

A Fully-Dynamic Closed-Form Solution for Δ -Hedging with Market Impact

Tianhui Michael Li*

Princeton Bendheim Center for Finance and Operations Research and Financial Engineering

Robert Almgren

Quantitative Brokers, LLC and NYU Courant Institute

We present a closed-form Δ -hedging result for a large investor whose trades generate adverse market impact. Unlike in the complete-market case, the agent no longer finds it tenable to be perfectly hedged or even within a fixed distance away from being hedged. Instead, he may find himself arbitrarily mishedged and optimally trades towards the classical Black-Scholes Δ , with trading intensity proportional to the degree of mishedge and inversely proportional to illiquidity. When combined with a recent result of Garleanu and Pedersen (2009), this suggests that Δ -hedging can be thought of as a Merton problem where the Merton-optimal portfolio is the Black-Scholes Δ -hedge. Both the discrete-time and continuous-time problems are solved. We discuss a number of applications of our result, including the equilibrium implications of our model on intraday trading patterns and stock pinning at options' expiry. Finally, numerical simulations on TAQ data based on intraday hedging of call options suggest that this strategy is able to significantly minimize market impact cost without incurring a significant increase in mishedging error (with respect to a Black-Scholes Δ -hedging strategy).

1. Introduction

The solution to hedging an option in a complete market is well known (Black and Scholes 1973, Merton 1973). However, the ability to perfectly replicate in a frictionless complete market is far from reality. Practitioners who trade large positions are familiar with the concept of market impact when trading in real markets with finite liquidity. For instance, a buyer who crosses the spread and lifts orders from the other side eats up liquidity in the order book, pushing his execution price ever higher away from him. Furthermore, liquidity often dries up during financial crises such as the 1987 Crash, the Asian Financial Crisis, and the LTCM Collapse. The recent subprime debacle has only further heightened investor concerns about liquidity.

We are interested in the optimal hedging of an option by a large risk-averse investor in an illiquid market. There is a large literature on trading under transaction costs. Part of the literature is interested in super-replication (Cetin et al. 2010, Soner et al. 1995). Our paper relaxes this requirement by having a finite penalty for being mishedged. In this respect, our paper is more closely linked to those using a utility-based framework (Cvitanic and Wang 2001, Davis and Norman 1990, Janecek and Shreve 2004, Shreve and Soner 1994). However, results within this strand of the literature typically assume a fixed cost of trading (for example, a fixed brokerage fee) or a cost that is proportional to trade size (for example, crossing the bid-ask spread). Mathematically, these problems exhibit a scaling which allows the authors to greatly simplify their formulas. Our paper differs from these in that the source of incompleteness comes from a price impact model: the more aggressively one trades, the more impact is incurred proportionally. The cost is no longer linear in

*This research is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-0646086 and the Fannie and John Hertz Foundation Graduate Fellowship. The author would like to thank both organizations for their generous support.

the shares purchased and this non-scaling makes our problem significantly harder. Nonetheless, we are able to obtain a closed-form solution for trading.

We separate market impact in terms of temporary and permanent impact effects in the framework developed by Almgren and Chriss (2001). The temporary impact affects only the execution price \tilde{P}_t but has no effect on the “fair value” or fundamental price P_t . In contrast, the permanent impact directly affects the fair value of the security P_t while having no direct effect on the execution price \tilde{P}_t . Thus we can think of the temporary impact as connected to the liquidity cost faced by the agent while the permanent impact as linked to information transmitted to the market by the agent’s trades (see, for example, Back (1992), Kyle (1985)).

Consider the trading model

$$X_t = X_0 + \int_0^t \theta_u du$$

where X_t is the number of shares held by the agent and θ_t is the intensity of trading. The fundamental price is given by

$$P_t = P_0 + \nu(X_t - X_0) + \sigma W_t \quad (1)$$

where $\nu > 0$ is the coefficient of permanent impact, $\sigma > 0$ is the absolute volatility of the fair value, and W_t is a standard one-dimensional Brownian Motion. Using a linear Brownian Motion rather than a geometric one is appropriate over the short time horizons considered in the paper and leads to dramatic simplifications in our results.

To model temporary impact, we assume that the θ_t ’th share will cost a premium $\lambda\theta_t$ over the fair value

$$\tilde{P}_t(\theta_t) - P_t = \lambda\theta_t$$

where $\lambda > 0$ is the market impact parameter. Therefore, the total premium from temporary impact of buying θ_t shares is

$$\int_0^{\theta_t} [\tilde{P}_t(\xi) - P_t] d\xi = \lambda \int_0^{\theta_t} \xi d\xi = \lambda \frac{\theta_t^2}{2}.$$

We can think of temporary impact as coming from a limit-order book with constant depth $1/\lambda$ and instant resilience (see, for example, Alfonsi and Schied (2010), Predoiu et al. (2010)). In such a model, purchasing θ_t shares would consume all the shares priced from P_t to $P_t + \lambda\theta_t$ on the book, thus pushing the execution cost of the last share up by $\lambda\theta_t$. The limit orders that were eaten up are instantly replaced immediately after execution. Thus our coefficient $\lambda = 2\eta$ where η is the temporary market impact of Almgren and Chriss (2001).

Previous work using this model has involved trading in the presence of Dark Pools (Kratz and Schöneborn 2010), and competitive liquidation (Schoeneborn and Schied 2010). Rogers and Singh (2007) also examine the Δ hedging problem but for agents with different risk preferences and were only able to obtain analytic expressions asymptotically. Our work differs from theirs in that we derive closed-form expressions by using simpler risk preferences. The paper most closely related to ours is Garleanu and Pedersen (2009). They solve the infinite-horizon ‘Merton Problem’ under only temporary market-impact assumptions. As in our setup, they use a mean-variance objective rather than the traditional expected utility setup of the classical Merton Problem. They find that trading intensity at time t (Proposition 5) is given by

$$\theta_t = -\kappa h(X_t - \text{target}_t) \quad \kappa \propto 1/\sqrt{\lambda}$$

where X_t is the number of shares, κ is an urgency parameter with units of inverse time, $h > 0$ is a dimensionless constant of proportionality related to the convexity of the continuation value, and the “target portfolio” target_t is the frictionless Merton-optimal portfolio. That is, with market-impact costs, it is no longer optimal to hold the Merton-optimal portfolio but instead, the agent

‘chases’ it with his trading. The intensity of trading θ_t is proportional to the distance between the current holdings and the Merton-optimal portfolio ($X_t - \text{target}_t$) and is inversely proportional to the square-root of the illiquidity parameter λ .

We use an analogous finite-horizon setup on $[0, T]$ with only temporary impact and obtain that the trading intensity is

$$\theta_t = -\kappa h(\kappa(T-t))(X_t - \text{target}_t)$$

where the “target portfolio” target_t is now the frictionless Black-Scholes Δ -hedge and $h(\cdot)$ is a positive dimensionless function, which comes from the finite-horizon nature of the setup (compare to (8) with $\nu = 0$ or $K = 1$). Hence, as in Garleanu and Pedersen (2009), an agent facing market illiquidity no longer maintains the zero-liquidity-cost optimal portfolio target_t but instead trades towards it to correct this ‘misholding’. Furthermore, also as in Garleanu and Pedersen (2009), the agent’s trading intensity θ_t is proportional to the difference between his current holdings and the optimal no market-impact portfolio ($X_t - \text{target}_t$) and inversely related to the market-impact cost λ . The similarity suggests that we can think of Δ -hedging in an illiquid market as a Merton optimal investment problem where the Merton portfolio is the Black-Scholes hedge portfolio.

The rest of the paper is as follows. We motivate our assumptions and setup the problem in Section 2. The solution is presented in Section 3. In Section 4, we give a number of applications and equilibrium implications of our result. In Section 5, we give a discrete-time formulation of the problem which we show converges to the continuous-time solution. We explain how this is necessary for performing discretized numerical simulations on TAQ data. The discrete-time solution is used in the hedging simulation in Section 6. Finally, we conclude in Section 7.

2. Problem Setup

We think of hedging a European contingent claim over a finite-time horizon $[0, T]$. In this section, we first give a heuristic justification for our objective (5). The resulting optimal policy for this objective is then rigorously proved in Section 3.

For a hypothetical small trader whose execution has no price-impact and trades in a complete market, the value of the option’s price $g(t, P_t)$ is a function of the time t and the price P_t of the underlying asset,

$$g : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}.$$

Hence, the agent’s total portfolio value at time t is given by $X_t P_t + g(t, P_t)$. We assume the price at $t = T$ is given by $g_0 : \mathbb{R} \rightarrow \mathbb{R}$ and the value for $t \in [0, T)$ is prescribed by Feynman-Kac to be the solution of the PDE

$$\dot{g}(t, x) + \frac{\sigma^2}{2} g''(t, x) = 0 \quad \text{for } t, x \in [0, T] \times \mathbb{R} \quad \text{and} \quad g(T, \cdot) = g_0. \quad (2)$$

We can view this as the Black-Scholes option-pricing PDE in our setting so that g has the interpretation of the option price in the corresponding (no impact-cost) complete-market.

For a large trader who does face market-impact, the terminal wealth is the sum of the option’s value, the stock’s value, and the cost of acquiring the position,

$$R_T = R_0 + \underbrace{g(T, P_T)}_T \text{ option value} + \underbrace{X_T P_T}_T \text{ stock value} - \int_0^T \underbrace{\int_0^{\theta_t} \tilde{P}_t(\xi) d\xi}_{\text{acquisition cost}} dt.$$

Using integration by parts and the self-financing condition for the portfolio, we may rewrite the terminal wealth as

$$R_T = R_0 + \int_0^T [X_t + g'(t, P_t)] dP_t - \frac{\lambda}{2} \int_0^T \theta_t^2 dt$$

where we have used the fact that the option's price g satisfies Feynman-Kac (2). On the short-time scales under consideration, it is appropriate to make a Taylor-approximation of the option's value. We make the following approximation,

$$g'(t, P_t) \approx g'(t, P_0) + \Gamma(P_t - P_0), \quad \text{with } \Gamma \in \mathbb{R} \text{ constant.} \quad (3)$$

REMARK 1. The approximation says that for small intraperiod fluctuations of $P_t - P_0$, the option's Δ varies linearly with the stock price. In the derivatives jargon, this is equivalent to assuming the option has a constant Γ . While this assumption is not necessary for our model—it could be solved with only technical conditions on the integrability of g —the assumption considerably simplifies the problem by reducing state-variables (see below) and allows for a simple analytic closed-form solution.

It is now convenient to introduce the variable

$$Y_t = X_t + g'(t, P_t)$$

with the interpretation of the portfolio's **net Δ exposure**. This will allow us to reduce the state variables from the pair shares-owned and price (X_t, P_t) to a single variable representing the net Δ position, Y_t . Y_t also has the interpretation of **distance from being Δ -hedged** as the ideal net Δ -exposure is 0. Hence, we may decompose our wealth process very intuitively into three parts

$$R_T = R_0 + \int_0^T \underbrace{Y_u \sigma dW_u}_{\text{Intraperiod Fluctuation}} + \underbrace{Y_u \nu \theta_u du}_{\text{Permanent Impact}} - \underbrace{\frac{\lambda}{2} \theta_u^2 du}_{\text{Temporary Impact}}. \quad (4)$$

If we interpret the interval $[0, T]$ as the trading day, we have the wealth as the sum of the fluctuation during the trading day and the liquidity cost from permanent and temporary impacts.

We define our mean-variance objective to be $J(0, Y_0)$ where the function J has the interpretation of the continuation value and is given by

$$J(t, y) = \inf_{\theta_s: s \geq t} \mathbb{E} \left[\frac{\gamma \sigma_T^2}{2} Y_T^2 + \int_t^T \frac{\gamma \sigma^2 Y_u^2}{2} - Y_u \nu \theta_u + \frac{\lambda}{2} \theta_u^2 du \mid Y_t = y \right], \quad (5)$$

and the state variable has dynamics

$$dY_t = K \theta_t dt + \Gamma \sigma dW_t \quad K = 1 + \nu \Gamma. \quad (6)$$

The new constant σ_T^2 is a measure of the risk associated with being mishedged at the terminal time $t = T$ (see Remark 2). Hence, every share purchased by the agent increases his net Δ -position Y_t by $K = 1 + \nu \Gamma$: 1 for the increase in the stock position and $\nu \Gamma$ for the permanent impact's effect on the option's Δ . Observe that K is invisible without permanent impact ($\nu = 0$ implies $K = 1$).

Equation (5) has a natural interpretation of balancing the temporary and permanent liquidity costs with a penalty for variance of the intraperiod and terminal mishedging errors.

REMARK 2. We may interpret the penalty σ_T^2 as follows. Assume the price P_t is continuous and satisfies (1) for $t \in [0, T)$ but makes a mean-zero, normally-distributed jump at time T ,

$$\Delta P_T = P_T - P_{T-} \sim N(0, \sigma_T^2) \quad \sigma_T > 0.$$

A direct interpretation is to think of the trading period $[0, T]$ as one trading day. We can think of the jump ΔP_T as being an overnight jump, during which the agent is unable to trade (i.e.

$X_{T-} = X_T$). In this interpretation, it is clear that the agent wishes to be hedged at the terminal time $t = T$ to avoid exposure to overnight fluctuation. In addition to Approximation 3, we make the addition approximation

$$g(T, P_T) \approx g(T-, P_{T-}) + g'(T-, P_{T-})\Delta P_T.$$

That is, for small overnight jumps ΔP_T , the jump in the option value is linear. Then we may write the wealth with the new jump term,

$$R_T = R_0 + Y_T \Delta P_T + \int_0^{T-} Y_u \sigma dW_u + Y_u \nu \theta_u du - \frac{\lambda}{2} \theta_u^2 du. \quad (7)$$

The expected terminal utility under an exponential (constant absolute risk aversion) utility function with coefficient of risk aversion $\gamma > 0$ is

$$\sup_{\theta} E[-\exp(-\gamma R_T)] = \sup_{\theta} -E \left[\exp \left(\frac{\gamma^2 \sigma_T^2}{2} Y_T^2 + \int_0^T \frac{\gamma^2 \sigma^2 Y_u^2}{2} - \gamma Y_u \nu \theta_u + \frac{\gamma \lambda}{2} \theta_u^2 du \right) \right].$$

This new continuation value is very similar to J (see (5)) and is the expectation of the exponential of the sum of four terms which correspond to the four terms in the terminal wealth (7). It turns out that both objectives give similar results except that mean variance is better behaved for technical reasons related to integrability.

3. Continuous-Time Solution

The setup is that of a linear-quadratic optimization problem. The solution is then

$$\theta(t, y) = -\kappa h(\kappa K(T-t))y. \quad \kappa = \sigma \sqrt{\frac{\gamma}{\lambda}} \quad (8)$$

where κ has dimension of inverse time. In Almgren and Chriss (2001) κ is an ‘urgency parameter’ which dictates the speed of liquidation: the higher κ the faster the initial liquidation. In Garleanu and Pedersen (2009), this term gives the relative intensity of trading¹, that is, the higher κ , the faster the agent trades towards the Merton-optimal portfolio.

The agent trades towards minimizing his hedging penalty but is prevented from holding the exact Black-Scholes Δ -hedge by the convex temporary impact or liquidity cost. Hence, trading intensity θ_t is proportional to the degree of mishedge Y_t and the urgency parameter κ . There is a greater penalty to being mishedged with higher underlying volatility σ and risk aversion γ so these parameters increase urgency. Similarly, a more illiquid market (higher λ) makes trading more costly, which decreases trading intensity.

The solution falls into three cases depending on whether $h(\cdot)$, the trading intensity proportion, is decreasing, flat, or increasing in time, i.e. whether the agent trades less intensely, is flat, or more intensely (respectively) towards the end of the trading period. This, in turn, depends on the value of a constant d which gives the relative size of ΔP_T and dP_t ,

$$h(\tau) = \begin{cases} \tanh(\tau + \operatorname{arctanh}(d)) & d < 1 \\ 1 & d = 1 \\ \coth(\tau + \operatorname{arccoth}(d)) & d > 1 \end{cases} \quad d = \frac{\gamma K \sigma_T^2 - \nu}{\lambda \kappa}. \quad (9)$$

Intuitively, intraperiod trading can be thought of as primarily hedging out the dP_t fluctuations while trading near the end of the period is primarily for hedging the terminal jump ΔP_T . When

¹ The a/λ in Proposition 5 of Garleanu and Pedersen (2009) is equivalent to κh in our setup

$d < 1$, the majority of the fluctuations occur during the day, and the solution is marked by $t \mapsto h(\kappa K(T-t))$ decreasing as t approaches T . That is, the agent is more concerned about hedging intraperiod fluctuations dP_t and relaxes hedging intensity to reduce hedging cost as the end of the period approaches. The second case is when $d = 1$ and hedging risk for the intraperiod and end-of-day are weighted equally. Then the agent's trading intensity A is constant in time, The final case is when $d > 1$ and the terminal jump is large compared to the daily fluctuations. So the agent trades more intensely with time and $t \mapsto h(\kappa K(T-t))$ increases towards d as t approaches T ,

If we consider the case when $[0, T]$ is the trading day, the last case (when $d > 1$) is the most realistic: options market-makers typically increase their hedging towards the close of trading to minimize their overnight exposure. We are now in a position to state the theorem.

THEOREM 1. *Assume that we restrict our admissible θ so that*

$$\mathbb{E} \int_0^T \theta_s ds < \infty.$$

The optimal trading intensity $\theta_t = \theta(t, Y_t)$ for the objective (5) is given by (8) where the non-negative trading intensity proportion h is given by (9). Under the optimal trading strategy, $Y_T \neq 0$ a.s. That is, because of the market-impact costs, the position is not perfectly Δ -hedged, even at the terminal time.

Proof Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the complete probability space with \mathcal{F}_t being the filtration generated by W_t . The principle of stochastic optimal control tells us that

$$M_t := \int_0^t \frac{\sigma^2 \gamma}{2} Y_u^2 du - Y_u \nu \theta_u + \frac{\lambda}{2} \theta_u^2 du + J(t, Y_t)$$

is a submartingale for all θ_t and a martingale under the optimal control. Therefore, the proof follows if we can show that M_t is a true martingale. The HJB equation can be written as

$$\begin{aligned} 0 &= \frac{\gamma \sigma^2}{2} y^2 + j + \frac{(\Gamma \sigma)^2}{2} J'' - \frac{1}{2\lambda} (y\nu - KJ')^2 \\ \theta(t, y) &= \frac{1}{\lambda} (y\nu - KJ'). \end{aligned} \tag{10}$$

Making the Ansatz

$$J(t, y) = A_2(T-t) \frac{y^2}{2} + A_0(T-t) \tag{11}$$

and separating by powers of y , we have that our HJB PDE becomes the ODE pair

$$\begin{aligned} \dot{A}_2 &= \sigma^2 \gamma - \frac{1}{\lambda} (KA_2 - \nu)^2 & A_2(0) &= \gamma \sigma_T^2 \\ \dot{A}_0 &= \frac{(\Gamma \sigma)^2}{2} A_2 & A_0(0) &= 0. \end{aligned}$$

When $0 < d < 1$, the solution is given by

$$\begin{aligned} A_2(T-t) &= \frac{1}{K} \left[\nu + \lambda \kappa \tanh(\kappa K(T-t) + \operatorname{arctanh}(d)) \right] \\ A_0(T-t) &= \frac{(\Gamma \sigma)^2}{2K} \left\{ \nu(T-t) + \frac{\lambda}{K} \left[\log \circ \cosh(\kappa K(T-t) + \operatorname{arctanh}(d)) + \frac{1}{2} \log(1-d^2) \right] \right\} \end{aligned} \tag{12}$$

when $d = 1$,

$$\begin{aligned} A_2(T-t) &= \frac{\nu + \lambda \kappa d}{K} = \sigma_T^2 \gamma \\ A_0(T-t) &= \frac{\Gamma^2 \sigma^2}{2} \sigma_T^2 \gamma (T-t) \end{aligned} \tag{13}$$

and when $d > 1$,

$$\begin{aligned} A_2(T-t) &= \frac{1}{K} [\nu + \lambda \kappa \coth(\kappa K(T-t) + \operatorname{arccoth}(d))] \\ A_0(T-t) &= \frac{(\Gamma\sigma)^2}{2K} \left\{ \nu(T-t) + \frac{\lambda}{K} [\log \circ \sinh(\kappa K(T-t) + \operatorname{arccoth}(d)) + \frac{1}{2} \log(d^2 - 1)] \right\}. \end{aligned} \quad (14)$$

Thus M_t is a local martingale. To show that M_t is a true-martingale, it suffices to show that the process $J(\cdot, Y)$ has finite expected quadratic variation. Observe that the function $A_2(\cdot)$ is uniformly bounded by some $C > 0$ independent of t and $\omega \in \Omega$ so that for some other constant $C' > 0$,

$$\mathbb{E} \int_0^T |J'(t, Y_t)|^2 dt \leq C' \mathbb{E} \int_0^T |Y_t|^2 dt.$$

Using Jensen's inequality and the fact that for some constant $C'' > 0$, $|\theta_t| \leq C''|Y_t|$ for all $t \in [0, T]$, we obtain

$$|Y_t|^2 \leq 3(|P_0|^2 + K^2|X_t|^2 + \sigma^2|W_t|^2) \leq 3 \left(|P_0|^2 + K^2 C'' \int_0^t |Y_u|^2 du + \sigma^2|W_t|^2 \right). \quad (15)$$

Hence, $\mathbb{E}|Y_t|^2$ is bounded for $t \in [0, T]$ by a standard application of Gronwall's Lemma and this bound can be made to depend on T and not t . Hence, we have (8) is the optimal policy.

The second part also follows from the fact that since Y_t is given by a linear SDE, it has a solution

$$Y_t = e^{-\int_0^t \kappa K h(\kappa K(T-s)) ds} \left(Y_0 + \Gamma \sigma \int_0^t e^{\int_0^s \kappa K h(\kappa K(T-u)) du} dW_s \right)$$

whose density is non-singular with respect to the Lebesgue measure on \mathbb{R} as h is bounded.

We give some comparative statics for our solution. The effect of the temporary impact term λ is the easiest to understand. As the cost of hedging decreases, it becomes optimal to trade more aggressively to minimize mishedging error (see Figure 1c). The effect of risk aversion γ is similar. As γ increases, so does the penalty for being mishedged and there is a stronger incentive to trade more aggressively (see Figure 1e). To understand the effect for intraperiod volatility σ , recall that the agent is trying to hedge both intraperiod volatility σ (via intraperiod trading) and volatility from the terminal jump σ_T (via trading near the end). As intraperiod volatility σ increases while the terminal jump size σ_T remains constant, the agent optimally trades more aggressively initially while still ending at the same trading intensity at time $t = T$ (see Figure 1b). The effect of increasing terminal jump volatility σ_T while keeping intraperiod volatility constant is roughly the opposite (see Figure 1f): it increases trading near $t = T$ while trading during the rest of the period remains relatively unaffected. Note that there is a need to trade slightly more intensely even at the beginning of the period $t = 0$ as this impacts the mishedging at the end of the period $t = T$.

The option's Γ and the permanent impact ν have more complex effects, which are qualitatively similar to each other. There are two competing effects at play, whose importance differs with the time to maturity. Firstly, there is a **adverse-impact effect** stemming from the permanent impact. To see this, let $Y_t < 0$ so the agent optimally wishes to purchase shares. This purchase will push up each share's value by ν , hence resulting in a decrease of $\nu Y_t < 0$ in the stock value of the agent's portfolio. In general, rebalancing under permanent impact always hurts his existing position, giving the agent an incentive to trade less aggressively in the hope that subsequent price-fluctuations will automatically rehedge his portfolio. The second effect is dubbed the **leverage-effect**. With $\nu = 0$, the purchase of a single share increases the trader's net Δ position Y by 1. However, if $\nu > 0$ then his net Δ position actually increases by $K = 1 + \nu\Gamma$ due to the permanent impact's effect on the option's price $\nu\Gamma$. When $\Gamma > 0$, the effect of a purchase on the portfolio's net Δ -position is

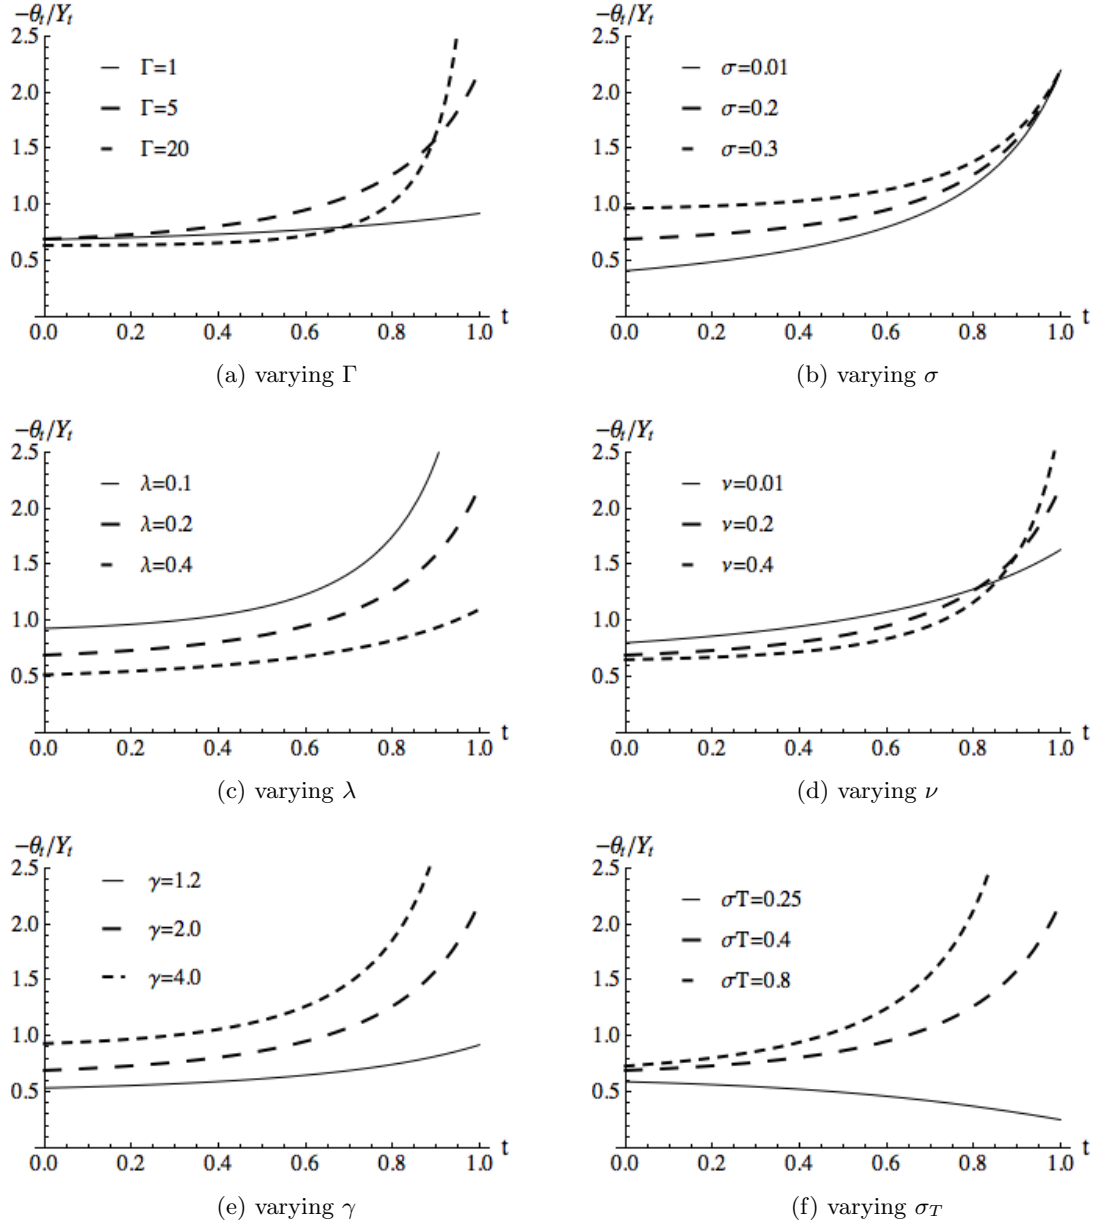


Figure 1 Comparative statics of $(KA_2(T-t) - \nu)/\lambda$ as a function of t . Parameters (unless otherwise specified): $\Gamma = 5.0, \sigma = 0.2, \lambda = 0.2, \nu = 0.2, \gamma = 2.0, \sigma_T = 0.4, T = 1.0$

Units of Fundamental Quantities		
Y_t sh	γ 1/\$	ν \$/sh ²
P \$/sh	λ sec · \$/sh ²	Γ sh ² /\$
θ sh/sec	σ \$/sh/ $\sqrt{\text{sec}}$	J 1

Units of Derived Quantities		
K 1	d 1	κ 1/sec

Figure 2 Units for constants: sh is a share, sec is a second, and \$ is a dollar.

leveraged by the permanent impact on the option's Δ . We can think of this effect as effectively lowering the cost of trading, (i.e. lowering the cost of achieving a fixed change in Δ). Hence, the agent optimally increases the intensity of trading. When $\Gamma < 0$, the effect of a sale or purchase is diminished, if not reversed. We can think of this as reducing the effectiveness of trading on the net Δ position or increasing the cost of rebalancing. Either way, the agent optimally decreases the intensity of trading.

The relative importance of the two effects can be seen in the expression

$$\theta_t = \frac{1}{\lambda} \cdot [\nu Y_t - KJ'(t, Y_t)].$$

The leverage-effect acts through KJ' while the adverse impact acts through νY_t . In Figure 1, $d > 1$ and the pressure to be hedged (KJ') increases towards the end of the period. Hence, near the beginning of the period, the adverse impact is the overriding concern and trading intensity is less aggressive with higher ν because of the adverse-impact effect (Figure 1d) Near the end, the agent is more concerned about hedging the terminal jump. This combined with the ease of hedging from the leverage effect acting through $KJ'(t, Y_t)$ implies the agent trades more aggressively with higher ν near the end. The effect of Γ is similar except that it only directly affects the leverage effect. With higher Γ , the increase in the leverage effect allows the agent to optimally trade more aggressively, particularly near the end of the day (Figure 1a). However, like varying σ_T , because Γ affects mishedging near the closing bell, it has an indirect effect on the trading intensity during the middle of the day.

4. Applications and Equilibrium Implications

4.1. Small Market-Impact Cost Limit

We are interested in the limit as both the temporary and permanent impacts vanish. The trading intensity θ_t is proportional to the distance from being Δ_t -hedged $Y_t = X_t + \Delta_t$ as given in

$$\theta_t = -\kappa \coth(\kappa K(T-t) + \operatorname{arccoth}(d)) Y_t.$$

We observe in Figure 1d that the profile of trading intensity flattens as a function of time with decreasing ν and from Figure 1c, we see that the trading intensity shifts upward with decreasing λ . Hence, we observe the trading intensity both increasing and flattening as market-impact decreases.

In the limit as $\lambda \downarrow 0$ and $\nu \downarrow 0$, we have $K \downarrow 1$, $\kappa \uparrow \infty$, and $d \uparrow \infty$. Therefore $\theta_t/Y_t \uparrow \infty$ and the process Y_t is driven back to zero very quickly. That is, the agent aggressively drives X_t towards Δ_t . The following theorem makes this idea precise.

PROPOSITION 1. *In the limit of vanishing market-impact costs, we recover the Black-Scholes Δ hedge. That is $\|Y\| \rightarrow 0$ as $\lambda, \nu \rightarrow 0$ where $\|\cdot\|$ is the $L^2(\mathbb{P} \times \mu)$ -norm and μ is the Lesbegue measure on $[0, T]$.*

Proof We adopt the notation of the proof of Theorem 3. In our limit, we are concerned with $d > 1$, and so we may take

$$C''' = \frac{\gamma K \sigma_T^2 + \nu}{\lambda}.$$

Since $C''' \rightarrow 0$ in the limit, the result follows from (15) and Gronwall's Lemma.

4.2. Small Intra-period Mismatching Penalty Limit

We consider the case when the penalty for the terminal jump is large compared to the intra-period hedging penalty. Mathematically, this is the case when $\sigma \rightarrow 0$ and $\sigma\Gamma$ is held constant. The objective in this limit is given by the continuation value

$$J(t, y) = \operatorname{ess\,inf}_{\theta_s: s \geq t} \mathbb{E} \left[\frac{\gamma\sigma_T^2}{2} Y_T^2 - \int_t^T Y_u \nu \theta_u + f(\theta_u) du \mid Y_t = y \right] \quad (16)$$

where we have removed the intra-period term. We include this result because (17) clearly demonstrates the increasing nature of the functions A_2 as $t \rightarrow T$ in terms of algebraic, rather than transcendental, functions.

THEOREM 2. *For small intra-period mismatching, the optimal trading intensity $\theta_t = \theta(t, Y_t)$ can be written as*

$$\theta(t, y) = -\frac{\kappa d}{1 + \kappa K d(T-t)} y. \quad (17)$$

Proof Again, we make the Ansatz (11), and separate the resulting HJB by powers of y . This yields

$$\theta(t, y) = \frac{1}{\lambda} (\nu - K A_2(T-t)) y$$

and

$$\begin{aligned} A_2' &= -\frac{1}{\lambda} [(1 + \nu\Gamma)A_2 - \nu]^2 & A_2(0) &= \gamma\sigma_T^2 \\ A_0' &= \frac{\Gamma^2\sigma^2}{2} A_2 & A_0(0) &= 0. \end{aligned}$$

By a similar argument as in Theorem 3, we obtain that

$$A_2(T-t) = \frac{1}{K} \left[\nu + \frac{\lambda\kappa d}{1 + \kappa d(T-t)} \right] \quad A_0(T-t) = \frac{\Gamma^2\sigma^2}{2K} \left[\nu(T-t) + \lambda \log(1 + \kappa d(T-t)) \right]. \quad \square$$

For a sense of how the execution compares to the full case, see Figure 3. Observe that without an intra-period Δ -hedging penalty, the trading intensity ratio $-\theta_t/Y_t$ (see (8)) is lower than with the penalty (see Figure 3), especially during the beginning of the period. The extra trading comes from hedging out fluctuations during the period. Intuitively, the optimal investor trades less aggressively when only faced with a terminal and no running mismatching penalty. As in the $d > 1$ case (14), trading becomes more aggressive towards the end of the period.

If we take the interpretation of $[0, T]$ as the trading day, then this approximation is valid when the overnight jump is large compared to the daily fluctuations. This may be the case in advance of a major closing-bell announcement.

4.3. Restriction on the direction of trading

For a broker, the Δ -hedging problem is significantly more difficult. Regulatory policy mandates that the broker can only buy or sell for any given client order. This is to prevent market-manipulation and to protect clients from potentially unscrupulous dealers. We solve the optimal Δ hedging problem under these constraints.

Without loss of generality, we impose a buy-only restriction on trading. That is, we add the constraint $\theta_t \geq 0$ a.s. to the minimization of our objective (5). While the problem is still Markovian, a closed-form solution is no-longer readily available. We use a policy-improvement algorithm that assumes a Markovian optimal policy $\theta_t = \theta(t, Y_t)$ to solve this problem. As a check, the same policy-improvement code was used to solve the Simplified Model in Section 4.2 and compared against

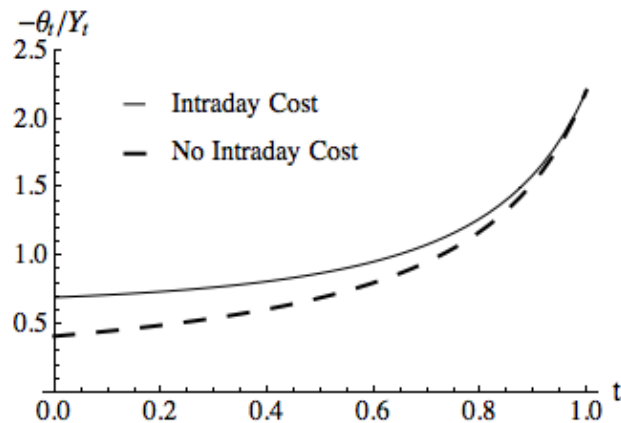


Figure 3 Comparison of trading intensity with and without intraperiod mishedging penalty. The parameter values are defined as in Figure 1

the analytic solution found in Section 4.2. It obtained continuation-value functions J accurate to within $\sim .1\%$.

A comparison of the restricted and unrestricted cases are plotted in Figure 4. Observe the asymmetry in trading policy $\theta(t, y)$ in Figure 4c. For negative values of y , $\theta(t, y)$ is positive to reduce the Δ -hedging error, that is the agent still purchases stock as in the unrestricted case (compare with Figure 4a). For positive values of y , an unrestricted agent would sell shares but a restricted agent cannot do so, so θ must be zero (see Figure 4c). There is a no-trade curve $y(t) \leq 0$ such that the agent trades if and only if he his current net Δ position is below this level, i.e. $\theta(t, y) = 0$ when $Y_t \geq y(t)$ and $\theta(t, y) > 0$ when $Y_t < y(t)$. This curve is marked in the figure in Figure 4c. Observe that $y(0) < 0$ and $y \rightarrow y(t)$ is an increasing function. Hence, far from the terminal time ($t \ll T$), the agent does not purchase stocks even when he is net short Δ . The result is interpreted thus: a selling-restricted agent is hesitant to purchase stock only to see his Δ position become positive before T due to stock-price fluctuations. However, $y(T) = 0$ so this effect disappears as $t \rightarrow T$: as the time remaining for Y_t to fluctuate above 0 runs out, the agent's rule becomes buy if I am net short Δ .

In the restricted case, the asymmetry in execution strategy for θ translates into an asymmetry for the value function J (compare the restricted case Figure 4d with the unrestricted one Figure 4b). For negative values of Y_t , J behaves similarly in both the selling-restricted and non-restricted cases as the optimal strategies are similar. For positive values of Y_t , J is significantly higher in the selling-restricted case as the control cannot be exercised to reduce the hedging error. The difference between the two cases represents the premium of being able to sell.

4.4. Stock Pinning

The increase in trading intensity for high Γ near the expiry is the mechanism behind stock pinning. A put or call option that is about to expire with a large outstanding institutional interest and with a stock price near its strike level will induce a so-called 'pinning effect' whereby the stock price at the close of trading on expiry is 'pinned' to the strike level. Empirically, this manifests itself in the observation that the distribution of end-of-day stock price on these days deviates substantially from the distribution of end-of-day stock price on any other days. That is, the stock price clumps around the strike level at the close of such an option expiry (Avellaneda and Lipkin 2003).

A call or put option near expiry with the underlying prices near its strike exhibits a large positive Γ . When institutional investors are net short Γ (perhaps because they collectively short the put or call) option market makers are net long Γ . The market-makers will hedge their position by buying




Figure 4 Plots of the continuation-value function J and the optimal policy θ for the unrestricted case (Figure 4b and Figure 4a) and the case when trading is restricted to purchases $\theta \geq 0$ (Figure 4d and Figure 4c). In the restricted case, we also plot the no-trade boundary in black, above which $\theta(t, y) = 0$ (see Figure 4c). Parameters: $\Gamma = 5.$, $\sigma = .4$, $\lambda = .2$, $\nu = .2$, $\gamma = 2.$, $\sigma_T = .4$, $T = 1$.

stock as the price dips and selling it as the price increases, thus stabilizing the price. Since the Γ peaks near expiry with the underlying price near the strike, the effect manifests itself as pinning near the strike.

Our model gives a micro-foundations explanation of this story. The agent is now the option market maker, and there is a large outstanding institutional short- Γ interest on an option expiring at T with the underlying price P_t near the option's strike. If the intraday fluctuations P_t are small, then the options position of the market maker effectively has a large flat Γ . Then as Figure 1a shows, a large Γ implies the agent trades intensely as expiry approaches and as his trades are stabilizing, the permanent impact pins the stock price at the strike. Indeed, the effect is self-stabilizing, even with moderately large fluctuations in P_t , the fact that market-makers are stabilizing the price limits price excursions and keeps the option within the band of high Γ for the option.

4.5. Intraday Trading Patterns

If we continue with the interpretation of the trading period as the trading day, we can give a microfoundations account of intraday trading patterns. The typical daily profile is U -shaped, with

the most intense trading occurring during the opening and close (see Figure 5a). Compare Figure 5a with the graph of $f(t)$ where

$$f(t) = \mathbb{E}[\theta_t^2] = \mathbb{E}\left[\kappa^2 \coth^2(\kappa(T-t) + \operatorname{arccoth}(d)) Y_t^2\right]$$

with $Y_0 \sim N(0, (\Gamma\sigma_T)^2)$. The distributional assumption on Y_0 corresponds to being perfectly hedged at the previous close ($Y_{0-} = 0$) but experiencing a change in the option's Δ due to an overnight price movement of the underlying stock. Near the open, the high $f(t)$ comes from the high initial variance in Y_t as the trader trades down his net Δ position from the previous evening's jump. Near the close, the high $f(t)$ comes from the increase in h near $t = T$, which represents aggressive trading in anticipation of the coming evening's overnight jump.

PROPOSITION 2. *The daily trading volume profile generated by Δ -hedging is*

$$f(t) = \kappa^2 \cosh^2(\kappa K(T-t) + \operatorname{arccoth}(d)) \left[\frac{(\Gamma\sigma_T)^2}{\sinh^2(\kappa K T + \operatorname{arccoth}(d))} + \frac{(\Gamma\sigma)^2}{2\kappa K} \left(\coth(\kappa K(T-t) + \operatorname{arccoth}(d)) - \coth(\kappa K T + \operatorname{arccoth}(d)) \right) \right]$$

Proof Observe that from (6) and (8) that we can write

$$\frac{d}{dt} \mathbb{E} Y_t^2 = -2\kappa K h(\kappa K(T-t)) + \frac{1}{2}(\Gamma\sigma)^2$$

whose solution is given by

$$\mathbb{E} Y_t^2 = \sinh^2(\kappa K(T-t) + \operatorname{arccoth}(d)) \left[\frac{\mathbb{E} Y_0^2}{\sinh^2(\kappa K T + \operatorname{arccoth}(d))} + \frac{(\Gamma\sigma)^2}{2\kappa K} \left(\coth(\kappa K(T-t) + \operatorname{arccoth}(d)) - \coth(\kappa K T + \operatorname{arccoth}(d)) \right) \right]. \quad \square$$

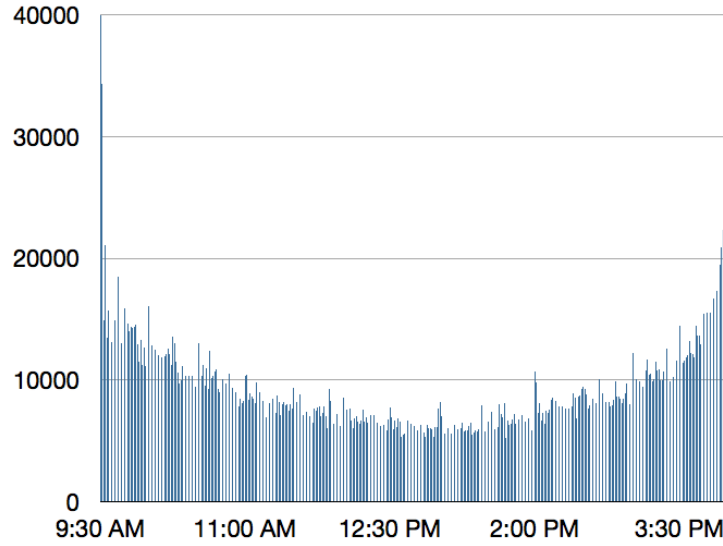
5. Discrete-Time Formulation and Solution

We need a discrete-time formulation and solution to the hedging problem in order to perform the simulations on TAQ data in Section 6. The simplest discrete strategy would be to evaluate the continuous-time strategy at discrete time points, buying $\theta_t \Delta t$ shares over the interval $[t, t + \Delta t]$. In effect, this is a forward Euler discretization of the underlying dynamics, and exhibits the well-known overshooting instability of that method for reasonable parameter values: when temporary impact λ is small the problem is “stiff” (see, for example, Ascher and Petzold (1998)). In order to obtain well-behaved discrete time solutions we must pose address the discrete-time problem directly.


5.1. Discrete-Time Formulation

We will discretize the dynamics by imposing a lattice of N points of width $\Delta t = T/N$, $\mathcal{T}_N = \{\frac{T}{N}(\mathbb{Z} \cap [0, N])\}$. At each time $t \in \mathcal{T}_N$, the agent sets his strategy for the entire period $[t, t + \Delta t]$. Under such conditions, it is only necessary to consider the implied discrete-time process. Hence, the agent's stock position is given by

$$X_t = X_0 + \sum_{u \in \mathcal{T}_N, u < t} \theta_u$$



(a) Average Minutely Trading Volume of Boeing (NYSE:BA) in the period April 3rd, 2000 to March 30th, 2001.



(b) Intensity of trading $f(t)$ for one trading day with $\lambda = 10^{-6}\sigma$, $\sigma = .0008$, $\nu = 10^{-6}$, $\Gamma = 2000$, $\sigma_T = .1$, $\gamma = 10^{-10}$.

Figure 5 Intraday trading intensity, actual, and from delta hedging.

while the dynamics for the stock price becomes

$$P_t = P_0 + \nu X_t + \sum_{u \in \mathcal{T}_N, u < t} \sigma \Delta W_u$$

where $\Delta W_{kT/N} = W_{(k+1)T/N} - W_{kT/N}$. Then the new discrete-time value function becomes

$$J(t, y) = \inf_{\theta} \mathbb{E} \left[\sum_{u \in \mathcal{T}_N, u \geq t} \left(\frac{\gamma \sigma^2}{2} Y_u^2 - Y_u \nu \theta_u + \frac{\lambda}{2} \theta_u^2 \right) \Delta t + \frac{\gamma \sigma_T^2}{2} Y_T^2 \middle| Y_t = y \right]. \quad (18)$$

	P_T	σ	σ_T	Γ
BA	53.27084	0.0433451	0.8270762	0.0515023
BAC	49.5763872	0.0243846	0.7866773	0.0780991
MSFT	64.9043047	0.0122172	1.4496147	0.1173921
PFE	43.5416817	0.0164636	0.6695676	0.1008441
WMT	52.5306878	0.0208141	0.9813519	0.0773531

Figure 6 The average closing stock price P_T , absolute volatility in dollars per share per square-root seconds σ , the average overnight volatility in dollars per share σ_T , and the average Γ of a call-option that matures in 5 days and was At-The-Money at the previous close.

5.2. Problem Solution

Again, we make a quadratic Ansatz that for $T - t \in \mathcal{T}_N$,

$$J(T - t, y) = A_2(T - t + \Delta t) \frac{y^2}{2} + A_0(T - t + \Delta t). \quad (19)$$

THEOREM 3. *The discrete-time optimal trading intensity $\theta_t = \theta(t, Y_t)$ is given by*

$$\theta(t, y) = \frac{\nu - K A_2(T - t)}{\lambda + K^2 A_2(T - t) \Delta t} y \quad (20)$$

and the continuation value J is given by (19) where

$$\begin{aligned} A_2(T - t + \Delta t) &= A_2(T - t) + \gamma \sigma^2 \Delta t + \frac{(\nu - K A_2(T - t))^2}{\lambda + K^2 A_2(T - t) \Delta t} \Delta t & A_2(0) &= \gamma \sigma_T^2 \\ A_0(T - t + \Delta t) &= A_0(T - t) + \frac{1}{2} \Gamma^2 \sigma^2 A_2(T - t) \Delta t & A_0(0) &= 0. \end{aligned}$$

Proof The discrete-time HJB equation becomes

$$J(T - t, y) = \inf_{\theta_t} \left[\frac{\gamma \sigma^2}{2} y^2 - y \nu \theta_t + \frac{\lambda}{2} \theta_t^2 \right] \Delta t + \frac{A_2(T - t)}{2} \left[(y + K \theta_t \Delta t)^2 + (\Gamma \sigma)^2 \Delta t \right] + A_0(T - t) = 0.$$

Separating both sides by powers of y and using (19) yields the desired result.

6. Simulation Using TAQ Data

To test the hedging strategy, we use stock TAQ Data to simulate Δ -hedging a call option using our trading strategy and the benchmark Black-Scholes strategy. We choose five companies: Boeing (NYSE:BA), Bank of America (NYSE:BAC), Microsoft (NASDAQ:MSFT), Pfizer (NYSE:PFE), and Wal-Mart (NYSE:WMT), each representing a different sector of the economy. Daily NBBO data was collected for the 251 trading days from April 3rd, 2000 to March 30th, 2001 (inclusive) and the mid-price prevailing at each minute of each trading afternoon (noon to 4:00 PM) was struck (computed). (The results would apply equally well had we used trading data from the entire trading day. We explain this later-on in this section. Using only afternoon data was done to reduce the computational load.) Summary statistics of the data are given in Figure 6.

We use the price data to calibrate our the discrete-time model (see Section 5) under minutely rehedging. We then use the data to simulate Δ -hedging a derivatives' position with a flat Γ . We chose the Γ to corresponded to that of 10,000 call options that were At-The-Money at the previous trading day's close and mature 5 days after that day's close. The price and greeks of the option are computed using the Black-Scholes formula.

We can think of the flat Γ derivative as approximating that of an At-The-Money call option under Approximation 3. A call-option's Γ profile becomes very narrow in moneyness near the option's expiry so that small fluctuations in the stock price will greatly affect the Γ , thus violating the assumptions of the model. This feature is not unique to our setup but stems from the 'kink' in the 'hockey-stick shape' payoff of the call and the associated difficulty of hedging At-The-Money call options near expiry is a widely recognized problem among practitioners. We sidestep this issue in our simulation by choosing a call option that expires in 5 days.

We set the initial position X_0 so that the option is initially Δ -hedged at the previous close. This (optimistically) simulates the position of a trader who hedges an option across multiple days. His initial position in the morning may not be hedged due to the overnight fluctuation. The absolute risk-aversion coefficient is chosen to be of the order $\gamma \sim 1 \times 10^{-9}$. This can be thought of as corresponding to a relative risk aversion of 1 for an agent with one billion dollars of wealth. We assume no permanent impact $\nu = 0$ and a temporary impact of the form $\lambda = \lambda_p \sigma$ where $\lambda_p > 0$ is a proportionality constant and σ is the absolute volatility of the stock. This accounts for the well-known stylized fact that, caeteris paribus, market impact is higher for stocks with greater volatility. We choose $\lambda_p = 1 \times 10^{-6}$. This has the interpretation that for a stock with $\sigma = .01$, purchasing 1000 shares of stock per second would have a market impact of one penny (see Figure 2 for units). For reference, the average σ in our sample ranged from .016 for PFE to .043 for BA (see Figure 6). The characteristic time-scale at which trading intensity peaks near the market close is $1/\kappa \approx 100$ seconds. In other words, for our liquid stocks, the agent trades with a constant proportional intensity until the last few minutes of trading when he trades more aggressively. Hence, our results are not affected by only hedging in the afternoon.

We track the performance of two different agents, one using the proposed hedging intensity of Section 5 (Opt) and another whose strategy is to maintain a Black-Scholes Δ -hedged (BS). The Black-Scholes trader will not be perfectly hedged due to fluctuations in the subsequent minute interval before rehedging. However, (BS) will likely maintain a tighter hedge on average as he would be perfectly hedged if the stock did not fluctuate in the subsequent interval whereas (Opt) is not even hedged at the start of the interval to begin with.

A summary of the simulation results is given in Figure 7. For an agent with $\gamma = 2.0 \times 10^{-9}$, (Opt) is able to save $\sim 80\%$ of the impact cost (relative (BS)) while only incurring $\sim 20\%$ increase in mishedging penalty. Figure 7, we give the average terminal mishedging (Terminal), running mishedging (Running), and temporary market impact (Impact) cost in equation (18). With higher risk aversion, (Opt)'s trading becomes more aggressive and more closely tracks that of (BS). This is seen in terminal and running mishedging fractional spreads between the two agents (see Terminal and Running columns in Figures 7a, 7c, and 7e). This difference decreases as risk-aversion increases, demonstrating that the Opt agent is maintaining a tighter hedge. However, this comes at an increased cost to (Opt), whose fractional temporary market-impact cost savings with respect to (BS) decreases with increasing risk-aversion (see Impact columns in Figure 7a, 7c, and 7e).

7. Conclusion

Market-impact from hedging a non-trivial outstanding position can have an important effect on costs. This paper develops a highly-tractable framework for analyzing optimal hedging of options for large agents who face market impact. We use a mean-variance framework and find that the optimal solution for option's hedging is for the agent to trade towards a "target portfolio." The target portfolio is the the optimal frictionless Black-Scholes Δ -hedge. The trader is prohibited from exactly holding the Black-Scholes hedge portfolio due to the market-impact costs and can only trade towards it to minimize the turnover cost. We also show that our results can be interpreted as a kind of Merton problem where the Merton-optimal portfolio is the Black-Scholes hedge portfolio.

	Terminal	Running	Impact
BA	1.178	1.751	-0.890
BAC	1.543	3.195	-0.900
MSFT	0.283	3.842	-0.903
PFE	1.014	2.816	-0.922
WMT	1.067	3.048	-0.897

(a) Fractional, $\gamma = 0.5 \times 10^{-9}$

	Terminal	Running	Impact
BA	0.293	0.816	-0.837
BAC	0.355	0.717	-1.089
MSFT	0.246	1.083	-1.729
PFE	0.187	0.923	-1.637
WMT	0.242	0.851	-1.334

(b) Per Std, $\gamma = 0.5 \times 10^{-9}$

	Terminal	Running	Impact
BA	0.580	1.158	-0.844
BAC	0.674	2.146	-0.862
MSFT	0.090	2.671	-0.870
PFE	0.494	1.891	-0.890
WMT	0.435	2.066	-0.857

(c) Fractional $\gamma = 1.0 \times 10^{-9}$

	Terminal	Running	Impact
BA	0.207	0.733	-0.833
BAC	0.331	0.735	-1.093
MSFT	0.157	1.093	-1.738
PFE	0.142	0.945	-1.641
WMT	0.203	0.883	-1.334

(d) Per Std $\gamma = 1.0 \times 10^{-9}$

	Terminal	Running	Impact
BA	0.266	0.736	-0.783
BAC	0.268	1.404	-0.811
MSFT	0.029	1.807	-0.826
PFE	0.217	1.234	-0.847
WMT	0.164	1.362	-0.804

(e) Fractional, $\gamma = 2.0 \times 10^{-9}$

	Terminal	Running	Impact
BA	0.157	0.629	-0.836
BAC	0.274	0.743	-1.097
MSFT	0.099	1.103	-1.748
PFE	0.113	0.958	-1.642
WMT	0.162	0.915	-1.333

(f) Per Std, $\gamma = 2.0 \times 10^{-9}$

Figure 7 Simulation results are given for the terminal, running, and temporary impact costs for various values of the risk-aversion γ . Fractional gives ‘Fractional spread’ or the average value of $(\text{Opt} - \text{BS})/\text{BS}$ over the trading period for the terminal mishedging (Terminal), running mishedging (Running), and temporary market-impact (Impact) costs in (18). Per Std gives the sample average $(\text{Opt} - \text{BS})$ in terms of the sample standard deviation. For 250 trading days in our dataset, multiply by $\sqrt{250} \approx 15.8$ for the z -score.

We show that in the limit of vanishing market-impact costs, we recover that the agent engages in Black-Scholes hedging. We are able to give a micro-foundations account of stock-pinning as market-makers hedging a large outstanding option position driving the price towards an option strike at expiry. Our strategy can also partially account for the U -shaped intraday trading patterns as the optimal response to hedging an option position with overnight jumps. Finally, we present the results of numerical simulations which demonstrate the practicality of this technique even for real-world intraday stock-price fluctuations.

References

- A. Alfonsi and A. Schied. Optimal trade execution and absence of price manipulations in limit order book models. *SIAM J. Financial Math.*, 1:490–522, 2010.
- R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, Apr 2001.
- U. M. Ascher and L. R. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, Philadelphia, PA, 1998.
- M. Avellaneda and M. D. Lipkin. A market-induced mechanism for stock pinning. *Quantitative Finance*, 3: 417–425, Sep 2003.
- K. Back. Insider trading in continuous time. *Review of Financial Studies*, Jan 1992.
- F. Black and M. Scholes. The pricing of options and corporate liabilities. *The Journal of Political Economy*, Jan 1973.
- U. Cetin, H. M. Soner, and N. Touzi. Option hedging for small investors under liquidity costs. *Finance and Stochastics*, 14:317–341, Jan 2010.
- J. Cvitanic and H. Wang. On optimal terminal wealth under transaction costs. *Journal of Mathematical Economics*, 35:223–231, Jan 2001.
- M. H. A. Davis and A. R. Norman. Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15(4):676–713, Nov 1990.
- N. Garleanu and L. Pedersen. Dynamic trading with predictable returns and transaction costs. *NBER Working Paper*, Mar 2009.
- K. Janecek and S. E. Shreve. Asymptotic analysis for optimal investment and consumption with transaction costs. *Finance and Stochastics*, 8:181–206, Jan 2004.
- P. Kratz and T. Schöneborn. Optimal liquidation in dark pools. *papers.ssrn.com*, Mar 2010.
- A. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, Jan 1985.
- R. Merton. Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4 (1):141–183, Jan 1973.
- S. Predoiu, G. Shaikhet, and S. Shreve. Optimal execution in a general one-sided limit-order book. *math.cmu.edu*, Feb 2010.
- L. C. G. Rogers and S. Singh. The cost of illiquidity and its effects on hedging. *Preprint*, Aug 2007.
- T. Schoeneborn and A. Schied. Liquidation in the face of adversity: stealth vs. sunshine trading. *papers.ssrn.com*, Mar 2010.
- S. Shreve and H. Soner. Optimal investment and consumption with transaction costs. *The Annals of Applied Probability*, 4(3):609–692, Aug 1994.
- H. M. Soner, S. E. Shreve, and J. Cvitanic. There is no nontrivial hedging portfolio for option pricing with transaction costs. *The Annals of Applied Probability*, 5(2):327–355, Jan 1995.