

INFORMATION-NEUTRAL HEDGING OF DERIVATIVES UNDER MARKET IMPACT AND MANIPULATION RISK

Behzad Alimoradian, Karim Barigou, Anne Eyraud

REPRINT | 2026 / 10

ISBA

Voie du Roman Pays 20 - L1.04.01

B-1348 Louvain-la-Neuve

Email : lidam-library@uclouvain.be

<https://uclouvain.be/en/research-institutes/lidam/isba/publication>



Article

Information-Neutral Hedging of Derivatives Under Market Impact and Manipulation Risk

Behzad Alimoradian¹ , Karim Barigou^{2,*} and Anne Eyraud³

¹ Independent Researcher; behzad.alimoradian@gmail.com

² Institute of Statistics, Biostatistics and Actuarial Science—ISBA, Louvain Institute of Data Analysis and Modeling—LIDAM, UCLouvain, Voie du Roman Pays 20, 1348 Louvain-La-Neuve, Belgium

³ Laboratoire de Sciences Actuarielle et Financière, Institut de Science Financière et d'Assurances, Université Claude Bernard Lyon 1, 50 Avenue Tony Garnier, F-69007 Lyon, France; anne.eyraud@univ-lyon1.fr

* Correspondence: karim.barigou@uclouvain.be

Abstract

The literature on derivative pricing in illiquid markets has mostly focused on computing optimal hedging controls, but empirical microstructure studies show that large order flow generates persistent and predictable price effects. Therefore, these controls can themselves induce endogenous market manipulation because traders can internalize the impact of their own trades. We identify the key shortcoming as the absence of a formal separation between a large trader's informational advantage and the mechanical price impact and temporary cost-of-hedging. To address this gap, we introduce a counterfactual informed observer—an agent who knows the large trader's strategy but does not face trading frictions—and use this device to isolate informational order-flow effects from mechanical price impact, a distinction explicitly observed in microstructure data. We prove the existence of information-neutral probability measures under which the discounted asset is a martingale for this observer and derive a hedging framework that jointly accounts for transaction costs and permanent market impact. Numerical experiments show that because price pressure and order-flow effects create non-linear execution costs, the optimal hedge for an out-of-the-money call can deviate substantially from the Black–Scholes hedge, with implications for risk management and regulatory monitoring.

Keywords: option pricing; market impact; illiquid markets; transaction costs; stochastic optimal control



Academic Editor: Vassilios
Papavassiliou

Received: 4 October 2025

Revised: 1 December 2025

Accepted: 16 December 2025

Published: 1 January 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The traditional option pricing framework, as outlined in the seminal paper by [Black and Scholes \(1973\)](#), assumes a frictionless and liquid financial market. However, this framework does not hold when dealing with large-notional derivatives on illiquid assets that are subject to market impact. Empirical microstructure studies show that when large trades or block hedges are executed, liquidity is often depleted, producing immediate price adjustments ([Hasbrouck, 1991](#)). A portion of this impact can persist, as documented in empirical analyses showing that large institutional metaorders leave lasting effects on prices ([Bouchaud et al., 2010](#); [Tóth et al., 2011](#)). Consequently, the spot price of the underlying—and therefore any derivative hedge—may shift as a direct result of the hedging trades themselves.

A first approach to account for such frictions, most notably the Leland model ([Leland, 1985](#)), incorporates transaction costs by inflating volatility. However, this method captures

only the temporary cost of adjusting the hedge and ignores liquidity pressure or any feedback of the hedge on the underlying price. To incorporate such feedback effects, two principal approaches have emerged. The first is a full-hedging approach, introduced by [Liu and Yong \(2005\)](#) and further developed in [Loeper \(2018\)](#), [Said \(2020\)](#), [Bordag and Frey \(2008\)](#), and [Abergel and Loeper \(2017\)](#), which seeks to eliminate risk entirely but can become impractical because continuously adjusting the hedge can generate prohibitively high option prices, particularly near maturity or when the spot is close to the strike. The second is a partial-hedging approach, exemplified by [Jonsson and Sircar \(2002\)](#) and extended in utility- or gain-maximization frameworks such as [Guéant and Pu \(2017\)](#), which trade off hedging costs against residual risk but can inadvertently create incentives for the trader to profit from their own impact.

These types of model-driven price manipulation have been highlighted in several theoretical studies, which show that objective-function-based hedging frameworks can inadvertently create incentives for traders to influence prices through their own trades ([Horst & Naujokat, 2011](#); [Kraft & Kühn, 2011](#); [Roch, 2022](#)). Such algorithmic or model-induced distortions fall under the broader category of new market manipulation described by [Lin \(2017\)](#), where price effects arise not from deliberate intent but as unintended consequences of the trading rules themselves. Empirical research also shows that derivatives markets can exhibit settlement-related distortions or manipulation incentives ([Griffin & Shams, 2018](#); [R. Jarrow et al., 2018](#); [Ni et al., 2005](#)). Beyond these derivatives-specific patterns, microstructure research establishes that order flow contains both informational and liquidity-driven components ([Barclay & Warner, 1993](#); [Easley & O'Hara, 1996](#); [Huang & Stoll, 1997](#)), and manipulation risks are more acute in settings where price impact includes a permanent component. Empirical studies show that such permanent components arise in institutional trading ([Almgren et al., 2005](#); [Bouchaud et al., 2010](#); [Tóth et al., 2011](#)), making lasting price distortions more exploitable than purely temporary execution costs. Together, these observations reveal a conceptual gap: existing models lack a way to capture informational drivers of market impact while ensuring non-manipulable pricing.

To address this gap, we introduce an insider observer: an agent who shares the large trader's informational advantage but exerts no influence on prices. As [R. A. Jarrow \(1994\)](#) explains, in the presence of derivatives the market effectively consists of two venues—the underlying and the derivative—and misalignment between them can create opportunities for manipulation. By analyzing the market from the insider's perspective, we capture how informational asymmetry modifies the martingale pricing structure and thereby alters the hedging dynamics faced by the large trader. This interpretation—where the large trader's trading rate reflects informed order flow in an analogous sense to classical microstructure models—motivates a conceptual distinction between informational and mechanical components of permanent impact. Unlike empirical decompositions of order flow into informed and liquidity-driven parts ([Barclay & Warner, 1993](#); [Easley & O'Hara, 1996](#); [Huang & Stoll, 1997](#)), our separation is implemented through an insider filtration combined, which isolates only the informed part of permanent impact in the reference measure. This formulation aligns with literature on asymmetric information and filtration enlargement ([Amendinger et al., 1998](#); [Eyraud-Loisel, 2013](#); [Gorud & Pontier, 1998](#); [Hillairet, 2004](#)), although—unlike some of these studies—we do not study the perspective of the uninformed trader.

A key part of our methodology relies on specifying a market-impact structure. The microstructure literature distinguishes between temporary and permanent price impact, and more elaborate formulations also include transient, decaying impact ([Almgren & Chriss, 2001](#); [Gatheral, 2010](#)). Temporary impact reflects the instantaneous liquidity cost of a trade, whereas permanent impact captures the informational or mechanical footprint left on the mid-price, and both components are well documented empirically ([Hasbrouck, 1991](#);

Madhavan et al., 1997). In our framework, temporary impact is modelled entirely through a nonlinear transaction-cost function à la Almgren–Chriss, capturing only the mechanical execution cost of adjusting the hedge. Permanent impact is represented by a drift shift in the underlying price process and is decomposed through a Poisson mechanism: the trader knows the expected magnitude of the permanent price displacement (informational component), whereas the random jump arrival times reflect liquidity-driven uncertainty (mechanical component). The insider observer filters out the informational part and treats the mechanical component as exogenous noise, allowing only the latter to enter price formation under the reference measure. In contrast, we do not incorporate transient or decaying impact (Gatheral, 2012; Klock, 2014), as our objective is not to model order-book resilience but to isolate informational asymmetry from liquidity-driven price pressure within a tractable continuous-time setting.

This paper is structured as follows. Section 2 discusses the regulatory notes, challenges and example of unintended market manipulation within derivatives pricing and hedging. It provides a case example and introduces our problem reduction approach by hypothesizing a market with one large trader and an insider trader. Section 3 describes the stock dynamics and the market impact model used in this paper. Based on this dynamics, Section 4 explores the pricing equation for the insider trader and elaborates on information-neutral probabilities for asset valuation as a martingale. Section 5 analyzes the portfolio dynamics of the large trader and formulates the mathematics of its hedging portfolio. Section 5.2 explores the optimal hedging strategy for the large trader, addressing market manipulation within the stochastic control framework. Section 6 presents a numerical implementation based on finite-difference methods, comparing the strategies of the large trader, the insider trader, and the classical Black–Scholes hedge. The model’s limitations, potential areas for future research, and the concluding remarks are then discussed in Sections 7 and 8.

2. Challenges in Traditional Hedging Models and Motivation for Our Approach

Empirical microstructure studies show that large order flow generates measurable price pressure even in the absence of news (Hasbrouck, 1991; Madhavan et al., 1997). Consequently, large hedging trades in derivative markets can influence the underlying asset price, creating distortions that may appear as spurious price movements or unintended changes in the payoff structure. Understanding and mitigating such effects is essential for traders managing sizeable derivative positions in markets with limited liquidity. This section therefore outlines the forms of market manipulation that are most relevant to large traders, highlights the challenges faced by pricing models that may unintentionally induce these distortions, and introduces our methodological approach for addressing these concerns.

2.1. Regulatory Notes

There are two types of market manipulative strategies, relevant for our ecosystem that is discussed in the regulations.

Abusive Squeeze: An abusive squeeze occurs when a trader deliberately creates a market condition that amplifies their ability to profit from the subsequent price movements of an asset. This typically involves sequentially buying or selling a large quantity of the asset to influence its price in a manner that benefits their trading position. Such actions can lead to artificial price levels. In the context of derivatives, an abusive squeeze might involve a large trader influencing the underlying asset’s price of a derivative contract tied to it. Regulatory bodies recognize this behavior as a form of market abuse (Financial Conduct Authority, 2005). Related distortions have also been characterized within the

market microstructure literature, where conditions for profitable manipulation are formally derived (Huberman & Stanzl, 2004; R. A. Jarrow, 1992).

Payoff Manipulation: Payoff manipulation refers to scenarios in which a trader's actions directly affect the payoff of a derivative instrument through their market impact. By strategically buying or selling the underlying asset, a trader can influence the settlement price of options or other derivatives to their advantage, particularly when holding sufficiently large positions. Such practices are prohibited under regulations designed to safeguard market integrity (European Parliament and of the Council, 2014; Financial Conduct Authority, 2005). Theoretical price-impact models likewise show that hedging flows or liquidity shocks can distort settlement levels and derivative payoffs (Kraft & Kühn, 2011; Roch, 2022).

One important consideration is that payoff manipulation can also arise unintentionally, for example when delta hedging is performed under models such as Black–Scholes that do not account for the price effects of the hedge itself. Empirical evidence of strike pinning and expiration-day distortions (Ni et al., 2005) illustrates how such unintended effects can arise in practice.

2.2. Practical Example of Unintended Manipulative Strategies

As a result, the standard delta-hedging strategies derived from these models can inadvertently lead to market conditions that result in payoff manipulation. To further explain this, the following example illustrates how such unintentional manipulation can happen in the case of a KO (Knock-out) call option.

Example 1 (Payoff Manipulation for an Out-of-the-Money Knock-out Call Option). Consider a large trader who holds a large position of a short-term, out-of-the-money Knock-out (KO) call option with a knock-out barrier set at 105, while the current stock price is 100. A KO call option is a type of barrier option that ceases to exist if the underlying asset's price reaches or exceeds a specified level, known as the knock-out barrier. If an investor holds a KO call option with a barrier at 105 and the stock price hits 105, the option becomes void and worthless. This characteristic makes KO options particularly sensitive to price movements of the underlying asset near the barrier level, exposing them to potential manipulation more than other types of derivatives.

Figure 1 illustrates the relationship between the spot price of the underlying asset and the price of the Knock-out (KO) call option. In the figure, the horizontal axis represents the spot price of the underlying asset, while the vertical axis denotes the price of the option. The solid line shows how the price of the KO call option changes as the spot price varies. As the spot price increases but remains below the KO barrier, the option price increases, reflecting its growing value. However, as the spot price approaches the KO barrier, the potential for the option to become worthless at the barrier exerts downward pressure on its price. The dashed line represents the payoff of the option at expiry. This payoff is contingent on the spot price relative to the KO barrier. If the spot price remains below the KO barrier at expiry, the option payoff reflects its intrinsic value. Conversely, if the spot price reaches or exceeds the KO barrier, the option is knocked out, resulting in no payoff. In this figure, the current spot price, depicted via a red dot, is in the money (greater than the strike price) but still far from the KO barrier.

Under a delta hedging strategy using the classical Black–Scholes model, the seller of the option must purchase more of the underlying asset to hedge their position as the option's sensitivity to the spot price, i.e., the delta, increases. As delta increases, it reflects the growing sensitivity of the option's price to changes in the spot price. To maintain a delta-neutral position and mitigate risk, the trader buys additional stock. This strategy offsets the option's price changes with corresponding movements in the underlying asset.




Figure 1. The spot price is in the money but still far from the knock-out (KO) strike. The solid curve shows the price of the KO barrier option as a function of the spot price, while the dashed lines indicate the option strikes and the payoff discontinuity at the KO barrier. The red dot represents the current spot price (x-axis) and the corresponding barrier option price (y-axis).

Now, let's assume the stock price jumps closer to the barrier, but not too close, as illustrated in Figure 2. In this scenario, represented by the red dot moving further up, the stock price increases but still remains below the knock-out barrier.




Figure 2. As the spot moves further up, the delta increases. The solid curve shows the price of the KO barrier option as a function of the spot price, while the dashed lines indicate the option strikes and the payoff discontinuity at the KO barrier. The red dot represents the current spot price (x-axis) and the corresponding barrier option price (y-axis).

However, this aggressive hedging can drive the stock price closer to the knock-out barrier, potentially breaching it. When the barrier is breached, as shown in Figure 3, the KO condition is triggered, and the option becomes worthless, resulting in no payoff.




Figure 3. Large hedging activity drives the spot into the KO barrier. The solid curve shows the price of the KO barrier option as a function of the spot price, while the dashed lines indicate the option strikes and the payoff discontinuity at the KO barrier. The red dot represents the current spot price (x-axis) and the corresponding barrier option price (y-axis).

Example 1 demonstrates how a trader's hedging actions, dictated by models that do not account for market impact, can unintentionally manipulate the payoff of the option. Recognizing this oversight, risk managers and traders often adapt their strategies to mitigate the inefficiencies of these models.

In practice, traders and risk managers are aware of the potential for unintentional payoff manipulation. To mitigate this, a common heuristic method involves modifying the payoff structure by setting the barrier level slightly above the actual knock-out (KO) strike and shaping the option to behave like a reverse straddle, peaking at the KO strike. This

adjustment smooths the payoff near the boundaries and reduces the option's sensitivity to small movements in the underlying asset, thereby limiting the amplitude of daily delta adjustments required for hedging. By using this approach, traders effectively price and hedge a hypothetical payoff that is more likely to trigger a payment than the true payoff, resulting in selling the option at a higher premium. Although a detailed explanation of this heuristic lies beyond the scope of this paper, it illustrates that the challenges discussed earlier are well recognized in practice, and that practitioners routinely employ such adjustments to market manipulation risks.

2.3. Our Approach and the Ecosystem Studied in This Paper

The example above underscores the necessity for models that account for the market impacts of trading activities and sets some criteria for preventing market manipulation. Central to this is how a large trader's knowledge of their market impact influences their hedges. To explore how this knowledge influences their actions, we introduce the concept of a hypothetical insider trader. This insider trader possesses the same information as the large trader but does not impact the market because their trades are too small to influence prices. By studying the option pricing and delta hedging strategies of this insider, we gain insights into how information alone, without the market impact, affects pricing and hedging strategies in derivatives trading. This construction parallels the informational component of order-flow dynamics documented in microstructure research, where informed and liquidity-driven trades are distinguished (Easley & O'Hara, 1996; Huang & Stoll, 1997), while deliberately excluding the liquidity-driven (mechanical) component of price impact in order to isolate informational effects.

Therefore, we hypothesize a market characterized by limited liquidity, with two types of actors:¹

- The *large trader* who has market impact and is aware of the effect of its own actions on the market.
- The *insider trader* who possesses the same knowledge as the large trader but does not generate any market impact because its share execution size is small.

3. The Market Impact Model

We consider a finite, deterministic time horizon $T \in \mathbb{R}^+$. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space. Let $\{W_t\}_{0 \leq t \leq T}$ be a standard Brownian motion and $\{N_t\}_{0 \leq t \leq T}$ be a homogeneous Poisson process with intensity $\phi > 0$. We assume that the processes $\{W_t\}_{0 \leq t \leq T}$ and $\{N_t\}_{0 \leq t \leq T}$ are independent. The probability measure $\mathbb{P}$ is traditionally called a real-world probability measure. Let $\mathbb{F} = \{\mathcal{F}_t\}_{t \in [0, T]}$ be the filtration jointly generated by the processes $\{W_t\}_{0 \leq t \leq T}$ and $\{N_t\}_{0 \leq t \leq T}$.

While leveraging a homogeneous Poisson process for its simplicity and tractability, we acknowledge the limitations this assumption may impose, particularly in more nuanced market contexts where a Cox process, with an intensity that adapts to market conditions and is $\mathbb{F}$ -adapted, could offer a more refined model. Nonetheless, our focus remains on the homogeneous Poisson process to maintain model simplicity and computational feasibility.

The core of our model revolves around three $\mathbb{F}$ -adapted processes: the underlying spot price process $\{S_t\}_{0 \leq t \leq T}$, the process $\{\chi_t\}_{0 \leq t \leq T}$ denoting the annualized rate of shares executed by a large trader, and the process $\{V_t\}_{0 \leq t \leq T}$ representing the evolving volume of shares held by the large trader. These processes are governed by the following dynamics:

$$\begin{cases} \frac{dS_t}{S_{t-}} = \mu_t dt + \sigma dW_t + \Lambda(\chi_t) dN_t, \\ dV_t = \chi_t dt, \\ \chi_t = \tilde{\chi}(t, S_t, V_t), \end{cases} \quad (1)$$

with $S_0 > 0$ and $V_0 \geq 0$. The asset's drift, represented by μ_t , is assumed to be a deterministic real function, and the constant volatility $\sigma \in \mathbb{R}^+$ is also deterministic.

In Equation (1), $\tilde{\chi} : \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \mathbb{R}$ is a piecewise continuous and bounded function that specifies the *execution speed* of the large trader's orders, denoted by the trading strategy $\{\chi_t\}_{t \geq 0}$, where χ_t represents the annualized rate of share execution at time t . Here we assume that, to conceal its order size, the large trader keeps this strategy hidden from the market; hence, χ_t is known only to the large trader and to the insider trader.

While this goes beyond the scope of this paper, there are various algorithmic methods to conceal order size. The simplest is the price-taker approach, which places market orders that are immediately filled at the best available price. Another common method employs iceberg-type limit orders. See Christensen and Woodmansey (2013) for more details.

Moreover, $\Lambda : \mathbb{R} \rightarrow \mathbb{R}$, a non-decreasing, square-integrable function, encapsulates the permanent market impact of the large trader's actions, satisfying $\Lambda(0) = 0$. We postulate that during a brief interval $[t, t + dt)$, the market's liquidity is either sufficient to absorb the trade without any lasting impact or insufficient, leading to a permanent market impact. In scenarios where market liquidity can accommodate the trade volume, the executed orders do not result in a permanent change to the asset's price. Conversely, when market depth is inadequate, the trading activity induces a permanent market impact, quantified by the function $\Lambda(\chi_t)$.² This impact is realized at the occurrence of each Poisson jump, where the asset's price adjustment is directly tied to the instantaneous trading rate.

Note that the market impact could be positive or negative depending on whether the shares are bought or sold. In the case of no trading activity by the large trader, Equation (1) reduces to the standard Black-Scholes dynamics.

4. Option Pricing for the Insider Trader

The insider trader has access to the information about the trading policy $\{V_t, \chi_t\}_{0 \leq t \leq T}$ of the large trader over time. Let $\{r_t\}_{0 \leq t \leq T}$ be a deterministic interest rate process. We also assume that the holder of the stock benefits from a dividend yield of $q \geq 0$, as well as the presence of a securities lending market and a future market for the underlying stock. The holder of the stock bears a cost of carry b , which is the adjustment rate between the future market and the current value of the stock. These markets, although subject to liquidity issues, enable the holder to enter into short selling positions.

In the following theorem, we demonstrate the existence of a set of probability measures under which the asset price and the option value are martingales. This result is the main outcome of this section and formulates mathematically how the insider trader's information affects the pricing of derivatives and the hedging strategies.

Theorem 1. *Under the regularity condition:*

$$\Lambda(\chi_t)(\mu_t + (q - b) - r_t) < \sigma^2, \quad \mathbb{P} - a.s., 0 \leq t \leq T, \quad (2)$$

the following results hold:

1. *There exists a set of probability measures $\mathcal{I}$ where for any measure $I \in \mathcal{I}$ the returns of the discounted asset*

$$S_t = \mathbb{E}^I \left[e^{-\int_t^T r_s ds} e^{(q-b)(T-t)} S_T \middle| \mathcal{F}_t \right] \quad (3)$$

are martingales under the measure I .

2. *The option price $\mathcal{O}$ for the insider trader*

$$\mathcal{O}_t = \mathbb{E}^I \left[e^{-\int_t^T r_s ds} \mathcal{O}_T \middle| \mathcal{F}_t \right]$$

is also a martingale under probability $I \in \mathcal{I}$.

3. We call $\mathcal{I}$ the set of information-neutral probabilities. Then, among all the information-neutral measures, there exists a probability measure $\mathcal{I}^*$ that minimizes the quadratic variation of the difference between the hedging portfolio and option value. We call $\mathcal{I}^*$ the minimal variation information-neutral probability measure.
4. The discounted return of the asset is a martingale under measure $\mathcal{I}^*$. Moreover, the dynamic of the spot price under $\mathcal{I}^*$ is described by

$$\frac{dS_t}{S_{t-}} = (r_t - (q - b) - \Lambda(\chi_t)\omega_t)dt + \sigma dW_t^{\mathcal{I}^*} + \Lambda(\chi_t)dM_t, \quad (4)$$

where $\{M_t\}_{0 \leq t \leq T}$ is a doubly stochastic Poisson process whose stochastic intensity process $\{\omega_t\}_{0 \leq t \leq T}$ verifies

$$\omega_t = \phi \left(\frac{\sigma^2 - \Lambda(\chi_t)(\mu_t + (q - b) - r_t)}{\sigma^2 + \phi \Lambda(\chi_t)^2} \right),$$

for each $0 \leq t \leq T$.

5. Furthermore, if the insider trader has entered into a transaction with terminal payoff $\mathcal{O}(S_T)$, we call $\mathcal{O}(t, S, V)$, the price of this option at time t and spot level S and the net volume of shares hold by the large trader V . Then, for all $0 \leq t \leq T$, the option price $\mathcal{O}$ can be computed via the following partial differential equation:

$$\begin{aligned} r_t \mathcal{O} = \frac{\partial \mathcal{O}}{\partial t} + (r_t - (q - b) - \Lambda(\chi_t)\omega_t)S \frac{\partial \mathcal{O}}{\partial S} + \frac{\partial \mathcal{O}}{\partial V} \chi_t \\ + \frac{1}{2} \frac{\partial^2 \mathcal{O}}{\partial S^2} \sigma^2 S^2 + (\mathcal{O}(t, S + S\Lambda(\chi_t), V) - \mathcal{O}(t, S, V))\omega_t. \end{aligned} \quad (5)$$

Proof is available in Appendix A. The demonstration of employs quadratic hedging and local risk-minimization techniques, akin to the approach of Cont and Tankov (2003); Föllmer and Schweizer (1990). Central to this proof is the application of the Girsanov Theorem to both the jump and Brownian processes. This application results in an adjustment of the drift term from μ_t to $r_t - (q - b)$ for the Brownian motion and modifies the intensity of jumps for the Poisson process. It's noteworthy that under the conditions where $\chi_t = 0$ throughout the interval $[0, T]$ or $\phi = 0$, the Equation (5) simplifies to the classical Black-Scholes partial differential equation.

It is worth noting that due to market incompleteness, the condition for the uniqueness of the martingale measure is not satisfied. As further elaborated in Tankov and Voltchkova (2009), when jumps coexist with a Brownian motion, the market becomes incomplete. This means that one cannot perfectly hedge all sources of variability using a conventional hedging portfolio comprising the underlying asset and a risk-free instrument. This also indicates the existence of numerous martingale measures.

We refer to Theorem 1 as defining an *information-neutral* measure because, for an insider, the market-impact process χ_t is not an exogenous source of uncertainty but a process known to them. Theorem 1 extends the classical risk-neutral change of measure to a probability measure under which, conditional on the insider's information, the discounted asset price remains a martingale by which the informational advantage is incorporated and then filtered from the drift of the stock dynamics in Equation (4). This result isolates the informational component of price dynamics from genuine market risk and provides the theoretical foundation for the subsequent pricing and hedging results developed under this framework. By analyzing the market from the insider's perspective, we capture how informational asymmetry alters the martingale pricing structure.

5. Physical Delivery Option for the Large Trader

In this section, we consider a large trader hedging an option with a terminal payoff of cash and physical delivery of shares. Physical delivery is the most prevalent transaction type for stock options, wherein the buyer receives the underlying asset upon expiration. This differs from cash settlement, where the payoff is settled in cash. For further details on option types and settlement methods, we refer to [Cboe Inc. \(2022\)](#).

Definition 1. We define the hedging strategy account by the couplet $\{H_t, V_t\}_{0 \leq t \leq T}$, with $\{V_t\}_{0 \leq t \leq T}$ representing the number of shares holding account, and $\{H_t\}_{0 \leq t \leq T}$ an $\mathbb{F}$ -adapted process representing the cash account, with V_0 as a constant and $H_0 = 0$. The dynamics of these processes are described as follows:

$$\begin{aligned} dV_t &= \chi_t dt, \\ dH_t &= H_t r_t dt + V_t S_t (q - b) dt - S_t dV_t - dC_t, \end{aligned} \quad (6)$$

where the process $\{C_t\}_{0 \leq t \leq T}$ represents the transaction costs process by the large trader and is defined by $C_0 = 0$ and $dC_t = g(\chi_t) dt$, where g is a sufficiently regular function (we provide further details on the possible parametric modeling of g in Section 6.2).

5.1. Hedges Trading P&L

Definition 2. We define the P&L of trading activity, noted as $P\&L_{t,T}^{\text{trading}}$, as the product-sum of V_u , which is the number of shares held at time u , and the discounted asset return over the interval du :

$$P\&L_{t,T}^{\text{trading}} = \int_t^T V_u d\left(e^{-\int_t^u (r_s - (q-b)) ds} S_u\right). \quad (7)$$

Essentially, $P\&L^{\text{trading}}$ represents the profit-and-loss resulting from delta rebalancing trading activities. Under any information-neutral measure $I \in \mathcal{I}$, the expectation is given by

$$\mathbb{E}^I \left[P\&L_{t,T}^{\text{trading}} \mid \mathcal{F}_t \right] = 0, \quad (8)$$

because the discounted asset price is a martingale under I .

However, under the historical measure $\mathbb{P}$ we may have

$$\mathbb{E}^{\mathbb{P}} \left[P\&L_{t,T}^{\text{trading}} \mid \mathcal{F}_t \right] \neq 0. \quad (9)$$

If the $P\&L_{t,T}^{\text{trading}}$ is positive in expectation in the historical measure, it suggests that there are dynamic arbitrage opportunities in the market that can be exploited with the right knowledge. This kind of practice is flagged as “abusive squeeze,” as discussed in Section 2, which manifests when an individual or entity exerts notable control over the supply, demand, or delivery processes of a qualifying or related investment.

5.2. Option and Hedges P&L

Proposition 1. At any time $0 \leq s \leq T$, we define $H_s + V_s S_s$ as the accounting value of the trading activity. The discounted accounting value at the terminal date T satisfies

$$e^{\int_t^T -r_s ds} (H_T + S_T V_T) = (H_t + V_t S_t) + P\&L_{t,T}^{\text{trading}} - \int_t^T e^{-\int_t^u r_s ds} S_u dC_u. \quad (10)$$

Proof is available in Appendix A. The left-hand side of (10) is the discounted value of the portfolio of shares and cash between t and T . The right-hand side decomposes this

into trading P&L net of execution costs, consistent with the economic structure of a large trader's hedging program.

Definition 3. We assume a physical delivery European option with an expiry date of T and a settlement date $T + \tau$. Hereafter, $Q(S_T)$ denotes the terminal number of shares to be either delivered or received, and $\mathcal{H}(S_T)$ is the terminal cash transaction.

It is important to recall that V_T represents the quantity of shares held by the large trader at time T . As such, the collective sum of $Q(S_T)$ and V_T reflects the remaining shares in the large trader's account at the terminal date. For effective hedging, $Q(S_T)$ and V_T should have opposite signs. Specifically, if $Q(S_T)$ is positive, then V_T should be negative, and conversely.

Any residual shares, represented as $Q(S_T) + V_T$, need to be executed between the expiry and settlement dates. This execution cost is denoted as $\bar{C}(Q(S_T) + V_T, \tau)$. Notably, at time T , the cost associated with executing these shares between T and $T + \tau$ is considered known. Hence, $\bar{C}(Q(S_T) + V_T, \tau)$ is $\mathcal{F}_T$ -adapted.

To hedge their position, the large trader maintains a portfolio consisting of H_t in cash and V_t shares at time t , using the same notation as in the previous section. Given the high hedging costs, the large trader may opt to leave some positions unhedged, anticipating favorable market movements to offset any potential issues arising from the unhedged portion. The final profit and loss at maturity seen from t is then given by

$$P\&L_{t,T} = \overbrace{\mathcal{H}(S_T) + S_T Q(S_T)}^{\text{Payoff}} + \overbrace{H_T + S_T V_T}^{\text{Accounting value of hedges}} - \overbrace{S_T \bar{C}(Q(S_T) + V_T, \tau)}^{\text{Terminal execution cost of missing shares}}. \quad (11)$$

Equation (11) has three main components: the option payoff, the accounting value of the hedging portfolio, and the terminal cost of executing any missing shares between the option expiry time T and the settlement date $T + \tau$ for physical delivery options.

By combining Equations (10) and (11), we can obtain the following formulation for the terminal $P\&L_{t,T}$:

$$P\&L_{t,T} = \overbrace{e^{\int_t^T -r_s ds} (\mathcal{H}(S_T) + S_T Q(S_T))}^{\text{Payoff}} + \overbrace{H_t + S_t V_t}^{\text{Hedge portfolio at time } t} + \underbrace{\left(\int_t^T e^{-\int_t^u r_s ds} S_u dC_u + e^{\int_t^T -r_s ds} S_T \bar{C}(Q(S_T) + V_T, \tau) \right)}_{\text{Transaction costs}} + P\&L_{t,T}^{\text{trading}}. \quad (12)$$

On the right-hand side of Equation (12), the first component represents the option payoff, and the second term is the profit and loss from trading activities, as defined in Equation (7), due to changes in the value of the underlying asset during the option's lifetime. The third term represents the temporary transaction costs of the delta-hedging program, while the last term represents the accounting value of the hedging portfolio at expiry.

5.3. Manipulating the Payoff Constraints and Problem Reduction

The trader's objective is to ensure that the P&L is non-negative. Without considering market abuse restrictions, the large trader would therefore seek to maximize expected P&L by solving $\sup_{\{\chi_s\}_{s \geq t}} \mathbb{E}[P\&L_{t,T} | \mathcal{F}_t]$. Two key considerations arise.

First, we must determine under which probability measure the expectation of the P&L should be computed. A natural choice is the historical (real-world) measure, as it reflects actual probabilities. However, as discussed in Section 2, the expected trading P&L, denoted $P\&L_{t,T}^{\text{trading}}$, is not zero under the historical measure. This implies that using the historical

measure could introduce an incentive for manipulative strategies, such as squeeze arbitrage. In contrast, under the information-neutral measure, the expected $P\&L_{t,T}^{\text{trading}}$ is zero, which is desirable because it ensures that neither the trader nor the model has an informational incentive to influence prices for a higher expected trading profit.

Second, to prevent payoff manipulation, we refine the stochastic optimization problem by excluding the payoff component from Equation (12). This modification removes any direct incentive to alter the option's payoff and thereby eliminates one potential source of manipulation.

It is worth noting that these two points are necessary conditions but not sufficient to avoid market abuse. While these measures do not guarantee a market manipulation-proof model, they address essential aspects to minimize the risk of falling into model-induced market abuse.

Specifically, among the information-neutral probability measures, we deploy the minimum variance probability measure, which we developed and defined in detail in the previous section. Consequently, we reduce the problem to Equation (13), to an optimization of the cost function per share J acting as its solution:

$$J(t, V_t, S_t) = \frac{1}{S_t} \inf_{\{\chi_s\}_{s \geq t}} \mathbb{E}^{I^*} \left[e^{-\int_t^T r_s ds} S_T \bar{C}(Q(S_T) + V_T, \tau) + \int_t^T e^{-\int_t^s r_u du} S_s dC_s \middle| \mathcal{F}_t \right]. \quad (13)$$

Equation (13) consists of two parts: the transaction costs of executing missing shares between the expiry date T and the settlement date and the transaction costs of executing shares before the expiry date. Both terms act against each other, and the trader needs to strike a balance between buying shares too early and getting closer to a full delta hedge with fewer missing shares to recover in the future but paying extra costs because of executing the shares too quickly, or leaving its position partially unhedged. In the favorable scenario, the market will move in the trader's direction, and the partial hedge will be closer to the full hedge. Otherwise, the trader will have some time between the expiry date T and the settlement date $T + \tau$ to execute the missing shares. Note that at the expiry date, the number of required shares is fixed and is no longer uncertain.

Theorem 2 establishes the relationship between the cost function J and the partial differential equation (PDE) that governs it. The theorem states that J can be deduced from the PDE expressed in Equation (18), subject to the terminal condition in Equation (19).

Theorem 2. Let $J(t, V_t, S_t)$ denote the cost-per-share value function associated with (13).

1. Define the Radon–Nikodym density and introduce the probability measure $\mathbb{S}$ by

$$\frac{d\mathbb{S}}{d\mathbb{I}^*} \bigg|_{\mathcal{F}_t} = \frac{e^{-\int_0^t r_s ds} S_t}{S_0}. \quad (14)$$

Under $\mathbb{S}$, the Brownian motion and jump intensity satisfy

$$dW_t^{\mathbb{S}} = dW_t^{\mathbb{I}^*} - \sigma dt, \quad \omega_t^{\mathbb{S}} = (1 + \Lambda(\chi_t)) \omega_t, \quad (15)$$

and the stock price dynamics become

$$\frac{dS_t}{S_{t-}} = \left(r_t + \sigma^2 - \Lambda(\chi_t) \omega_t \right) dt + \sigma dW_t^{\mathbb{S}} + \Lambda(\chi_t) dM_t^{\mathbb{S}}, \quad (16)$$

where $M_t^{\mathbb{S}}$ is a doubly stochastic Poisson process with intensity $\omega_t^{\mathbb{S}}$.

Changing numéraire from $\mathcal{I}^*$ to the discounted stock measure $\mathcal{S}$ transforms (13) into the reduced cost-per-share representation

$$J(t, V_t, S_t) = \inf_{\{\chi_s\}_{s \geq t}} \mathbb{E}^{\mathcal{S}} \left[\bar{C}(Q(S_T) + V_T, \tau) + \int_t^T dC_s \middle| \mathcal{F}_t \right]. \quad (17)$$

2. The value function $J(t, V, S)$ satisfies the Hamilton–Jacobi–Bellman equation

$$0 = \frac{\partial J}{\partial t} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 J}{\partial S^2} + (r_t + \sigma^2) S \frac{\partial J}{\partial S} + \mathcal{J}^*(t, V, S), \quad (18)$$

with terminal condition

$$J(T, V, S) = \bar{C}(Q(S) + V, \tau). \quad (19)$$

For any admissible control χ , the running cost contribution is

$$\begin{aligned} \mathcal{J}_\chi(t, V, S) = & -\Lambda(\chi) \omega_t S \frac{\partial J}{\partial S} + \chi \frac{\partial J}{\partial V} \\ & + \left(J(t, S(1 + \Lambda(\chi)), V) - J(t, S, V) \right) \omega_t^{\mathcal{S}} + \frac{dC_t(\chi)}{dt}, \end{aligned} \quad (20)$$

and the optimal running term is

$$\mathcal{J}^*(t, V, S) := \inf_{\chi} \mathcal{J}_\chi(t, V, S). \quad (21)$$

The optimal execution rate is therefore

$$\chi_t^* = \operatorname{arginf}_{\chi} \mathcal{J}_\chi(t, V, S). \quad (22)$$

Proof is available in Appendix A. It is important to note that $\omega_t, \omega_t^{\mathcal{S}}$ are dependent on χ .

6. Numerical Results

In this section, we present numerical results focusing on the large trader hedging an out-of-the-money call option. We analyze the option costs, hedging dynamics for the large trader, and the implications for the insider trader. Additionally, we perform a sensitivity analysis on key model parameters.

6.1. Solving the Partial Differential Equations

To numerically solve the two-dimensional partial differential Equation (18), we can use the dynamic programming approach described in Algorithm 1. First, we initialize the cost function $J(T, S, V)$ to the terminal condition $\bar{C}(Q(S) - V, \tau)$ at the expiration time T . Then, we discretize time into Δt intervals and iterate backward in time from T to 0, updating $\mathcal{J}$ at each time step using Equations (18), (20) and (21). We repeat this process for n iterations to ensure convergence.

Once we have obtained the optimal hedging strategy χ^* , we can use it to solve for the insider's trader option price $\mathcal{O}$ using a finite difference scheme for the partial differential in Equation (5). Finally, we obtain the values of $\mathcal{J}$ and $\mathcal{O}$ at time 0, denoted by $J(0, S_0, 0)$ and $\mathcal{O}(0, S_0, V_0)$, respectively. More details on the finite difference scheme used to solve Equation (18) are provided in the Appendix A.

Algorithm 1 Methodology for solving the large trader's cost function

```

1:  $J(T, S, V) \leftarrow \bar{C}(Q(S) + V, \tau)$ ;
2: Discretize the time to  $\Delta t$  intervals;
3: for  $t = T; t \leftarrow t + \Delta t; t = 0$  do
4:   Using  $J(t + \Delta t, \cdot, \cdot)$  apply Equation (21) and solve  $\chi_t^*$ ;
5:   Using  $\chi^*$ , apply Equation (20) and calculate  $\mathcal{J}^*(t + \Delta t, \cdot, \cdot)$ ;
6:   Using  $\mathcal{J}^*(t + \Delta t, \cdot, \cdot)$ , solve Equation (18) and obtain  $J(t, \cdot, V)$ ;
7:   for  $iter = 0; iter \leftarrow iter + 1; iter < n = \text{numiter}$  do
8:     Using  $J(t, \cdot, \cdot)$  apply Equation (21) and solve  $\chi_t^*$ ;
9:     Using  $\chi^*$ , apply Equation (20) and recalculate  $\mathcal{J}^*(t + \Delta t, \cdot, \cdot)$ ;
10:    Using  $\mathcal{J}^*(t + \Delta t, \cdot, \cdot)$ , solve Equation (18) and obtain  $J(t, \cdot, V)$ ;
11:   end for
12:   Given  $\chi_t^*$ , solve the insider's trader option price  $\mathcal{O}$  using a finite difference scheme
    of the partial differential in Equation (5);
13: end for
14: Obtain  $J(0, S_0, 0)$  and  $\mathcal{O}(0, S_0, V_0)$ 

```

Note that, to solve $\mathcal{O}$ in Equation (5) as well as in Equation (18), we employ a numerical scheme similar in structure to those used for path-dependent (e.g., Asian) options. The algorithm executes a set of two-dimensional finite-difference schemes in parallel—each corresponding to a discretized level of the grid, the net volume of shares held by the large trader. In each step, the PDE for $\mathcal{O}(t, S, \cdot)$ is first solved independently across all grid-levels, and a subsequent semi-Lagrangian remapping stage integrates the impact of changes across the grid. This approach alternates between the diffusion dynamics in (t, S) and the advection dynamics within the grid. A detailed presentation of this well-established method can be found in Rogers and Shi (1995).

From a computational point of view, the overall complexity of the numerical scheme scales linearly in each of the three discretization dimensions. At every time step, the implicit finite-difference solve in the (t, S) dimension requires, for each of the N_V values of V , the construction of a tri-diagonal matrix and the solution of a tri-diagonal linear system of size $N_S \times N_S$. Since a tridiagonal matrix algorithm is used, each such solve costs $\mathcal{O}(N_S)$ operations, yielding a total PDE cost of $\mathcal{O}(N_S N_V)$ per time step. The optimization step that computes κ^* involves a directional search at each grid node (S_i, V_j) ; because this search is local and terminates after a small number of updates, its total cost per time step remains $\mathcal{O}(N_S N_V)$. The semi-Lagrangian reorganization of the grid, which updates the solution according to the shift induced by V , also loops once over the entire grid and therefore contributes another $\mathcal{O}(N_S N_V)$ operations per step. Summing all stages over the N_T time steps of the backward induction, the total computational cost of the method is $\mathcal{O}(N_S N_V N_T)$, which reflects linear scaling in the spot discretization, the volume discretization, and the number of time steps.

6.2. Choice of Transaction Costs Parametrization

We first need to define a parameterization for the transaction costs function. Let $\mathcal{C}_t$ be a stochastic process with the dynamics defined by Equation (23). The process $\{\mathcal{C}_t\}_{0 \leq t \leq T}$ represents the transaction costs beared by the large trader :

$$\begin{aligned} \mathcal{C}_0 &= 0 \\ d\mathcal{C}_t &= \eta \left(\frac{\chi_t}{\chi_t + A} \right)^\alpha \chi_t dt, \end{aligned} \quad (23)$$

where $A \in \mathbb{R}$ is the (annualized) average volume, $\eta \in \mathbb{R}^+$ is the impact factor, and $\alpha > 0$ is the exponential factor. The form is consistent with the empirical microstructure literature, in

which marginal execution cost grows sublinearly with trade size, commonly approximated by square-root or other concave laws (Almgren et al., 2005; Grinold & Kahn, 2000; Zarinelli et al., 2015). For longer periods between the expiry and settlement date, an optimal purchase/liquidation algorithm can be applied, as explored in the literature (e.g., Almgren and Chriss (2001)). However, in this study, we use a constant rate program for simplicity. In this case, the expected value of function $\bar{C}(x, \tau)$ —the terminal cost of executing x at time τ between the expiry and settlement—can be expressed as Equation (24) below.

Proposition 2. Assume the trader needs to execute X shares within a period of τ and chooses to do so at a constant execution speed of $\frac{X}{\tau}$.³ Moreover, we assume that there are no dividends and lending costs during the time period from T to $T + \tau$. Under these conditions, the expected value of the cost under an information-neutral measure $I \in \mathcal{I}$ is described by the following equation:

$$\mathbb{E}^{\mathcal{I}}[\bar{C}(x, \tau) | \mathcal{F}_T] = \eta \times \frac{\left(\frac{x}{\tau}\right)^{\alpha}}{\left(\frac{x}{\tau} + A\right)^{\alpha}} x. \quad (24)$$

Proof is available in Appendix A. Overall, our cost model allows to capture the trade-off between the size of the trade (x) and the time until the settlement date (τ). Specifically, as x increases, so does the cost of execution, but at a decreasing rate governed by the exponent α . In contrast, as τ increases, the cost of execution decreases due to the larger time frame for trading. These properties parallel empirical observations that execution cost rises with urgency and notional size, but at a diminishing marginal rate.

6.3. Out of the Money Call Option

In this section, we present numerical results for a large trader who is the seller of a call option defined in Table 1.

Table 1. Large Trader’s option.

T	τ	Call Strike	Notional
120 days	15 days	110	1×10^6

We also make the hypothetical parameter assumptions as per Table 2. These assumptions, though hypothetical, fall within the realistic ranges commonly used in financial modeling and empirical studies. The annual drift rate, $\mu = 0.02$, is within the range of U.S. annual returns of small-cap stocks as documented by Damodaran (2025). The volatility parameter, $\sigma = 0.25$, aligns with the historical range of the CBOE Volatility Index (VIX), which typically fluctuates between 10 and 20 during stable markets and exceeds 30 during periods of high market stress VIX (2025). The risk-free rate, $r = 0.01$, reflects the low-interest-rate environment observed in the U.S. market during 2010–2020. Also, further in this section, we will provide sensitivity testing to the choice of parameters ϕ , ν , and λ . Finally, for simplicity, we assume that the stock does not pay any dividends ($q = 0.0$) and its lending yield is zero ($b = 0.0$).

Table 2. Market parameters assumptions.

μ	σ	η	ϕ	λ	S_0	A	r	α	q	b
0.02	0.25	0.1	20.0	0.2 to 2.0	100	5×10^6	0.01	2.0	0.0	0.0

Figure 4 depicts the terminal cost function J for the large trader’s call option. The graph illustrates different scenarios depending on the final stock price and the number of shares V_T to be delivered at expiration. When the option is in the money, i.e., the stock

price is above the strike price, the cost function is zero since the large trader delivers the shares at no additional cost. Conversely, when the option is out of the money, i.e., the stock price is below the strike price, the large trader needs to purchase the notional amount of shares before the settlement date, leading to higher costs. As shown in the graph, the cost function varies across different corners but gradually flattens as we move backward in time, as observed in Figure 5. These results are consistent with our cost model in Equation (19), which assumes a linear relationship between the cost and the spot price of the stock. These properties parallel empirical findings that market impact increases with order size and urgency, but with a concave (diminishing marginal) dependence (Almgren et al., 2005; Zarinelli et al., 2015).




Figure 4. The terminal cost function $J \times S$ for a hedged option. V_T shall be read as a multiple of the notional. Colors encode the magnitude of the terminal cost function on the surface, with a continuous monotonic mapping from lower to higher values.




Figure 5. The cost function $J \times S$ at time 0 for a hedged call option. Data source: Matlab simulation. Colors encode the magnitude of the cost function on the surface, with a continuous monotonic mapping from lower to higher values.

Figures 6 and 7 show the optimal amount of shares to be executed (χ^*) at time 0 and just before expiry, respectively. Our findings suggest that the optimal χ^* is much lower than the Black-Scholes Delta due to the high cost of a full hedge strategy. As shown in Figure 8, the difference in Delta between the two strategies is significant. The large trader executes

only a small percentage of χ^* at time 0, whereas the Black-Scholes Delta is 100% for the same spot levels. However, as the expiry date approaches, the large trader's optimal strategy gets closer to the Black-Scholes Delta, as illustrated in Figure 9. This muted hedging intensity aligns with empirical evidence that large institutional traders hedge less aggressively than theoretical deltas because of market impact constraints (Abergel & Loeper, 2017; Kraft & Kühn, 2011).




Figure 6. The optimal shares strategy function at time 0 for a call option. Data source: Matlab simulation. Colors encode the magnitude of the cost function on the surface, with a continuous monotonic mapping from lower to higher values.




Figure 7. The optimal shares strategy function at time $T - dt$ for a call option. Data source: Matlab simulation. Colors encode the magnitude of the cost function on the surface, with a continuous monotonic mapping from lower to higher values.

Examining the impact of the large trader's hedging strategy on the insider trader's option price, as defined in Table 3, provides valuable insights. Figure 10 compares the insider trader's option price with the Black-Scholes model for the same call option. The graph highlights the option's increased convexity around the money and the higher prices than the Black-Scholes model when the option is in the money, surpassing even the intrinsic value. This phenomenon is due to the directional market impact of the large trader. If the spot price is in the money, the large trader will purchase shares, driving the spot price

deeper into the money. The insider trader's option price, therefore, reflects the market impact of the large trader, resulting in a higher option price than predicted by the standard Black-Scholes model. This upward deviation is consistent with empirical price-pressure effects documented in transaction-level studies (Hasbrouck, 1991).




Figure 8. Comparison between BS delta and optimal shares for the large trader at inception, assuming $V_0 = 0$.




Figure 9. Comparison between BS delta and optimal shares for the large trader at expiry, assuming $V_T = 0$. Data source: Matlab simulation.

We would also like to highlight the fact that the insider option price is sensitive to parameters ϕ , η , λ as shown in Figures 11 and 12. Economically speaking, an increase in the market impact parameters ϕ and λ results in a higher call option price for the insider trader (for the option defined in Table 1). Recall that λ is the market impact factor, and ϕ is the frequency of the market impact function. Higher values of these two parameters result in a higher market impact function and, accordingly, more divergence from the Black & Scholes price. These sensitivities resemble empirical relationships between participation rate, execution frequency, and permanent impact reported in equity microstructure (Almgren et al., 2005), with the additional touch that in here we study the price of a convex payoff, i.e. a call option (Cetin et al., 2006).




Figure 10. Comparison between Black and Scholes price and insider trader price for a call option. Data source: Matlab simulation.




Figure 11. Call option price for insider trader at $t = 0$, (Assuming the large trader holds no shares at the time ($V_0 = 0$), and $\phi = 20$), sensitivity to λ parameter. Dots correspond to simulated option prices obtained from the numerical model, while the solid curves are interpolations between the simulated points and are shown for visual guidance only.




Figure 12. Call option price for insider trader at $t = 0$, (Assuming the large trader holds no shares at the time ($V_0 = 0$), and $\lambda = 0.1$), sensitivity to ϕ parameter. Dots correspond to simulated option prices obtained from the numerical model, while the solid curves are interpolations between the simulated points and are shown for visual guidance only.

Table 3. Small trader's option.

T	Call Strike
120 days	110

7. Discussion, Limitations and Further Research

While the model presented here is conceptually sound, it is worth acknowledging several shortcomings and directions for further research. These limitations also suggest how the framework may speak to broader questions in market microstructure, particularly concerning the interaction between informed trading, mechanical price impact, and liquidity fluctuations.

First and foremost, we have not proven mathematically that the model presented in this paper is non-manipulative. This challenge exists because, as of the date of this paper, market manipulation—and specifically payoff manipulation—has no formal mathematical definition to refer from which such a proof could be derived. This is potentially because any mathematical definition of manipulation would need to depend on the legal definition. The legal definition, however, has turned out to be subject to judicial interpretation, as illustrated by Fletcher (2018), which provides examples of differing legal interpretations. Curato et al. (2017) also note that defining manipulation in a mathematically consistent and empirically valid way remains an open problem. A similar challenge arises in the microstructure literature, where empirical evidence of settlement distortions and price-pressure effects (Griffin & Shams, 2018; Ni et al., 2005) has not yet produced a unified theoretical definition of manipulation. In short, before we can claim that a model is non-manipulative, two prerequisites are needed: (i) a legal definition that minimizes subjective interpretation and (ii) a general mathematical theory of market manipulation for derivatives. Developing such a general theoretical framework is outside the scope of this paper, but with the proposed method—which models the problem from the standpoint of an informed trader—we can assert that, through this framework, the large trader becomes effectively blind to the market impact of its own hedging activity. This method may prove useful for future research on formalizing market manipulation mathematically. In particular, the information-neutral construction proposed here may serve as a building block for microstructure models that explicitly separate informational asymmetry from mechanically induced price impact, in the spirit of Glosten and Milgrom (1985) and Hasbrouck (1991).

For completeness, we note that our market-impact specification also relates to a second strand of the market microstructure literature concerned with dynamic no-arbitrage and round-trip trading strategies. Classical results by Huberman and Stanzl (2004) and Gatheral (2010) show that permanent price-impact functions must satisfy linearity conditions to preclude profitable round trips. Furthermore, Guéant (2013) proved that other forms of market impact, which are not linear, also conform with the no-arbitrage condition. Although we do not include a full technical development here for reasons of conciseness and to maintain focus on the information-neutral framework, we verified that—under linear permanent impact—the Poisson permanent-impact specification used in this paper is consistent with these no-dynamic-arbitrage principles in the sense that round-trip trading does not yield positive expected P&L under certain condition.

Also, some of the intricacies of our model relate to the specification of the market-impact function, as summarized below:

- The permanent-impact model presented in Equation (1) operates at *random times*, contrasting with deterministic formulations such as Almgren and Chriss (2001).⁴ This randomness is plausible in markets with stochastic liquidity, as proposed in recent studies (Almgren, 2012; Barger & Lorig, 2019; Klöck, 2012; Ma et al., 2020), and our

model follows this line of research while leaving empirical validation of stochastic market-impact functions for future research.

Another advantage of a stochastic market-impact specification is that this component of the underlying spot process becomes part of its *quadratic-variation* term. If the impact were deterministic, it would enter the model at the same order as the drift. By introducing randomness, however, the market-impact term contributes directly to the volatility of the underlying asset. This modeling choice motivates a potential line of investigation: if market volatility arises primarily from the market impact of demand–supply imbalances, then the impact function itself should be of the same stochastic order as volatility. This idea is consistent with empirical evidence that short-term volatility is closely tied to order-flow imbalances and liquidity fluctuations (Bouchaud et al., 2004; Hasbrouck & Seppi, 2001). A natural extension would be to examine whether the information-neutral decomposition proposed here is borne out in transaction-level data.

- Apart from the stochastic nature of market impact, our model assumes a Poisson jump process. This choice introduces discontinuities, representing periods with no jumps as times when market depth is sufficient to absorb trades without permanent impact. In practice, however, large transactions rarely exhibit such intervals. Alternative continuous-time formulations, such as exponential or Lévy-based specifications, could capture this behavior more realistically. Exploring these alternatives constitutes a promising avenue for future mathematical research.
- Another limitation of a pure-jump specification is that when the jump intensity ϕ is large—as in our model, where frequent market-impact events are intended—the Poisson process begins to approximate a one-sided continuous motion. This behavior increases numerical error in the PDE simulations, since the jump component is handled separately from the tridiagonal scheme used to solve the drift–diffusion part of the equation, see Abeille et al. (2023). From a microstructure perspective, understanding such numerical effects may be relevant for modelling high-frequency impact dynamics where order flow arrives in rapid bursts (Curato et al., 2017).
- Another assumption in our market impact model is the absence of a correlation between the stochastic liquidity process $\{N_t\}_{0 \leq t \leq T}$ and other external sources of spot volatility $\{W_t\}_{0 \leq t \leq T}$. In practice, days with higher volatility are often associated with greater challenges in executing stock trades without significantly impacting the market. A more realistic model might assume some degree of correlation between $\{W_t\}_{0 \leq t \leq T}$ and $\{N_t\}_{0 \leq t \leq T}$. This correlation is consistent with empirical evidence that periods of elevated volatility are associated with reduced liquidity supply, as documented in Chordia et al. (2001).

Lastly, one may observe that this paper provides a theoretical price for the insider trader but does not explicitly define a price for the large trader. In martingale asset-pricing theory, the price of a contingent claim is obtained as the expected value of its discounted payoff under a martingale measure. The martingale measure used here is the *minimal-variance information-neutral* measure, constructed based on the insider's hedge δ_t , which is accessible to the insider because such a trader can execute it.

The situation differs for the large trader. The large trader hedges through the process $dV_t = \chi_t dt$, not δ_t . Therefore, the trading strategy of the large trader does not fully replicate the payoff, and no arbitrage-free “large-trader price” can be derived. Instead, we evaluated the expectation of temporary cost function, denoted J_t , under $\mathbb{I}^*$ as perceived by the insider, and determine an admissible hedging policy χ_t based on that. The existence of an information-neutral martingale measure that is also accessible to the large trader

remains an open question and is left for future research. If such measure exists, it can be then used also for obtaining the admissible hedging policy χ_t .⁵

8. Conclusions

In conclusion, this paper examines the implications of a large trader's activities on option pricing and hedging. We addressed an important but overlooked problem in the literature: how a large trader's knowledge of their own trading activity can influence their actions. Therefore, a key contribution of this paper is its methodology to disentangle this knowledge from other factors influencing options' hedging strategies. This separation between informational and mechanical components of order flow places the analysis squarely within the market microstructure literature, where this distinction plays a central role. To achieve this, we introduced the concept of an insider trader, who has the same information but does not impact the market.

Our analysis established the existence of information-neutral probability measures, under which the discounted asset prices are martingales for the insider trader. Using these information-neutral measures, along with necessary conditions for mitigating market manipulation, we derived a cost optimization function to establish a hedging policy for the large trader. This framework allows us to develop a hedging policy that accounts for transaction costs and market impact while taking into account some conditions for non-manipulative practices. These findings offer a practical and effective approach to pricing derivatives under market impact, which can be readily implemented by practitioners. In this respect, the framework contributes to the microstructure literature on large-trader models, permanent impact, and manipulation-resistant execution.

The paper concludes with numerical results that showcase the optimal delta-hedging strategy for an out-of-the-money call option. Our analysis indicates that the model proposes a more relaxed delta-hedging strategy compared to the Black and Scholes model. Additionally, the price of the call option for the insider trader is observed to be higher under market impact than the Black and Scholes model. These numerical patterns align with empirical findings in market microstructure, on price-pressure effects, and drift induced by directional order flow, and complement how this could impact option prices.

Although the analysis was developed under the assumption of a single large trader, several of our findings extend to environments with multiple strategic traders, provided that the insider trader can observe or infer their aggregate impact. Likewise, the construction of the information-neutral measure is not tied to the specific hedging objective studied here and can be adapted to other trading strategies. Future research may explore these extensions in greater depth, especially in settings where systemic behavior arise among different classes of large traders pursuing non-derivative trading strategies. In addition, the framework suggests several further directions for market microstructure research, including extension to multi-agent settings where several large or informed traders interact through liquidity constraints.

Author Contributions: Conceptualization, B.A. and K.B.; Methodology, B.A., K.B. and A.E.; Software, B.A.; Validation, B.A., K.B. and A.E.; Formal analysis, B.A.; Investigation, B.A., K.B. and A.E.; Data curation, B.A.; Writing—original draft, B.A.; Writing—review & editing, B.A., K.B. and A.E.; Visualization, B.A.; Supervision, K.B. and A.E. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: There is no conflict of interest related to this paper from the authors.

Appendix A. Proofs

Proof of Theorem 1. By Itô formula applied to a process with jumps, the option price dynamics are given by

$$d\mathcal{O}_t = \frac{\partial \mathcal{O}}{\partial t} dt + \frac{\partial \mathcal{O}}{\partial S} (\mu_t dt + \sigma dW_t) S_{t-} + \frac{1}{2} \frac{\partial^2 \mathcal{O}}{\partial S^2} \sigma^2 S_{t-}^2 dt + \frac{\partial \mathcal{O}}{\partial V} \chi_t dt + \left[\mathcal{O}(t, S_{t-}(1 + \Lambda(\chi_t)), V_t) - \mathcal{O}(t, S_{t-}, V_t) \right] dN_t.$$

where $\mathcal{O}(t, S_{t-}(1 + \Lambda(\chi_t)), V_t) - \mathcal{O}(t, S_{t-}, V_t)$ is the jump in the price of the option due to large traders' activity χ_t , and $S_t(\mu_t dt + \sigma dW_t)$ is the continuous part of the spot movement.

On the other hand, the insider trader's option hedging portfolio at time t is given by

$$P_t = \delta_t S_t + \overbrace{(P_t - \delta_t S_t)}^{\text{Cash Account}},$$

where the insider trader has δ_t in shares and $P_t - \delta_t S_t$ in cash amount. Given that the notional is small, the insider trader—contrary to the large trader—can hedge its position without additional transaction costs. The dynamics of the portfolio P_t are as follows:

$$dP_t = \delta_t (dS_t + S_{t-} q dt - S_{t-} b dt - S_{t-} r_t dt) + r_t P_t dt.$$

In our incomplete market, the option cannot be perfectly hedged by a self-financing portfolio. Therefore, we search for a hedging strategy that satisfies the following properties:

- $\delta_t = \arg \inf_{\delta} d\langle (P - \mathcal{O}) \rangle_t$: If the market is complete, it is possible to find δ_t such that the quadratic variation $d\langle (P - \mathcal{O}) \rangle_t$ is always zero. However, this is impossible due to the jump component of the market impact. One can only find a strategy that minimizes the quadratic variation, but it will not neutralize it.

The stochastic part of $\mathcal{O}_t - P_t$ can be formulated as:

$$d(\mathcal{O}_t - P_t) = \dots dt + S_{t-} \left(\frac{\partial \mathcal{O}}{\partial S} - \delta_t \right) \sigma dW_t + \left[\mathcal{O}(t, S_{t-}(1 + \Lambda(\chi_t)), V_t) - \mathcal{O}(t, S_{t-}, V_t) - \delta_t S_{t-} \Lambda(\chi_t) \right] dN_t.$$

Therefore, the quadratic variation of $\mathcal{O}_t - P_t$ reads as follows:

$$d\langle P_t - \mathcal{O}_t \rangle = \left(\frac{\partial \mathcal{O}}{\partial S} - \delta_t \right)^2 S_{t-}^2 \sigma^2 dt + \phi S_{t-}^2 \Lambda(\chi_t)^2 \left[\delta_t - \frac{\mathcal{O}(t, S_{t-}(1 + \Lambda(\chi_t))) - \mathcal{O}(t, S_{t-})}{S_{t-} \Lambda(\chi_t)} \right]^2 dt.$$

The insider trader minimizing the quadratic variation will obtain the following formula:

$$\delta_t S_{t-} = \frac{\sigma^2}{\sigma^2 + \phi \Lambda(\chi_t)^2} \frac{\partial \mathcal{O}}{\partial S} S_{t-} + \frac{\phi \Lambda(\chi_t)}{\sigma^2 + \phi \Lambda(\chi_t)^2} (\mathcal{O}(t, S_{t-}(1 + \Lambda(\chi_t)), V_t) - \mathcal{O}(t, S_{t-}, V_t)). \quad (\text{A1})$$

- $\mathbb{E}^{\mathbb{P}}[dP_t|\mathcal{F}_t] = \mathbb{E}^{\mathbb{P}}[dO_t|\mathcal{F}_t]$: The average (first order) option price movement is fully hedged by the portfolio.

We can derive the following partial differential equation:

$$\frac{\mathbb{E}^{\mathbb{P}}[dP_t|\mathcal{F}_t]}{dt} = r_t P_t + (\mu_t + q - b - r_t + \Lambda(\chi_t)\phi)\delta_t S_{t-}.$$

On the other hand, we have

$$\begin{aligned} \frac{\mathbb{E}^{\mathbb{P}}[dO_t|\mathcal{F}_t]}{dt} &= \frac{\partial O}{\partial t} + \frac{\partial O}{\partial S} S_{t-} \mu_t + \frac{1}{2} \frac{\partial^2 O}{\partial S^2} \sigma^2 S_{t-}^2 + \frac{\partial O}{\partial V} \chi_t \\ &\quad + \phi[\mathcal{O}(t, S_{t-}(1 + \Lambda(\chi_t)), V_t) - \mathcal{O}(t, S_{t-}, V_t)]. \end{aligned}$$

Combining the two equations above with Equation (A1) we obtain that:

$$\begin{aligned} r_t P_t &= \frac{\partial O}{\partial t} + \left(\mu_t - \frac{\sigma^2(\mu_t + q - b - r_t + \Lambda(\chi_t)\phi)}{\sigma^2 + \phi\Lambda(\chi_t)^2} \right) S_t \frac{\partial O}{\partial S} + \frac{1}{2} \frac{\partial^2 O}{\partial S^2} \sigma^2 S_t^2 + \frac{\partial O}{\partial V} \chi_t dt \\ &\quad + \phi \left(1 - \frac{\Lambda(\chi_t)(\mu_t + q - b - r_t + \Lambda(\chi_t)\phi)}{\sigma^2 + \phi\Lambda(\chi_t)^2} \right) [\mathcal{O}(t, S_{t-}(1 + \Lambda(\chi_t)), V_t) - \mathcal{O}(t, S_{t-}, V_t)]. \end{aligned} \quad (A2)$$

Let's define the variable ω_t , which represent a jump intensity, as

$$\omega_t = \phi \frac{\sigma^2 - \Lambda(\chi_t)(\mu_t + q - b - r_t)}{\sigma^2 + \phi\Lambda(\chi_t)^2}. \quad (A3)$$

We also assume the regularity condition (2) holds. This condition is necessary as otherwise ω_t will not be positive almost surely, and the Feynman-Kac theorem could not be applied as the intensity of the Poisson process will not be positive.

- Since this is not a self-financing portfolio, the trader can intervene at each time step by injecting or depleting cash to ensure that the hedging portfolio's value matches the theoretical value of the option. Mathematically, the trader will inject $P_t - O_t$ in cash into the hedging portfolio, ensuring that $P_t = O_t$. Therefore, due to these cash injections, the hedging portfolio's value is continuously adjusted to equal the option price. Based on the definition of ω_t and the assumption that the hedging portfolio is adjusted to satisfy $P_t = O_t$, we obtain:

$$\begin{aligned} r_t O &= \frac{\partial O}{\partial t} + (r_t - q + b - \Lambda(\chi_t)\omega_t) S_{t-} \frac{\partial O}{\partial S} + \frac{\partial O}{\partial V} \chi_t dt \\ &\quad + \frac{1}{2} \frac{\partial^2 O}{\partial S^2} \sigma^2 S_{t-}^2 + \omega_t [\mathcal{O}(t, S_{t-}(1 + \Lambda(\chi_t)), V_t) - \mathcal{O}(t, S_{t-}, V_t)]. \end{aligned} \quad (A4)$$

Having the partial differential equation above, the most common technique would be to use the Feynman-Kac formula. However, the stochastic process derived is not an exponential Levy, so one has to be careful in using the Feynman-Kac technique for non-exponential-Levy processes. Using the Feynman-Kac formula, the equivalent stochastic equation for the spot price S_t under a probability measure that we call minimal-variance information neutral probability and note it as $\mathcal{I}^*$ is given by

$$\frac{dS_t}{S_{t-}} = (r_t + b - q - \chi_t \Lambda(\chi_t) \omega_t) dt + \sigma dW_t + \Lambda(\chi_t) dM_t, \quad (A5)$$

where $\{M_t\}_{0 \leq t \leq T}$ is a doubly stochastic Poisson process with stochastic intensity process $\{\omega_t\}_{0 \leq t \leq T}$. Also, the following equations hold under the I^* probability:

$$S_t = \mathbb{E}^{I^*} \left[e^{-\int_t^T r_s ds} e^{(q-b)(T-t)} S_T \middle| \mathcal{F}_t \right] \quad (\text{A6})$$

$$\mathcal{O}(t, S_t, V_t) = \mathbb{E}^{I^*} \left[e^{-\int_t^T r_s ds} \mathcal{O}(T, S_T, V_T) \middle| \mathcal{F}_t \right]. \quad (\text{A7})$$

To validate our use of the Feynman-Kac technique, we prove that having the stochastic process of Equation (A5) and the function $\mathcal{O}$ of Equation (A7) would result in the partial differential equation in (A4).

On the one hand, based on Equation (A7), we have

$$\mathbb{E}^{I^*} [d\mathcal{O}_t | \mathcal{F}_t] = r_t \mathcal{O}_t dt. \quad (\text{A8})$$

On the other hand, applying Itô's lemma to the function $\mathcal{O}$:

$$\begin{aligned} \mathbb{E}^{I^*} [d\mathcal{O}_t | \mathcal{F}_t] &= \frac{\partial \mathcal{O}}{\partial t} dt + \frac{\partial \mathcal{O}}{\partial S} \mathbb{E}^{I^*} [dS_t^C] + \frac{\partial^2 \mathcal{O}}{\partial S^2} \mathbb{E}^{I^*} [d\langle S^C \rangle_t] + \frac{\partial \mathcal{O}}{\partial V} \chi_t dt \\ &\quad + S_t \mathbb{E}^{I^*} [dM_t] \left[\mathcal{O}(t, S_{t-}(1 + \Lambda(\chi_t)), V_t) - \mathcal{O}(t, S_{t-}, V_t) \right]. \end{aligned} \quad (\text{A9})$$

where S^C is the continuous part of the diffusion of S in the dynamics (A5). Now, expanding the dynamics of (A5) in (A8) and (A9), we obtain again the partial differential Equation (A4). Therefore, our usage of the Feynman-Kac technique was valid.

Note that we have proved the existence of at least one element in the set of information-neutral probability measures, which is the I^* . The existence of I^* is sufficient proof for the existence of the set of information neutral measures $\mathcal{I}$ and the fact that it is not empty. $\square$

Proof of Proposition 1. Let $b_u = e^{\int_t^u -r_s ds}$ denote the discount factor at time t with maturity T . or the zero coupon bond price, $\bar{H}_u$ the discounted cash $\bar{H}_u = b_u H_u$, and $\bar{S}_u = b_u S_u$ the discounted asset. Therefore:

$$\begin{aligned} d\bar{H}_u &= -r_u \bar{H}_u du + b_u dH_u = V_u \bar{S}_u (q - b) du - \bar{S}_u dV_u - b_u S_u dC_u \\ d\bar{S}_u &= -r_u \bar{S}_u du + b_u dS_u. \end{aligned} \quad (\text{A10})$$

On the other hand, given the fact that $\{V_t\}_{0 \leq t \leq T}$ —defined in Equation (6)—is a finite variation process, we do have the following equation:

$$d(V_u \bar{S}_u) = V_u d\bar{S}_u + \bar{S}_u dV_u. \quad (\text{A11})$$

Therefore, we have:

$$d(V_u \bar{S}_u + \bar{H}_u) = V_u d\bar{S}_u + V_u \bar{S}_u (q - b) du - b_u dC_u(\chi_u). \quad (\text{A12})$$

Let's now call $\hat{S}_t = e^{\int_t^u (q-b)ds} S_u$, then one can write:

$$d(V_u \bar{S}_u + \bar{H}_u) = V_u d\hat{S}_u - b_u dC_u. \quad (\text{A13})$$

Taking the integration, we obtain:

$$e^{-\int_t^T r_s ds} (S_T V_T + H_T) = H_t + S_t V_t + \int_t^T V_u d \left(e^{-\int_t^u (r_s - (q-b))ds} S_u \right) - \int_t^T e^{-\int_t^u r_s ds} dC_u,$$

which ends the proof of Proposition 1. $\square$

Proof of Theorem 2. We assume $q - b = 0$. Under the measure $\mathcal{I}^*$, the stock price S_t satisfies

$$\frac{dS_t}{S_{t-}} = r_t dt + \sigma dW_t^{\mathcal{I}^*} + \Lambda(\chi_t) d\tilde{M}_t^{\mathcal{I}^*},$$

where $\tilde{M}_t^{\mathcal{I}^*} = M_t^{\mathcal{I}^*} - \int_0^t \omega_s ds$ is the $\mathcal{I}^*$ —compensated doubly stochastic Poisson process with intensity ω_t .

The discounted stock

$$\tilde{S}_t := e^{-\int_0^t r_s ds} S_t$$

is a strictly positive local martingale under $\mathcal{I}^*$. Hence

$$Z_t := \frac{\tilde{S}_t}{\tilde{S}_0} = \frac{e^{-\int_0^t r_s ds} S_t}{S_0}$$

is a positive $\mathcal{I}^*$ —martingale, which defines the measure $\mathbb{S}$ via (14). By the Girsanov–Esscher theorem for semimartingales with jumps, the Brownian motion and intensity under $\mathbb{S}$ are given by (15), and substituting these into the $\mathcal{I}^*$ —dynamics yields:

$$\frac{dS_t}{S_{t-}} = \left(r_t + \sigma^2 - \Lambda(\chi_t) \omega_t \right) dt + \sigma dW_t^{\mathbb{S}} + \Lambda(\chi_t) dM_t^{\mathbb{S}},$$

where $M_t^{\mathbb{S}}$ has intensity $\omega_t^{\mathbb{S}} = (1 + \Lambda(\chi_t)) \omega_t$.

Starting from the original definition of J in (13) and using the standard change-of-numéraire identity with the stock as numéraire,

$$\mathbb{E}^{\mathcal{I}^*} \left[e^{-\int_t^T r_s ds} S_T f(S_T) \mid \mathcal{F}_t \right] = S_t \mathbb{E}^{\mathbb{S}} [f(S_T) \mid \mathcal{F}_t],$$

we divide the original objective by S_t and obtain the representation (17) under the stock-numéraire measure $\mathcal{S}$:

$$J(t, V_t, S_t) = \inf_{\{\chi_s\}_{s \geq t}} \mathbb{E}^{\mathcal{S}} \left[\tilde{C}(Q(S_T) + V_T, \tau) + \int_t^T dC_s \mid \mathcal{F}_t \right].$$

Now fix t and consider a small time step $dt > 0$. From (17), we have:

$$J(t, V_t, S_t) = \inf_{\chi_t} \left\{ \mathbb{E}^{\mathcal{S}} [J(t+dt, V_t + \chi_t dt, S_{t+dt}) \mid \mathcal{F}_t] + dC_t(\chi_t) \right\}, \quad (\text{A14})$$

where $V_{t+dt} = V_t + \chi_t dt$.

We apply Itô's formula with jumps to $J(t, V_t, S_t)$ under the stock dynamics (16). Over $[t, t+dt]$ and for a fixed control χ_t we obtain

$$\begin{aligned} J(t+dt, V_t + \chi_t dt, S_{t+dt}) &= J(t, V_t, S_t) + J_t dt + J_S dS_t + \frac{1}{2} J_{SS} (dS_t)^2 \\ &\quad + J_V \chi_t dt + [J(t, S_{t-}(1 + \Lambda(\chi_t)), V_t) - J(t, S_{t-}, V_t)] dM_t^{\mathbb{S}}, \end{aligned}$$

where we use the shorthand $J_t = \partial J / \partial t$, $J_S = \partial J / \partial S$, $J_{SS} = \partial^2 J / \partial S^2$, and $J_V = \partial J / \partial V$. From (16) and taking $\mathbb{E}^{\mathcal{S}}[\cdot \mid \mathcal{F}_t]$ in (A14), we obtain

$$\begin{aligned} \mathbb{E}^{\mathcal{S}} [J(t+dt, V_t + \chi_t dt, S_{t+dt}) \mid \mathcal{F}_t] &= J(t, V_t, S_t) + J_t dt + \frac{1}{2} \sigma^2 S^2 J_{SS} dt \\ &\quad + (r_t + \sigma^2 - \Lambda(\chi_t) \omega_t) S J_S dt + \chi_t J_V dt \\ &\quad + [J(t, S(1 + \Lambda(\chi_t)), V) - J(t, S, V)] \omega_t^{\mathbb{S}} dt, \end{aligned}$$

where we have written $S = S_t$ and $V = V_t$ for brevity.

Substituting this into (A14), subtracting $J(t, V_t, S_t)$ from both sides, dividing by dt , and letting $dt \downarrow 0$ yields

$$0 = J_t + \frac{1}{2}\sigma^2 S^2 J_{SS} + (r_t + \sigma^2) S J_S + \inf_{\chi} \left\{ -\Lambda(\chi) \omega_t S J_S + \chi J_V + [J(t, S(1 + \Lambda(\chi)), V) - J(t, S, V)] \omega_t^S + \frac{dC_t(\chi)}{dt} \right\},$$

with the terminal payoff $\bar{C}(Q(S_T) + V_T, \tau)$, which completes the proof. $\square$

Proof of Proposition 2. Having Equation (23) for the short-term cost, and with the strong assumptions that all shares requested will be executed, the long-term cost of execution of X shares in τ time is

$$\bar{C}(X, \tau) = \int_0^\tau S_t dC_t.$$

Note that in our case, $\frac{X}{\tau}$ is the constant execution rate. In case we have Equation (23) for the short-term cost, the long-term cost of execution is as follows:

$$\begin{aligned} \int_T^{T+\tau} S_t dC_t &= \eta \times \frac{\left(\frac{X}{\tau}\right)^{\alpha+1}}{\left(\frac{X}{\tau} + A\right)^\alpha} \int_T^{T+\tau} e^{-rt} \left(1 + \Lambda\left(\frac{X}{\tau}\right)\right)^t S_t dt, \\ E^I \left[\int_T^{T+\tau} S_t dC_t \right] &= \eta \frac{\left(\frac{X}{\tau}\right)^{\alpha+1}}{\left(\frac{X}{\tau} + A\right)^\alpha} S_T \tau. \end{aligned}$$

$\square$

Notes

- ¹ There are other types of traders such as *average traders* who are too small to influence the market and are unaware of the large trader's activities. However, this group of traders is not the focus of this paper.
- ² In Equation (1), the size of the jump is denoted by $\Lambda(\chi_t)$, as opposed to $\Lambda(\chi_{t-})$. The notation t^- is typically used to ensure the validity of the closed-form exponential representation of the process $\{S_t\}_{0 \leq t \leq T}$, which is a product of the impact of all jumps and the exponential Gaussian process, as detailed in Doléans-Dade (1970). In simpler terms, for the exponential formulation to hold, jumps need to be applied first to the spot price early at t^- , followed by other sources of variation. However, in our model, an exponential expansion is not employed. From a financial perspective, this implies that while the jump increments are applied to the price at S_{t-} before other variations, its actual magnitude is determined at time t , not at t^- .
- ³ While optimal purchase/liquidation algorithms for longer purchase periods are well-researched topics, as seen in works like Almgren and Chriss (2001), our focus here is on the constant rate program. There are compelling reasons for selecting a constant purchase program over a more intricate one. Notably, most derivatives desks do not prioritize optimizing their execution cost, placing greater emphasis on volatility instead. Thus, even if optimal algorithms are available, they might not be employed in derivatives trading. Our goal here is not to suggest how the purchasing should be executed but to understand the behavior in the most prevalent scenario. This reasoning drives our choice for a constant rate program.
- ⁴ Later, Gatheral (2010) proved that linearity of the permanent-impact function is a sufficient condition for no-arbitrage, while Guéant (2013) and others explored nonlinear and stochastic extensions.
- ⁵ A practitioner, however, may form a *proxy valuation* by combining the insider's informational price with the large trader's expected cost, but this is a heuristic and does not constitute an arbitrage-free price.

References

- Abeille, M., Bouchard, B., & Croissant, L. (2023). Diffusive limit approximation of pure-jump optimal stochastic control problems. *Journal of Optimization Theory and Applications*, 196(1), 147–176. [CrossRef]
- Abergel, F., & Loeper, G. (2017). Option pricing and hedging with liquidity costs and market impact. In *Econophysics and sociophysics: Recent progress and future directions* (pp. 19–40). Springer.
- Almgren, R. (2012). Optimal trading with stochastic liquidity and volatility. *SIAM Journal on Financial Mathematics*, 3(1), 163–181. [CrossRef]

- Almgren, R., & Chriss, N. (2001). Optimal execution of portfolio transactions. *Journal of Risk*, 3, 5–40. [CrossRef]
- Almgren, R., Thum, C., Hauptmann, E., & Li, H. (2005). Direct estimation of equity market impact. *Risk*, 18(7), 58–62.
- Amendinger, J., Imkeller, P., & Schweizer, M. (1998). Additional logarithmic utility of an insider. *Stochastic Processes and Their Applications*, 75(2), 263–286. [CrossRef]
- Barclay, M. J., & Warner, J. B. (1993). Stealth trading and volatility: Which trades move prices? *Journal of Financial Economics*, 34(3), 281–305. [CrossRef]
- Barger, W., & Lorig, M. (2019). Optimal liquidation under stochastic price impact. *International Journal of Theoretical and Applied Finance*, 22(2), 1850059. [CrossRef]
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654. [CrossRef]
- Bordag, L. A., & Frey, R. (2008). Pricing options in illiquid markets: Symmetry reductions and exact solutions. In *Nonlinear models in mathematical finance: New research trends in option pricing* (pp. 103–129). Nova Science Publishers, Inc.
- Bouchaud, J.-P., Farmer, J. D., & Lillo, F. (2010). How markets slowly digest changes in supply and demand. In T. Hens, & K. R. Schenk-Hoppé (Eds.), *Handbook of financial markets: Dynamics and evolution* (pp. 57–160). Elsevier.
- Bouchaud, J.-P., Gefen, Y., Potters, M., & Wyart, M. (2004). Fluctuations and response in financial markets: The subtle nature of “random” price changes. *Quantitative Finance*, 4(2), 176–190. [CrossRef]
- Cboe Inc. (2022). *Equity options product specifications*. Available online: https://www.cboe.com/exchange_traded_stock/equity_options_spec (accessed on 30 January 2022).
- Cetin, U., Jarrow, R., Protter, P., & Warachka, M. (2006). Pricing options in an extended Black–Scholes economy with illiquidity. *Review of Financial Studies*, 19(2), 493–529. [CrossRef]
- Chordia, T., Roll, R., & Subrahmanyam, A. (2001). Market liquidity and trading activity. *Journal of Finance*, 56(2), 501–530. [CrossRef]
- Christensen, H. L., & Woodmansey, R. (2013). Prediction of hidden liquidity in the limit order book of GLOBEX futures. *The Journal of Trading*, 8(3), 68–95. [CrossRef]
- Cont, R., & Tankov, P. (2003). *Financial modelling with jump processes*. Chapman and Hall/CRC.
- Curato, G., Gatheral, J., & Lillo, F. (2017). Optimal execution with non-linear transient market impact. *Quantitative Finance*, 17(1), 41–54. [CrossRef]
- Damodaran, A. (2025). *Historical returns on stocks, bonds and bills—United States*. Available online : http://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/histretSP.html (accessed on 14 January 2025).
- Doléans-Dade, C. (1970). Quelques applications de la formule de changement de variables pour les semimartingales. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 16, 181–194. [CrossRef]
- Easley, D., & O’Hara, M. (1996). Liquidity, information, and infrequently traded stocks. *Journal of Finance*, 51(4), 1405–1436. [CrossRef]
- European Parliament and of the Council. (2014). *Regulation (EU) No 596/2014 on market abuse (market abuse regulation)*. European Parliament and of the Council.
- Eyraud-Loisel, A. (2013). Quadratic hedging in an incomplete market derived by an influential informed investor. *Stochastics and International Journal of Probability and Stochastic Processes*, 85(3), 412–430. [CrossRef]
- Financial Conduct Authority. (2005). *FCA handbook: MAR 1.6.15 market abuse*. Financial Conduct Authority.
- Fletcher, G. G. (2018). Legitimate yet manipulative: The conundrum of open-market manipulation. *Duke Law Journal*, 68, 479.
- Föllmer, H., & Schweizer, M. (1990). *Hedging of contingent claims under incomplete information*. Rheinische Friedrich-Wilhelms-Universität Bonn.
- Gatheral, J. (2010). No-dynamic-arbitrage and market impact. *Quantitative Finance*, 10(7), 749–759. [CrossRef]
- Gatheral, J. (2012). Transient linear price impact and Fredholm integral equations. In J.-P. Fouque, & J. Langsam (Eds.), *Representing, measuring and managing systemic risk* (pp. 309–321). World Scientific.
- Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100. [CrossRef]
- Griffin, J. M., & Shams, A. (2018). Manipulation in the VIX? *Review of Financial Studies*, 31(4), 1377–1417. [CrossRef]
- Grinold, R. C., & Kahn, R. N. (2000). *Active portfolio management*. McGraw Hill New York.
- Gorod, A., & Pontier, M. (1998). Insider trading in a continuous time market model. *International Journal of Theoretical and Applied Finance*, 1(03), 331–347. [CrossRef]
- Guéant, O. (2013). Permanent market impact can be nonlinear. *arXiv*, arXiv:1305.0413.
- Guéant, O., & Pu, J. (2017). Option pricing and hedging with execution costs and market impact. *Mathematical Finance*, 27(3), 803–831. [CrossRef]
- Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance*, 46(1), 179–207. [CrossRef]
- Hasbrouck, J., & Seppi, D. J. (2001). Common factors in prices, order flows, and liquidity. *Journal of Financial Economics*, 59(3), 383–411. [CrossRef]
- Hillairet, C. (2004). *Equilibres sur un marché financier avec asymétrie d’information et discontinuité des prix* [Ph.D. thesis, Université Paul Sabatier-Toulouse III].

- Horst, U., & Naujokat, F. (2011). On derivatives with illiquid underlying and market manipulation. *Quantitative Finance*, 11(7), 1051–1066. [CrossRef]
- Huang, R. D., & Stoll, H. R. (1997). The components of the bid-ask spread: A general approach. *Review of Financial Studies*, 10(4), 995–1034. [CrossRef]
- Huberman, G., & Stanzl, W. (2004). Price manipulation and quasi-arbitrage. *Econometrica*, 72(4), 1247–1275. [CrossRef]
- Jarrow, R., Fung, S., & Tsai, S.-C. (2018). An empirical investigation of large trader market manipulation in derivatives markets. *Review of Derivatives Research*, 21(3), 331–374. [CrossRef]
- Jarrow, R. A. (1992). Market manipulation, bubbles, corners, and short squeezes. *Journal of Financial and Quantitative Analysis*, 27(3), 311–336. [CrossRef]
- Jarrow, R. A. (1994). Derivative security markets, market manipulation, and option pricing theory. *Journal of Financial and Quantitative Analysis*, 29(2), 241–261. [CrossRef]
- Jonsson, M., & Sircar, K. R. (2002). Partial hedging in a stochastic volatility environment. *Mathematical Finance*, 12(4), 375–409. [CrossRef]
- Klock, F. (2014). Price manipulation in a market with interacting investor and venue choice. *SIAM Journal on Financial Mathematics*, 5(1), 607–649.
- Klöß, F. (2012). *Regularity of market impact models with stochastic price* (Working paper, May 14, 2012). Available online: <https://ssrn.com/abstract=2057610> (accessed on 26 November 2025).
- Kraft, H., & Kühn, C. (2011). Large traders and illiquid options: Hedging vs. manipulation. *Journal of Economic Dynamics and Control*, 35(11), 1898–1915. [CrossRef]
- Leland, H. E. (1985). Option pricing and replication with transaction costs. *Journal of Finance*, 40(5), 1283–1301. [CrossRef]
- Lin, J. (2017). The new market manipulation. *Emory Law Journal*, 66(6), 1253–1308.
- Liu, H., & Yong, J. (2005). Option pricing with an illiquid underlying asset market. *Journal of Economic Dynamics and Control*, 29(12), 2125–2156. [CrossRef]
- Loeper, G. (2018). Option pricing with linear market impact and nonlinear black–scholes equations. *Annals of Applied Probability*, 28(5), 2664–2726. [CrossRef]
- Ma, G., Siu, C. C., Zhu, S.-P., & Elliott, R. J. (2020). Optimal portfolio execution problem with stochastic price impact. *Automatica*, 112, 108739. [CrossRef]
- Madhavan, A., Richardson, M., & Roomans, M. (1997). Why security prices change. *Review of Financial Studies*, 10(4), 1035–1064. [CrossRef]
- Ni, S. X., Pearson, N. D., & Poteshman, A. M. (2005). Stock price clustering on option expiration dates. *Journal of Financial Economics*, 78(1), 49–87. [CrossRef]
- Roch, A. F. (2022). Hedging of American options in illiquid markets with price impacts. *International Journal of Theoretical and Applied Finance*, 25(01), 2250001. [CrossRef]
- Rogers, L. C. G., & Shi, Z. (1995). The value of an Asian option. *Journal of Applied Probability*, 32(4), 1077–1088. [CrossRef]
- Said, E. (2020). *Market impact in systematic trading and option pricing* [Ph.D. thesis, Université Paris-Saclay].
- Tankov, P., & Voltchkova, E. (2009). Jump-diffusion models: A practitioner’s guide. *Banque et Marchés*, 99(1), 24.
- Tóth, B., Lemperiere, Y., Deremble, C., de Lataillade, J., Kockelkoren, J., & Bouchaud, J.-P. (2011). Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2), 021006. [CrossRef]
- VIX, CBOE. (2025). *CBOE volatility index (VIX) historical data*. Available online: https://www.cboe.com/tradable_products/vix/vix_historical_data (accessed on 14 January 2025).
- Zarinelli, E., Treccani, M., Farmer, J. D., & Lillo, F. (2015). Beyond the square root: Evidence for logarithmic dependence of market impact on size and participation rate. *Market Microstructure and Liquidity*, 1(02), 1550004. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.