Hedging through a limit order book with varying liquidity

Rossella Agliardi* Ramazan Gençay†

Abstract

We relax the classical price-taking assumption and study the impact of orders of arbitrary size on price when the availability of liquidity is a concern in hedging. Our paper extends the earlier literature where illiquidity is modeled both in terms of depth and resilience. Some of our results hold for more general stochastic processes for the underlying, for example, a generic Lévy process and the model applies to a generic contingent claim. Furthermore, we offer an insight into the trade-off between hedging and price manipulation.

Keywords: hedging, large traders, limited liquidity, resilience, limit order book.

JEL No: G13, C22.

*Corresponding author: Dipartimento di Matematica, Università di Bologna and IMATI - CNR, Pavia, Italy. Email: rossella.agliardi@unibo.it

†Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, British Columbia, V5A 1S6, Canada and RCEA, Rimini, Italy. Email: rgencay@sfu.ca. Ramazan Gençay gratefully acknowledges financial support from the Natural Sciences and Engineering Research Council of Canada and the Social Sciences and Humanities Research Council of Canada.

1 Introduction

As computerized trading has been steadily growing, an effective management of the price impact effects, and the other frictions related to high-frequency orders, have become a critical focus for trading firms. Thus, price impact models have become central to decisions on how to trade optimally and they pose new challenges when more complex financial problems are at hand, such as the hedging and pricing of derivatives.

We explicitly solve the portfolio problem of a large trader (e.g. a corporate investor such as a hedge fund) that needs to hedge its derivative position in an illiquid environment and incurs a price impact as the one arising when trading through a limit order book. Most models/algorithms handle the problem of a stock trader who has a large block order to unwind and needs to minimize her/his market impact (optimal trade execution), while far less research work investigates how to implement an optimal hedging strategy in such an illiquid environment. Our contribution is to examine the costs associated with optimal hedging in an illiquid market with respect to the depth and the resilience of the limit order book.

The interplay between option hedging and the dynamics of the limit order book, which governs the underlying equity, has become more relevant with the introduction of options with shorter maturities.[1] Weekly options have the same contract specifications as standard options except for their weekly Thursday-to-Friday maturity. Since the introduction of weekly options in 2005, the average daily volume of weekly options has grown, and their coverage has been extended to exchange-traded funds in addition to equities and indices. One-, two-, three- or four-day options are also becoming more visible in the marketplace. Such options provide market participants with a tool for hedging overnight risk, weekend risk, or the risk associated with special events such as earnings announcements and economic reports.[2] In this paper, we investigate the costs associated with hedging a European option when both the depth and the resilience speed of the limit order book for the underlying equity are associated with the illiquidity costs.

We use a discrete setting for trading times and provide an explicit solution that indicates the optimal strategy for a hedger who aims to minimize the illiquidity costs incurred and the hedging error. The classical models of option pricing and hedging assume that agents are price takers. In practice, however, large traders face adverse market impact when trading with limited liquidity. Thus, in many models of optimal block execution, the agent trades patiently to minimize the

[1]The Chicago Board Options Exchange (CBOE) launched weekly options, "Weeklys," on Friday, October 28, 2005.

[2]NYSE Euronext has a short maturity index option, and CBOE plans to offer option maturities shorter than a week.

costs associated with price impact. A block trade by a large hedger can significantly affect asset prices, which makes perfect hedging a non-optimal strategy. In particular, a Black-Scholes hedge usually results in higher replication costs than a more patient strategy that is intended to minimize the impact cost without increasing significantly the distance of the replicating portfolio from the target payoff. Price impact decreases the effectiveness of the classical hedging strategies, and what Bouchaud and Potters (2003) refer to as the "Black-Scholes miracle" does not take place.[3]

The challenge of including liquidity risk in option pricing and hedging (namely, the relaxation of the classical price-taking assumption and the study of the impact of orders of arbitrary size on price) is considered in a vast body of literature. A stream of research presents "large traders" models in which the price of an asset depends on certain fundamentals and on a large trader's activity based on the reaction function for his current holdings. The equilibrium stock price is a function of a fundamental value and the current stock position of the large trader. A reaction function can be viewed as a reduced-form representation of an economic equilibrium model like those proposed by Frey and Stremme (1997), Frey (1998), and Platen and Schweizer (1998) and of the general framework offered in Bank and Baum (2004). These models represent the feedback between asset prices and the competing market price manipulation generated by adverse price adjustments due to large trades. One main aim of this research is to quantify the price distortions induced by feedback effects from hedging strategies, such as a level-dependent volatility. Although the size of the trade is an important component of liquidity risk, other features also play a role, including the trading rate of individual traders and the speed of the market clearing by all agents. Thus, it has been argued that such models should be considered models of feedback effects rather than models of liquidity effects.

A theoretical framework for a theory of option pricing under liquidity risk is offered in Cetin, Jarrow and Protter (2004) and other related works (see also Cetin, Jarrow, Protter and Warachka (2006) and the comprehensive survey in Jarrow and Protter (2007)), in which the concept of liquidity risk is integrated into the classical theory for an exogenously given supply curve. An arbitrary trader acts as a price taker with respect to a supply curve for the shares. As such, the supply curve depends only on the order size and is otherwise independent of the trader's past actions, endowments and risk aversion; the price effect of an order is limited to the moment when the order is placed, and the order has no lasting impact on the price process. In our paper, the persistent effects of orders that affects future price dynamics are taken into account. As Roch and

[3] In essence, in the "Black-Scholes miracle", the residual risk is zero, and Black and Scholes' Δ-hedge allows one to completely eliminate the risk associated with an option, but this occurs only in special cases. As Bouchaud and Potters (2003) note, "zero-risk is rather an exception than the rule".

Soner (2011) note, "due to its exogenous nature, the supply curve does not take into account the impact of trading on its future evolution".[4] Matsumoto (2009) considers general trade times and focuses on a scenario in which the investor does not know the success of his trade in advance. However, other liquidity effects such as the costs incurred, price impacts and the presence of large investors are not considered in Matsumoto's paper.

These studies do not explicitly consider the shape and dynamics of the order book. Explicit references to the limit order book (hereafter LOB) literature within a hedging problem are offered in Rogers and Singh (2010) and in Almgren and Li (2010), which examine the hedging of options via the purchase of the underlying assets through LOBs.[5] In Rogers and Singh (2010), only the temporary impact is considered; illiquidity is modeled as a transaction cost based on the depth of the book and the trade size. In Almgen and Li (2010), the permanent impact of trade size on stock prices is modeled; there is a component of the impact that affects all current and future stock prices equally. The aim is to minimize the hedging error or to maximize the terminal value of the replication portfolio. Both models employ a continuous time setting. These models do not consider the transient nature of the price impact (i.e., its decay over time) or the resilience of the LOB (i.e., the speed at which the limit order book rebuilds itself after being consumed by a trade). In contrast, we show that the time needed to restore a book after the execution of an order plays a critical role in determining the optimal execution strategy.

In our paper, an extension of this model is presented that takes these features into account. Resilience is modeled in terms of a parameter that measures the speed at which the spread between the fundamental price and the best ask (bid) price (which is generated by past trades) is resolved. A model with a fully permanent impact can be viewed as a special case with zero resilience speed, whereas a model with an exclusively temporary impact can be viewed as a limit case with infinite resilience speed. Furthermore, our paper differs from previous research in that our main results in Section 3 also hold for more general stochastic processes for the underlying assets (for example, a generic Lévy process), thus allowing for jumps in the price paths and non-Gaussian return distribution.[6] Because of the central limit theorem, the distribution of the stock returns can be approximated using a Gaussian distribution when data from a long period are used and aggregation thins out the tails; however, this does not hold for dynamic hedging, which depends

[4] See also Cetin, Soner and Touzi (2010).

[5] A different approach is taken in Kraft, C. Kühn (2011); there, the price impact is assumed to affect the drift and volatility of the stock dynamics.

[6] We refer to Geman (2002) for an explanation of the motivations that have led to introduce Lévy processes in finance.

on high-frequency data.[7]

Finally, we provide conditions under which the agent has in incentive to stray away from the optimal hedging strategy and to engage in a manipulation pursuit. Thus, we gain an insight into the issue of the trade-off between hedging and manipulating which is usually overlooked in the literature.

We stress that the results of this paper are by no means straightforward or self-evident, due to the several conflicting objectives of the hedger's problem in an illiquid environment. First, the usual trader's dilemma applies where if a trade is too aggressive it incurs high market impact costs, while if execution is too passive, it is exposed to significant timing risk due to adverse price movement. Therefore, the trader has to balance the trade-off between cost and risk. Second, in the case of a lasting price impact the hedger may benefit from her/his stock holdings, but the hedging activity negatively affects the option payoff. At the same time, the illiquidity costs that arise from the hedging activity can be offset by the price manipulation strategy. Finally, in the case of a quickly dissipating price impact there is no scope for price manipulation and the only hedger's concern is to minimize the illiquidity costs (by trading less aggressively) while not increasing significantly the mis-hedging error. Thus, there is a complicated interplay between hedging, exploiting the price impact and minimizing incurred costs, which makes the problem non trivial, and the resilience plays a crucial role in dictating the optimal strategy.

Section 2 presents the setup. Section 3 studies the special case of an entirely temporary impact, which can be considered as a discrete version of the case presented by Rogers and Singh (2010) where we allow for more general price dynamics for the underlying equity. Section 4 provides an explicit solution to the general problem with finite resilience. Section 5 offers some comments on sensitivity analysis and the robustness of the model to possible extensions. Section 6 concludes.

2 The setup

Consider an economy that consists of a risky asset (a stock) that is traded through a limit order book (LOB), an associated contingent claim and a risk-free asset. The LOB is exposed to repeated market orders by a large trader. For the sake of simplicity, the risk-free interest rate is assumed to be zero. We adopt an LOB-based market impact model as in Obizhaeva and Wang (2005): a model in which the size and decay of the price impact are interpreted within the context of an order placed in an electronic LOB and in which it is assumed that the impact is linear and that its

[7] See Taleb (1997) for a discussion of this point.

time decay is exponential.[8]

We adopt a discrete time schedule in which $t_0 < t_1 < .... < t_N$ denotes the times at which trade occurs and a fixed filtered probability space with the triplet $(\Omega, \Im, P, \{\Im_k\}_{k=0,...,N})$. Let the existence of such a filtered probability space be given, and let S_t^0 denote the unaffected share price: its value in the long term when large traders are inactive. We assume that S_t^0 is a martingale with respect to the given filtration. Usually, in the financial models of LOBs, the unaffected stock price follows either a Bachelier or a geometric Brownian motion process, i.e., $dS_t^0 = \sigma dW_t$ or $dS_t^0 = \sigma S_t^o dW_t$ where W_t is a Wiener process with respect to the fixed filtration that represents the information available to market investors. Here, S_t^0 may be any martingale. To simplify the notation, $S_{t_k}^0$ will be denoted as S_k^0, the expectations $E(. \mid \Im_k)$ will be denoted as E_k and the price process S_k^0 is a martingale with respect to the fixed filtration, i.e., $E_k(S_{k+1}^0) = S_k^0$.

Let x_0, x_1, ... x_N denote the sizes of the trade orders at time t_0, $t_1, ...t_N$. We do not assume equidistant trading times - although equidistant portfolio re-balancing is a common practice - because the convergence of a hedging scheme can in some cases be improved by using unequally spaced dates (see Geiss and Geiss (2006)) and shrinking the interval width near the option expiry.

Consider a block-shaped limit order book, that is, a continuous price distribution with a constant height $q > 0$. When a buy (sell) order x_k is placed at time t_k, the ask (bid) price is shifted up (down) to S_k^+, and the order consumes the shares between the best ask (bid) price S_k where $S_k^+ = S_k + \frac{x_k}{q}$. In other words, q is the size of the order flow needed to change the stock price by one unit, i.e., it measures the depth of the LOB. The smaller the parameter q, the lower the liquidity of the stock. Let V_k denote the volume of potential orders between S_k^0 and S_k and assume that the gap between the two prices that is generated by previous trading is gradually mitigated according to $dV_t = -\rho V_t dt$. In other words, $V_t = e^{-\rho(t-s)} V_s$ in the time interval between the two trades.

The dynamics are as follows:

$$S_k = S_k^0 + \frac{1}{q} V_k \tag{1}$$

$$S_k^+ = S_k + \frac{x_k}{q} = S_k^0 + \frac{1}{q} V_k^+ \tag{2}$$

where $V_k^+ = V_k + x_k$ and

[8] This model is consistent with the no-dynamic arbitrage principle as stated in Gatheral (2010).

$$V_{k+1} = e^{-\rho(t_{k+1}-t_k)} V_k^+ \Rightarrow V_k = \sum_{j=0}^{k-1} e^{-\rho(t_k - t_j)} x_j \tag{3}$$

if one assumes $V_0 = 0$. Thus, the permanent impact is

$$S_{k+1} - S_k = S_{k+1}^0 - S_k^0 + \frac{1}{q}[e^{-\rho(t_{k+1}-t_k)} V_k^+ - V_k] \tag{4}$$

and the temporary impact is

$$S_k^+ - S_{k+1} = -(S_{k+1}^0 - S_k^0) + \frac{1}{q}[1 - e^{-\rho(t_{k+1}-t_k)}] V_k^+ \tag{5}$$

In Figure 1, we plot a trajectory for a stock price S_t that is affected by two subsequent purchases. In panel (a), the impact is only temporary, whereas in panel (b), there is a full permanent impact. In panel (c), a resilience effect mitigates the impact over time.

Because x_k may be positive or negative in a hedging problem, one must model both the buy side and the sell part of the LOB. Here, the simplification presented by Alfonsi, Fruth, Schied (2010, page 153) is adopted such that

$$V_k^+ = V_k + x_k \quad \text{and} \quad S_k^+ = S_k + \frac{x_k}{q}$$

holds regardless of the sign of x_k, i.e., for both purchases and sales.

Consider the problem of hedging a European option over a time horizon $[0, T]$ where $T = t_N$. The option is described by a square-integrable random variable G, as is usual in the hedging literature, where the closeness of a replicating portfolio to the target option payoff is measured in a mean-squared sense. In the absence of a price impact, the value of the option's price is a function, g, of the time t and the price S of the underlying asset; $G_k = E_k[G] = g(t_k, S_k)$, and G_k is also a martingale. We assume that even in the illiquid case, the option value is described by a function, g, of the time and the affected stock price S. We also assume that $g(t, S_t)$ is $\Im_t$-adapted.

Let R_0 denote the initial wealth of an agent who wants to hedge the option. Let $x_0, x_1, \ldots, x_{N-1}$ be the quantity of shares of the underlying asset that he purchases/sells at time $t_0, t_1, \ldots, t_{N-1}$, and let $X_k = \sum_{j=0}^{k-1} x_j$, $k = 1, \ldots, N$, denote the number of stocks held by the agent in the period $[t_{k-1}, t_k]$. (Define $X_0 = 0$). Assume that x_k is $\Im_k$-measurable, and hence, that X_k is $\Im_{k-1}$-measurable. In other words, the two stochastic processes X and x are predictable and adapted, respectively, with respect to the given filtration.

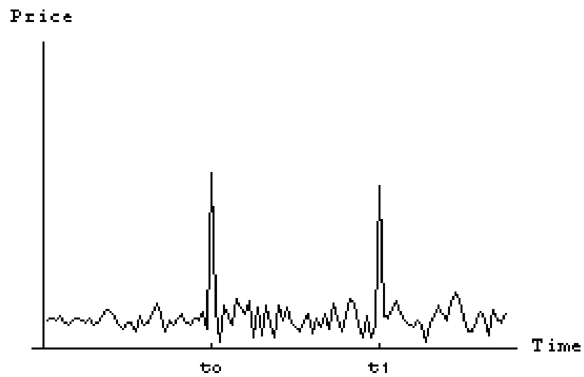

(a) Temporary Impact

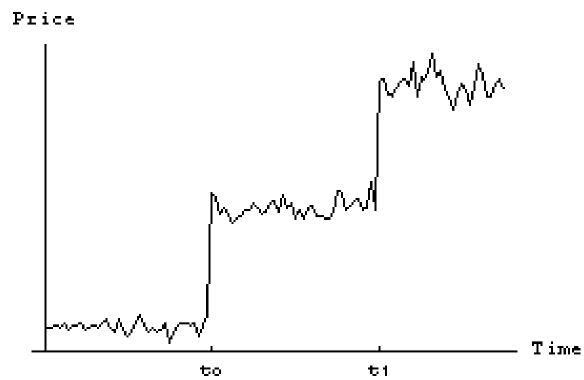

(b) Permanent Impact

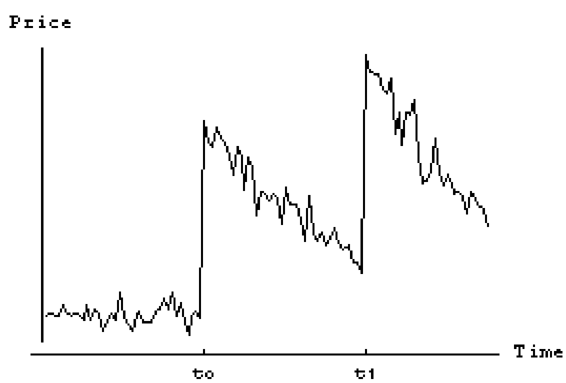

(c) Impact with Resilience

Figure 1: These figures illustrate the trajectory of a stock price S_t that is affected by two subsequent purchases. In panel (a), the impact is only temporary, whereas in panel (b), there is a full permanent impact. In panel (c), a resilience effect mitigates the impact over time.

Assuming a zero interest rate, the wealth dynamics are such that at time t_0^- is R_0, and soon after trading x_0 stocks, it is

$$R_0^+ = R_0 + (S_0^+ - S_0)x_0 - \frac{1}{2q}x_0^2 \tag{6}$$

at time t_1^- :

$$R_1 = x_0(S_1 - S_0^+) + R_0^+ = x_0[S_1^0 - S_0^0] - \frac{1}{2q}(1 - 2e^{-\rho(t_1 - t_0)})x_0^2 + R_0$$

Arguing recursively, we can state that

$$\begin{aligned} R_{k+1} = {} & \sum_{j=0}^{k} x_j(S_{k+1}^0 - S_j^0) - \frac{1}{q}\sum_{j=0}^{k}(\frac{1}{2} - e^{-\rho(t_{k+1}-t_j)})x_j^2 \\ & - \frac{1}{q}\sum_{\substack{i,j=0,\dots,k \\ i>j}} (e^{-\rho(t_i - t_j)} - e^{-\rho(t_{k+1}-t_i)} - e^{-\rho(t_{k+1}-t_j)})x_i x_j + R_0 \end{aligned} \tag{7}$$

In an illiquid environment, a large hedger must accept the trade-off between targeting a low hedging error and incurring a high impact cost. In what follows, we adopt a criterion similar to the one that Rogers and Singh (2010) use in minimizing the sum of the mean-squared hedging error and incurring illiquidity costs. Such an approach yields a tractable problem and makes it possible to "balance low illiquidity costs against poor replication".

The terminal wealth associated with the portfolio of the stock holding and a short option is

$$R_N - g(t_N, S_N) + g(t_0, S_0) \tag{8}$$

where R_N is given by Equation (7). Another expression for Equation (8) is $\sum_{k=0}^{N-1} r_k + R_0$ with

$$\begin{aligned} r_k = {} & X_{k+1}(S_{k+1}^0 - S_k^0) - g(t_{k+1}, S_{k+1}^0 + \frac{V_{k+1}}{q}) + g(t_k, S_k^0 + \frac{V_k}{q}) \\ & + \frac{1}{q}[X_{k+1}(V_{k+1} - V_k) - \frac{x_k^2}{2}] \end{aligned} \tag{9}$$

The expected value of the terminal wealth in Equation (8) can be written as

$$-\sum_{k=0}^{N-1} E_0[r_k^{(e)}] + R_0 \tag{10}$$

where

$$r_k^{(e)} = \frac{1}{q}[X_{k+1}(V_k - V_{k+1}) + \frac{x_k^2}{2}] + E_k[g(t_{k+1}, S_{k+1}^0 + \frac{V_{k+1}}{q})] - g(t_k, S_k^0 + \frac{V_k}{q}) \tag{11}$$

The sum $\sum_{k=0}^{N-1} r_k^{(e)}$ represents the cumulated effect of the price impact on the total wealth and can be viewed as a type of liquidity cost associated with the trading activity. On the other hand, $r_k^{(f)} = r_k - r_k^{(e)}$ represents the contribution due to the price fluctuation S_t^0 between the two trading times, t_k and $t_{k+1.}$

In this paper, the optimal hedging strategies are obtained by minimizing a quantity that reflects the need to protect the portfolio against market fluctuations in risky assets and simultaneously to optimize the expected costs due to the impact risk generated by the trading strategy. For example, in the case considered in Section 3, the quantity to be minimized is

$$\inf_{x_0,\ldots,x_{N-1}} E_0[\frac{1}{2q}\sum_{j=0}^{N-1} x_j^2] + \frac{1}{2}var_0[\sum_{j=0}^{N-1} x_j(S_N^0 - S_j^0) - G_N]$$

that reflects the need to protect the portfolio against market fluctuations in risky assets and simultaneously to optimize the expected costs due to the impact risk generated by the trading strategy. Thus, we depart from the traditional quadratic criterion for hedging error, which is usually employed when the market is frictionless but would not entirely account for the risk that is intrinsic to our setting. Indeed, the aim of eliminating or reducing volatility risk is only part of the goal, and a pure hedging approach may yield prohibitively high costs.

A trader in an illiquid environment faces the usual trader's dilemma, that is, "trading too aggressively incurs too much market impact while trading too passively incurs too much risk" (Kissel et al. (2004)). Here the goal is to minimize a function consisting of two terms. The first one represents the expected loss within the portfolio due to illiquidity, whereas the second term is taken from the minimum-variance hedge literature.[9]

[9] A discussion on the use of mean quadratic variation versus variance to represent minimization of portfolio volatility is offered in Forsyth et al (2012) in the context of the optimal trade execution literature.

The idea of splitting the objective function of a hedging problem into two parts under illiquidity was previously introduced by Rogers and Singh (2010) and is motivated mainly by tractability issues. We note that a hedging problem is usually more difficult to resolve than the problem of optimally liquidating an initial position because in the latter case, deterministic trajectories (in terms of time and initial position) are allowed or even desirable, whereas hedge paths also depend on the stochastic price of the risky asset. This simplifying approach does not investigate the possible complicated interactions between the two components of risk. However, this approach does highlight the different roles of those components in the hedging strategy. An additional advantage of this method is that one may also consider combinations of the two terms (see expression (12)) to represent the different weights that an agent attaches either to the illiquidity risk or to the market risk of mis-hedging.[10] In the next section, an optimal solution is obtained in the special case involving only a temporary impact. The effect of finite resilience is then studied in Section 4.

3 Hedging under a perfectly resilient market

In this section, a hedging problem is solved when the price impact is only temporary. The illiquidity costs that arise from the hedging strategy make perfect replication no longer advisable. The optimal strategy is the one that makes it possible to achieve two objectives: to finish close to the option value and to avoid incurring excessively large impact costs during the process by using a less aggressive strategy than in the liquid case. Unlike Rogers and Singh (2010), we allow for any stock dynamics and for any level of illiquidity; however, we work under a discrete-time formulation rather than a continuous one.

The problem to be solved takes the following form:

$$\inf_{x_0,\ldots,x_{N-1}} \quad E_0[\frac{1}{2q}\sum_{j=0}^{N-1} x_j^2] + \frac{\gamma}{2} var_0[\sum_{j=0}^{N-1} x_j(S_N^0 - S_j^0) - G_N] \tag{12}$$

Here the parameter $\gamma > 0$ captures the hedger's attitude towards the risk of market price fluctuation. If a level of $\gamma = 1$ is interpreted as representing an equal concern about the cost of hedging and the risk to be hedged, a level of $\gamma > 1$ ($\gamma < 1$) depicts a minor (major, respectively) concern about cost than the risk of market price fluctuations.

[10]The further discussion of this issue is in Section 5.

Proposition 1. *The optimal hedging strategy under criterion (12) is*

$$X^*_{k+1} = \frac{\gamma cov_k(S_{k+1}, G_{k+1}) + \frac{1}{q} X_k - \overline{b}_{k+1}}{\gamma var_k(S_{k+1}) + \overline{c}_{k+1}}$$

where

$$c_k = \frac{2}{q} - \frac{1}{q^2[\gamma var_k(S_{k+1}) + \overline{c}_{k+1}]}, \qquad c_N = \frac{1}{q},$$

$$b_k = \frac{\overline{b}_{k+1} - \gamma cov_k(S_{k+1}, G_{k+1})}{q[\gamma var_k(S_{k+1}) + \overline{c}_{k+1}]}, \qquad b_N = 0,$$

and $\overline{c}_{k+1} \equiv E_k[c_{k+1}]$ *and* $\overline{b}_{k+1} \equiv E_k[b_{k+1}]$.

Proof: See Appendix.

Example 1: The classical liquid case is obtained by letting $q \to \infty$; in other words, the optimal hedge is $X^*_{k+1} = \frac{cov_k(S_{k+1}, G_{k+1})}{var_k(S_{k+1})}$. See Schweizer (1995) or Föllmer and Schweizer (1989).

Example 2: Consider small trading intervals $\Delta t = t_k - t_{k-1}$. If S^0_t exhibits a geometric Brownian motion with volatility σ, then the expression for $var_k(S_{k+1})$ can be replaced with $\sigma^2 S_k^2 \Delta t$, and the expression for $cov_k(S_{k+1}, G_{k+1})$ can be rewritten as $\sigma^2 S_k^2 \partial_S g(t_k, S_k)\Delta t$. Obviously, the limit case $q \to \infty$ yields the Black-Scholes hedge.

Example 3: Consider small trading intervals $\Delta t = t_k - t_{k-1}$. Assume that the stock price $S^0_t = e^{L_t}$ where $(L_t)_{t \geq 0}$ is a Lévy process. Let ψ denote the characteristic exponent with respect to the martingale measure P, i.e., $E(e^{i\xi X_t}) = e^{-t\psi(\xi)}$, and $\psi(-i) = 0$ if the interest rate is 0. Then, $var_k(S_{k+1})$ can be approximated using $-\psi(-2i)S_k^2 \Delta t$ and $cov_k(S_{k+1}, G_{k+1})$ can be obtained as $e^{L_k}[\psi(D_x) - \psi(D_x - i)]g(t_k, e^{L_k})\Delta t$, where $\psi(D_x)$ is the pseudo-differential operator whose symbol is the characteristic exponent $\psi(\xi)$.[11]
If $L_t = -\frac{\sigma^2 t}{2} + \sigma W_t$ where W_t is a Wiener process, the Black-Scholes framework applies again (with a zero, risk-free interest rate), and in this case, $\psi(\xi) = \frac{i\sigma^2}{2}\xi + \frac{\sigma^2 \xi^2}{2}$. Thus,

$$var_k(S_{k+1}) \approx -\psi(-2i)S_k^2 \Delta t = \sigma^2 S_k^2 \Delta t$$

and

$$cov_k(S_{k+1}, G_{k+1}) \approx S_k i\sigma^2 D_x g(t_k, e^{L_k})\Delta t = \sigma^2 S_k^2 \partial_S g(t_k, S_k)\Delta t$$

[11] See Boyarchenko and Levendorskiĭ (2002) for the notation for pseudo-differential operators and Lévy processes.

as $D_x = -i\partial_x$. An alternative solution to the variance-optimal hedging problem for stochastic processes with stationary independent increments is offered in Hubalek et al. (2006) in the fully liquid case.

Example 4: If S_t^0 follows a Bachelier process with volatility σ, then $var_k(S_{k+1}) = \sigma^2(t_{k+1} - t_k)$, and c_k is deterministic. The expression for the hedges simplifies to

$$x_k^* = [1 - \frac{1}{q(\gamma\sigma^2(t_{k+1} - t_k) + c_{k+1})}](\partial_S g(t_k, S_k) - X_k)$$

which clearly shows that the intensity of the trading depends linearly on the distance from being perfectly hedged, i.e., $X_k - \partial_S g(t_k, S_k)$. A detailed proof of this approximate expression is offered in the Appendix.

Proposition 1 explicitly represents the optimal trading strategy under a temporary price impact caused by hedging through a limit order book whose depth is described by the parameter q. To show the effect of the optimal strategy x^* and to compare it with the classical Black-Scholes Δ-hedge (which does not take into account the price impact), we simulate both strategies in hedging a short position on a call. The path of the price of the underlying stock is simulated assuming a geometric Brownian motion and 100 trades. The case of an in-the-money call option near expiry is considered. The parameter γ is set equal to 1. Both hedging strategies – the Black-Scholes strategy and the optimal strategy – are plotted in the same graph (see Figure 2(a) and Figure 2(b)) for a comparison. Moreover, Figure 2(a) and Figure 2(b) are obtained by employing two different liquidity parameters, q, to show the sensitivity of the optimal hedging strategy to the depth of the LOB.

Figure 2(a) is obtained from $\frac{1}{q} = 0.001$, whereas in Figure 2(b), $\frac{1}{q} = 0.005$. $\frac{1}{q} = 0.001$ (0.005) indicates that a trade of 1000 (200) shares of stock would have a market impact of one unit of currency on the stock price. The optimal hedging strategy is generally less aggressive than the Black-Scholes strategy. As a result, the incurred cost due to the price impact is much lower. Of course, this effect is enhanced in the presence of a shallower LOB. This finding confirms that a hedger trading in an LOB should give up the perfect hedge. For example, in the case where $\frac{1}{q} = 10^{-3}$, the optimal hedge makes it possible to eliminate 80% of the impact cost that would be accrued using the Black-Scholes hedge while leaving acceptable terminal mis-hedging.

It is interesting to note that although a Black-Scholes strategy flattens out as it approaches the end, in the optimal strategy, the agent may trade more aggressively than a Black-Scholes trader when he is near the close of the trading period. Essentially, the intensity of the trading

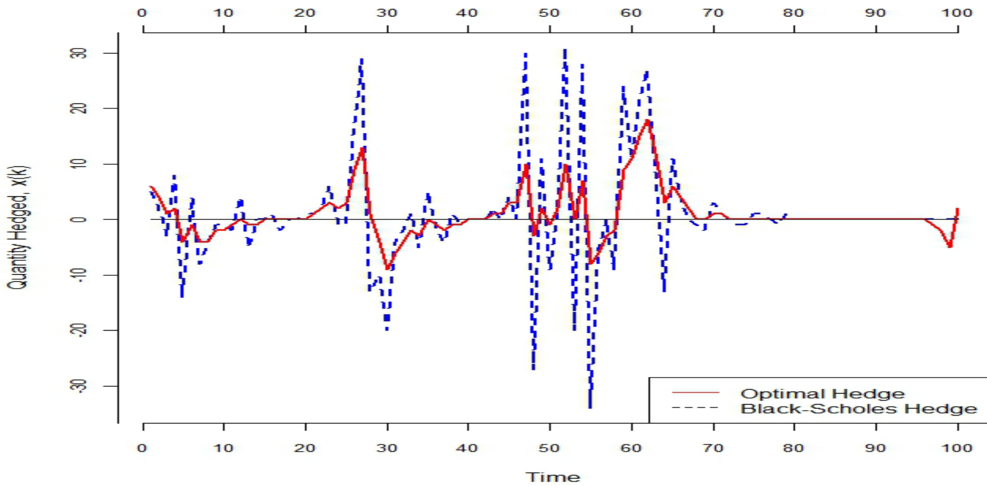

(a)

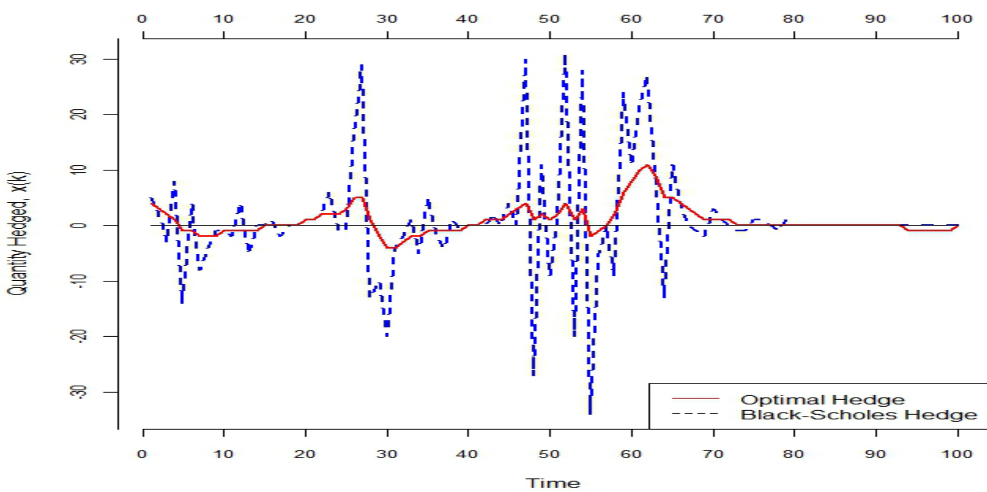

(b)

Figure 2: Black-Scholes hedging (which neglects the price impact) and the optimal strategy are plotted for two different liquidity parameters, q, which are related to the depth of the LOB. Figure 2(a) is obtained from $\frac{1}{q} = 0.001$, whereas in Figure 2(b), $\frac{1}{q} = 0.005$. $\frac{1}{q} = 0.001$ (0.005) indicates that a trade of 1000 (200) shares of stock would have a market impact of one unit of currency on the stock price.

increases slightly in the terminal period of the trading time interval. In Almgen and Li (2011), such shapes for the trading pattern are obtained by introducing a penalty that incorporates the agent's concerns about movements beyond the trading interval, for example, an overnight jump if the trading period is one day. In contrast, we obtain a moderate increase in the final trading activity without employing this type of device or even considering the issue of liquidating the terminal wealth. Obviously, the adoption of a final condition incorporating this concern would increase the intensity of the phenomenon. However, this effect is seemingly intrinsic to the optimal strategy under the price impact.

4 Hedging with price impact and resilience

Consider the hedging problem under the general setting described in Section 2. We use the notation from Section 2 and write the terminal value of the hedger's stock portfolio as $\sum_{k=0}^{N-1} r_k + R_0$ where r_k is given by Equation (9). In what follows, we suppose that $\frac{1}{q}$ is a small parameter in comparison to the intra-period variance of S^0_t. Thus, the concern regarding market price fluctuation is not negligible with respect to the effect of the price impact. In the opposite case, the hedging problem would become meaningless. (See Remark 4).

Because the time scales are usually short and because $\frac{1}{q}$ is a small parameter, the following approximation is reasonable:

$$g(t_{k+1}, S^0_{k+1} + \frac{V_{k+1}}{q}) - g(t_k, S^0_k + \frac{V_k}{q}) \approx \partial_S g(t_k, S^0_k + \frac{V_k}{q})[S^0_{k+1} - S^0_k + \frac{V_{k+1} - V_k}{q}] \tag{13}$$

Introducing the variable $Y_k = X_k - \partial_S g(t_k, S_k)$, which measures the distance of the stock holdings, X_k, from the traditional Δ-hedge, we can represent the terminal wealth as follows:

$$R_0 + \sum_{k=0}^{N-1}[(Y_k + x_k)(S^0_{k+1} - S^0_k) + \frac{1}{q}[(Y_k + x_k)(V_{k+1} - V_k) - \frac{x_k^2}{2}] \tag{14}$$

The hedger's problem is

$$\inf_{x_0,\ldots,x_{N-1}} E_0[\frac{1}{q}\sum_{k=0}^{N-1}((Y_k + x_k)(V_k - V_{k+1}) + \frac{x_k^2}{2})] + \frac{\gamma}{2} var_0[\sum_{k=0}^{N-1}(Y_k + x_k)(S^0_{k+1} - S^0_k)] \tag{15}$$

The dynamics of the state variables are as follows:

$$V_{k+1} = \beta_k(V_k + x_k) \tag{16}$$
$$Y_{k+1} = Y_k + x_k[1 - \frac{\Gamma}{q}\beta_k] + \frac{\Gamma}{q}[1 - \beta_k]V_k - \Gamma(S^0_{k+1} - S^0_k) \tag{17}$$

where $\beta_k = e^{-\rho(t_{k+1}-t_k)}$ and $\Gamma = \partial^2_S g(t_0, S_0)$.

The following proposition provides an explicit solution for the case in which $S^0_t = S_0 + \sigma W_t$.

Proposition 2. *The optimal hedging strategy under criterion (15) is*

$$x^*_k = \frac{L_k Y_k - H_k V_k}{D_k}$$

where the following notation is used:

$$D_k = \gamma\sigma^2(t_{k+1} - t_k) + \frac{1 - 2\beta_k}{q} + a_{k+1}\beta_k^2 + 2b_{k+1}\beta_k(1 - \frac{\Gamma}{q}\beta_k) + c_{k+1}(1 - \frac{\Gamma}{q}\beta_k)^2 \tag{18}$$
$$H_k = \frac{1 - \beta_k}{q} + a_{k+1}\beta_k^2 + 2b_{k+1}\beta_k(1 - \frac{\Gamma}{q}\beta_k) + c_{k+1}(1 - \frac{\Gamma}{q}\beta_k)^2 \tag{19}$$
$$K_k = \frac{1 - \beta_k}{q} + a_{k+1}\beta_k^2 + b_{k+1}\beta_k(1 - \frac{2\Gamma}{q}\beta_k) - c_{k+1}\frac{\Gamma}{q}\beta_k(1 - \frac{\Gamma}{q}\beta_k)^2 \tag{20}$$
$$L_k = \frac{\beta_k}{q} - \gamma\sigma^2(t_{k+1} - t_k) - c_{k+1}(1 - \frac{\Gamma}{q}\beta_k) - b_{k+1}\beta_k \tag{21}$$

and

$$
\begin{aligned}
a_k &= (\frac{H_k}{D_k})^2[\gamma\sigma^2(t_{k+1}-t_k)+\frac{1-2\beta_k}{q}+c_{k+1}(1-\frac{\Gamma}{q}\beta_k)] \qquad (22)\\
&+\frac{H_k}{D_k}[\frac{2\beta_k-2}{q}+\frac{2c_{k+1}\Gamma(1-\beta_k)}{q}(\frac{\Gamma\beta_k}{q}-1)+\frac{b_{k+1}\beta_k(\gamma\sigma^2(t_{k+1}-t_k)-\frac{\beta_k}{q})(\frac{\Gamma\beta_k}{q}-1)}{D_k}]\\
&+c_{k+1}(\frac{\Gamma}{q})^2(1-\beta_k)^2+\frac{2b_{k+1}\beta_k}{D_k}[\gamma\sigma^2(t_{k+1}-t_k)-\frac{\beta_k}{q}]\frac{\Gamma(1-\beta_k)}{q}\\
&+\frac{a_{k+1}\beta_k^2}{D_k}[\gamma\sigma^2(t_{k+1}-t_k)-\frac{\beta_k}{q}]^2, \quad a_N=0\\
b_k &= \frac{H_k}{D_k^2}[\frac{1-2\beta_k}{q}L_k-\gamma\sigma^2(t_{k+1}-t_k)K_k]+\frac{1-\beta_k}{q}+\frac{L_k(\beta_k-1)}{qD_k} \qquad (23)\\
&+\frac{H_k}{D_k}[\frac{\beta_k}{q}+c_{k+1}(\frac{\Gamma\beta_k}{q}-1)(1+\frac{L_k}{D_k})]\\
&+c_{k+1}\frac{\Gamma(1-\beta_k)}{q}(1+\frac{L_k}{D_k})-a_{k+1}\beta_k^2\frac{L_k[\gamma\sigma^2(t_{k+1}-t_k)-\frac{\beta_k}{q}]}{D_k^2}\\
&+\frac{b_{k+1}\beta_k}{D_k}[(1+(\frac{\Gamma\beta_k}{q}-1)\frac{L_k}{D_k})(\gamma\sigma^2(t_{k+1}-t_k)-\frac{\beta_k}{q})-(\frac{\Gamma(1-\beta_k)}{q}+(\frac{\Gamma\beta_k}{q}-1)\frac{H_k}{D_k})L_k], \quad b_N=0\\
c_k &= \frac{\gamma\sigma^2(t_{k+1}-t_k)K_k^2+\frac{1-2\beta_k}{q}L_k^2+a_{k+1}\beta_k^2L_k^2-2b_{k+1}\beta_kL_k^2(\frac{\Gamma\beta_k}{q}-1)}{D_k^2} \qquad (24)\\
&+\frac{L_k}{D_k}[\frac{2\beta_k}{q}-2b_{k+1}\beta_k]+c_{k+1}[1+(\frac{\Gamma\beta_k}{q}-1)\frac{L_k}{D_k}]^2, \quad c_N=0
\end{aligned}
$$

Proof: See Appendix

Remark 3. *The case of Section 3, Example 4, is obtained as a limit case for Proposition 2 in which* $\rho\to\infty$*.*

Remark 4. *The optimal solution is obtained using some restrictions on the relevant parameters. For example, in the case involving equally spaced time intervals,* $t_k-t_{k-1}=\Delta t$*, we need to assume*

$$
\gamma\sigma^2\Delta t>\frac{2e^{-\rho\Delta t}-1}{q}.
$$

In this scenario, the need to hedge against market price fluctuation is not irrelevant given the incurred illiquidity effects. If this is not the case, the trade-off between hedging and manipulation makes the second strategy more favorable, and an infinite demand for stocks may be optimal. Note *that when there is merely a temporary impact, no sufficient condition is required because*

$$\gamma\sigma^2\Delta t > -\frac{1}{q}$$

is trivially satisfied.

The case of a fully permanent impact is obtained when one takes $\rho = 0$, that is, $\beta_k = 1$ in Proposition 2. In this case, one obtains $V_k = X_k$; thus, there is a single state variable, and the expression for the optimal hedges is less involved. The optimal strategy is given as follows:

$$x_k^* = (-1 + \frac{c_{k+1}(\frac{\Gamma}{q} - 1)\frac{\Gamma}{q}}{\gamma\sigma^2(t_{k+1} - t_k) - \frac{1}{q} + c_{k+1}(1 - \frac{\Gamma}{q})^2})Y_k \tag{25}$$

where

$$c_k = \frac{1}{q} + c_{k+1}(\frac{\Gamma}{q})^2 + 2\frac{\Gamma}{q}(1 - \frac{\Gamma}{q})\frac{c_{k+1}^2}{\gamma\sigma^2(t_{k+1} - t_k) - \frac{1}{q} + c_{k+1}(1 - \frac{\Gamma}{q})^2}, \quad c_N = 0. \tag{26}$$

The numerical simulation for the case of a permanent impact shows that the optimal hedge is close to the corresponding strategy without a price impact and that the hedger slightly overhedges the option at the beginning. The reason is that although the hedger can exploit the permanent price impact by benefiting from his stock holdings, the hedging activity also negatively affects the option payoff. Thus, the illiquidity costs that arise from the hedging activity are offset by the price manipulation strategy. As a result, when the price impact does not dissipate immediately, the optimal strategy is closer to that of a perfect hedge than in the case of Section 3. In Figure 3, a case involving a permanent impact is compared to one involving an exclusively temporary impact when the stock price follows a Bachelier process and the hedger has a short position on a call option. Figure 3 shows that when the price impact does not vanish immediately the hedger trades more aggressively than in the case of a merely temporary impact. The main reason is that, in the case of a persistent price impact, the loss in portfolio due to the execution costs is partly compensated by the systematic increase in value of the stock holding due to the hedging activity. As a result, the persistence of price impact allows for a more aggressive strategy which is also closer to the classical (liquid) hedge than the optimal strategy obtained in Section 3.

In the extreme case of infinite resilience the hedger only faces execution costs and cannot benefit from stock price manipulation. However, even in the case of a permanent price impact, the hedger will not switch to a pure price manipulation strategy as far as his trading activity moves prices in a direction which is unfavorable for his derivative position. This latter effect is weak when the stock volatility is low relatively to the parameters modeling the price impact.

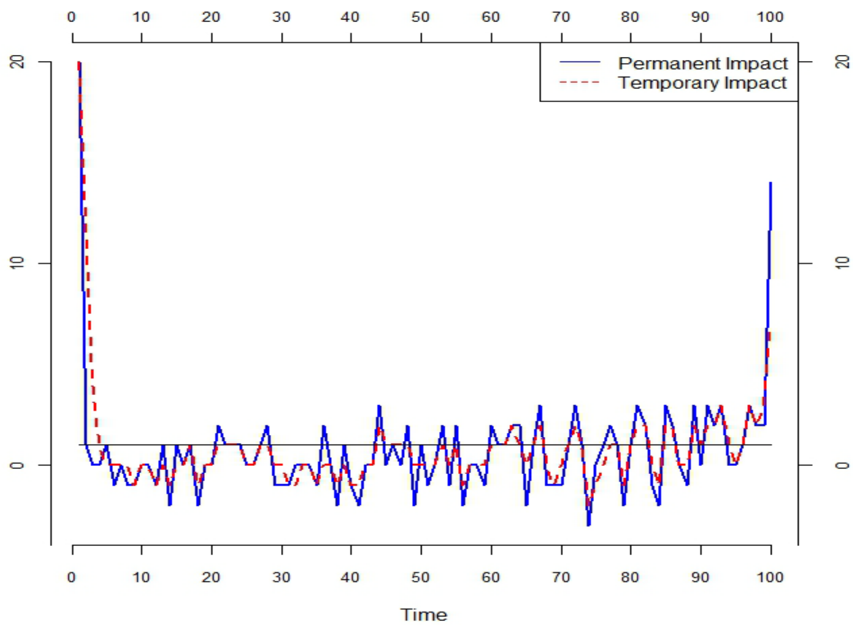

Figure 3: The case involving a permanent impact is compared to that of an exclusively temporary impact when the stock price follows a Bachelier pattern and the hedger has a short position on a call option.

However, our model does not take into account the issue of how to liquidate the terminal portfolio; that is, it is implicitly assumed that the terminal portfolio can be sold at the spot price with negligible liquidity costs. If this is not the case, the terminal conditions of our optimization problem must be changed, and the solution is modified accordingly. This issue is tackled in Section 5, where we discuss the robustness of the model to the introduction of liquidation costs in the terminal stage. More importantly, a thorough analysis of non-transient price impacts would require a market model in which the affected price would be a long-term effect of market clearing. Finally, a realistic model should allow the trader to cancel and possibly re-enter an order or to route the order to another book. However, all of these features would make the problem mathematically intractable.

5 Further remarks and possible extensions

5.1. Volatility

The solution to a hedging problem in an illiquid environment depends crucially on the interplay between the volatility of the underlying risky asset and the parameters modeling the liquidity of the limit order book. The liquidity of the limit order book is modeled in terms of the LOB's depth, q, and the speed at which the price impact decays with time as the result of new orders coming in the order book, that is, the resilience (ρ) of the LOB. The sensitivity of the hedge policy to the volatility of the underlying asset has the same direction as in the liquid case, i.e. the absolute value of the hedge is increasing in the volatility level. Numerical simulations (not included for brevity) confirm that the larger the volatility of the underlying asset, the larger is the number of stocks which are required to hedge.

The model assumes that the volume of the order book consumed by the large trader recovers exponentially or, equivalently for a block-shaped LOB, that the extra spread decays exponentially, where the rate of exponential decay is ρ. In comparison with the infinitely-liquid case ($q = \infty$) it is generally optimal to adopt a less aggressive hedging strategy in order to reduce the cost induced by the price impact. Furthermore, the deeper the LOB, the closer is the optimal strategy to the classical hedge in the liquid case, and vice versa, because the added costs coming from the finite liquidity are proportional to the level of illiquidity described by $\frac{1}{q}$, as expression (12) reveals.

However, the resilience of the LOB plays a crucial role in dictating the optimal behavior. At one extreme, when ρ is very large ($\rho \to \infty$) the hedger's concern is to minimize the illiquidity costs and he/she just needs to trade the more patiently the less deep is the LOB. The hedger always pursues this strategy, no matter what the level of market riskiness of the underlying asset is. At the other extreme, when $\rho = 0$ the impact of trading is at its fullest and the hedger's trades have a permanent impact on the quoted price which is proportional to the level of illiquidity. When the resilience is low the interplay between hedging and exploiting the price impact becomes more complicated and can switch to the latter pursuit, i.e. increasing the value of the stock holding by exploiting the persistence of the price impact, in case of low volatility of the underlying asset, as Remark 4 reveals.

5.2. Types of price decay

It is also interesting to note that although we model the resilience throughout a constant parameter, ρ, the solution to the hedging problem depends on $\beta_k = e^{-\rho(t_{k+1}-t_k)}$ as a whole, and not on

ρ explicitly. On one hand, this means that the timing of the hedging strategy is an important determinant of the optimal solution. For example, if the choice is a frequent rebalancing of the hedging portfolio, that is $\Delta t \simeq t_{k+1} - t_t$ is small, it is not advisable to pursue the hedging strategy with $\Delta t < 1/(q\sigma^2 + 2\rho)$ (under the setting of Proposition 2 and $\gamma = 1$), because the condition in Remark 4 yields $(q\sigma^2 + 2\rho)\Delta t > 1 + \rho^2(\Delta t)^2 + o((\Delta t)^2)$ which is incompatible with the previous inequality.

Alternatively, one can allow for a variable resilience speed in the model just interpreting β_k as $e^{-\int_{t_k}^{t_{k+1}} \rho(t)dt}$. The results in Section 4 are robust to extensions adopting a more general time decay function, say d. Throughout the paper d is an exponential function, while one might consider a different decay factor and the notation β_k could be safely interpreted as $d(t_{k+1} - t_k)$, that is, the main result holds under any decay factor. An increasingly expanding empirical literature promotes the use of alternative decay factors for the temporary impact, for example a power-law decay in combination with a non linear price impact. We point out that however the setting we adopted here - a linear price impact and an exponential time decay - is motivated by the findings in Gatheral (2010) showing that the model is compatible with the no-dynamic arbitrage principle.

5.3. Degree of risk aversion

Our model allows for a whole range of weights that different agents may attach to the risk components of their target function (12) or in (15). For example, in (12):

$$\inf_{x_0,\ldots,x_{N-1}} E_0[\frac{1}{q}\sum_{k=0}^{N-1}\frac{x_k^2}{2}] + \frac{\gamma}{2}var_0[\sum_{k=0}^{N-1} x_k(S_N^0 - S_k^0) - G_N]$$

the parameter γ captures the hedger's attitude towards the risk of market price fluctuations. If $\gamma = 1$ represents an equal concern about the cost of hedging and the risk to be hedged, then $\gamma > 1$ ($\gamma < 1$) can be interpreted as a major (minor, respectively) risk aversion to market price fluctuations affecting the option position. It is easy to conclude that the larger γ, the more aggressive the hedging strategy can be, that is, when the agent is strongly averse to the risk of being mis-hedged his/her strategy follows the stock price closely, while, in the opposite case, his/her main concern is to avoid high liquidity costs. A similar remark applies in the case of a persistent price impact. However, in this case, the optimal strategy is pursued insofar as the coefficient γ is sufficiently large. For example, the condition in Remark 4 reads: $\gamma\sigma^2\Delta t > \frac{2e^{-\rho\Delta t} - 1}{q}$. When γ is small, the agent is better off if he/she switches from a hedging strategy to price manipulation exploiting the price impact effect.

5.4. Costly liquidation of the portfolio terminal position

All the discussion above implicitly assumes that the terminal portfolio can be liquidated with negligible costs. Here we show how the argument can be generalized to incorporate the cost of liquidating the terminal position. To fix ideas, let us focus on the extreme case of a full impact and let us assume a quadratic liquidation cost of the form ℓY^2, $\ell \geq 0$, which is consistent with the shape of the liquidity costs that are modelled throughout the paper. Under this scenario the hedger's problem takes the form:

$$\inf_{x_0,\ldots,x_{N-1}} E_0[\frac{1}{q}\sum_{k=0}^{N-1}((Y_k+x_k)x_k+\frac{x_k^2}{2})+\ell Y_{N-1}^2]+\frac{\gamma}{2}var_0[\sum_{k=0}^{N-1}(Y_k+x_k)(S_{k+1}^0-S_k^0)]$$

where the dynamics of the state variable is $Y_{k+1} = Y_k + x_k[1-\frac{\Gamma}{q}] - \Gamma(S_{k+1}^0 - S_k^0)$. The optimal hedge is given by the formulas (25), (26) with $c_N = 2\ell$. Moreover, the optimal solution is obtained under a weaker condition. For example, in the case of equally spaced time intervals Δt and a not too small parameter q, whenever $\gamma\sigma^2\Delta t + 2\ell > \frac{1}{q}$.

The conclusion is that, if the cost of liquidating the terminal portfolio is taken into account in the form of a quadratic function, the agent has a minor incentive to stray far from his/her hedging activity and to engage in a manipulating strategy. In particular, the smaller is that incentive, the higher is the terminal liquidation cost.

6 Conclusions

Our results illustrate the hedging costs accrued in the presence of a price impact through a LOB. In the presence of this type of price friction, the optimal hedging strategy involves trading less aggressively than would be dictated by the Black-Scholes model to minimize the costs associated with such price impacts. Obviously, the magnitude of the latter is contingent on the shape of the LOB and on how much liquidity is available relative to the desired trade size. If the LOB is shallow, the influence of the price impact is magnified, thus making perfect hedging expensive relative to a patient strategy of trading according to the resilience speed of the book.

The illiquidity costs associated with hedging are more pronounced given a pure temporary impact than under a pure permanent price impact. This is because, although the hedger can exploit the permanent price impact, on one hand, the hedging activity negatively affects the option payoff, on the other hand, the illiquidity costs that arise from the hedging activity are offset by the price manipulation strategy. Under both temporary and permanent price impact, the cost of a terminal mishedge may be less than the eliminated price impact costs. Our work also contributes towards the understanding of the issue of risk hedging versus price manipulating behavior.

References

[1] R. Almgren, T. M. Li (2010), A fully-dynamic closed-form solution for $\Delta-$hedging with market impact, Operations Research (submitted).

[2] A. Alfonsi, A. Fruth, A. Schied (2010), Optimal execution strategies in limit order books with general shape functions, Quantitative Finance, 10, 2, 143-157.

[3] P. Bank, D. Baum (2004), Hedging and portfolio optimization in financial markets with a large trader, Mathematical Finance, 14, 1-18.

[4] J. P. Bouchaud, M. Potters (2003), Theory of Financial Risk and Derivative Pricing: from Statistical Physics to Risk Management, CUP.

[5] S- I. Boyarchenko, S. Z. Levendorskiĭ (2002), Non-Gaussian Merton-Black-Scholes Theory, World Scientific.

[6] U. Cetin, R. A. Jarrow, P. Protter (2004), Liquity risk and arbitrage pricing theory, Finance and Stochastics, 8, 311-341.

[7] U. Cetin, R. Jarrow, P. Protter and M. Warachka (2006), Pricing Options in an Extended Black Scholes Economy with Illiquidity: Theory and Empirical Evidence, Review of Financial Studies, 19, 493-529

[8] U. Cetin, H. M. Soner, N. Touzi (2010), Options hedging for small investors under liquidity costs, Finance and Stochastics, 14, 317-341.

[9] H. Föllmer, M. Schweizer (1989), Hedging by sequential regression: An introduction to the mathematics of option trading, ASTIN Bulletin, 19, 29-42.

[10] R. Frey (1998), Perfect option hedging for a large trader, Finance and Stochastics, 2, 115-141.

[11] R. Frey, A. Stremme (1997), Market volatility and feedback effects from dynamic hedging, Mathematical Finance, 7, 351-374.

[12] P. A. Forsyth, J. S. Kennedy, S. T. Tse, H. Windcliff (2012), Optimal trade execution: a mean-quadratic-variation approach, Journal of Economic Dynamics and Control, forthcoming.

[13] J. Gatheral (2010), No-dynamic-arbitrage and market impact, Quantitative Finance, 10, 749-759.

[14] C. Geiss, S. Geiss (2006), On an approximation problem for stochastic integrals where random time nets do not help, Stochastic Processes and Applications, 116, 407-422.

[15] H. Geman (2002), Pure jump Lévy processes for asset price modeling, Journal of Banking and Finance, 26, 1297-1316.

[16] F. Hubalek, J. Kallsen, L. Krawczyk (2006), Variance-optimal hedging for processes with stationary independent increments, Annals of Applied Probability, 16, 853-885.

[17] R. A. Jarrow, P. Protter (2007), Liquidity risk and option pricing theory, in Handbooks in Operations Research and Management Science: Financial Engineering (eds. T. R. Birge and V. Linetski), Elsevier.

[18] R. Kissel, M. Glantz, R. Malamut (2004), A practical framework for estimating transaction costs and developing optimal trading strategies to achieve best execution, Finance Research Letters, 1, 35-46.

[19] K. Matsumoto (2009), Mean-variance hedging with uncertain trade execution, Applied Mathematical Finance, 16, 3, 219-252.

[20] H. Kraft, C. Kühn (2011), Large traders and illiquid options: Hedging versus manipulation (http://ssrn.com/abstract=01482469).

[21] A. Obizhaeva, J. Wang (2005), Optimal trading strategy and supply/demand dynamics, Technical report MIT.

[22] E. Platen, M. Schweizer (1998), On feedback effects from hedging derivatives, Mathematical Finance, 8, 67-84.

[23] A. Roch, H. M. Soner (2011), Resilient price impact of trading and the cost of illiquidity (http://ssrn.com/abstract=1923840).

[24] L. C. G. Rogers, S. Singh (2010), The cost of illiquidity and its effects on hedging, Mathematical Finance, 20, 4, 597-615.

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

[26] N. Taleb (1997), Dynamic Hedging, J. Wiley & Sons, N.Y.

Appendix

Proof of Proposition 1: By the law of total variance,

$$var_t(X) = E_t[var_{t+\Delta t}(X)] + var_t[E_{t+\Delta t}(X)]$$

for $\Delta t > 0$. Arguing recursively, one obtains

$$var_0[\sum_{j=0}^{N-1} x_j(S_N - S_j) - G_N] = \sum_{k=0}^{N-1} var_k[X_{k+1}S_{k+1} - G_{k+1}].$$

The value function for Equation (12) then takes the form:

$$J_0(X) \ = \inf_{x_0,\ldots,x_{N-1}} E_0[\sum_{k=0}^{N-1} \frac{1}{2q}x_k^2 + \frac{\gamma}{2}var_k[(X + x_k)S_{k+1} - G_{k+1}]].$$

By standard argument in stochastic dynamic programming, J_0 is the last step of:

$$\begin{aligned} J_k(X) \ &= \ \inf_{x_k} \ E_0[\frac{1}{2q}x_k^2 + \frac{\gamma}{2}var_k[(X + x_k)S_{k+1} - G_{k+1}] \\ &\quad +E_k[J_{k+1}(X + x_k)]], \quad k = 0, \ldots, N-1, \quad J_N(X) = 0. \end{aligned} \tag{27}$$

Solving the minimization problem in Equation (27), we obtain the optimal hedge x_k as the solution of:

$$x_k(\frac{1}{q} + \gamma(var_k(S_{k+1})) + Xvar_k(S_{k+1}) - cov_k(S_{k+1}, G_{k+1})) + E_k[J'_{k+1}(X + x_k)] = 0.$$

Making the Ansatz:

$$J_k(X) = a_k + b_kX + \frac{\widetilde{c}_k}{2}X^2$$

one obtains the expressions for a_k, b_k and $\widetilde{c}_k$ by comparing the two sides of Equation (27) and separating both sides by powers of X. The result follows if we denote $c_k = \widetilde{c}_k + \frac{1}{q}$. □

Proof of Proposition 2: The value function for Equation (15) takes the form:

$$J_0(Y,V) \;=\inf_{x_0,\ldots,x_{N-1}} \;E_0[\sum_{k=0}^{N-1}\frac{1}{q}[(Y+x_k)((1-\beta_k)V-\beta_k x_k)+\frac{x_k^2}{2}]+\frac{\gamma}{2}var_k[(Y+x_k)S^0_{k+1}].$$

By standard argument in stochastic dynamic programming, J_0 is the last step of:

$$\begin{aligned} J_k(Y,V) \;&=\; \inf_{x_k} \;\; E_0[\frac{-1}{q}[Y(\beta_k x_k+(\beta_k-1)V)+(\beta_k-1)Vx_k \\ &\quad +x_k^2(\beta_k-\frac{1}{2})]+\frac{\gamma}{2}(Y+x_k)^2 var_k[S^0_{k+1}] \\ &\quad +E_k[J_{k+1}(Y+x_k(1-\frac{\Gamma}{q}\beta_k)+\frac{\Gamma}{q}(1-\beta_k)V \\ &\quad -\Gamma(S^0_{k+1}-S^0_k),(\beta_k-1)V+\beta_k x_k)]], \quad k=0,\ldots,N-1, \quad J_N(Y,V)=0. \end{aligned} \tag{28}$$

Solving the minimization problem in Equation (28), we obtain the optimal hedge x_k as the solution of:

$$\begin{aligned} &(Y+x_k)\gamma var_k(S_{k+1})-\frac{1}{q}[Y\beta_k+(\beta_k-1)V+x_k(2\beta_k-1)] \\ &+E_k[J'_{k+1}(Y+x_k(1-\frac{\Gamma}{q}\beta_k)+\frac{\Gamma}{q}(1-\beta_k)V-\Gamma(S^0_{k+1}-S^0_k),(\beta_k-1)V+\beta_k x_k)]=0. \end{aligned} \tag{29}$$

Making the Ansatz:

$$J_k(Y,V) \;=\frac{a_k}{2}V^2+b_kVY+\frac{c_k}{2}Y^2+d_k$$

the expressions for a_k, b_k and c_k are obtained by comparing the two sides of Equation (29). □

Proof of Example 4: If S_t^0 follows an arithmetic Brownian motion with volatility σ then Proposition 1 yields:

$$x_k^* = \frac{\gamma\sigma^2(t_{k+1}-t_k)\partial_S g_k + \frac{1}{q}X_k - \overline{b}_{k+1}}{\gamma\sigma^2(t_{k+1}-t_k)+c_{k+1}} - X_k$$

where $\overline{b}_{k+1}$ and c_{k+1} are defined in Proposition 1 and $\partial_S g_k = \partial_S g(t_k, S_k)$. One can rewrite x_k^* as:

$$x_k^* = \frac{[\gamma\sigma^2(t_{k+1}-t_k)+c_{k+1}-\frac{1}{q}](\partial_S g_k - X_k) - \overline{b}_{k+1} + [\frac{1}{q}-c_{k+1}]\partial_S g_k}{\gamma\sigma^2(t_{k+1}-t_k)+c_{k+1}}.$$

Then we just need to show that $E_k[b_{k+1} + (c_{k+1} - \frac{1}{q})\partial_S g_k] = 0$ for any $k = 0, \ldots, N-1$ and with small trading intervals. First we note that $b_N = 0$ and $c_N = \frac{1}{q}$. Arguing recursively and denoting $t_{k+1} - t_k$ by Δt_k one gets:

$$b_k = \frac{\overline{b}_{k+1} - \gamma\sigma^2\Delta t_k \partial_S g_k}{q[\gamma\sigma^2\Delta t_k + c_{k+1}]} = -\partial_S g_k \frac{[c_{k+1}+\gamma\sigma^2\Delta t_k - \frac{1}{q}}{q[\gamma\sigma^2\Delta t_k + c_{k+1}]}$$

which is equal to $(\frac{1}{q} - c_k)\partial_S g_k$ by definition of c_k. Thus $E_{k-1}[b_k + (c_k - \frac{1}{q})\partial_S g_{k-1}] = 0$ follows. $\square$