

Discrete Delta Hedging under Stochastic Volatility and Jumps: A Monte Carlo Cost–Risk Frontier

Yue Li

*College of Arts and Sciences, The Ohio State University, Columbus, USA
li.13232@buckeyemail.osu.edu*

Abstract. Delta hedging is a fundamental strategy in options risk management, relying on continuous adjustment of a replicating portfolio to eliminate risk. However, real markets exhibit features such as stochastic volatility and jumps that violate the assumptions of the Black–Scholes model, rendering perfect replication impossible and the market incomplete. In such cases, hedging can only reduce risk at best, and frequent rebalancing incurs significant transaction costs. This article investigates discrete delta hedging under stochastic volatility and jump-diffusion dynamics, quantifying the trade-off between hedging cost and risk reduction via Monte Carlo simulation. We construct a cost–risk frontier, analogous to an efficient frontier, that shows the minimal achievable risk for a given cost (and vice versa). The results demonstrate that increasing the hedge frequency (trading more often) generally lowers the variance of hedging errors but at a rapidly diminishing rate and with higher accumulated costs. Even with very frequent rebalancing, a residual risk remains due to jumps and unhedgeable volatility fluctuations. We discuss how this frontier can inform optimal hedging policies, balancing transaction costs against risk appetite, and we compare our findings with prior theoretical and empirical studies.

Keywords: Stochastic Volatility, Transaction Costs, Jump-Diffusion.

1. Introduction

In their seminal work, Black and Scholes showed that under certain idealized conditions (constant volatility, frictionless trading, continuous time), an option's payoff can be perfectly replicated by continuously adjusting a portfolio of the underlying asset and risk-free asset [1]. This continuous delta hedging strategy leads to a riskless portfolio and yields the celebrated Black–Scholes option pricing formula. Merton extended this theory, and together they demonstrated that, if the model assumptions hold, the cost of hedging an option (i.e., the initial option premium) equals the option's arbitrage-free price [2]. The necessary conditions for perfect hedging include a single source of uncertainty (the underlying's price following a geometric Brownian motion with constant volatility) and no market frictions. Under these conditions, the market is complete and admits a unique risk-neutral pricing measure.

However, financial markets in reality violate several Black–Scholes assumptions. Two important deviations are stochastic volatility and jump discontinuities in asset prices. Empirical evidence shows that asset return distributions exhibit excess kurtosis (fat tails) relative to the lognormal, and

volatility is time-varying and stochastic. To address these phenomena, more complex models have been developed. Notably, jump–diffusion models introduce sudden price jumps [2], and continuous-time stochastic volatility models introduce additional randomness in volatility [3]. These richer models can better fit market option price patterns such as implied volatility skews and smiles. However, they also make the market incomplete – there are now multiple sources of uncertainty (e.g. a volatility factor or jump shocks) but only one primary asset available for hedging. In an incomplete market, a contingent claim cannot be perfectly hedged using only the underlying asset. As a result, there is no unique arbitrage-free price; instead, a continuum of fair prices exists, and hedging strategies can only minimize risk rather than eliminate it.

Practitioners have responded to incompleteness in various ways. One approach is to augment the hedging portfolio with additional instruments—for example, using traded options to hedge jump risk or volatility risk in addition to the underlying stock. It has been shown that holding a short-term option in the hedge portfolio can significantly mitigate jump-induced losses [4]. Such multi-instrument hedging can, in theory, restore completeness if a sufficient set of contingent claims is available (e.g. using a vanilla option to hedge Vega exposure in an SV model). Another approach is to accept that some risk is unhedgeable and to seek an optimal hedge that balances risk and cost. Strategies like mean-variance hedging and quadratic risk minimization explicitly choose the hedge ratio that minimizes the variance of the hedging error, rather than using the Black–Scholes delta. These optimal hedge ratios generally differ from Black–Scholes delta and do not fully eliminate risk, but they minimize it under a given criterion.

Even setting aside model incompleteness, discrete-time hedging (in contrast to continuous rebalancing) introduces additional imperfection. In practice, hedging adjustments occur at finite intervals (e.g., daily) rather than continuously. Studies have examined the discretization error associated with infrequent rebalancing [5-7]. They found that discretely rebalanced hedges lead to a distribution of final hedging costs or errors, rather than a single deterministic cost equal to the option’s price. In a Black–Scholes world (no jumps, constant vol), the average outcome of a high-frequency discrete hedge still equals the Black–Scholes price, but there is appreciable variance around this mean cost. This variance grows as the rebalancing interval lengthens (i.e. hedging less frequently), and it represents the residual risk of a discrete hedge.

Another crucial consideration is transaction costs. Continuous hedging is infeasible and would imply infinite trading costs under frictions. Even a very small proportional cost per trade can drastically alter hedging strategy and option pricing [6]. Frequent trading becomes expensive, prompting the question of an optimal rebalancing frequency that balances risk reduction against cost. Modified option-pricing equations under transaction costs have been derived, effectively widening the no-arbitrage bounds to account for the cost–risk trade-off [8, 9]. In general, with transaction costs present, hedging strategies that are too aggressive (trading on every tiny price move) may over-hedge, spending more on transactions than the incremental risk reduction is worth. On the other hand, hedging too infrequently leaves substantial risk. Thus, there is an intuitive trade-off between cost and risk in choosing a hedging policy.

In an incomplete market setting with jumps or SV (or both) and with transaction costs, the problem becomes finding a strategy that optimally compromises between risk and cost. Recent research has addressed this problem from various angles. For instance, stochastic control and model predictive control approaches have been applied to hedging with cost penalization, yielding a Pareto-optimal frontier of strategies [10]. More directly, “restricted but optimal” delta hedging under scenarios including jumps, stochastic volatility, and transaction costs has been investigated [7]. In his analysis, the hedge is only adjusted at discrete intervals and can incorporate realistic costs; the

optimal strategy is one that minimizes a combination of hedging error variance and cost. Such studies confirm that incorporating jumps and volatility risk leads to a non-zero minimum attainable risk – even the best strategy cannot eliminate risk completely, because some sources of uncertainty remain unhedged. They also underscore that the relationship between how often one hedges and the resulting cost and risk is nonlinear: beyond a certain point, more frequent hedging yields only marginal risk reduction but incurs significantly higher costs.

This paper contributes to the literature by explicitly quantifying the cost–risk frontier for discrete delta hedging in a model that includes both stochastic volatility and jumps. It uses Monte Carlo simulation to evaluate hedging performance across a range of rebalancing frequencies, from very frequent (approaching continuous) to very sparse. For each frequency, it computes the expected transaction cost incurred and the residual risk (measured by the standard deviation of final hedging P&L). Plotting these as a trade-off curve yields the frontier: strategies on the frontier are Pareto-optimal in the sense that you cannot reduce risk without increasing cost, or reduce cost without increasing risk. This provides a concrete illustration of the principle that “there is no free lunch” in incomplete markets – any reduction in risk must be paid for, either via higher cost or via accepting some other form of risk.

The remainder of the paper is organized as follows. In Section 2, it describes the market model (stochastic volatility with jumps) and the discrete delta hedging strategy, including how we incorporate transaction costs. Section 3 details the Monte Carlo simulation approach and parameter choices. In Section 4, it presents the results, including the cost–risk frontier chart obtained and discussion of its shape. Section 5 concludes with insights on practical hedging policy implications and potential extensions of this work.

2. Model and methodology

Underlying Asset Dynamics: It assumes the underlying asset S_t follows a stochastic volatility jump-diffusion process. Specifically, we adopt a variant of the Bates model [11]. Under the risk-neutral measure, the dynamics are:

- **Stochastic volatility:** $dV_t = \kappa(\theta - V_t)dt + \eta\sqrt{V_t}dW_t^V$, where V_t is the instantaneous variance, κ is the mean-reversion rate, θ the long-run mean variance, θ vol-of-vol, and W_t^V a Wiener process. This is the Heston SV model for volatility. We assume the Feller condition for V_t to stay positive. The correlation between the volatility shock dW_t^V and the asset price shock dW_t^S is ρ (which may be zero or negative as commonly estimated in equity markets).

- **Jump component:** In addition to the diffusive price shock, we include jumps via a Poisson process N_t with intensity λ (jumps per year). When a jump occurs, the asset price experiences a multiplicative jump $J = e^Y$, where Y is a random jump size (we assume $Y \sim N(\mu_J, \sigma_J^2)$, i.e., log-normal jumps). Thus, S_t satisfies:

$$dS_t = S_t (r dt + \sqrt{V_t} dW_t^S + (J - 1) dN_t) \quad (1)$$

where r is the risk-free rate. We choose μ_J such that $E[J - 1] = 0$ under the risk-neutral measure (so on average jumps do not drift the price), and σ_J controls jump volatility. Jumps arriving with intensity λ make the price process exhibit sudden discontinuities, contributing to heavier tails and skewness in returns.

This model incorporates two independent risk factors: the Brownian diffusive risk (with stochastic variance) and the jump risk. With only the underlying stock available to trade (and

perhaps a money market account), the market is incomplete. There is no trading strategy in S_t alone that can hedge a derivative perfectly, even with continuous trading. Intuitively, continuous delta-hedging can neutralize the instantaneous diffusion risk (the dW_t^S term) but cannot foresee jumps — any jump will cause a discontinuous change in S that leads to a hedging error. Likewise, unpredictable changes in volatility V_t will alter the option's value in ways that cannot be offset by just trading S_t . This inherent limitation means a delta-hedged portfolio will still exhibit random P&L fluctuations.

Option and Hedging Setup: For concreteness, we consider hedging a short position in a European call option on S (strike K , maturity T). The initial option premium (theoretical fair price) is denoted C_0 . In our risk-neutral simulation, C_0 is computed as the expected discounted payoff $e^{-rT} \mathbb{E}[(S_T - K)^+]$. At time 0, the hedger sells the call for C_0 and begins a delta-hedging strategy. The strategy is discrete delta hedging, meaning we choose a sequence of rebalancing times $0 = t_0 < t_1 < \dots < t_N = T$ at which the hedge ratio is updated. Between these times, the holdings are fixed. At each hedge time t_i , we set the position in the underlying stock to $\Delta(t_i)$, where $\Delta(t)$ is an estimate of the option's delta (sensitivity to S). In a complete-market setting, the ideal delta would be the partial derivative $\partial C / \partial S$ given the model. In our setting, one could use the theoretical delta from a model (if available in closed form via Heston's formula or simulation). For simplicity, we use an approximate delta given by the Black–Scholes formula, but updated at each time with the current observed volatility level. Specifically, we compute:

$$\Delta(t_i) = N(d_1(S_{t_i}, V_{t_i}, T - t_i)) \quad (2)$$

where $N(\cdot)$ is the standard normal CDF (the Black–Scholes delta for a call), and $d_1(S, \sigma, \tau) = \frac{\ln(S/K) + \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}}$,

with $\sigma = \sqrt{V_{t_i}}$ interpreted as the current volatility and $\tau = T - t_i$ the remaining time. This approach effectively treats the local volatility as $\sigma_{loc} = \sqrt{V_{t_i}}$ and computes a Black–Scholes delta with that volatility. While not exact (since the true optimal delta in stochastic volatility differs), this mimics common practice where traders update greeks based on current implied or local volatilities. It also ensures $\Delta(t)$ evolves in a plausible manner between 0 and 1 depending on moneyness and time. At maturity, we do not rebalance (the option payoff is realized).

Trading Costs: Each time the hedge is adjusted, the change in the stock position incurs a transaction cost. We assume a simple proportional cost model: whenever Δ is changed, the trader pays a cost equal to a fraction α of the dollar value of stock traded. For example, if at time t_{i-1} the position was Δ_{i-1} and at t_i the new target is Δ_i , and the stock price at t_i is S_{t_i} , then:

$$\text{transaction cost} = \alpha, |\Delta_i - \Delta_{i-1}|, S_{t_i}. \quad (3)$$

This can represent a bid–ask spread or commission. In our simulations, we take a modest α (e.g., 0.1% per trade), but accumulated over many trades this can meaningfully erode profits. The presence of transaction costs discourages extremely frequent rebalancing — the incremental risk reduction from an extra rebalance must be weighed against the certain cost incurred.

Hedging P&L: The performance of the hedging strategy is evaluated by the profit-and-loss (P&L) at option expiration. We track the replicating portfolio consisting of $\Delta(t)$ shares of stock and a cash account (financing the stock purchases or generated from stock sales). Starting with initial cash C_0

(from selling the option) and $\Delta(0)$ shares, the portfolio is rebalanced at each t_i . Between hedges, the stock holding $\Delta(t_{i-1})$ grows in value as S moves, and the cash earns interest at rate r (here we set r low or zero for simplicity). At t_i , we adjust the stock holding by $\Delta_i - \Delta_{i-1}$, and the cash account changes by the negative of the stock trade plus the transaction cost. At T , the short call position is settled by paying the payoff $(S_T - K)^+$ to the option holder. The hedging error (final P&L) is:

$$\Pi = (\text{value of stock} + \text{cash}) - (\text{option payoff liability}) \quad (4)$$

If $\Pi = 0$, the hedge perfectly replicated the payoff; $\Pi > 0$ indicates a profit for the hedger (the strategy over-hedged), and $\Pi < 0$ a loss (under-hedged). In an ideal frictionless complete-market world, we would have $\Pi = 0$ deterministically for the correct strategy. In our incomplete, discrete, costly setting, Π is a random variable. We focus on its mean and standard deviation as key metrics. The mean $E[\Pi]$ represents any systematic bias (which in our case comes mostly from transaction costs — we expect $E[\Pi] < 0$, i.e., an average loss equal to the total cost paid by the hedger, since the initial option price was fair). The standard deviation σ_Π measures the risk — the volatility of the hedging outcome. A well-chosen hedge will have a small σ_Π , indicating outcomes tightly clustered around the mean (ideally around zero or a small negative cost). We will use σ_Π as our metric of residual risk, and the average transaction cost $E[\text{cost}] = -E[\Pi]$ as our cost measure (since any shortfall in P&L for the hedger is essentially the cost expended, assuming zero drift under risk-neutral pricing).

3. Simulation and implementation

To investigate the cost–risk trade-off, we perform a Monte Carlo simulation for various hedging frequencies, and the corresponding results are shown in Table 1. It fixes a set of model parameters for illustration: an initial stock price $S_0 = 100$, option strike $K = 100$ (at-the-money), maturity $T = 1$ year. Volatility starts at $V_0 = 0.04$ (so 20% initial vol). It chooses Heston parameters $\kappa = 1.0$, $\theta = 0.04$, $\eta = 0.2$, and correlation $\rho = 0$ (uncorrelated for simplicity). Jump parameters are $\lambda = 1.0$ (on average one jump per year), $\sigma_J = 0.25$ (so about 25% jump size volatility), and $\mu_J = -0.5\sigma_J^2$ (approximately -0.03125) to ensure mean jump size is 1 (no drift). The transaction cost rate is set to $\alpha = 0.001$ (0.1% of traded value each hedge).

Table 1. Delta hedging frequency and rebalancing intervals

Hedging frequency	Δt (interval)	Notes
Daily hedging	$\Delta t = 1$ day	~252 hedges per year; effectively as often as trading days allow.
Weekly hedging	$\Delta t \approx 5$ days	
Monthly hedging	$\Delta t \approx 21$ days	
Quarterly hedging	$\Delta t \approx 63$ days	
Semi-annual hedging	$\Delta t \approx 126$ days	
Annual hedging	$\Delta t = T$	No interim adjustments; only set initial hedge and hold to expiration.

For each strategy, it simulates a large number of price paths (on the order of 50,000 or more) to accurately estimate the distribution of P&L. Each path is generated by discretizing the SDEs for S_t and V_t on a daily time grid ($\Delta t = \frac{1}{252}$). We use an Euler discretization for the volatility process and incorporate jumps by a simple Poisson sampling each day (if a jump occurs during a day, we multiply S by the jump factor J). The delta hedge is applied by sampling the state at the designated hedge times and rebalancing accordingly.

The entire simulation and analysis were implemented in MATLAB. Figure 1 reports the cost–risk frontier generated by our Monte Carlo engine for the parameter set described in the Model and Methodology section. The algorithm simulates Bates dynamics (Heston stochastic volatility with Merton-type jumps) on a daily grid, applies discrete delta hedging at the stated rebalancing intervals with proportional transaction costs, and records terminal hedging P&L across simulated paths. For reproducibility, it initializes the MATLAB random number generator to a fixed state (seed 123456, twister stream). The fully annotated script and run instructions are provided in the Supplementary Material (and mirrored at the public repository listed in “Data and Code Availability”).




Figure 1. Cost-risk frontier for discrete delta hedging

The simulation evolves price and variance path-by-path and computes the P&L for each hedging frequency; we use 10^5 paths for the final estimates to control Monte Carlo error. Mean transaction cost and the standard deviation of the hedging P&L are computed for each frequency, and the resulting pairs are plotted in Figure 1 to form the empirical cost–risk frontier. To improve efficiency and comparability across frequencies, we employ vectorization and common random numbers; Figure 1 was exported from MATLAB in vector format for publication quality.

4. Results and discussion

As Figure 1 shows, Monte Carlo estimated cost–risk frontier for discrete delta hedging under stochastic volatility and jumps. The x-axis is the risk (standard deviation of hedging P&L) and the y-axis is the cost (average total transaction costs incurred). Each labeled point corresponds to a different hedging frequency (e.g. daily, weekly, monthly, etc.). The curve illustrates the trade-off: achieving lower risk requires higher cost. Notably, the frontier is nonlinear, showing diminishing risk reduction returns for increasing cost.

The cost–risk frontier obtained from the simulation is shown in Figure 1. Each point on the plot represents one hedging strategy parameterized by rebalancing frequency. The rightmost point (highest risk, lowest cost) corresponds to a buy-and-hold strategy with essentially no rebalancing (in

our case, an initial delta hedge set at $t = 0$ and never adjusted, which is analogous to annual hedging for a 1-year option). The leftmost point (lowest risk, highest cost) corresponds to daily hedging. As expected, more frequent hedging moves the outcome towards lower risk: the standard deviation of P&L for daily hedging is substantially smaller than for monthly or quarterly hedging. Meanwhile, the cumulative transaction costs for daily hedging are much higher than for infrequent hedging.

Here the frontier delivers four big takeaways. First, even very frequent rebalancing cannot eliminate risk: jumps (gap risk) and stochastic volatility (unspanned vega risk) leave an irreducible hedging error when you hedge only with the underlying. Second, increasing hedge frequency exhibits strong diminishing returns—moving from very infrequent to moderate rebalancing slashes risk, but beyond that the curve flattens and extra trades mostly buy small variance reductions. Third, transaction costs reshape the optimum: because every adjustment burns cash and adds cost variability, extremely frequent hedging can raise total risk after costs; the efficient choice is an interior cadence on the flat part of the frontier where the marginal risk reduction roughly equals the marginal cost. Fourth, the mean P&L is negative across strategies, effectively matching average transaction costs, so delta hedging behaves like paying an insurance premium to compress tails rather than a profit center.

Managerially, the slope of the frontier at a candidate cadence is the “price of risk reduction.” Choose the frequency where that price aligns with risk appetite or capital constraints. Parameter shifts move the whole curve: more frequent/larger jumps or higher vol-of-vol push it right (more residual risk), while higher trading frictions push it up (greater cost). Thus, a jumpier asset or wider spreads warrant less aggressive rebalancing; quieter markets with tight spreads can justify more. If deeper risk cuts are required, adding a second instrument (e.g., a short-dated option for vega/jump exposure) doesn’t change the trade-off logic but can shift the entire frontier inward, achieving lower risk for the same cost.

Overall, the cost–risk frontier quantifies the intuitive idea that you can buy reduced risk by paying higher transaction cost. Each point on the frontier could be seen as an efficient hedging strategy for a particular risk tolerance: a risk-averse hedger might choose a high-frequency, high-cost strategy to minimize variance, whereas a cost-sensitive (or less risk-averse) hedger might accept more risk in exchange for lower costs.

It is important to note that the exact shape and numbers on the frontier will depend on model parameters (volatility of volatility, jump intensity, etc.) and on the option’s characteristics (maturity, strike). For instance, a higher jump intensity λ or larger jump size σ_J would increase the residual risk for all strategies (shifting the frontier rightwards – more risk) because jumps are harder to hedge. A higher transaction cost rate α would make the high-frequency strategies even more costly (shifting the frontier upwards – more cost). The frontier we present is thus illustrative, but the methodology can be applied to specific situations to determine an appropriate hedging frequency or to estimate the efficient frontier for that scenario.

Our findings are consistent with previous analytical results in the literature. As Merton (1976) pointed out, in a jump-diffusion setting a delta-hedged portfolio will still experience jumps in value leading to losses. Our simulation shows those losses manifest as a baseline level of risk that cannot be diversified away by faster trading. Similarly, the transaction cost effects we observe align with Leland’s theory that effective volatility is increased by transaction costs – in other words, the presence of costs makes it as if the option is riskier, because one cannot chase the deltas as closely. What we see in the frontier’s flattening is essentially that phenomenon: beyond a point, increasing

hedge frequency yields minimal risk reduction because effectively the strategy is “running in place” – incurring costs to shave off ever smaller risk components.

One interesting aspect worth discussing is that in our simulation we used a simple Black–Scholes delta estimate. One could ask: would using a more optimal delta (for example, the mean-variance optimal hedge ratio given the model) improve the trade-off? Potentially yes – an optimal strategy could achieve a lower variance for the same cost. In that case, our frontier would shift towards the origin (improvement). However, the general shape would remain – there would still be a convex, diminishing-return curve. Our focus here was on delta hedging as it is a common baseline. Future work could incorporate optimal hedging rules or non-linear instruments to see how the frontier shifts (for example, adding a second instrument like a short-term put might allow further risk reduction, effectively pushing the frontier downward at the low-risk end by completing more of the market).

We also note that the distribution of hedging P&L (not just its variance) is of interest. In our simulation, the unhedged position (short call without any hedge) would have a very dispersed P&L distribution, effectively the negative of the call’s payoff distribution. Hedging narrows this distribution. We observed that with daily hedging, the P&L distribution becomes more concentrated around a small loss (the cost), whereas with monthly hedging the distribution is wider. Hedging particularly trims the extreme tail outcomes – for example, scenarios where the stock price skyrockets at expiration (which would be a huge loss for an unhedged short call) are substantially mitigated by delta hedging, since the hedger would have acquired a lot of stock in those scenarios. However, jump risk implies that a sudden large move just before a hedge adjustment can still produce a significant loss. This highlights that while delta hedging addresses continuous price risk effectively, it is less effective for gap risk. In practice, risk managers often complement delta hedging with other tools (like stop-loss rules or options positions) to handle jump risk.

5. Conclusion

Here we examined discrete-time delta hedging in an equity market with stochastic volatility and jumps, using Monte Carlo simulation to map the trade-off between hedging cost and residual risk. Unlike the idealized Black–Scholes setting—where continuous rebalancing with constant volatility yields exact replication—the presence of volatility randomness and discontinuous price jumps makes the market intrinsically incomplete. Even with very frequent rebalancing, a nontrivial, irreducible variance of the hedging error persists because trading the underlying alone cannot span volatility and jump risks. Discreteness compounds this limitation: in a Black–Scholes world, discrete hedging produces a distribution of replication outcomes with zero mean and nonzero variance; in our richer setting, proportional transaction costs push the mean outcome negative while leaving a wider dispersion of P&L. As trading frequency increases, variance typically falls, but at a diminishing rate, while costs climb quickly; beyond moderate frequencies, additional rebalancing primarily purchases marginal risk reduction at disproportionate expense. The resulting cost–risk frontier makes this trade-off explicit, tracing the Pareto set of efficient strategies from low-cost/high-risk to high-cost/low-risk. This frontier offers a practical decision tool: risk managers can select a rebalancing cadence consistent with their risk appetite and cost budget, or assess whether a targeted variance reduction justifies the incremental transaction cost. The framework is readily extensible. The same simulation engine can be adapted to different maturities and strikes, to alternative instruments (e.g., American-style contracts, bearing in mind early-exercise complexity), and to multi-asset portfolios where cross-greeks matter. Likewise, the risk axis need not be the standard deviation of hedging P&L; Value-at-Risk or Expected Shortfall can be substituted to produce frontiers aligned with regulatory or internal risk metrics. Overall, our findings formalize a central

practical insight: in incomplete markets with frictions, hedging precision is bought—not assumed—and the efficient purchase happens along a quantifiable cost–risk frontier.

In practice, traders often use a mix of strategies to manage risks from jumps and volatility. This might include holding residual positions in options (vega hedges or jump hedges), or structuring trades such that extreme moves are limited (through stop-loss orders). Those approaches lie outside the scope of pure delta hedging and were not addressed here. However, the framework we developed could be expanded to include multiple instruments – effectively, that would likely shift the cost–risk frontier downward (achieving lower risk for a given cost) by enlarging the hedger’s toolkit.

In summary, discrete delta hedging in an incomplete market comes with an inherent inefficiency – it cannot eliminate risk, and whatever risk remains can only be curtailed at a proportional cost. Understanding this cost–risk frontier is crucial for realistic expectations of hedging performance and for making informed decisions about hedging strategies. Our results underscore the importance of calibrating hedging frequency and trades to a firm’s risk tolerance and cost constraints, an insight that is valuable for both option market-makers and end-users employing dynamic hedging for risk management.

References

- [1] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654.
- [2] Merton, R.C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2), 125–144.
- [3] Hull, J., & White, A. (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42(2), 281–300.
- [4] He, T., et al. (2006). Dynamic hedging under jump diffusion by using short-term options.
- [5] Boyle, P., & Emanuel, D. (1980). Discretely adjusted option hedges. *Journal of Financial Economics*, 8(3), 259–282.
- [6] Leland, H. (1985). Option pricing and replication with transactions costs. *Journal of Finance*, 40(5), 1283–1301.
- [7] Wilmott, P. (2022). A note on hedging: restricted but optimal delta hedging, mean, variance, jumps, stochastic volatility, and costs. *Wilmott Journal* 2022, 11 pages (DOI: 10.1002/WILJ.9).
- [8] Figlewski, S. (1989). Option arbitrage in imperfect markets. *Journal of Finance*, 44(5), 1289–1311.
- [9] Hoggard, T., Whalley, A.E., & Wilmott, P. (1994). Hedging option portfolios in the presence of transaction costs. *Advances in Futures and Options Research*, 7, 21–35.
- [10] Noorian, F., & Leong, P.H.W. (2014). Dynamic hedging of foreign exchange risk using stochastic model predictive control. 2014 IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr), 439–446.
- [11] Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2), 327–343.