risk management is concerned with controlling three financial risks: market risk, credit risk and liquidity risk.¹ Starting with the Black-Scholes-Merton option pricing formula, both market and credit risk have been successfully modeled with Duffie (1996) and Bielecki and Rutkowski (2002) offering excellent summaries of these literatures. In contrast, our understanding of liquidity risk is still preliminary.

This paper defines liquidity risk as the increased variability in realized returns from forming a replicating portfolio or implementing a hedging strategy due to the price impact of random transactions. In particular, the corresponding price impacts for the series of transactions required to hedge an option are stochastic since they depend on the evolution of the stock price. Consequently, the liquidity cost associated with replicating an option is random. This liquidity cost is the manifestation of the liquidity risk inherent in the hedging strategy's performance.²

The approach in Çetin, Jarrow and Protter (2004) hypothesizes the existence of a stochastic supply curve for a security's price as a function of order flow.³ Specifically, a second argument incorporates the size (number of shares) and direction (buy versus sell) of a transaction to determine the price at which the trade is executed. For a given supply curve, traders act as price takers. The greater an asset's liquidity, the more horizontal its supply curve. In the context of continuous trading, they characterize necessary and sufficient conditions on the supply curve's evolution to ensure there are no arbitrage opportunities in the economy. Furthermore, conditions for an approximately complete market are also provided.

In the most general setting with unrestricted predictable trading strategies, three primary conclusions regarding liquidity risk are available. First, all liquidity costs are avoidable when

(approximately) replicating a derivative's payoff using continuous trading strategies of finite variation. Second, as a consequence of the previous statement, the value of a derivative security is identical to its price in the classical theory which assumes markets are perfectly liquid. Third, there are no implied bid-ask option spreads that are attributable to illiquidities in the underlying asset. It is important to emphasize that these conclusions assume continuous trading of arbitrarily small quantities.

However, not all predictable trading strategies are possible to implement in practice. In particular, one cannot trade continuously with arbitrarily small quantities. To accommodate these limitations, we define *discrete trading strategies* as those simple trading strategies where the minimum time between successive trades is greater than a given constant $\delta > 0$. Although one may trade at any point in time, subsequent trades occur after at least δ time units have elapsed.⁴ This situation is distinct from the classical approach in which “discrete” trading strategies are not constrained to have a positive time step between trades, enabling them to approximate any predictable trading strategy, with the cost of approximately replicating a contingent claim given in Duffie and Protter (1988). In contrast, by imposing a minimum time between trades, the classical theory is no longer valid. Specifically, one cannot approximate (arbitrarily closely) a derivative's payoff even in the absence of illiquidity. Furthermore, with illiquidity, we have an additional complication as the liquidity costs do not converge to zero even for highly liquid assets. Consequently, option bid-ask spreads are partially attributable to illiquidities present in the underlying asset, along with the number of options being hedged.

The purpose of this paper is to investigate the pricing of derivatives using discrete trading

strategies when the underlying asset is not assumed to be perfectly liquid. Specifically, we study the pricing and hedging of a European call option on a stock in an extended Black Scholes economy with illiquidity. In this economy, transaction prices reflect the price impact of order flow.

After calibrating the model's parameters to market data, we investigate optimal hedging strategies in the context of these illiquidities. The optimal hedging strategy results from a dynamic program which super-replicates the option payoffs. Two non-optimal discrete time trading strategies based on the Black Scholes hedge are implemented for comparison. In particular, we implement Black Scholes hedges at random as well as fixed time points. These Black Scholes hedging strategies are studied for two reasons. First, they are used in practice (see Jarrow and Turnbull (2004)) given the infeasibility of continuous hedging (due to market frictions). Therefore, it is instructive to investigate whether these strategies are nearly optimal in our setting. Second, because of their popularity in practice, they provide a useful standard for comparison for understanding the optimal hedging strategy. Not surprising, our empirical results confirm the non-optimality of the Black Scholes hedging strategies. Furthermore, our empirical results demonstrate that even under the optimal hedging strategy, liquidity costs comprise a significant component of an option's price.

It is important to relate our paper to the literature on transactions costs, including Leland (1985), Boyle and Vorst (1992), as well as Edirisinghe, Naik and Uppal (1993). Although transactions costs also increase option prices, liquidity risk is endogenous to the trading process. Moreover, liquidity risk is characterized by a continuous supply curve that is differentiable at the

origin, implying a well defined limit exists, even for continuous trading. Çetin (2003) contains further details regarding the distinction between illiquidity and transaction costs.

Our proposed framework is also similar to the feedback effects on stock prices generated by option hedging demands as well as the literature on large traders. These issues are studied in Schonbucher and Wilmott (2000), Platen and Schweizer (1998) and Frey (1998). However, in this context, it is important to emphasize that our specification is a “reduced-form” illiquidity model since the supply curve is independent of transactions by other agents. Furthermore, our framework focuses on temporary price impacts.⁵ Permanent price impacts associated with very large transactions or a sequence of trades with the same direction are not typical properties of hedging strategies.

The remainder of this paper begins with a description of the general model in the next section while Section 2 introduces the extended Black Scholes economy. Estimation of the liquidity parameter using a sample of five NYSE firms is conducted in Section 3. The derivation of optimal discrete time hedging strategies is the subject of Section 4. For comparative purposes, two non-optimal discrete time hedging strategies are also implemented in Section 5. Empirical results follow in Section 6 with Section 7 offering our conclusions.

1 The Model

This section summarizes the relevant portions of Çetin, Jarrow and Protter (2004) used in the subsequent analysis. We are given a filtered probability space $(\Omega, F, (F_t)_{0 \leq t \leq T}, \mathbf{P})$ satisfying the usual conditions where T is a fixed time and \mathbf{P} represents the statistical or empirical probabil-

ity measure for a stock that pays no dividends. Also traded is a money market account that accumulates value at the spot rate of interest denoted r .

Let $S(t, x)$ represent the stock price, *per share*, at time $t \in [0, T]$ that a trader pays/receives for order flow x normalized by the value of a money market account. A positive order ($x > 0$) represents a buy, a negative order ($x < 0$) signifies a sale and $x = 0$ corresponds to the *marginal trade*.

1.1 Trading Strategies and Liquidity Costs

A *trading strategy* is summarized as $(X_t, Y_t : t \in [0, T], \tau)$ where X_t represents the trader's aggregate holding of stock at time t , and Y_t the aggregate position in the money market account, while τ denotes the liquidation time of the stock in the replicating portfolio. The trading strategy is subject to the following restrictions:⁶

1. $X_{0-} \equiv Y_{0-} \equiv 0$,
2. $X_T = 0$ and
3. $X = H1_{[0, \tau)}$ for some process $H(t, \omega)$ where $\tau \leq T$ is a stopping time.⁷

These restrictions ensure that the trading strategy is liquidated prior to time T which ensures that round-trip liquidity costs are incurred. The stopping time τ allows the portfolio to be liquidated prior to time T .

A *self-financing trading strategy (s.f.t.s)* generates no cash flows for all times $t \in (0, T)$ after the initial purchase. More formally, a self-financing trading strategy is represented as

$(X_t, Y_t : t \in [0, T], \tau)$ where:

1. X_t is càdlàg with finite quadratic variation ($[X, X]_T < \infty$),
2. $Y_0 = -X_0 S(0, X_0)$ and
3. for $0 < t \leq T$,

$$Y_t = Y_0 + X_0 S(0, X_0) + \int_0^t X_{u-} dS(u, 0) - X_t S(t, 0) - L_t, \quad (1)$$

where L_t is the liquidity cost, defined as

$$L_t = \sum_{0 \leq u \leq t} \Delta X_u [S(u, \Delta X_u) - S(u, 0)] + \int_0^t \frac{\partial S}{\partial x}(u, 0) d[X, X]_u^c \geq 0 \quad (2)$$

with $L_{0-} = 0$.

The expression $[X, X]_t^c$ denotes the quadratic variation of the continuous component of X at time t (see Protter (2004)).

Observe that the liquidity cost consists of two components. The first is due to discontinuous changes in the share holdings while the second results from continuous rebalancing. For a continuous trading strategy, the first term in expression (2) equals

$$L_0 = X_0 [S(0, X_0) - S(0, 0)],$$

even after time zero. If the trading strategy is also of finite variation, then the second term in equation (2) is zero because $[X, X]_t^c = 0$. Thus, for a trading strategy that is both continuous and of finite variation, the entire liquidity cost is due to forming the initial position and manifested in L_0 . Furthermore, if one chooses a continuous trading strategy of finite variation with an

initial stock position equal to zero ($X_0 = 0$) that quickly “approaches” the desired level, then the liquidity cost of this approximating s.f.t.s. is also zero. This insight is required to extend the fundamental theorems of asset pricing.

1.2 Fundamental Theorems of Asset Pricing with Illiquidity

As is standard in the literature, an *arbitrage opportunity* is any s.f.t.s. (X, Y, τ) such that $\mathbb{P}\{Y_T \geq 0\} = 1$ and $\mathbb{P}\{Y_T > 0\} > 0$. A modified first fundamental theorem of asset pricing is available with illiquidity. For $\beta \geq 0$, define

$$\Theta_\beta \equiv \{\text{s.f.t.s } (X, Y, \tau) \mid (X_- \cdot s)_t \geq -\beta \text{ for all } t \text{ almost surely}\} .$$

The modified first fundamental theorem of finance with illiquidity states that if there exists a probability measure $\mathbb{Q} \sim \mathbb{P}$ such that $S(\cdot, 0)$ is a \mathbb{Q} -local martingale, then there is no arbitrage for $(X, Y, \tau) \in \Theta_\beta$ for any β .

For pricing derivatives, we assume the existence of such a \mathbb{Q} -local martingale for $S(\cdot, 0)$. Next, a market is defined to be *approximately complete* if given any contingent claim C , defined as a \mathbb{Q} -square integrable random variable, there exists a sequence of self-financing trading strategies (X^n, Y^n, τ^n) such that $Y_T^n \rightarrow C$ as $n \rightarrow \infty$ in $L^2(d\mathbb{Q})$.⁸ This definition parallels the standard definition of a complete market. The difference is that the payoff of any contingent claim is only approximately attained.

Given this definition, a modified second fundamental theorem of asset pricing also holds in this setting. The modified second theorem states that the existence of a unique probability measure $\mathbb{Q} \sim \mathbb{P}$ such that $S(\cdot, 0) = s$ is a \mathbb{Q} -local martingale implies the market is approximately

complete.

It is perhaps surprising that in an approximately complete market, a continuous and finite variation trading strategy is always available to approximate any contingent claim that starts with zero initial holdings in the underlying stock. Indeed, given any contingent claim C , there exists a predictable process X such that $C = c + \int_0^T X_u ds_u$ and a sequence of s.f.t.s. (X^n, Y^n, τ^n) , where X^n is continuous and of finite variation with the properties $X_0^n = 0$, $X_T^n = 0$ and $Y_0^n = c$ such that $Y_T^n \rightarrow C$ in $L^2(d\mathbb{Q})$. The liquidity cost of this sequence of s.f.t.s. equals zero (since the first trade is of zero magnitude and the s.f.t.s. is continuous and of finite variation) with the contingent claim's payoff at time T approximated by

$$Y_T^n = c + \int_0^T X_u^n dS(u, 0). \quad (3)$$

Consequently, the arbitrage free value for the contingent claim is its classical value,

$$E^{\mathbb{Q}}[C]. \quad (4)$$

This summarizes the three primary conclusions of Çetin, Jarrow and Protter (2004) and motivates the extended Black Scholes economy with illiquidity presented in the next section.

2 An Extended Black Scholes Economy

To value a European call option in an extended Black Scholes economy, we assume the stock's supply curve satisfies

$$S(t, x) = e^{\alpha x} S(t, 0) \quad \text{with} \quad \alpha > 0 \quad (5)$$

where

$$S(t, 0) \equiv \frac{s_t}{e^{rt}} = \frac{s_0 e^{\mu t + \sigma W_t}}{e^{rt}}, \quad (6)$$

for constants μ and σ , with W_t denoting a standard Brownian motion.

Equation (6) states that the marginal stock price $S(t, 0)$ follows a geometric Brownian motion, while the extended Black Scholes economy's supply curve is given in equation (5). The $e^{\alpha x}$ functional form for the supply curve is chosen for simplicity and is easily generalized. It is important to emphasize that the supply curve given in expression (5) is stochastic. After a trade is executed, a new supply curve $S(t, x)$ is generated for subsequent trades.

Under the supply curve in equation (5), there exists a unique martingale measure \mathbb{Q} for $S(t, 0)$ as explained in Duffie (1996). Hence, applying the extended first and second fundamental theorems of asset pricing with illiquidity, the market is arbitrage-free and approximately complete.

2.1 Pricing a European Call Option

Consider a European call option on the stock with a strike price of K and maturity T , with the corresponding payoff $C_T = \max[S(T, 0) - K e^{-rT}, 0]$. Expression (4) implies the European call value equals

$$\begin{aligned} E^{\mathbb{Q}}[C_T] &= e^{-rT} E^{\mathbb{Q}}(\max[s_T - K, 0]) \\ &= s_0 N(h_0) - K e^{-rT} N(h_0 - \sigma\sqrt{T}), \end{aligned} \quad (7)$$

where $N(\cdot)$ denotes the standard cumulative normal distribution function whose argument is

$$h_t \equiv \frac{\log(s_t) - \log K + r(T - t)}{\sigma\sqrt{T - t}} + \frac{\sigma}{2}\sqrt{T - t}. \quad (8)$$

In this setting, the Black Scholes hedging strategy $X_t = N(h_t)$ is continuous but not of finite variation. Therefore, although the Black Scholes formula remains valid given liquidity costs, the standard hedging strategy does not attain this value. Indeed, equation (2) implies the Black Scholes hedging strategy results in a positive liquidity cost⁹

$$L_T = X_0(S(0, X_0) - S(0, 0)) + \int_0^T \frac{\alpha (N'(h_u))^2 s_u}{T - u} du, \quad (9)$$

as proved in Appendix A.

An example of s.f.t.s. that is continuous and of finite variation, yet approximates the call option's payoff is

$$X_t^n = \begin{cases} n 1_{[\frac{1}{n}, T-\frac{1}{n})} (t) \int_{(t-\frac{1}{n})^+}^t N(h(u)) du, & \text{if } 0 \leq t \leq T - \frac{1}{n} \\ n \left(T X_{(T-\frac{1}{n})}^n - t X_{(T-\frac{1}{n})}^n \right), & \text{if } T - \frac{1}{n} \leq t \leq T. \end{cases} \quad (10)$$

This strategy starts with $X_0^n = 0$ and quickly approaches an “average” of the Black-Scholes hedge at time t . Then, just before maturity, liquidation transfers the accumulated value into the money market account. The above trading strategy is seen to have zero liquidity costs with

$$Y_T^n = E^Q[C_T] + \int_0^T X_{u-}^n dS(u, 0) \quad (11)$$

$$\rightarrow C_T = \max[S(T, 0) - K e^{-rT}, 0]$$

in $L^2(d\mathbb{Q})$. To summarize, the above trading strategy is a “smoothed” version of the Black Scholes hedging strategy that eliminates liquidity risk.¹⁰

2.2 Discrete Hedging Strategies

As previously noted, the continuous hedging strategy in expression (10) cannot be implemented in practice. A class of feasible trading strategies are the *discrete trading strategies* defined as any simple s.f.t.s. X_t where

$$X_t \in \left\{ x_{\tau_0} 1_{\{\tau_0\}} + \sum_{j=1}^N x_{\tau_j} 1_{(\tau_{j-1}, \tau_j]} \left| \begin{array}{l} 1. \tau_j \text{ are } \mathbb{F} \text{ stopping times for each } j, \\ 2. x_{\tau_j} \text{ is in } \mathcal{F}_{\tau_{j-1}} \text{ for each } j \text{ (predictable),} \\ 3. \tau_0 \equiv 0 \text{ and } \tau_j > \tau_{j-1} + \delta \text{ for a fixed } \delta > 0. \end{array} \right. \right\}$$

These trading strategies are discontinuous because once a trade is executed, the subsequent trade is separated by a minimum of $\delta > 0$ time units, as in Cheridito (2003). For the remainder of the paper, lower case values x and y denote discrete trading strategies.

By restricting the class of trading strategies, we retain an arbitrage-free environment although the minimum distance δ between trades prevents the market from being approximately complete. In an incomplete (not approximately complete market), the cost of replicating an option depends on the chosen trading strategy.

For any discrete trading strategy, the liquidity cost equals

$$L_T = \sum_{j=0}^N [x_{\tau_{j+1}} - x_{\tau_j}] [S(\tau_j, x_{\tau_{j+1}} - x_{\tau_j}) - S(\tau_j, 0)]. \quad (12)$$

For a discrete trading strategy with $x_T = 0$, the hedging error is given by

$$C_T - Y_T = C_T - \left[y_0 + x_0 S(0, 0) + \sum_{j=0}^{N-1} x_{\tau_{j+1}} [S(\tau_{j+1}, 0) - S(\tau_j, 0)] \right] + L_T. \quad (13)$$

Thus, there are two components to this hedging error. The first quantity, with a sign reversal for ease of reference in later applications,

$$\left[y_0 + x_0 S(0, 0) + \sum_{j=0}^{N-1} x_{\tau_{j+1}} [S(\tau_{j+1}, 0) - S(\tau_j, 0)] \right] - C_T, \quad (14)$$

is the error in replicating the option's payoff C_T , and is consequently referred to as the *approximation error*. Thus, a positive approximation error signifies a surplus in the replicating portfolio relative to the liability of the contingent claim's payoff while a negative value represents a deficit. The second component in equation (13) is the *liquidity cost* L_T defined in equation (12).

Furthermore, since a perfect balance between long and short option positions offsets their liabilities and eliminates the need to hedge using the underlying asset, this paper offers a methodology to infer option spreads *conditional* on a specified imbalance in the number of long and short option positions (aggregated over strike prices and maturities). As emphasized repeatedly in the remainder of this paper, equation (12) implies that liquidity costs increase quadratically with the number of options being hedged while prices increase at a linear rate. This is best seen from equation (16) in the next section. Finally, we focus on replicating long (hedging short) call positions as they entail the possibility of negative cash flows at maturity.

To provide realistic illustrations of the impact of liquidity costs and hedging errors on the option's price, we first need to calibrate the α parameter and confirm the supply curve is upward sloping. This is the subject of the next section.

3 Supply Curve Estimation

To investigate the liquidity costs incurred when constructing options with discrete trading strategies, this section details the estimation of the supply curve liquidity parameter α using the TAQ database. For illustrative purposes, we select five well known companies trading on the NYSE with varying degrees of liquidity; General Electric (GE), International Business Machines (IBM), Federal Express (FDX), Reebok (RBK) and Barnes & Noble (BKS). Our empirical analysis is conducted over a four year period with 1,011 trading days, from January 3, 1995 to December 31, 1998. These five firms represent a cross-section of stocks with respect to daily trading volume which have options on the Chicago Board of Options Exchange (CBOE).¹¹

3.1 Estimation Procedure

A simple regression methodology is employed to estimate the liquidity parameter α in equation (5). Although the true price $S(t, 0)$ is unobservable, this term is eliminated by considering two consecutive intra-day transactions. Let τ_i denote the time index with corresponding order flow x_{τ_i} and stock price $S(\tau_i, x_{\tau_i})$ for every transaction $i = 1, \dots, N$ in a given day. Thus, we are led to the following regression specification

$$\ln \left(\frac{S(\tau_{i+1}, x_{\tau_{i+1}})}{S(\tau_i, x_{\tau_i})} \right) = \alpha [x_{\tau_{i+1}} - x_{\tau_i}] + \mu [\tau_{i+1} - \tau_i] + \sigma \epsilon_{\tau_{i+1}, \tau_i}. \quad (15)$$

The error $\epsilon_{\tau_{i+1}, \tau_i}$ equals $\epsilon \sqrt{\tau_{i+1} - \tau_i}$ with ϵ being distributed $\mathcal{N}(0, 1)$. Observe that the left side of equation (15) is the percentage return between two consecutive trades and this expression reduces

to a standard geometric Brownian motion when α is identically zero. Transaction prices instead of the bid-ask spread are utilized for three reasons. First, trades may be executed “inside the spread” which implies that quotes potentially overestimate liquidity costs. Second, the bid-ask spread is a commitment to a specific volume which may change according to market conditions. Third, quotes may be “stale” for infrequently traded stocks.

Given a series of transaction prices, the first issue is to sign the trade volume as either buys or sells. To accomplish this task, the Lee and Ready (1991) algorithm is employed. Since our analysis concerns the frequent hedging of options in small quantities, we limit our attention to transaction sizes (absolute value of order flow) that are less than or equal to ten lots.

Our estimation procedure generates daily estimates for α over the sample period. Therefore, a total of 1,011 regressions are performed for each of the five firms. The average number of observations per day for each firm is reported in Table 1.

3.2 Estimation Results

Table 1 below displays the regression results from equation (15). The ninth and tenth columns record the number of significant α and μ coefficients. Observe that the estimated α parameters are almost always significantly positive at the 5% level, in contrast to the μ estimates. Thus, much of the variation in intraday stock prices is attributable to order flow. Furthermore, the statistically positive α estimates confirms the existence of our hypothesized upward-sloping supply curve in equation (5).

However, the α estimates are somewhat “noisy” which suggests that alternative supply curve

formulations may be worth considering, two of which are explored in the next subsection as robustness tests. Furthermore, it is important to emphasize that the standard errors are only valid goodness-of-fit measures for an individual day since the estimation procedure is performed daily.

Figure 1 displays the estimated time series of α parameters over the sample period for IBM, FDX and BKS which represent high, medium and low liquidity securities respectively. Figure 2 plots the corresponding estimates for GE and RBK. As expected, the α estimates for IBM and GE are lower than those of BKS and RBK, confirming our intuition regarding differences in their liquidity.

It is important to emphasize that the liquidity cost of a transaction depends on both α and the marginal stock price $S(\cdot, 0)$. In particular, for a small α , a Taylor series expansion of $\exp\{\alpha(x_{\tau_{j+1}} - x_{\tau_j})\}$ in equation (12) implies the terms in the summation are approximately

$$[x_{\tau_{j+1}} - x_{\tau_j}] S(\tau_j, 0) [\exp\{\alpha(x_{\tau_{j+1}} - x_{\tau_j})\} - 1] \approx \alpha S(\tau_j, 0) [x_{\tau_{j+1}} - x_{\tau_j}]^2 . \quad (16)$$

Observe the symmetry between purchases and sales as equation (16) indicates the sign of order flow is irrelevant. Furthermore, we find evidence in Figure 2 that α and the marginal stock price move inversely to each other. For example, as GE's stock price increases over this sample period, its α parameter declines. Intuitively, this suggests that market makers strive to obtain a constant dollar-denominated fee per lot transacted. Overall, the liquidity cost is a function of both the α parameter and the frictionless stock price $S(\cdot, 0)$.

3.3 Robustness Tests

As a robustness check of our α estimates, we alter our original estimation procedure in three ways. The first modification continues with equation (15) but excludes transactions larger than five lots. Although the empirical evidence in Hausman, Lo and MacKinlay (1992) suggests decreasing marginal price impacts, frequent hedging of an option implies small transactions and this renders many of the transactions recorded in the TAQ database irrelevant to our analysis.¹² The second and third robustness checks consist of two alternative supply curves with diminishing marginal price impacts,¹³

$$\text{sgn}(x) \sqrt{|x|} \tag{17}$$

and

$$\text{sgn}(x) \ln(1 + |x|). \tag{18}$$

These two formulations are upward sloping and satisfy the property that $S(\cdot, 0) = 0$. Furthermore, they only require the estimation of one parameter using the regression procedure in equation (15). In particular, after applying a logarithm to the ratio of transaction prices, the exponential functional form facilitates a simple regression analysis for the calibration of α . More complicated supply curves involving additional parameters are left for future research.

The modified marginal price impacts in equations (17) and (18) are both calibrated with order

flow up to and including ten lots using the following two regressions which parallel equation (15),

$$\ln \left(\frac{S(\tau_{i+1}, x_{\tau_{i+1}})}{S(\tau_i, x_{\tau_i})} \right) = \alpha \operatorname{sgn} ([x_{\tau_{i+1}} - x_{\tau_i}]) \sqrt{|x_{\tau_{i+1}} - x_{\tau_i}|} + \mu [\tau_{i+1} - \tau_i] + \sigma \epsilon_{\tau_{i+1}, \tau_i}$$

and (19)

$$\ln \left(\frac{S(\tau_{i+1}, x_{\tau_{i+1}})}{S(\tau_i, x_{\tau_i})} \right) = \alpha \operatorname{sgn} ([x_{\tau_{i+1}} - x_{\tau_i}]) \ln (1 + |x_{\tau_{i+1}} - x_{\tau_i}|) + \mu [\tau_{i+1} - \tau_i] + \sigma \epsilon_{\tau_{i+1}, \tau_i}.$$

The results of our three robustness tests are now summarized.¹⁴ The first robustness test results in the average number of available daily transactions declining from those reported as n in Table 1 to 720, 432, 181, 226 and 122 respectively for GE, IBM, FDX, BKS and RBK. Thus, a loss of information is incurred.

All three robustness tests produce very significant α estimates. Overall, the α estimates are smallest for equation (15) with a ten lot upper bound (the original procedure) followed by those generated by transactions of five lots or less. Next in magnitude are the square root estimates from equation (17), with the log formulation in equation (18) producing the largest α coefficients.

However, the original formulation in equation (15) produces α estimates with the lowest standard errors. This result is consistent across all five firms. Thus, our subsequent option pricing investigation is justified in using the α estimates from Table 1.

To further validate the supply curve formulation of liquidity in Section 1, we perform an additional experiment whose results are contained in Table 2. The “hypotheses” on the left side of Table 2 offer the implications of our upward sloping supply curve when both time (t) and information (ω) are fixed. The price inequalities indicate the ordering of transaction prices

realized in a given sequence of trades. These statements amount to identifying different locations on the supply curve. To approximate this comparative static in the data, we examine consecutive transactions in which the influence of μ is negligible. Despite the noise introduced by a change in time (and information), if the upward sloping supply curve formulation is valid, then we expect the inequalities to hold for the majority of transaction sequences studied.

The results on the right side of Table 2 confirm the validity of our liquidity model. It is important to emphasize that our model assumes transitory price impacts and the comparative static implication assumes no change in information. Hence, only transactions less than 10 lots are likely to be consistent with those arising from frequent hedging and non-information based trades (with ω fixed). These situations are best reflected in the first two rows of Table 2. The last four rows of Table 2 consider situations in which small buy (sell) orders are subsequent to large buy (sell) orders. In these four rows, information may not be constant across the two transactions. Therefore, we expect a decline in consistency relative to the first two rows. This pattern is observed although the decline is slight. More importantly, the bottom four rows in Table 2 continue to support the implications of our liquidity framework.

3.4 Stochastic Liquidity

It is important to emphasize that variability in the α estimates over time does not imply a misspecified liquidity model. The standard errors of the α coefficients reported in Table 1 attest to its accuracy. Moreover, a “stochastic” α process is unnecessary if this parameter varies inversely with the marginal price, implying a roughly constant dollar-denominated liquidity cost

per transaction.¹⁵ Figures 3 and 4 offer visual evidence consistent with this property.

Statistically, with t representing daily estimates, an AR(1) model is applied to the product $\alpha(t)S(t, 0)$

$$\alpha(t)S(t, 0) = \beta \alpha(t-1)S(t-1, 0) + \epsilon_t \quad (20)$$

$$\alpha(t)S(t, 0) - \alpha(t-1)S(t-1, 0) = [\beta - 1] \alpha(t-1)S(t-1, 0) + \epsilon_t. \quad (21)$$

We then estimate the β coefficient to investigate whether it is statistically different from one. If not, then the model

$$\alpha(t)S(t, 0) - \alpha(t-1)S(t-1, 0) = \epsilon_t, \quad (22)$$

cannot be rejected and variation in the product over time is merely noise. Several time intervals, such as 30, 90 and 180 days are examined with equation (22) offering similar results. For example, 95% confidence intervals for the β parameters over the last 90 days of the sample period are 0.9943 ± 0.0293 , 0.9856 ± 0.0876 , 0.9693 ± 0.0902 , 0.9947 ± 0.0811 and 0.9744 ± 0.0801 for GE, IBM, FDX, BKS and RBK respectively. Consequently, time variation in α is not a serious concern when calibrating liquidity costs. Over the entire sample, a slight downward bias is detected with β being statistically less than one. However, this time period greatly exceeds available equity option maturities.

In our later empirical implementation of the discrete option hedging strategies we consider two calculations of the total liquidity cost; the first using equation (12) and the second derived from the approximate liquidity cost in equation (16). In the second instance, we assume the $\alpha(\tau_j)S(\tau_j, 0)$ terms are fixed at their initial value $\alpha S(0, 0)$. In particular, we define the approx-

imate liquidity cost as

$$\begin{aligned}\bar{L}_T &= \sum_{j=0}^N \alpha(\tau_j) S(\tau_j, 0) [x_{\tau_{j+1}} - x_{\tau_j}]^2 \\ &\equiv \alpha S(0, 0) \sum_{j=0}^N [x_{\tau_{j+1}} - x_{\tau_j}]^2.\end{aligned}\tag{23}$$

The economic importance of a time-varying α to option pricing is examined empirically by comparing the liquidity cost estimates in equation (12) versus equation (23). This comparison is conducted for each of the three discrete time trading strategies analyzed, with details contained in Section 6.

For emphasis, although the estimates of α and $S(t, 0)$ appear inversely related, one cannot estimate α after normalizing the $[x_{\tau_{i+1}} - x_{\tau_i}]$ term in equation (15) by $S(t, 0)$. Indeed, the marginal price $S(t, 0)$ is implied from transaction prices $S(t, x)$ conditional on an estimated α parameter. Specifically, a time series of marginal prices corresponding to transactions with zero order flow is unobservable. Instead, only transaction prices in a less than perfectly liquid market are available. Consequently, inferences regarding $S(t, 0)$ as well as its volatility σ (and expected return μ) utilize transaction prices $S(t, x)$. These transaction prices are then “inverted” to determine marginal prices under an assumed supply curve formulation along with estimates for its parameters such as α .

Therefore, an estimation procedure which attempts to calibrate α' in the regression formulation

$$\ln \left(\frac{S(\tau_{i+1}, x_{\tau_{i+1}})}{S(\tau_i, x_{\tau_i})} \right) = \frac{\alpha'}{S(t, 0)} [x_{\tau_{i+1}} - x_{\tau_i}] + \mu [\tau_{i+1} - \tau_i] + \sigma \epsilon_{\tau_{i+1}, \tau_i}.\tag{24}$$

instead of equation (15) is misspecified.

4 Optimal Discrete Option Hedging Strategies

This section derives optimal discrete time hedging strategies for super-replicating an option. Super-replication is often invoked in the incomplete markets literature due to its independence from investor preferences and probability beliefs.

4.1 Super-Replication of Options

Our hedging analysis utilizes the discrete trading strategies from Subsection 2.2 with the property that $\tau_N = T$. Let $E_t[\cdot]$ denote an expectation with respect to the martingale measure and define $Z_t = X_t S(t, 0) + Y_t$ as the time t marked-to-market value of the replicating portfolio.

For super-replicating a call option, the optimization problem is:

$$\min_{(X,Y)} Z_0 \quad \text{s.t.} \quad Z_T \geq C_T = \max\{S(T, 0) - Ke^{-rT}, 0\} \quad (25)$$

where

$$Z_T = y_0 + x_0 S(0, 0) + \sum_{j=0}^{N-1} x_{\tau_{j+1}} [S(\tau_{j+1}, 0) - S(\tau_j, 0)] - L_T.$$

At an intermediate time $t \geq 0$, this problem is written as¹⁶

$$\min_{(X,Y)} Z_t \quad \text{s.t.} \quad Z_T \geq C_T = \max\{S(T, 0) - Ke^{-rT}, 0\}. \quad (26)$$

4.2 Solution Methodology

The solution to the super-replication problem exists since an investor can always hedge by purchasing one unit of the underlying stock. Given the super-replication problem in equation (25),

the following lemma demonstrates that liquidity costs are minimized by trading as frequently as possible with the smallest possible quantities.

Lemma 1. *The optimal hedging strategy trades whenever possible with the smallest possible transactions. Thus, this strategy has the minimum expected liquidity cost among all super-replicating portfolios with trades at δ -intervals.*

Proof: Under the martingale measure,

$$E[C_T] \leq E[Z_T] = y_0 + x_0 S(0, 0) - E[L_T]. \quad (27)$$

Now, consider trading at two different time points 0 and t with quantities $a > 0$, $b > 0$ versus one combined transaction at time t comprised of $a + b > 0$. The expected liquidity cost of trading twice is

$$S(0, 0)[e^{\alpha a} - 1] + E_0[S(t, 0)][e^{\alpha b} - 1] \quad (28)$$

versus the expected liquidity cost of trading once, $E_0[S(t, 0)][e^{\alpha(a+b)} - 1]$. Under the martingale measure (with $\mu = r - \frac{1}{2}\sigma^2$), the discounted stock price is a martingale with the property

$$E_0[S(t, 0)] = E_0 \left[s_0 e^{-\frac{\sigma^2}{2}t + \sigma W_t} \right] = \frac{s_0}{e^{r \cdot 0}} = S(0, 0). \quad (29)$$

Using this equation, it is seen that

$$[e^{\alpha a} - 1] + [e^{\alpha b} - 1] < [e^{\alpha(a+b)} - 1]. \quad (30)$$

Thus, two trades incur a lower total liquidity cost than one, implying more frequent transactions are optimal. Indeed, not trading at the first available instant leaves the option position unhedged

for a period of time. Consequently, an extra transaction is needed at a later date which increases the expected total liquidity cost and yields a suboptimal portfolio. \square

To clarify and interpret the above lemma, we emphasize that the optimal hedge minimizes liquidity costs with respect to all trading strategies that super-replicate the option. However, the optimal hedge does not produce the smallest liquidity cost among all possible trading strategies. Indeed, imagine a trivial strategy that does not trade at all. This strategy yields a zero liquidity cost but does not attempt to control the approximation error. Furthermore, the intuition why more frequent trading reduces liquidity costs may be drawn from Subsection 1.1 where it is seen that a continuous strategy of finite variation eliminates the liquidity cost, with an example provided in equation (10). In our discrete time context, more frequent hedging offers a better approximation to a continuous hedge strategy of finite variation.

With the previous lemma, we proceed to solve the problem as a constrained discrete time dynamic program for a fixed Δ time step. Notationally, δ is reserved for the minimum time between trades of the same investor. In other words, lower case δ is a market-based parameter that signifies the smallest duration between consecutive market orders executed for the *same* trader. The upper case Δ represents an input for a binomial option pricing solution and is simply a computational specification.¹⁷

We implement a numerical procedure to solve this problem based on a binomial approximation to the geometric Brownian motion in expression (6). A two step illustration of this numerical procedure is discussed in the next subsection. Unlike the transaction costs literature whose limit implies infinite option prices when trading is continuous, our methodology has $\Delta \rightarrow 0$ being

well defined. In particular, we are discretizing a liquidity framework that allows for dynamic continuous time rebalancing of the replicating portfolio, as illustrated in the smoothed Black Scholes hedging strategy of equation (10). In contrast, the traditional super-replication approach with transaction costs cannot appeal to continuous trading. Instead, a static hedge strategy for a call option that purchases one unit of stock at time zero is their result. Overall, we impose an exogenous constraint on trading strategies to conform with market practice, not because the underlying mathematics limits our analysis to discrete trading. Consequently, the known convergence of the binomial process to a geometric Brownian motion justifies our numerical solution, and is consistent with applications to exchange-traded American equity option.

Note that the super-replication procedure ensures the approximation error in equation (14) is non-negative, even if intermediate stock price movements between hedge portfolio rebalancings occur. In this instance, super-replicating the option adds the largest (worst case) approximation errors to the constrained optimization. Therefore, after being translated into higher liquidity costs, these positive errors are interpreted as forcing the investor to confront more illiquidity.

4.3 Implementation

The following offers a brief summary of the steps required to implement the dynamic programming procedure using a binomial stock price process. Additional details regarding its general solution are provided in Appendix B. Consider a two-period binomial tree with an initial stock price denoted S . Up and down factors are signified by U and D respectively and are identical to those in the standard binomial tree literature. At time 1 ($\Delta = 1$), the stock price is either SU

or SD while at the option's maturity, three stock prices are available; SUU , SDU and SDD . By construction, $SUD = SDU$ since the binomial tree is recombining.

At time 1, in the up state, our objective is to solve for y_U and x_U , the amount in the money market account and stock respectively. The minimization problem involves two constraints

$$\begin{aligned} \min Z_1^U &= y_U + x_U SU + \alpha[x_U - x_1]^2 SU & (31) \\ \text{such that} \quad & y_U + x_U SUU \geq \max\{SUU - K, 0\} \\ & y_U + x_U SDU \geq \max\{SDU - K, 0\}. \end{aligned}$$

Similarly, in the down node, the quantities y_D and x_D are obtained as the solution to

$$\begin{aligned} \min Z_1^D &= y_D + x_D SD + \alpha[x_D - x_1]^2 SD & (32) \\ \text{such that} \quad & y_D + x_D SDU \geq \max\{SDU - K, 0\} \\ & y_D + x_D SDD \geq \max\{SDD - K, 0\}. \end{aligned}$$

Denote these optimal solutions as x_U^* , x_D^* , y_U^* and y_D^* which appear in the time 0 constraints. At time 0

$$\begin{aligned} \min Z_0 &= x_1 S + y_1 + \alpha x_1^2 S \quad \text{such that} \\ x_1 S + y_1 + x_1(SU - S) &= y_U^* + x_U^* SU + \alpha[x_U^* - x_1]^2 SU & (33) \end{aligned}$$

$$x_1 S + y_1 + x_1(SD - S) = y_D^* + x_D^* SD + \alpha[x_D^* - x_1]^2 SD \quad (34)$$

for which the optimal x_1 and y_1 values are solved. The super-replication price of the call option thus equals $x_1^* S + y_1^* + \alpha(x_1^*)^2 S$. Equations (33) and (34) each pertain to one of the two possible stock price paths from time 0 to time 1. For example, the left side displays the initial portfolio

value plus the gain (or loss) on stock position, while the right side is the optimal value of the replicating portfolio at either the up or down node plus the associated liquidity cost of rebalancing.

5 Discrete Hedging Strategies for Comparison

Besides the optimal hedge in the previous section, we also investigate non-optimal discrete trading strategies that employ the Black-Scholes hedge at fixed time intervals as well as random time points.¹⁸ These alternative hedging strategies are implemented via simulation for the geometric Brownian motion stock price process. These Black Scholes hedging strategies are studied for two reasons. First, these strategies are implemented in practice (see Jarrow and Turnbull (2004)) because of the infeasibility of continuous hedging. Second, their popularity in practice provides an appropriate benchmark for comparison with the optimal hedging strategy.

Empirically, we analyze 10 European call options, each on 100 shares of the underlying stock. The average implied volatility for call options (from Bloomberg)¹⁹ and the average closing stock price (from TAQ) during the sample period serve as inputs with their values reported in Table 1. Option moneyness is adjusted according to \$5 intervals. For example, in-the-money (out-of-the-money) options have a corresponding stock price that is \$5 higher (lower) than the strike price. Throughout our analysis, results are presented for 30 day option prices although the errors are similar for other maturities. For simplicity, the riskfree rate is set equal to zero as this has a negligible impact on the 30 day option's price (average interest rate below 5% during our sample period).

The fixed time trading strategy is represented as $x_u = x_t$ for $u \in [t, t + \Delta]$ for a specified Δ where x_t is the Black Scholes delta hedge parameter. Observe that the amount transacted is random. As alluded to earlier, the dependence on the random stock price necessitates simulation to price the options. We consider hedging frequencies of one and two days along the stochastic price path.

The second trading strategy hedges at random time points with $x_u = x_t$ provided $|x_u - x_t| \leq \theta$ for a given $\theta > 0$. In particular, trades only occur when the previously executed Black Scholes delta hedge differs by more than θ from the replicating portfolio's current requirement. Thus, transactions are induced by the need to rebalance the hedge portfolio. Moreover, with a transaction executed the instant the θ barrier is breached, we refer to the quantity traded as being fixed. The control width θ is chosen to coincide with transaction sizes of 5 and 10 lots.²⁰

The two non-optimal hedging strategies are evaluated using 10,000 simulations, each over a 30 day period. For the random time strategy, rebalancing of the hedge portfolio may occur at any of 300 points along the stock price path once the control width is breached. However, our results are not sensitive to this figure. Indeed, as many as 5,000 potential hedge time points are examined with similar results.²¹

6 Empirical Results

Our results are contained in Tables 3 to 5 for 10 options, each on 100 shares. Naturally, larger positions imply greater percentage price impacts for illiquidity as equation (16) illustrates that the liquidity cost increases quadratically with transaction sizes. Our empirical results are reported

for components of the hedging error; liquidity costs and approximation errors. For ease of comparison, the absolute value of the average approximation error is reported since super-replication of the option implies only non-negative errors for the optimal hedge strategy.

Using the optimal trading strategy detailed in Subsection 4.3, Table 3 reports summary statistics for the dollar-denominated liquidity costs²² and as well as their percentage impacts. The dynamic programming procedure produces an option price $x_1 S(0, 0) + y_1 + \alpha S(0, 0)x_1^2$ that includes the initial cost of forming the replicating portfolio. To proxy for the frictionless option price, we resolve the dynamic program under the constraint that $\alpha = 0$ and utilize the output $x_1^0 S(0, 0) + y_1^0$ where the 0 superscript refers to the constraint on the liquidity parameter. This price serves as the basis for computing the percentage impact of illiquidity.

The last column of Table 3 (as well as Tables 4 and 5), is concerned with the approximate liquidity cost defined in equation (23). When compared with the results in earlier columns derived from equation (12), only minor discrepancies are reported. Thus, the economic repercussions of stochastic liquidity appear to be minor.

6.1 Non-Optimal Hedges

According to Tables 4 and 5, liquidity costs are not sensitive to the rebalancing frequency when compared with the option's moneyness. Based on the approach which hedges at fixed time intervals, approximation errors are reported in Table 4 and experience a significant decrease after reducing Δ from two days to one, while the liquidity costs are almost identical.²³ Consequently, more frequent rebalancing yields smaller hedging errors. As recorded in Table 5, the random

time Black Scholes hedge also produces similar liquidity costs across the two control bands, equal to 5 and 10 lots. In contrast, the approximation errors are larger for the wider 10 lot control band. As a result, a smaller value of θ is desirable for reducing the hedging error.

Overall, for the chosen parameters we study, hedging more frequently (smaller Δ or θ) reduces the approximation errors, although the liquidity costs are less sensitive to these parameters. However, even for the smaller values of Δ and θ , the approximation error remains considerable.

Observe that the optimal trading strategy produces similar, and often lower, liquidity costs relative to the Black Scholes implementations executed at more frequent intervals. Not surprisingly, for out-of-the-money options, the safety offered by super-replicating an option occasionally results in higher liquidity costs as more of the stock is purchased to guard against the possibility of a “bad” scenario coinciding with an increase in the stock price.

As emphasized previously, there are fundamental differences between the optimal and non-optimal hedge strategies which complicates a direct comparison. Specifically, the non-optimal hedges add an associated liquidity cost to a frictionless option price, while illiquidity is intrinsic to the optimal strategy since the hedging error implications of a transaction are accounted for in the dynamic program’s constrained minimization. Thus, the liquidity costs associated with a transaction are manifested in the optimal trading strategy as inputs. However, in economic terms, imposing the constraint that the replicating portfolio’s value at maturity is at least as much as the option’s payoff does not appear to result in significantly higher liquidity costs.

6.2 Summary of Empirical Results

First, liquidity costs are a significant component of an option's price. These costs increase quadratically with the imbalance in the number of short or long positions being hedged. Second, the optimal hedging strategy provides similar and often reduced liquidity costs relative to the Black-Scholes hedge, but has the advantage of avoiding negative approximation errors. Third, liquidity risk is primarily generated by random transactions since the cost per transaction (a function of $\alpha(t) \cdot S(t, 0)$) is relatively stable over time. In other words, conditional on a specified transaction size, stochastic liquidity is not a serious concern. This property is a consequence of $\alpha(t)$ and $S(t, 0)$ exhibiting inverse fluctuations over time.

Finally, our fourth result finds that employing the Black Scholes hedge induces a relationship between an option's moneyness and the percentage impact of illiquidity. Specifically, in-the-money options are subject to the lowest percentage impact from illiquidity, despite having the largest dollar-denominated liquidity costs. This large dollar-denominated liquidity cost is partially attributed to the high initial cost of forming the replicating portfolio as the option trader is assumed to start with zero shares of the underlying stock. With in-the-money options, most of hedge portfolio re-balancing occurs when the stock price decreases. Conversely, with the initial cost of the option being relatively large, a correspondingly small percentage price increase is incurred. However, for out-of-the-money options with low initial prices, the impact of illiquidity is very significant despite a small dollar-denominated liquidity cost. As expected, at-the-money options lie between these extremes.

Since the impact of illiquidity is related to an option's moneyness, assuming a frictionless

market yields biased implied Black Scholes volatilities when its associated hedging strategy is implemented. In other words, conditional on an observed option price, ignoring liquidity costs is tantamount to overestimating option prices, which implies implied volatilities are overestimated as a consequence. The extent of this bias depends on the strike price.

Table 6 reports implied volatilities for the two non-optimal Black Scholes hedge strategies, while Figure 5 plots a “smile” over five different strike prices and two maturity dates (30 and 90 days). In Figure 5, σ and α are chosen to be 30% and 1.25×10^{-4} respectively as these are representative parameter values. Furthermore, Figure 5 considers the replication of 25 options using a discrete Black Scholes strategy implemented every two days (with nearly identical results for its random time counterpart). With the Black Scholes implied volatilities computed under the assumption of a perfectly liquid market ($\alpha = 0$) for the underlying asset, while option prices are increased by the liquidity cost associated with their replication, an upward bias in implied volatilities is detected. Moreover, the only source of this upward bias is the liquidity cost of replicating the option. Both Table 6 and Figure 5 indicate that illiquidity causes the implied volatility of in-the-money options to have the largest upward bias.²⁴

However, it is important to emphasize that illiquidity is not the exclusive catalyst for having implied volatilities depend on strike prices. Other alternatives include jumps and stochastic volatility as well as feedback effects from a trader with a large stock or option position. Nonetheless, unless our optimal hedging strategy is implemented, liquidity costs may be partially responsible for generating the implied volatility “smile” documented by Rubinstein (1985).

7 Conclusion

This paper represents the first attempt to consider the impact of illiquidity in the underlying asset market on option pricing. Consequently, this paper serves to link the market microstructure and option pricing literatures.

An extended Black Scholes economy is utilized to illustrate the theory and provide initial estimates for the impact of illiquidity on option prices. Liquidity costs are modeled as a stochastic supply curve with the underlying asset price depending on order flow. Consistent with the market microstructure literature, purchases are executed at higher prices while sales are executed at lower prices.

In addition, optimal hedging strategies that super-replicate an option are derived by solving a dynamic program. For comparative purposes, two non-optimal but intuitive discrete hedging strategies are also implemented.

Empirical results document the importance of illiquidity to option pricing. In particular, liquidity costs are a significant component of the option's price, and increase quadratically in the number of options being hedged. Second, the standard Black-Scholes hedging strategies often have higher liquidity costs than the optimal hedging strategy, and admit the possibility that the replicating portfolio is worth less than the option's liability at maturity. Third, non-optimal Black-Scholes hedges cause the impact of illiquidity to depend on the option's moneyness, offering another explanation for the implied volatility smile.

Appendices

A Liquidity Cost of Black Scholes Hedge

The liquidity cost of the Black Scholes hedge is

$$\begin{aligned}
 L_T &= X_0(S(0, X_0) - S(0, 0)) + \int_0^T \alpha s_u d[N(h), N(h)]_u^c \\
 &= X_0(S(0, X_0) - S(0, 0)) + \int_0^T \alpha s_u (N'(h_u))^2 d[h, h]_u^c \\
 &= X_0(S(0, X_0) - S(0, 0)) + \int_0^T \alpha s_u (N'(h_u))^2 \frac{1}{s_u^2 \sigma^2 (T - u)} d[s, s]_u \\
 &= X_0(S(0, X_0) - S(0, 0)) + \int_0^T \frac{\alpha (N'(h_u))^2}{\sigma^2 s_u (T - u)} \sigma^2 s_u^2 du \\
 &= X_0(S(0, X_0) - S(0, 0)) + \int_0^T \frac{\alpha (N'(h_u))^2 s_u}{T - u} du. \tag{35}
 \end{aligned}$$

B Dynamic Programming Details

The dynamic programming technique is a two-stage process. First, self-financing trading strategies that minimize the terminal deficit, as in equation (36) below are found. The second step then selects the hedge with the lowest initial cost from the previous set of strategies. A recursive technique is introduced to solve the super-replication problem. Let

$$J(x, y) = \min_{X \in \mathcal{X}(x, y)} E [\max\{C_T - Z_T, 0\}] , \tag{36}$$

where $\mathcal{X}(x, y)$ is the set of s.f.t.s. whose initial value are $X_0 = x$ and $Y_0 = y$. Thus, $J(x, y)$ is the minimum expected terminal deficit starting with an initial position (x, y) . Clearly, $J(x, y) = 0$ for sufficiently large values of x and y since such initial positions ensure the contingent claim is

hedged. Thus, the following expression is well defined²⁵

$$z^* = \min\{y + xS(0, 0) : J(x, y) = 0\}. \quad (37)$$

The above is the minimal super-replication price for hedging the contingent claim. Thus, finding the minimal super-replication price is a two-step procedure. First, $J(x, y)$ is calculated for all x, y . Then, the price $y + xS(0, 0)$ is minimized over all x, y such that $J(x, y) = 0$. Note that the expected terminal deficit of a s.f.t.s. is zero if and only if the strategy super-replicates the claim. In addition, to uniquely identify (x, y) from the minimization procedure, the self-financing condition is required.

For notational simplicity, let $T = \Delta \cdot M$ for some $M > 0$ representing the total number of trades. To apply the recursive algorithm, define for $t \in \{n \cdot \Delta : n \in \{0, \dots, M\}\}$,

$$J_t(x, y, S(t, 0)) = \min_{X \in \mathcal{X}_t(x, y)} E_t [\max\{C_T - Z_T, 0\}], \quad (38)$$

where $\mathcal{X}_t(x, y)$ is the set of s.f.t.s. with $X_t = x$ and $Y_t = y$. For simplicity, $T - n$ represents $t = \Delta \cdot (M - n)$ throughout the remainder of this appendix. Consider the boundary condition,

$$J_T(x, y, S(T, 0)) = \max\{C_T - y - xS(T, 0), 0\}, \quad (39)$$

at time T where x and y are the positions in the stock and money market at T as defined after equation (38). The following is the recursive relation between time t and $t + 1$:

$$J_t(x, y, S(t, 0)) = \min_{\Delta x} E_t [J_{t+1}(x + \Delta x, y - \Delta x S(t, \Delta x), S(t + 1, 0))] . \quad (40)$$

Intuitively, x shares in the stock and y shares in the money market account at time t is equivalent to having $x + \Delta x$ shares and $y - \Delta x S(t, \Delta x)$ in the money market (implied by the self-financing condition) at time $t + 1$ for an arbitrary transaction size Δx .

The first order condition for the minimization problem in (40) is given by

$$\begin{aligned} & E_t \left[\frac{\partial J_{t+1}}{\partial x} (x + \Delta x, y - \Delta x S(t, \Delta x), S(t+1, 0)) \right] \\ = & E_t \left[\frac{\partial J_{t+1}}{\partial y} (x + \Delta x, y - \Delta x S(t, \Delta x), S(t+1, 0)) \left(\frac{d}{d\Delta x} (\Delta x S(t, \Delta x)) \right) \right]. \end{aligned} \quad (41)$$

Equation (41) states that the change in the value of the stock position equals, on average, the marginal cost of executing the transaction scaled by the sensitivity of the value function to movements in the money market account. The marginal cost of executing the transaction is $\frac{d}{d\Delta x} (\Delta x S(t, \Delta x))$ while the scaling factor equals $\frac{\partial J_{t+1}}{\partial y} (x + \Delta x, y - \Delta x S(t, \Delta x), S(t+1, 0))$.

B.1 Binomial Implementation of Dynamic Program

Recall that the self-financing condition at time t implies²⁶

$$y_t - y_{t-1} = -(x_t - x_{t-1})S(t-1, x_t - x_{t-1}) \quad (42)$$

$$= -(x_t - x_{t-1})S(t-1, 0) - (x_t - x_{t-1})[S(t-1, x_t - x_{t-1}) - S(t-1, 0)]. \quad (43)$$

Thus, for each t , we have

$$Z_{t-1} = x_{t-1}S(t-1, 0) + y_{t-1} \quad (44)$$

$$= x_t S(t-1, 0) + y_t + (x_t - x_{t-1}) [S(t-1, x_t - x_{t-1}) - S(t-1, 0)].$$

Recall that x_t is F_{t-1} -measurable along with y_t according to the self-financing condition. Proceeding backwards from time $T-1$,

$$\min Z_{T-1} = y_T + x_T S(T-1, 0) + (x_T - x_{T-1}) [S(T-1, x_T - x_{T-1}) - S(T-1, 0)],$$

such that $Z_T \geq C_T$. Observe that liquidity costs are minimized in the above formulation subject to the approximation errors being non-negative. Invoking the binomial approximation with $S(T, 0)$ being either $S(T - 1, 0)U$ or $S(T - 1, 0)D$, x_T and y_T should satisfy

$$y_T + x_T S(T - 1, 0)U \geq (S(T - 1, 0)U - K)^+ \quad (45)$$

$$y_T + x_T S(T - 1, 0)D \geq (S(T - 1, 0)D - K)^+. \quad (46)$$

Since U and D are constants, x_T and y_T are F_{T-1} -measurable as required. Among all values of x_T and y_T satisfying the above inequalities, the optimal values minimize $y_T + x_T S(T - 1, 0) + (x_T - x_{T-1}) [S(T - 1, x_T - x_{T-1}) - S(T - 1, 0)]$. Denote the optimal solution as

$$x_T^*(S(T - 1, 0), x_{T-1}, y_{T-1}). \quad (47)$$

At time $T - 2$ the super-replication problem becomes

$$\min Z_{T-2} = x_{T-1} S(T - 2, 0) + y_{T-1} + (x_{T-1} - x_{T-2}) [S(T - 2, x_{T-1} - x_{T-2}) - S(T - 2, 0)],$$

such that $Z_T \geq C_T$ which is equivalent to

$$\begin{aligned} \min \quad & y_{T-1} + x_{T-1} S(T - 2, 0) + (x_{T-1} - x_{T-2}) [S(T - 2, x_{T-1} - x_{T-2}) - S(T - 2, 0)] \\ \text{s.t.} \quad & y_{T-2} + x_{T-2} S(T - 2, 0) - (x_{T-1} - x_{T-2}) [S(T - 2, x_{T-1} - x_{T-2}) - S(T - 2, 0)] \\ & + x_{T-1} [S(T - 1, 0) - S(T - 2, 0)] \\ = \quad & y_T^* + x_T^* S(T - 1, 0) + (x_T^* - x_{T-1}) [S(T - 1, x_T^* - x_{T-1}) - S(T - 1, 0)]. \end{aligned} \quad (48)$$

In the above, dependency of x_T^* and y_T^* on $S(T - 1, 0)$, x_{T-1} and y_{T-1} is implicit. Thus, we begin at time $T - 2$ to end at $T - 1$ with the optimal allocation in the stock and bond which

is capable of super-replicating the claim at time T . The expression in equation (48) equals $y_{T-1} + x_{T-1}S(T-2, 0)$. Thus, the minimization problem becomes

$$\begin{aligned}
\min \quad & y_{T-1} + x_{T-1}S(T-2, 0) + (x_{T-1} - x_{T-2}) [S(T-2, x_{T-1} - x_{T-2}) - S(T-2, 0)] \quad (49) \\
\text{s.t.} \quad & y_{T-1} + x_{T-1}S(T-2, 0) + x_{T-1} [S(T-1, 0) - S(T-2, 0)] \\
& = y_T^* + x_T^*S(T-1, 0) + (x_T^* - x_{T-1}) [S(T-1, x_T^* - x_{T-1}) - S(T-1, 0)] .
\end{aligned}$$

The method continues until time 0. For any $t < T$, super-replication requires finding x_t and y_t that minimize

$$y_t + x_tS(t-1, 0) + (x_t - x_{t-1}) [S(t-1, x_t - x_{t-1}) - S(t-1, 0)] , \quad (50)$$

subject to the constraint

$$\begin{aligned}
& y_t + x_tS(t-1, 0) + x_t [S(t, 0) - S(t-1, 0)] \quad (51) \\
& = y_{t+1}^* + x_{t+1}^*S(t, 0) + (x_{t+1}^* - x_t) [S(t, x_{t+1}^* - x_t) - S(t, 0)] ,
\end{aligned}$$

given x_{t+1}^* from the previous solution.

References

- [1] Bielecki, T., and M. Rutkowski, 2002, *Credit Risk: Modeling, Valuation and Hedging*, Springer-Verlag, New York.
- [2] Blais, M., and P. Protter, 2005, “An Analysis of the Supply Curve for Liquidity Risk through Book Data,” working paper, Cornell University.
- [3] Boyle, P., and D. Emanuel, 1980, “Discretely Adjusted Option Hedges,” *Journal of Financial Economics*, 8, 259-282.
- [4] Boyle, P., and T. Vorst, 1992, “Option Replication in Discrete Time with Transaction Costs,” *Journal of Finance*, 48, 271-293.
- [5] Çetin, U., 2003, “Default and Liquidity Risk Modeling,” Ph.D. thesis, Cornell University.
- [6] Çetin, U., R. Jarrow, and P. Protter, 2004, “Liquidity Risk and Arbitrage Pricing Theory,” *Finance and Stochastics*, 8, 311-341.
- [7] Chen, Z., Stanzl, W., and M. Watanabe, 2001, “Price Impact Costs and the Limit of Arbitrage,” working paper, Yale University.
- [8] Cheridito, P., 2003, “Arbitrage in Fractional Brownian Motion Models,” *Finance and Stochastics*, 7, 417-554.
- [9] Duffie, D., 1996, *Dynamic Asset Pricing Theory* (2d ed.), Princeton University Press, New Jersey.

- [10] Duffie, D. and P. Protter, 1988, "From Discrete to Continuous Time Finance: Weak Convergence of the Financial Gain Process," *Mathematical Finance*, 2, 1-16.
- [11] Edirisinghe, C., V. Naik, and R. Uppal, 1993, "Optimal Replication of Options with Transactions Costs and Trading Restrictions," *Journal of Financial and Quantitative Analysis*, 28, 117-138.
- [12] Frey, R., 1998, "Perfect Option Hedging for a Large Trader," *Finance and Stochastics*, 2, 115-141.
- [13] Hausman, J. A., A. W. Lo, and C. MacKinlay, 1992, "An Ordered Probit Analysis of Transaction Stock Prices," *Journal of Financial Economics*, 31, 319-379.
- [14] Jarrow, R. and S. Turnbull, 2004, *Derivative Securities* (2d ed.), Southwestern Publishers, Cincinnati.
- [15] Lee, C. M. C., and M. J. Ready, 1991, "Inferring Trade Direction from Intraday Data," *Journal of Finance*, 46, 733-746.
- [16] Leland, H., 1985, "Option Pricing and Replication with Transaction Costs," *Journal of Finance*, 40, 1283-1301.
- [17] Lillo, F., Farmer, J.D., and R. N. Mantegna, 2002, "Single Curve Collapse of the Price Impact Function for the New York Stock Exchange," working paper, Sante Fe Institute.
- [18] Platen E., and M. Schweizer, 1998, "On Feedback Effects from Hedging Derivatives," *Mathematical Finance*, 8, 67-84.

- [19] Protter, P., 2004, *Stochastic Integration and Differential Equations* (2d ed.), Springer Verlag, New York.
- [20] Rubinstein, M., 1985, “Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 through August 3, 1978,” *Journal of Finance*, 40, 455-480.
- [21] Schonbucher, P., and P. Wilmott, 2000, “The Feedback Effect of Hedging in Illiquid Markets,” *SIAM Journal on Applied Mathematics*, 61, 232-272.

Footnotes

1. Two other risks are often discussed in the literature; operational and model risk. However, since these risks are due to the legal system and model usage respectively, they may be considered non-financial.
2. This notion of liquidity risk is valid even if the cost per transaction is stable over time since the sign and size of transactions generated by the hedge strategy remain functions of the random stock price.
3. Chen, Stanzl and Watanabe (2001) examine the impact of illiquidity on the price of subsequent transactions, but not the price impact of trades within the context of a coherent mathematical model. By means of a statistical analysis, Lillo, Farmer and Mantegna (2002) also construct a supply curve without a mathematical model.
4. These trading strategies have been previously studied by Cheridito (2003) in the context of fractional Brownian motions.
5. As a consequence, our model is not appropriate for applications involving the liquidation or acquisition of large positions in the stock relative to the number of outstanding shares.
6. X_t and Y_t are predictable and optional processes, respectively.
7. Here, $H(t, \omega)$ is a predictable process and τ is a predictable ($F_t : 0 \leq t \leq T$) stopping time.
8. The space $L^2(d\mathbb{Q})$ is the set of F_{T-} measurable random variables that are square integrable using the probability measure \mathbb{Q} .

9. Both L_T and Y_T^n are already normalized by the value of the money market account.
10. An integral of a continuous function with respect to the Lebesgue measure is of bounded variation and continuous.
11. Subsequent research by Blais and Protter (2005) provides additional support for the linear supply curve structure given in expression (15) below for liquid stocks. Our stocks, with options trading on the CBOE, would fall into this category. For illiquid stocks, the supply curve appears to be piecewise linear but with stochastic slopes and a jump discontinuity at zero.
12. In reality, there may exist economies of scale for option pricing since replication costs are not linear in the number of underlying contracts.
13. The function $\text{sgn}(x)$ is defined as
$$\left\{ \begin{array}{ll} 1 & \text{if } x > 0 \\ 0 & \text{if } x = 0. \\ -1 & \text{if } x < 0 \end{array} \right.$$
14. For brevity, detailed tables containing the results are unreported but available upon request.
15. Dramatic evidence of this property is found around stock splits. Note that α expresses illiquidity in percentage terms while L_T is a dollar-denominated quantity.
16. Since $Z_T = Y_T + x_T S(T, 0)$, the same principle applies at time T with x_T being non-zero.
17. Since the liquidity cost for the *optimal* hedge strategy declines with more frequent trading, in theory, Δ would be reduced by an infinitely powerful computer until it reached the lower

bound δ , while in practice, the option price converges (to two decimal places) for $\Delta > \delta$. To clarify, Δ may be further divided into Δ_H and Δ_S , corresponding to the time interval between hedge portfolio rebalancings and stock price movements respectively. This enables $\Delta_S \rightarrow 0$ (decline below δ), producing a geometric Brownian motion as in our comparative study of non-optimal hedging strategies. However, with regards to our optimal trading strategy in this section, $\Delta_H = \Delta_S$, and there is no distinction.

18. The trading strategy in equation (10) is not implemented as it is continuous and of finite variation by design.
19. Bloomberg implied volatilities are computed using closing prices of three options which are closest to being at-the-money (across possible maturities). Closing prices are the bid (ask) if the last transaction is below (above) the bid (ask) or the transaction price itself if it lies between the option's bid-ask spread.
20. Justification for these values follows from plotting the distribution of absolute transactions (unsigned order flow) although some dependence on the strike price is detected.
21. As a consequence of the numerical grid, a minimum distance between trades is enforced in the random time non-optimal hedging approach, although a Brownian motion process may generate $x_u - x_t$ values that exceed $\theta > 0$ for $u - t < \delta$ in theory.
22. Note that the optimal hedge need not produce lower liquidity costs than the non-optimal strategies which are not constrained to produce non-negative approximation errors.

23. The slightly smaller liquidity costs for the two day non-optimal fixed time hedge strategy, versus its one day counterpart, result from being able to avoid certain intermediate transactions. For example, an increase (decrease) in the geometric Brownian motion process followed by a decrease (increase) to a similar level implies an intermediate transaction may be ignored (at least reduced) by the non-optimal strategy with a larger Δ . This does not contradict Lemma 1 which applies to optimal hedge strategies that super-replicate the option as discussed after its proof.
24. It is inappropriate to infer implied volatilities for the optimal trading strategy using the Black Scholes formula. Indeed, if the marketplace was aware of the optimal strategy, then the Black Scholes formula would not be utilized to compute prices or infer implied volatilities. Instead, obtaining an implied volatility for the optimal strategy requires an iterative procedure that involves the dynamic program. In addition, the binomial structure underlying the dynamic program requires smaller time intervals to ensure the corresponding option prices have converged. Otherwise, their values may be below those computed using the Black Scholes model as a result of the numerical approximation.
25. To clarify, x and y in equation (37) are not equal to x_1 and y_1 respectively in Section 4. Instead, they refer to positions *before* the initial trade. Thus, y in equation (37) equals $x_1S(0,0) + y_1 + \alpha S(0,0)x_1^2$ and $x = 0$.
26. The only change in the portfolio holdings during the interval $(t - 1, t]$ occurs right after time $t - 1$.

Figure 1: Plot of estimated α parameters each day of sample period from January 3, 1995 to December 31, 1998 based on equation (15) for IBM, Federal Express (FDX) and Barnes & Noble (BKS). The dotted line denotes the average daily stock price of the firm. These three companies represent high, medium, and low liquidity firms with respect to NYSE stocks that have traded CBOE options.

Figure 2: Plot of estimated α parameters each day of sample period from January 3, 1995 to December 31, 1998 based on equation (15) for GE and Reebok (RBK). The dotted line denotes the average daily stock price of the firm. Note the inverse relationship between α and $S(t, 0)$.

Figure 3: Plot of the product $\alpha(t) \cdot S(t, 0)$ each day of sample period from January 3, 1995 to December 31, 1998 for IBM, FDX and BKS. The value of $S(t, 0)$ is the average stock price on a particular date according to transactions in the TAQ database.

Figure 4: Plot of the product $\alpha(t) \cdot S(t,0)$ each day of sample period from January 3, 1995 to December 31, 1998 for GE and RBK. The value of $S(t,0)$ is the average stock price on a particular date according to transactions in the TAQ database.

Figure 5: Plots of the implied volatility across five different strike prices and two distinct maturities when 25 options (each on 100 shares) are replicated. The plots above correspond to a representative stock with a current price of \$50 whose underlying volatility is 30% per annum. In \$5 increments, the strike price ranges from \$40 to \$60 (in-the-money to out-of-the-money) while maturities of 30 and 90 days are considered. The α parameter is 1.25×10^{-4} per lot transacted while the trading strategy consists of implementing the Black Scholes hedge parameters at fixed time intervals of $\Delta = 2$ days. The disparity between the 30% volatility and those implied from option prices is the result of illiquidity. Specifically, the Black Scholes implied volatilities are computed under the assumption of a perfectly liquid market ($\alpha = 0$) for the underlying asset, while option prices are increased by the liquidity cost of their replication. Consequently, higher option prices induced by illiquidity generate an upward bias in their implied volatilities.

Table 1: Summary statistics for the daily parameter estimates of α (in lots) and μ generated by the regression model in equation (15) for each of the five firms. The second column records the average number of daily transactions available. Under each of the α summary statistics is their standard error. The ninth and tenth columns detail the number of days for which the parameter estimates are significant at the 5% level, along with the corresponding percentage relative to the 1,011 day sample period. The last two columns record the average stock price and implied volatility of each firm during the sample period.

Company Ticker	Mean n	Parameter	1st Percentile	Median	99th Percentile	Mean	Standard Deviation	5% Level Days	5% Level Percentage	Stock Price	Volatility (annual)
GE	1,061	$\hat{\alpha} \times 10^{-4}$	0.20	0.44	1.63	0.59	0.38	1,011	99.99	\$75	23.25%
		$SE \times 10^{-4}$	0.012	0.050	0.106	0.061	0.038	-	-	-	-
IBM	656	$\hat{\mu} \times 10^{-7}$	-33.96	-3.62	13.03	-4.34	9.08	24	2.37	-	-
		$\hat{\alpha} \times 10^{-4}$	0.06	0.17	0.43	0.19	0.08	998	98.71	\$115	30.93%
FDX	249	$SE \times 10^{-4}$	0.018	0.033	0.094	0.037	0.016	-	-	-	-
		$\hat{\mu} \times 10^{-7}$	-26.22	-2.81	13.76	-3.71	7.93	31	3.07	-	-
BKS	318	$\hat{\alpha} \times 10^{-4}$	0.15	0.53	1.32	0.56	0.21	971	96.04	\$65	28.41%
		$SE \times 10^{-4}$	0.051	0.137	0.311	0.143	0.063	-	-	-	-
RBK	159	$\hat{\mu} \times 10^{-7}$	-13.12	-0.61	10.90	-0.66	4.83	50	4.95	-	-
		$\hat{\alpha} \times 10^{-4}$	0.34	1.18	2.84	1.28	0.50	962	95.15	\$35	37.39%
RBK	159	$SE \times 10^{-4}$	0.099	0.263	0.828	0.306	0.163	-	-	-	-
		$\hat{\mu} \times 10^{-7}$	-30.46	-0.73	14.58	-1.41	7.56	46	4.55	-	-
RBK	159	$\hat{\alpha} \times 10^{-4}$	0.11	0.99	2.50	1.08	5.00	957	94.66	\$35	36.22%
		$SE \times 10^{-4}$	0.152	0.278	0.684	0.317	0.119	-	-	-	-
RBK	159	$\hat{\mu} \times 10^{-7}$	-23.26	-2.70	14.24	-3.09	7.01	25	2.47	-	-

Table 2: Summary of order flow and transaction prices. Recorded below are the relationships between consecutive transaction prices as a function of order flow. The first two rows impose no constraints on the magnitude of the transaction. The next four rows have small (large) trades defined as those less than or equal to (greater than) 10 lots. Note that large trades are not included in our subsequent estimation procedure for α as these transactions are deemed irrelevant for the purposes of frequent hedging.

Trade Type and Hypothesis			Company Ticker				
First Trade (x_1)	Second Trade (x_2)	Transaction Price Hypothesis	GE	IBM	FDX	BKS	RBK
Sell	Buy	$S(t_1, x_1) \leq S(t_2, x_2)$	100.00%	100.00%	100.00%	100.00%	100.00%
Buy	Sell	$S(t_1, x_1) \geq S(t_2, x_2)$	100.00%	100.00%	100.00%	100.00%	100.00%
Small Buy	Large Buy	$S(t_1, x_1) \leq S(t_2, x_2)$	98.77%	98.58%	98.46%	98.58%	99.05%
Small Sell	Large Sell	$S(t_1, x_1) \geq S(t_2, x_2)$	98.10%	97.99%	98.61%	98.43%	98.95%
Large Buy	Small Buy	$S(t_1, x_1) \geq S(t_2, x_2)$	93.60%	90.28%	80.59%	80.53%	83.88%
Large Sell	Small Sell	$S(t_1, x_1) \leq S(t_2, x_2)$	90.33%	87.33%	78.87%	78.67%	81.89%

Table 3: Summary of optimal trading strategy for 10 options, each on 100 shares. Option prices in the third and fourth columns are reported for an individual option on 100 shares with the other entries representing quantities for an imbalance of 10 options. For emphasis, unlike option prices, the liquidity cost is not linear in the number of options (or shares) under consideration and is not rescaled as a consequence. For at-the-money options, the strike price equals the initial stock price recorded in Table 1. The stock price is then increased (decreased) by \$5 for in-the-money (out-of-the-money) options. The volatility parameter is also found in Table 1. The first option price $x_1^0 S(0, 0) + y_1^0$ utilizes hedge parameters x_1^0 and y_1^0 that are solved from a dynamic program with $\alpha = 0$. This “frictionless” option price proxy serves as the basis for the percentage impact of illiquidity. The second reported option price includes the liquidity cost at time zero of forming the replicating portfolio, $\alpha S(0, 0)x_1^2$, with x_1 and y_1 computed for an economy with $\alpha > 0$. Hence, the values x_1 versus x_1^0 and y_1 versus y_1^0 are not comparable. After time zero, we infer an implied liquidity cost using the optimal transactions at each node of the binomial procedure for comparison with the approximate liquidity cost. The last column is the approximate liquidity cost in equation (23) for transactions after time zero.

Option Characteristics		Individual Option		Costs Associated with Replicating Portfolio for 10 Options					
Company Name	Option Moneyness	Option Price with $\alpha = 0$	Option Price with $\alpha S(0, 0)x_1^2$	Liquidity Cost (at $t = 0$)	Liquidity Cost (after $t = 0$)	Total Liquidity Cost	Percentage Impact	Excess at T	Approximate (after $t=0$) Liquidity Cost
GE	In	548.91	551.40	19.31	5.77	25.08	0.46	0.00	5.79
	At	177.94	180.20	9.59	10.06	19.65	1.10	0.00	10.08
	Out	46.94	47.11	0.86	0.82	1.68	0.36	0.00	0.83
IBM	In	730.56	731.42	8.61	3.46	12.07	0.17	0.00	3.47
	At	362.91	364.04	4.81	4.46	9.27	0.26	0.00	4.48
	Out	229.90	230.37	2.44	2.29	4.73	0.21	0.00	2.30
FDX	In	562.42	527.99	15.70	4.96	20.66	0.37	0.00	4.98
	At	188.42	190.29	8.46	8.09	16.55	0.88	0.00	8.12
	Out	56.07	56.25	0.92	0.87	1.79	0.32	0.00	0.88
BKS	In	517.87	520.21	23.39	4.52	27.91	0.54	0.00	4.54
	At	133.50	135.82	11.07	10.18	21.25	1.59	0.00	10.25
	Out	12.25	12.28	0.12	0.11	0.23	0.19	0.00	0.11
RBK	In	513.40	515.40	20.02	3.67	23.69	0.46	0.00	3.69
	At	129.33	131.29	9.16	8.44	17.60	1.36	0.00	8.49
	Out	8.36	8.37	0.05	0.05	0.10	0.12	0.00	0.05

Table 4: Summary of hedging strategy at fixed time points with random quantities for 10 options, each on 100 shares. Option prices below are reported for an individual option on 100 shares with the other entries representing quantities for an imbalance of 10 options. For emphasis, unlike option prices, the liquidity cost is not linear in the number of options (or shares) under consideration. This strategy consists of hedging the options at predetermined discrete time points using Black Scholes hedge parameters with the stock price evolving as a geometric Brownian motion. The liquidity cost and approximation error are recorded below along with their percentage impact on the option price. For at-the-money options, the strike price equals the initial stock price recorded in Table 1. The stock price is then increased (decreased) by \$5 for in-the-money (out-of-the-money) options. The volatility parameter is also found in Table 1. The fifth column which states the combined frictionless option price with its associated liquidity cost is for ease of reference, and divides the total liquidity cost across the 10 options. However, as seen in equation (16), liquidity costs increase quadratically with the number of options hedged while the price increases at a linear rate. Thus, the figures below are not directly comparable with those in Table 3 but are provided for illustration.

Company Name	Option Characteristics			Individual Option			Hedging Error (10 options)			Approximate Liquidity	
	Option Money	Hedging Frequency	Black Scholes	With Liquidity	Liquidity Cost	Approximation Error	Liquidity %	Liquidity Cost	Liquidity %		
GE	In	$\Delta = 2$ days	546.56	550.29	37.31	38.16	0.68	38.24	0.70		
		$\Delta = 1$ day	546.56	550.34	37.84	24.32	0.69	38.54	0.71		
	At	$\Delta = 2$ days	200.78	202.68	18.98	72.92	0.95	20.12	1.00		
		$\Delta = 1$ day	200.78	202.80	20.16	31.94	1.00	20.93	1.04		
	Out	$\Delta = 2$ days	38.39	38.84	4.52	44.06	1.18	4.90	1.28		
		$\Delta = 1$ day	38.39	38.92	5.29	23.01	1.38	5.51	1.43		
IBM	In	$\Delta = 2$ days	715.01	716.43	14.23	131.00	0.20	14.82	0.21		
		$\Delta = 1$ day	715.01	716.49	14.75	58.14	0.21	15.20	0.21		
	At	$\Delta = 2$ days	409.50	410.45	9.47	144.66	0.23	9.98	0.24		
		$\Delta = 1$ day	409.50	410.50	10.04	67.46	0.25	10.44	0.25		
	Out	$\Delta = 2$ days	199.20	199.73	5.27	126.35	0.26	5.64	0.28		
		$\Delta = 1$ day	199.20	199.77	5.72	61.76	0.29	5.93	0.30		
FDX	In	$\Delta = 2$ days	555.16	558.19	30.34	48.16	0.55	31.15	0.56		
		$\Delta = 1$ day	555.16	558.25	30.89	27.04	0.56	31.53	0.57		
	At	$\Delta = 2$ days	212.61	214.19	15.82	68.57	0.74	16.73	0.79		
		$\Delta = 1$ day	212.61	214.28	16.67	34.00	0.78	17.29	0.81		
	Out	$\Delta = 2$ days	44.64	45.05	4.11	42.23	0.92	4.45	1.00		
		$\Delta = 1$ day	44.64	45.12	4.76	21.44	1.07	4.92	1.10		
BKS	In	$\Delta = 2$ days	520.96	525.38	44.23	21.11	0.85	45.02	0.86		
		$\Delta = 1$ day	520.96	525.44	44.78	14.67	0.86	45.48	0.87		
	At	$\Delta = 2$ days	150.64	152.60	19.62	51.75	1.30	20.74	1.38		
		$\Delta = 1$ day	150.64	152.73	20.85	21.70	1.38	21.61	1.43		
	Out	$\Delta = 2$ days	12.04	12.26	2.15	18.58	1.79	2.28	1.90		
		$\Delta = 1$ day	12.04	12.28	2.43	8.66	2.02	2.44	2.03		
RBK	In	$\Delta = 2$ days	518.68	522.45	37.70	17.21	0.73	38.33	0.74		
		$\Delta = 1$ day	518.68	522.49	38.09	12.39	0.73	38.65	0.75		
	At	$\Delta = 2$ days	145.93	147.59	16.59	48.21	1.14	17.56	1.20		
		$\Delta = 1$ day	145.93	147.68	17.52	24.65	1.20	18.17	1.24		
	Out	$\Delta = 2$ days	10.52	10.68	1.64	12.85	1.56	1.74	1.66		
		$\Delta = 1$ day	10.52	10.72	1.95	6.08	1.85	1.96	1.86		

Table 5: Summary of hedging strategy at random time points with fixed quantities for 10 options, each on 100 shares. Option prices below are reported for an individual option on 100 shares with the other entries representing quantities for an imbalance of 10 options. For emphasis, unlike option prices, the liquidity cost is not linear in the number of options (or shares) under consideration. This strategy consists of hedging the options at random time points using Black Scholes hedge parameters when needed to rebalance the replicating portfolio. The stock price evolves as a geometric Brownian motion. The liquidity cost and approximation error are recorded below along with their percentage impact on the option price. For at-the-money options, the strike price equals the initial stock price recorded in Table 1. The stock price is then increased (decreased) by \$5 for in-the-money (out-of-the-money) options. The volatility parameter is also found in Table 1. The fifth column which states the combined frictionless option price with its associated liquidity cost is for ease of reference, and divides the total liquidity cost across the 10 options. However, as seen in equation (16), liquidity costs increase quadratically with the number of options hedged while the price increases at a linear rate. Thus, the figures below are not directly comparable with those in Table 3 but are provided for illustration.

Company Name	Option Characteristics			Individual Option			Hedging Error (10 options)			Approximate Liquidity	
	Option Moneyness	Amount Hedged	Black Scholes	With Liquidity	Liquidity Cost	Approximation Error	Liquidity %	Liquidity Cost	Liquidity %		
GE	In	$\theta = 10$ lots	546.56	550.46	38.96	3.14	0.71	39.37	0.72	39.37	0.72
		$\theta = 5$ lots	546.56	550.46	39.03	1.72	0.71	39.47	0.72	39.47	0.72
	At	$\theta = 10$ lots	200.78	202.97	21.91	3.92	1.09	22.16	1.10	22.16	1.10
		$\theta = 5$ lots	200.78	202.97	21.88	2.78	1.09	22.11	1.10	22.11	1.10
	Out	$\theta = 10$ lots	38.39	39.02	6.29	2.46	1.64	6.17	1.61	6.17	1.61
		$\theta = 5$ lots	38.39	39.02	6.28	1.74	1.64	6.16	1.60	6.16	1.60
IBM	In	$\theta = 10$ lots	715.01	716.56	15.51	7.97	0.22	15.77	0.22	15.77	0.22
		$\theta = 5$ lots	715.01	716.56	15.46	6.51	0.22	15.71	0.22	15.71	0.22
	At	$\theta = 10$ lots	409.50	410.59	10.91	6.83	0.27	11.03	0.27	11.03	0.27
		$\theta = 5$ lots	409.50	410.59	10.90	3.80	0.27	11.01	0.27	11.01	0.27
	Out	$\theta = 10$ lots	199.20	199.85	6.51	5.52	0.33	6.46	0.32	6.46	0.32
		$\theta = 5$ lots	199.20	199.86	6.55	5.04	0.33	6.52	0.33	6.52	0.33
FDX	In	$\theta = 10$ lots	555.16	558.35	31.86	3.20	0.57	32.26	0.58	32.26	0.58
		$\theta = 5$ lots	555.16	558.35	31.86	0.78	0.57	32.27	0.58	32.27	0.58
	At	$\theta = 10$ lots	212.61	214.42	18.09	4.05	0.85	18.28	0.86	18.28	0.86
		$\theta = 5$ lots	212.61	214.42	18.10	2.60	0.85	18.31	0.86	18.31	0.86
	Out	$\theta = 10$ lots	44.64	45.20	5.59	1.39	1.25	5.45	1.22	5.45	1.22
		$\theta = 5$ lots	44.64	45.20	5.63	0.17	1.26	5.49	1.23	5.49	1.23
BKS	In	$\theta = 10$ lots	520.96	525.52	45.61	1.27	0.88	46.14	0.89	46.14	0.89
		$\theta = 5$ lots	520.96	525.53	45.66	0.64	0.88	46.21	0.89	46.21	0.89
	At	$\theta = 10$ lots	150.64	152.90	22.57	2.78	1.50	22.83	1.52	22.83	1.52
		$\theta = 5$ lots	150.64	152.91	22.68	2.60	1.51	22.93	1.52	22.93	1.52
	Out	$\theta = 10$ lots	12.04	12.36	3.15	1.29	2.62	2.92	2.42	2.92	2.42
		$\theta = 5$ lots	12.04	12.36	3.17	1.02	2.64	2.92	2.43	2.92	2.43
RBK	In	$\theta = 10$ lots	518.68	522.56	38.79	0.67	0.75	39.19	0.76	39.19	0.76
		$\theta = 5$ lots	518.68	522.55	38.74	0.26	0.75	39.15	0.75	39.15	0.75
	At	$\theta = 10$ lots	145.93	147.83	19.00	2.65	1.30	19.20	1.32	19.20	1.32
		$\theta = 5$ lots	145.93	147.83	19.02	2.63	1.30	19.24	1.32	19.24	1.32
	Out	$\theta = 10$ lots	10.52	10.76	2.43	1.72	2.31	2.24	2.13	2.24	2.13
		$\theta = 5$ lots	10.52	10.76	2.36	1.24	2.24	2.18	2.07	2.18	2.07

Table 6: Implied volatility of options under the non-optimal Black Scholes trading strategies. The Black Scholes implied volatility derived from option prices which account for the impact of illiquidity in the fifth columns of Tables 4 and 5 are reported below. In addition, the true underlying volatility from Table 1 utilized in generating the option prices is given for ease of reference. Results for $\Delta = 1$ day and $\theta = 5$ lots are nearly identical to those of 2 days and 10 lots reported below respectively.

Option Characteristics			Implied Black Scholes Volatility	
Company Name	Option Moneyness	True Volatility	Black Scholes Fixed Time $\Delta = 2$ days	Black Scholes Random Time $\theta = 10$ lots
GE	In	23.25%	23.91	23.94
	At	23.25%	23.47	23.50
	Out	23.25%	23.34	23.38
IBM	In	30.93%	31.05	31.06
	At	30.93%	31.00	31.01
	Out	30.93%	30.98	30.99
FDX	In	28.41%	29.00	29.02
	At	28.41%	28.62	28.65
	Out	28.41%	28.50	28.54
BKS	In	37.39%	39.51	39.57
	At	37.39%	37.88	37.95
	Out	37.39%	37.57	37.63
RBK	In	36.22%	38.10	38.17
	At	36.22%	36.63	36.69
	Out	36.22%	36.38	36.41




Figure 1:





Figure 2:




Figure 3:





Figure 4:




Figure 5: