

Dynamic Hedging under Jump Diffusion with Transaction Costs

J.S. Kennedy

Department of Applied Mathematics, University of Waterloo, Waterloo ON, Canada N2L 3G1, jskennedy@uwaterloo.ca

P.A. Forsyth

David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1,
paforsyt@uwaterloo.ca

K.R. Vetzal

Centre for Advanced Studies in Finance, University of Waterloo, Waterloo ON, Canada N2L 3G1, kvetzal@uwaterloo.ca

If the price of an asset follows a jump diffusion process, the market is in general incomplete. In this case, hedging a contingent claim written on the asset is not a trivial matter, and other instruments besides the underlying must be used to hedge in order to provide adequate protection against jump risk. We devise a dynamic hedging strategy that uses a hedge portfolio consisting of the underlying asset and liquidly traded options, where transaction costs are assumed present due to a relative bid-ask spread. At each rebalance time, the hedge weights are chosen to simultaneously (i) eliminate the instantaneous diffusion risk by imposing delta neutrality; and (ii) minimize an objective that is a linear combination of a jump risk and transaction cost penalty function. Since reducing the jump risk is a competing goal vis-à-vis controlling for transaction cost, the respective components in the objective must be appropriately weighted. Hedging simulations of this procedure are carried out, and our results indicate that the proposed dynamic hedging strategy provides sufficient protection against the diffusion and jump risk while not incurring large transaction costs.

Subject classifications: Finance: asset pricing

Area of review: Financial Engineering

History:

1. Introduction

Due to the incomplete nature of a market containing jumps, dynamically hedging an option under a jump diffusion process is far from a straightforward endeavour. Unlike the complete Black-Scholes model, a continuously rebalanced delta hedge will not result in an instantaneously risk-free portfolio (except in the degenerate case where the contract's payoff is linear in the underlying). For instance, executing the simple delta hedging strategy for an option with a convex payoff will lead to a loss if a jump occurs, regardless of the magnitude or direction of the jump. Therefore, any dynamic hedging procedure implemented within a jump diffusion framework must take into account the jump risk.

A dynamic hedging strategy that can be used under a jump diffusion model was explored in He et al. (2006). This method seeks to mitigate the jump risk by holding instruments in the hedge portfolio that protect against a sudden, extreme movement in the stock price. The weights of the hedging instruments are chosen to (i) enforce any desired constraints, including delta neutrality; and (ii) ensure that if a jump occurs, the change in the value of the entire portfolio is small for a suitable range of jump amplitudes. In He et al. (2006), this dynamic strategy is shown to provide good results when hedging a longer term European straddle and American put using short-term calls and puts. However, no consideration was given to the role of transaction costs: for the frequent rebalancing necessitated by dynamic hedging, these costs may make the procedure prohibitively expensive.

Most of the extant literature on hedging a target contract using other exchange-traded options focuses on static strategies, motivated at least in part by the desire to avoid the high costs of frequent trading. Examples of this type of approach include Derman et al. (1995) and Carr et al. (1998). As the strategy proposed in these papers involves a buy-and-hold portfolio of traded options, it does not incur significant transaction costs. However, this type of approach is not suitable for a wide variety of contracts, such as those with American early exercise provisions or path-dependent features.

Alternatively, a semi-static strategy may also be used to hedge under jump diffusion. With this approach, one chooses hedge portfolio weights that attempt to replicate the value of the target option at a future time (which usually corresponds to the expiry of the shorter term options used for hedging). The hedge will be infrequently rolled over before the target contract expires, thus limiting transaction costs. This procedure was explored in Carr and Wu (2004) and He et al. (2006) under two slightly different forms, and appears to be very effective when hedging vanilla options. Carr and Wu also show how their procedure can be extended to hedge discretely observed path-dependent options, although the technique becomes essentially dynamic when the monitoring frequency is high. Moreover, the application of semi-static methods to contracts with early exercise rights is not clear-cut.

A dynamic hedging strategy can handle contracts with path-dependent features. In the presence of transaction costs, however, the cumulative expense of the necessary updates may become large as the rebalancing frequency increases. The goal of this work is to devise a dynamic hedging strategy that protects against the diffusion and jump risk while not costing too much to maintain.

In this paper, we concentrate on hedging a single-factor jump diffusion process. Alternative approaches which have been suggested as improvements over the benchmark Black-Scholes diffusion model include diffusive stochastic volatility (Heston 1993) and stochastic volatility with jumps in both volatility and asset price (Duffie et al. 2000). Note that under a stochastic volatility model (without jumps), a perfect hedge can in principle be constructed with a dynamically rebalanced portfolio consisting of the underlying and one additional option. It is clear that hedging jumps is more challenging than hedging stochastic volatility. Once we have a method which is effective for hedging asset jumps, we can then extend this idea to the most general case of hedging under a model with stochastic volatility and jumps in both asset price and volatility.

The outline of the paper is as follows. Section 2 provides a brief summary of the dynamic strategy introduced in He et al. (2006) for hedging under jump diffusion in the absence of transaction costs. Section 3 describes how this strategy is related to the previous literature, and discusses existing studies connected to hedging under transaction costs. Section 4 then shows how the objective function in the dynamic strategy of He et al. (2006) can be augmented to include a component that takes into account transaction costs. After introducing the general hedging framework in Section 5, the behaviour of the optimization problem is explored in Section 6 for a specific rebalancing example. Section 7 looks at hedging simulations for both European and American-style claims; a constant relative bid-ask spread is used, as well as a more realistic bid-ask model drawn from market data. Section 8 concludes with a brief summary of our main results.

2. A Dynamic Hedging Strategy Under Jump Diffusion in the Absence of Transaction Costs

A simple delta hedge (i.e. containing the underlying asset and cash) carried out in a jump diffusion setting will eliminate the diffusion risk while ignoring all but the linear component of the jump risk. If there is a continuum of possible jump states, in principle an infinite number of hedging instruments would be needed to entirely eliminate the jump risk. Obviously, this is not possible in practice. The dynamic hedging strategy of He et al. (2006) aims to minimize a measure of the instantaneous jump risk, at each rebalance time, using a finite set of hedging instruments.

In a jump diffusion model with constant volatility, the evolution of the underlying asset S is governed by

$$\frac{dS_t}{S_{t^-}} = (\alpha - q - \kappa\lambda) dt + \sigma dZ_t + d\left(\sum_{i=1}^{\pi_t} (J_i - 1)\right), \quad (1)$$

where t^- denotes the instant immediately before time t , α is the instantaneous expected rate of return, q is the dividend yield, and σ is the diffusive volatility. In addition, π_t is a Poisson process with intensity $\lambda > 0$, and J_i are independent and identically distributed positive random variables representing the jump amplitudes, with distribution $g(\cdot)$ and mean $\kappa + 1$. To limit notational complexity, we will use the shorthand

$$\Delta F d\pi = d\left(\sum_{i=1}^{\pi_t} (F(J_i S_{t_i}) - F(S_{t_i}))\right) \quad (2)$$

in the remainder of this paper. Note that the general stochastic process (1) encompasses both the real-world (\mathbb{P} measure) process that represents how the market actually evolves, and the risk-adjusted (\mathbb{Q} measure) process used for no-arbitrage valuation. If required, the appropriate superscript (i.e. either \mathbb{P} or \mathbb{Q}) is appended to the above quantities to distinguish the measure with which it is associated. For those parameters that are invariant to changes of measure, such as σ , the superscript may be omitted without ambiguity. Furthermore under the risk-adjusted process, $\alpha^{\mathbb{Q}} = r$, the risk-free rate of interest (which is assumed to be non-negative). In practice, the \mathbb{Q} measure parameters may be obtained by calibrating to option prices in the market while, in general, the \mathbb{P} measure parameters are unobservable (but estimable from historical return data for the underlying asset).

Following standard arguments (e.g. Cont and Tankov 2004, Andersen and Andreasen 2000), the value of a European option is given by

$$\frac{\partial V}{\partial \tau} = \frac{\sigma^2 S^2}{2} \frac{\partial^2 V}{\partial S^2} + (r - q - \kappa^{\mathbb{Q}} \lambda^{\mathbb{Q}}) S \frac{\partial V}{\partial S} - rV + \lambda^{\mathbb{Q}} \left(\int_0^\infty V(SJ, \tau) g^{\mathbb{Q}}(J) dJ - V(S, \tau) \right), \quad (3)$$

where T is the expiry date of the contract, $\tau = T - t$, and $g^{\mathbb{Q}}(J)$ is the risk-adjusted distribution of jumps. Defining

$$\mathcal{L}V \equiv \frac{\partial V}{\partial \tau} - \left(\frac{\sigma^2 S^2}{2} \frac{\partial^2 V}{\partial S^2} + (r - q - \kappa^{\mathbb{Q}} \lambda^{\mathbb{Q}}) S \frac{\partial V}{\partial S} - (r + \lambda^{\mathbb{Q}}) V + \lambda^{\mathbb{Q}} \int_0^\infty V(SJ, \tau) g^{\mathbb{Q}}(J) dJ \right) \quad (4)$$

and letting V_e denote the early exercise payoff of an American claim, the price of an American option is given by (Barles 1997)

$$\min(\mathcal{L}V, V - V_e) = 0. \quad (5)$$

Assume a bank has sold a derivative V , and now holds a short position $-V$ in that contract. The bank establishes a hedge portfolio which contains an amount B in cash, is long e units of the underlying asset S , and long N additional hedging instruments $\vec{I} = [I_1, I_2, \dots, I_N]$ (written on the underlying) with weights $\vec{\phi} = [\phi_1, \phi_2, \dots, \phi_N]$. When combined with the short position in the target contract $-V$, the resulting overall hedged position has value

$$\Pi = -V + eS + \vec{\phi} \cdot \vec{I} + B.$$

To represent changes in the components of Π due to a jump of size J , we use the notation $\Delta V = V(JS) - V(S)$, $\Delta S = S(J - 1)$ and $\Delta \vec{I} = \vec{I}(JS) - \vec{I}(S)$.

By making the overall hedged position delta neutral, i.e.

$$-\frac{\partial V}{\partial S} + e + \vec{\phi} \cdot \frac{\partial \vec{I}}{\partial S} = 0, \quad (6)$$

it is shown in Appendix A that the instantaneous change in the value of the overall hedged position is

$$d\Pi = r\Pi dt + \lambda^{\mathbb{Q}} dt \mathbb{E}^{\mathbb{Q}} \left[\Delta V - (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right] + d\pi^{\mathbb{P}} \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right]. \quad (7)$$

Therefore, the value of the overall hedged position grows at the risk-free rate, but has additional terms due to the jump component:

$$\underbrace{\lambda^{\mathbb{Q}} dt \mathbb{E}^{\mathbb{Q}} \left[\Delta V - (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right] + d\pi^{\mathbb{P}} \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right]}_{\text{instantaneous jump risk}}. \quad (8)$$

The first constituent of the jump risk is deterministic, while the second part is stochastic as it depends on whether or not the Poisson event occurs over the instant dt .

The diffusion risk has been removed by the imposition of delta neutrality. When a jump occurs ($d\pi^{\mathbb{P}} = 1$), the change in the overall hedged position due to this jump is given by the random variable

$$\Delta H_J = -\Delta V + e\Delta S + \vec{\phi} \cdot \Delta \vec{I}. \quad (9)$$


Consider the expression

$$\int_0^{\infty} \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right]^2 W(J) dJ. \quad (10)$$

In Appendix B we consider an idealized trading environment, where $W(J)$ is a *proper weighting function* with respect to both $g^{\mathbb{P}}(J)$ and $g^{\mathbb{Q}}(J)$ (as defined in Appendix B). For example, $W(J) \geq g(J)$ guarantees that W is a proper weighting function with respect to g . In Appendix B we demonstrate that, by imposing delta neutrality and making the integral (10) sufficiently small at each instant, the variance of the terminal hedging error can be made small. Since only a moderate number of hedging instruments will be used in practice, it will be impossible to make the expression in (10) arbitrarily small. Consequently, the dynamic strategy of He et al. (2006) selects the weights e and $\vec{\phi}$ that minimize (10), while respecting delta neutrality and any other imposed constraints.

The weighting function $W(J)$ is set by the hedger. One possible choice for this function is the distribution of jumps observed in the market, but since this requires knowledge of the \mathbb{P} measure, it would often have to be approximated. In order to ensure that the jump risk can become small through bounding the integral in (10), $W(J)$ must be a proper weighting function as discussed above. However, since only a small number of hedging instruments are used in practice, $W(J)$ need not be a proper weighting function. Therefore a more practical weighting function would be a uniform density, set to non-zero for the range of jumps deemed likely.

The forthcoming numerical examples will employ the uniform-like weighting function plotted in Figure 1, which is constant between $\frac{1}{5} \leq J \leq \frac{9}{5}$ and extends linearly down to zero outside this range on either side.¹ This weighting function encapsulates a lack of knowledge pertaining to the jumps under the \mathbb{P} measure: with no information about which jump sizes are more likely than others, broad protection is sought for $J \in [0, 2]$. In general, when a uniform-like weighting function with support $[J_{\min}, J_{\max}]$ is used, J_{\min} should be chosen close to zero in order to protect against all downward jumps. Selecting J_{\max} is a bit more difficult. The weighting function of Figure 1 implies jumps of size $J > 2$ will not be taken into consideration when the hedge portfolio is formed/rebalanced—if it is suspected that $g^{\mathbb{P}}(J)$ has a slowly decaying right tail, a higher value of J_{\max} may be needed. The

Figure 1 Uniform-like weighting function used in the jump risk objective (10).

choice of $[J_{\min}, J_{\max}]$ can have a noticeable effect on the hedging performance: if the range is too large (e.g. $[0, 10]$), protection will be wasted on highly unlikely jump events, while a narrow band (e.g. $[0.8, 0.9]$) may ignore probable jump amplitudes. Tests reported in He et al. (2006) show that a uniform-like weighting function generally performs well, and is much better than a poor guess for the \mathbb{P} measure jump distribution.

In summary, the dynamic hedging procedure of He et al. (2006) computes the hedge weights $\{e^*, \vec{\phi}^*\}$ that solve the constrained optimization

$$\begin{aligned} & \arg \min_{\{e, \vec{\phi}\}} \int_0^\infty \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right]^2 W(J) dJ \\ & \text{subject to} \\ & e + \vec{\phi} \cdot \frac{\partial \vec{I}}{\partial S} = \frac{\partial V}{\partial S}. \end{aligned} \quad (11)$$

In the sequel, we will refer to this strategy as jump risk hedging. Other constraints, such as gamma neutrality, may also be imposed.

The above theoretical framework encompasses any jump diffusion process. Furthermore, since many infinite activity Lévy processes can be approximated as a jump diffusion (Asmussen and Rosiński 2001, Cont and Tankov 2004), this hedging strategy would also be applicable in these situations. For the numerical examples of this paper, we shall use a jump diffusion model where $\log J$ is normally distributed with constant mean μ and standard deviation γ .

3. Dynamic Hedging: Relationship to Previous Work

The idea of using a finite number of options as part of a dynamic hedging strategy to minimize jump risk has also been suggested by Bates (1988) and Andersen and Andreasen (2000). Note that neither of these papers provided any tests of the strategy. In Bates (1988), the hedge portfolio is selected so that (for infinitesimal hedging intervals) the diffusion risk is identically zero and the expected value of the local jump risk is minimized. In our notation, this amounts to

$$\{e^*, \vec{\phi}^*\} = \arg \min_{\{e, \vec{\phi}\}} \mathbb{E} \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right]^2, \quad (12)$$

subject to the delta-neutral constraint (6). Bates (1988) suggests that the expectation in equation (12) can be taken w.r.t. either \mathbb{P} or \mathbb{Q} . A similar procedure is suggested in Andersen and Andreasen (2000), but they recommend using $\mathbb{E}^{\mathbb{P}}$ in equation (12), along with the additional constraint

$$\mathbb{E}^{\mathbb{P}} \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right] = 0. \quad (13)$$

Both of these approaches are similar in spirit to the strategy of He et al. (2006): the jump risk is minimized in some sense, subject to the delta-neutral constraint that eliminates the instantaneous diffusion risk.

Cont et al. (2005) consider the quadratic program

$$\arg \min_{\{e_t, \vec{\phi}_t, V_0\}} \mathbb{E}^{\mathbb{Q}} [\Pi_T - V_T]^2,$$

where V_T is the payoff of the option being hedged and $\{e_t, \vec{\phi}_t\}$ is the trading strategy over $[0, T]$ that minimizes the objective. Also, V_0 is the initial capital. This program employs a global criterion, as its objective function is an expectation involving the terminal hedging error. However, since this expectation is taken w.r.t. the pricing measure \mathbb{Q} , the solution reduces to a local risk minimization (Cont and Tankov 2004). In our notation (assuming finite activity jumps), the optimal hedging weights at each instant are given by (Cont et al. 2005)

$$\{e^*, \vec{\phi}^*\} = \arg \min_{\{e, \vec{\phi}\}} \left[\lambda^{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}} \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right]^2 + \sigma^2 S^2 \left[-\frac{\partial V}{\partial S} + e + \vec{\phi} \cdot \frac{\partial \vec{I}}{\partial S} \right]^2 \right], \quad (14)$$

and V_0 is the option price. The jump risk hedging procedure of He et al. (2006) is similar to that of Cont et al. (2005), except that in (11) the diffusion risk is explicitly eliminated and a proper weighting function is employed to minimize the \mathbb{P} measure local risk. If the delta-neutral condition (6) is imposed in (14) and the weighting function in (11) is set to $g^{\mathbb{Q}}(J)$, then the two procedures are identical.

For infrequent jumps ($\lambda^{\mathbb{Q}} \ll 1$), the diffusion component in (14) may swamp the jump component, such that the protection against jumps may be lacking. On the other hand, by explicitly enforcing delta neutrality in (11), the jump protection is the same regardless of how often the jump events occur; no estimate of the jump intensity is required. Furthermore, only the weighting function has to be set: we can use $g^{\mathbb{P}}(J)$, $g^{\mathbb{Q}}(J)$, or even a uniform-like density as in Figure 1. Conversely, by not explicitly enforcing the delta-neutral constraint in (14), there exists an extra degree of freedom, which may be important if using a small number of hedging instruments. For example, if only the underlying is used in the hedge, the optimization (11) will trivially yield the delta-neutral position, while the weight in the underlying computed by solving (14) will take into account both the diffusion and jump risk. For infrequent jumps, however, the weight in the underlying will essentially be a small correction to the delta-neutral position. Note that in the case of only one possible jump size, both (11) and (14) will yield the perfect hedge that uses the underlying and one additional option.

None of the dynamic strategies discussed so far take into account transaction costs: since more than the underlying is required to hedge jumps, the effect of these costs becomes important. For complete-market one-factor diffusion models, there is an extensive body of literature that incorporates transaction costs into the pricing and hedging framework. For a market where the underlying evolves according to geometric Brownian motion (GBM) and transaction costs are proportional to the value of the transaction in the asset, a hedging argument similar to that used for deriving the Black-Scholes equation yields a nonlinear partial differential equation (PDE) (Hoggard et al. 1994). When pricing individual calls and puts, this nonlinearity is absent due to

the single-signed gamma, such that only a volatility adjustment is required in the Black-Scholes PDE. Leland (1985) considers this problem, and develops a hedging strategy based on the delta of the option values found using the augmented volatility. The hedging error in Leland's model as the rebalancing interval $\Delta t \rightarrow 0$ is a non-trivial function, and almost surely negative (Kabanov and Safarian 1997, Grandits and Schachinger 2001). For a jump diffusion model that uses Merton's original assumption of diversifiable jump risk, Mocioalca (2003) derives a volatility adjustment for pricing calls and puts that is analogous to Leland's. However, the delta hedge motivated by this analysis has the same drawback as in a jump model with no transaction costs, namely that it does not provide adequate protection against jump risk.

The approach outlined above is local in time. Global-in-time methods that use utility indifference pricing can also be employed. Davis et al. (1993) consider a GBM model, with proportional transaction costs, where the underlying can be traded continuously. Within this utility framework, a Hamilton-Jacobi-Bellman (HJB) system can be solved to establish the optimal hedging strategy. The solution defines a buy, sell and no-trade region: when the underlying enters either the buy or sell region, the hedger performs the necessary transaction that brings the position in the underlying back onto the boundary of the no-trade region.

An analogous stochastic control program could be set up for hedging under jump diffusion with transaction costs. Keppo and Peura (1999) consider a problem similar to this, only within a GBM model. The authors use approximations to yield an augmented formulation, not involving HJB equations, that can be treated numerically: a quadratic program is obtained, where the vector and matrix components of the objective are estimated using Monte Carlo simulation. The system of HJB equations would involve the values of the target option and hedging instruments, as well as the controls that yield the optimal hedging strategy. Clearly, any numerical solution involving the HJB equations would be computationally intractable for any more than a few hedging instruments.

4. Incorporating Transaction Costs

The jump risk hedging strategy (11) does not take into account transaction costs, so it may yield hedge portfolio weights that require expensive trading to implement. In this section we show how the objective function in (11) can be modified so that these costs are considered when rebalancing. For our purposes, transaction costs refer to the difference between the bid/ask price and the theoretical value of the security (underlying asset or option). Brokerage commissions and other fees are ignored.

We assume the following scenario: using the linear pricing equation (3), a hedger fits option pricing parameters to the midpoint option values observed in the market. Then, a hedging strategy is constructed using a simple market model of bid-ask spreads. This approach preserves the property that the prices are linear in the numbers bought/sold, and it makes minimal assumptions about a model of bid-ask spreads. In contrast, a nonlinear pricing equation (e.g. Hoggard et al. 1994, Mocioalca 2003) in effect attempts to predict the bid-ask spread for options. Furthermore, for nonlinear pricing equations, the value of the overall hedged position $-V + eS + \vec{\phi} \cdot \vec{I} + B$ is not the same as the sum of the values of the individual components. Linear pricing rules are the market standard for the simple contracts we use for hedging (Cont and Tankov 2004).

We incorporate transaction costs using a relative bid-ask spread. We will assume that the options to be used for hedging can be characterized by a single strike price K —this will include vanilla puts and calls.² This assumption will allow the relative bid-ask spread for a range of options to be modelled as a function of moneyness K/S . The dollar spread is assumed to be symmetric around the theoretical option value found using the pricing equation (3). Furthermore, the (relative) bid-ask spread will be quoted as a fraction of the option price. For example, with a relative bid-ask spread of 0.10 on an option with theoretical value \$5.00, the bid price is \$4.75 and the ask price is \$5.25. In other words, the hedger will have to pay \$5.25 to purchase the option with theoretical

value \$5.00, and would receive \$4.75 if the option were to be sold. If the weight of an instrument is $\rho(t_{n-1})$ before rebalancing and $\rho(t_n)$ after rebalancing, the total cost of the transaction is

$$|\rho(t_n) - \rho(t_{n-1})| \times \frac{\text{BA}}{2} \times \text{Instrument Value},$$

where BA denotes the relative bid-ask spread.

The quadratic objective (10) facilitates a straightforward application of Lagrange multipliers within the original optimization problem (11), so a similar quadratic representation for handling transaction costs is desirable. One suitable candidate is the sum of the squares of transaction costs

$$\left[\left(e(t_n) - e(t_{n-1}) \right) \times \frac{\text{BA}_S}{2} \times S \right]^2 + \sum_{j=1}^N \left[\left(\phi_j(t_n) - \phi_j(t_{n-1}) \right) \times \frac{\text{BA}_j}{2} \times \text{OptVal}_j \right]^2, \quad (15)$$

where ϕ_j is the weight in the j^{th} hedging option, which has value OptVal_j and relative bid-ask spread BA_j , and BA_S is the relative bid-ask spread for the underlying.

Reducing transaction costs and minimizing jump risk are competing goals, so this problem falls under the rubric of multi-objective optimization. One typical way of handling such a problem is to weight the objectives by a set of coefficients that sum to unity. If ξ is the weight on the jump risk exposure (10) and $1 - \xi$ is the weight on the transaction cost objective (15), the resulting optimization problem is

$$\begin{aligned} \arg \min_{\{e(t_n), \vec{\phi}(t_n)\}} & \left\{ \int_0^\infty \left[-\Delta V + (e(t_n)\Delta S + \vec{\phi}(t_n) \cdot \Delta \vec{I}) \right]^2 W(J) dJ \right\} \\ & + (1 - \xi) \left\{ \left[\left(e(t_n) - e(t_{n-1}) \right) \times \frac{\text{BA}_S}{2} \times S \right]^2 + \sum_{j=1}^N \left[\left(\phi_j(t_n) - \phi_j(t_{n-1}) \right) \times \frac{\text{BA}_j}{2} \times \text{OptVal}_j \right]^2 \right\} \end{aligned}$$

subject to

$$e(t_n) + \vec{\phi}(t_n) \cdot \frac{\partial \vec{I}}{\partial S} = \frac{\partial V}{\partial S}. \quad (16)$$

In Appendix B we demonstrate that for a properly chosen weighting function, if the overall position is delta neutral and the objective in (16) is made sufficiently small at each instant within a continuously rebalanced hedge, then the variance of the terminal hedging error can be made small. Since this will most likely be difficult in practice, we aim to make the objective as small as possible at each rebalancing time. Appendix B also shows that, under ideal conditions, ξ should be $\mathcal{O}(\Delta t^2)$, where Δt is the rebalancing interval. The optimization problem and the associated hedging simulations will be considered for a range of ξ values.

With the influence parameter ξ and the weighting function $W(J)$ specified by the hedger, the optimization problem (16) may be solved using Lagrange multipliers. This results in a linear system for the unknowns $e(t_n)$ and $\vec{\phi}(t_n)$, which are the weights that the hedge portfolio should have after rebalancing at t_n . The entries of the linear system which involve correlation-type integrals are precomputed using efficient FFT techniques (see d'Halluin et al. (2005) and He et al. (2006) for further details). The linear system may be poorly conditioned in certain situations, which manifests itself by unstable (as a function of time) portfolio weights e and $\vec{\phi}$. In order to avoid this behaviour, we use a Truncated Singular Value Decomposition (TSVD) (Hansen 1987): small singular values are set to zero—where the cutoff is imposed via a user controlled parameter—and the modified decomposition is used to solve the system in the standard way (Press et al. 1993).

Before considering a set of hedging simulations, we will introduce the general hedging procedure and investigate the behaviour of the optimization problem (16) that lies at the heart of the strategy.

5. The Hedging Procedure

To initiate the hedging strategy, an appropriate weighting function $W(J)$ must first be selected and the influence parameter ξ fixed. The initial hedge portfolio weights $e(0)$ and $\vec{\phi}(0)$ are chosen to solve the optimization (16), with the weights from the “previous” rebalancing set to 0. These trades must be financed by the bank account: at time zero the amount of cash $B(0)$ equals the aggregate cost of the long and short positions of the hedging instruments, plus the initial value of the option $V(S_0, 0)$ received by the hedger upon selling it.³ At each rebalance time t_n the hedge portfolio weights are recalculated by solving the optimization problem (16). The long position in the underlying asset is subsequently updated by purchasing $e(t_n) - e(t_{n-1})$ shares, where $e(t_n)$ is the new computed weight and t_{n-1} denotes the time of the last rebalancing. The long positions in the hedging options are updated similarly by purchasing $\vec{\phi}(t_n) - \vec{\phi}(t_{n-1})$ units. These trades are again financed by the cash account, which after rebalancing contains

$$B(t_n) = \exp\{r(t_n - t_{n-1})\}B(t_{n-1}) - \left[e(t_n) - e(t_{n-1}) \right] \left[1 + \operatorname{sgn}(e(t_n) - e(t_{n-1})) \frac{\text{BA}_S}{2} \right] S_{t_n} \\ - \sum_{j=1}^N \left[\phi_j(t_n) - \phi_j(t_{n-1}) \right] \left[1 + \operatorname{sgn}(\phi_j(t_n) - \phi_j(t_{n-1})) \frac{\text{BA}_j}{2} \right] I_j(S_{t_n}, t_n),$$

where it is assumed the hedging instruments have non-negative value. The above be written in the alternative form

$$B(t_n) = \exp\{r(t_n - t_{n-1})\}B(t_{n-1}) - \left[e(t_n) - e(t_{n-1}) \right] S_{t_n} - \left[\vec{\phi}(t_n) - \vec{\phi}(t_{n-1}) \right] \cdot \vec{I}(S_{t_n}, t_n) \\ - \underbrace{\left[\left| e(t_n) - e(t_{n-1}) \right| \left(\frac{\text{BA}_S}{2} \right) S_{t_n} + \sum_{j=1}^N \left| \phi_j(t_n) - \phi_j(t_{n-1}) \right| \left(\frac{\text{BA}_j}{2} \right) I_j(S_{t_n}, t_n) \right]}_{\text{transaction costs}}$$

to make explicit the cost of transactions due to the bid-ask spread. In our simulations, the stock value S_{t_n} is taken from a randomly generated asset price path, whose evolution is governed by the real-world measure. When liquidating the hedge portfolio to cover the short position $-V$, we must take into account transaction costs. The value of the overall hedged position at exercise/expiry T^* is given by

$$\Pi(T^*) = -V(S_{T^*}, T^*) + B(t') \exp\{r(T^* - t')\} + e(t') S_{T^*} + \vec{\phi}(t') \cdot \vec{I}(S_{T^*}, T^*) \\ - \left| e(t') \right| \left(\frac{\text{BA}_S}{2} \right) S_{T^*} - \sum_{j=1}^N \left| \phi_j(t') \right| \left(\frac{\text{BA}_j}{2} \right) I_j(S_{T^*}, T^*),$$

where t' is the time of the last rebalancing. When hedging a European option, this liquidation will most likely coincide with the expiry of shorter term options, such that the transaction costs will usually only come from disposing of the underlying.

We are interested in the value of the overall hedged position when V is exercised or expires. Ideally, the portfolio has a value of zero as this implies perfect replication. However due to transaction costs, the presence of jumps, and the discrete nature of the rebalancing, this obviously will not be the case. The value of the hedged portfolio upon liquidation is the hedging error. One common metric for the hedging error at the exercise/expiry time T^* is the relative profit and loss (P&L):

$$\text{Relative P\&L} = \frac{\exp\{-rT^*\} \Pi(T^*)}{V(S_0, 0)}. \quad (17)$$

Table 1 The pricing \mathbb{Q} measure and real-world \mathbb{P} measure that characterize the jump diffusion model, where $\log(J) \sim N(\mu, \gamma)$.

Probability Measure	λ	μ	γ	σ	α
Risk-adjusted (\mathbb{Q})	0.1000	-0.9200	0.4250	0.2000	0.0500
Objective (\mathbb{P})	0.0228	-0.5588	0.4250	0.2000	0.1779

The dividend yield $q = 0$ and $\alpha^{\mathbb{Q}} = r = 0.05$.

Table 2 Instruments in the overall hedged position.

Instrument	Initial		Value at	Weight at	Value at
	Maturity	Strike	$t = 0, S = \$100$	$t = 0, S = \$100$	$t = 0.05, S = \$106.5$
Straddle	1 year	\$100.00	\$21.41	-1.0000	\$24.05
Underlying	n.a.	n.a.	\$100.00	-0.6360	\$106.50
Put	0.25 years	\$80.00	\$0.91	1.2881	\$0.67
Put	0.25 years	\$90.00	\$1.53	-0.9367	\$0.91
Call	0.25 years	\$100.00	\$5.34	1.9197	\$9.38
Call	0.25 years	\$110.00	\$1.50	-0.9288	\$3.23
Call	0.25 years	\$120.00	\$0.28	0.6032	\$0.69

All options have European-style exercise rights.


6. A Representative Optimization Problem

At each rebalance time, our goal is simple: choose hedge portfolio weights that impose delta neutrality and reduce the jump risk, while keeping transaction costs as small as possible. The tradeoff between the minimization of jump risk and the reduction of transaction costs is controlled by the influence parameter ξ : for $\xi = 1$ we are only concerned with jump risk, while $\xi = 0$ corresponds to total concentration on transaction costs. A solution x^* is said to be Pareto optimal if any perturbation of x^* required to improve one of the component objectives can only be made at the expense of another objective. The collection of all such solutions is the Pareto optimal set, and the associated set of component objective values $\vec{F}(x^*)$ is the Pareto front. The most well-known example of a Pareto front in finance is the efficient frontier of the Markowitz model.

We consider a specific rebalancing example as a means to study the behaviour of the objective function in (16) for different values of the influence parameter ξ . Before doing so, however, it is necessary to provide some details regarding the real-world and pricing measures used for our tests. The values that characterize these measures are reported in Table 1. For the \mathbb{Q} measure, we use values quite similar to those reported by Andersen and Andreasen (2000), which were found by calibrating to observed prices of S&P 500 index options. To obtain the \mathbb{P} measure parameters, we transform the \mathbb{Q} measure using the power utility equilibrium model of Naik and Lee (1990) and Bates (1991). The linkage is based on the coefficient of relative risk aversion, which we assume is equal to 2. Details are provided in Appendix C. Note that the relation between \mathbb{Q} and \mathbb{P} is used simply as a means of obtaining the real-world parameters in a somewhat formal way. The invocation of power utility has no connection to our hedging criterion.

Assume a financial institution has sold an at-the-money one-year European straddle, where $S_0 = \$100$. To hedge its exposure, positions are taken in the underlying and five put and call options with three months until expiry—the instruments in the overall hedged position are given in Table 2. The initial hedge portfolio weights are found by solving (16) with $\xi = 1$. At $t = 0.05$ with the underlying at $S = \$106.5$, the hedge portfolio is to be rebalanced: the optimization problem (16) is solved for the range of influence parameter values $\xi \in [0, 1]$. The relative bid-ask spread is fixed at 0.10 for all hedging options and 0.002 for the underlying.

The hedge portfolio weights found from solving the optimization problem with varying ξ are used to compute the jump risk objective (10) and the transaction cost objective (15), and these are plotted together in Figure 2. The exposure to jump risk is smallest for $\xi = 1$, but gets large

Figure 2 The Pareto optimal front for the optimization (16) at $t = 0.05$, $S = \$106.5$.

as the influence parameter decreases. The opposite behaviour is observed for the transaction cost objective. This curve displays how the “best” solution is subjective: the interested party should opt for a solution from the Pareto optimal set, but the specific choice will be based on other considerations.

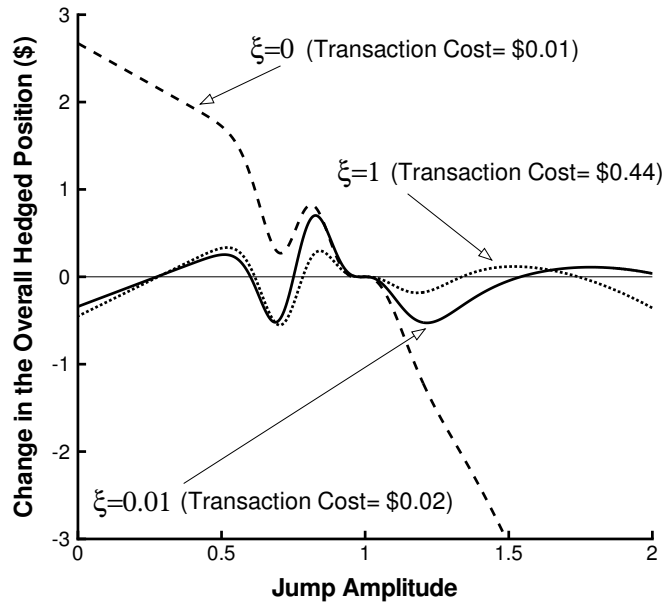
The value from the jump risk objective (10) is a rather blunt statistic, as it condenses the risk from a continuum of possible jumps into a single number. Nonetheless, the protection afforded against jump risk by a hedge portfolio may be visualized by considering a plot of ΔH_J . Recall that ΔH_J in (9) represents the change in the overall hedged position due to a jump of size J . Therefore, the desired behaviour of this curve is for it to remain very close to zero, as this corresponds to little change in the overall hedged position due to a jump. Figure 3 presents the jump risk profile for three values of the influence parameter ξ . The best possible curve in terms of minimizing jump risk is for $\xi = 1$, but the associated hedge portfolio weights are selected in a manner that ignores transaction costs. For $\xi = 0$ only transaction costs are considered; these costs are quite low, but the protection against jumps is not very good. The third curve, corresponding to $\xi = 0.01$, offers the middle ground we seek, namely low transaction costs with good protection against jumps. The conclusion that may be drawn from Figure 3 is that hedge portfolio weights can be found which do a good job of adequately satisfying both objectives.

7. Hedging Simulations

7.1. A Simple Hedging Example: Five Hedging Options

To provide a simple illustration of the hedging strategy, we extend the example of Section 6. A one-year European straddle is to be hedged over its lifetime. Initially, the underlying along with puts of strike $K = [80, 90]$ and calls with strike $K = [100, 110, 120]$ are used, where all of the hedging options have three months until maturity. These options are traded until they expire, at which time new options are purchased, and these new options have the same strikes and time to maturity as the initial set. The \mathbb{P} and \mathbb{Q} measures employed are in Table 1. The \mathbb{P} measure is used to simulate the path of the underlying, and the parameters are unobservable to the hedger. The hedger knows

Figure 3 Change in the overall hedged position due to a jump.



Note. Assumes a jump occurs an instant after rebalancing at $t = 0.05$, $S = \$106.5$. The curves correspond to the hedge portfolio weights found using three different values of the influence parameter ξ . The rebalancing cost of forming the hedge portfolio associated with each profile is given by “Transaction Cost”.

the \mathbb{Q} measure, and uses it to price. Each simulation set consists of 250,000 individual simulations, meaning a total of about 5,700 jumps are expected over each set. Note that, in general, there is little difference between the results for 100,000 and 250,000 simulations.

7.1.1. No Transaction Costs We first investigate a hedging example in which financial instruments may be traded without incurring transaction costs. Consider the case where only the underlying is used in a delta hedge, which is rebalanced every 0.025 years. The results, in the form of summary statistics for the relative P&L, are contained in the first row of Table 3. The outliers of the distribution are important, as they give an indication of the protection against jumps. The 0.02% and 0.2% percentiles are very negative in this case, corresponding to the large losses that often result when a jump occurs (recall that since the straddle has a convex payoff, every jump results in a loss). The fact that the mean is positive may seem surprising, but it is a simple consequence of the \mathbb{Q} measure being more “pessimistic” than the \mathbb{P} measure (i.e. the \mathbb{Q} measure parameters imply more frequent jumps with, on average, larger drops in the underlying). A further discussion of delta hedging under jump diffusion can be found in He et al. (2006). When five options are included in the hedge along with the underlying, the weights are chosen using the jump risk hedging strategy—represented by the optimization in (11)—that is employed when transaction costs are not present. Compared to the delta hedge, this procedure dramatically reduces the exposure to jump risk, as demonstrated by the results in the second row of Table 3.

7.1.2. Transaction Costs Present, but Ignored We next consider a set of simulations where transaction costs are present, with a constant relative bid-ask spread of 0.10 for the hedging options and 0.002 for the underlying. The results for the delta hedge, presented in the first row of Table 4, are very similar to those when no transaction costs are incurred (first row of Table 3), although there is a small negative change in the mean due to the cost of trading. Since trades involving the underlying are rather inexpensive, this small movement is not surprising. When five

Table 3 Relative P&L if there are no transaction costs.

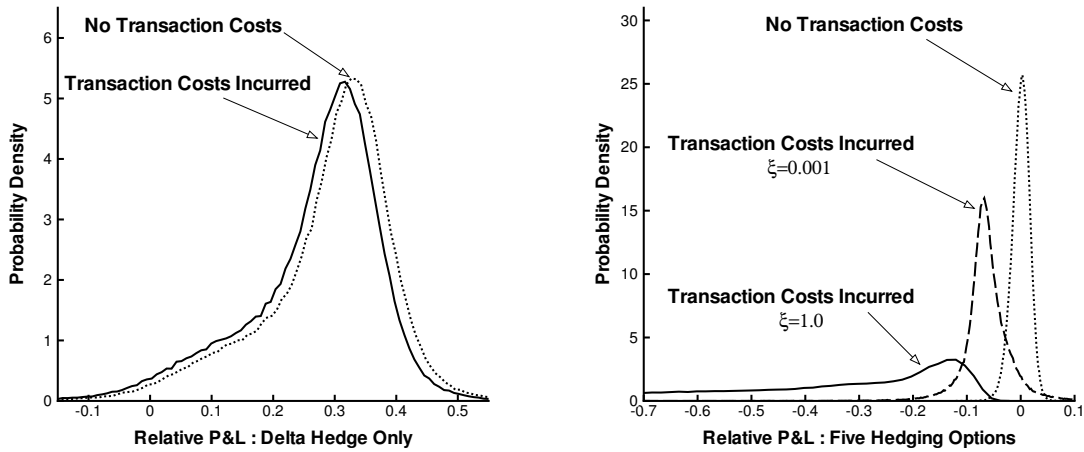
Hedging Strategy	Mean	Std. Dev.	Percentiles			
			0.02%	0.2%	99.8%	99.98%
Delta hedge	0.2452	0.3845	-5.6046	-3.8915	0.5503	0.6241
Five hedging options	0.0002	0.0166	-0.0792	-0.0577	0.0460	0.0695

The weights are chosen by solving the optimization (11). The option being hedged is a one-year European straddle.

Table 4 Relative P&L when transaction costs exist, but are ignored.

Hedging Strategy	Mean	Std. Dev.	Percentiles			
			0.02%	0.2%	99.8%	99.98%
Delta hedge	0.2244	0.3845	-5.6177	-3.9040	0.5289	0.6039
Five hedging options	-0.3822	0.2715	-2.0969	-1.4103	-0.0579	-0.0446

The weights are chosen by solving the optimization (16) with $\xi = 1$. The option being hedged is a one-year European straddle.

Figure 4 Distributions of relative P&L for hedging the one-year European straddle.

Note. When $\xi = 1.0$, the transaction costs are ignored in the optimization (16), while for $\xi = 0.001$ both transaction costs and jump risk are taken into account.

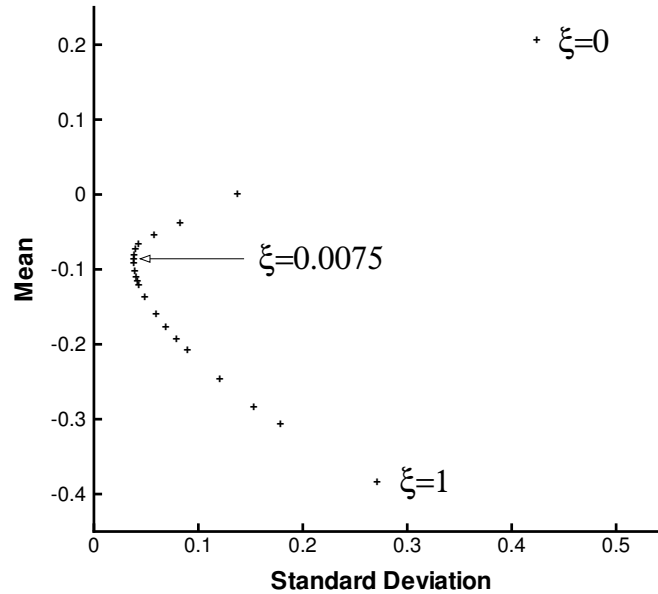
options are included in the hedge along with the underlying, the weights are chosen in a manner that ignores transaction costs. This strategy of using five hedging options yields results that are very poor—the delta hedge is better in this instance, as indicated by the statistics in Table 4. Even though the hedge will help protect the overall position when a jump occurs, the cost of the required transactions may be very high. For this example, the jump risk hedging strategy with five options performs quite well until transaction costs are introduced, at which point it becomes essentially useless.

7.1.3. Transaction Costs Present, and Taken into Account We now carry out the simulations again, only this time using the optimization problem (16) that takes into account both jump risk and transaction costs. The simulations are performed using the influence parameters

$$\xi \in [0, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 0.0025, 0.005, 0.0075, 0.01, 0.02, 0.03, 0.04, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.75, 0.9, 0.95, 1]. \quad (18)$$

For each of the 250,000 simulations, the same value of the influence parameter will be used throughout; for example, if forty rebalances are carried out over the entire hedging horizon, the same value

Figure 5 Mean and standard deviation of the relative P&L, with varying ξ , for hedging the one-year European straddle.



of ξ is used for each of the forty individual optimization problems. The criteria used for deciding which ξ is best for hedging will be the mean and standard deviation of the relative P&L.

For each value of the influence parameter ξ in (18), the mean and standard deviation of the relative P&L is plotted in Figure 5. The results for the extreme values of $\xi = 0$ and $\xi = 1$ are poor. If the standard deviation is used as the sole criterion to select the best value of the influence parameter, the choice is $\xi = 0.0075$. This corresponds to a standard deviation of 0.0384 and a mean of -0.0843. However, for $\xi = 0.001$ the standard deviation is only slightly higher at 0.0429, while the mean of -0.0632 is much better. The distribution of the relative P&L for this case is the middle density in the right panel of Figure 4. A mean of -0.0632 translates to \$1.35 in monetary terms: if the hedger charges this as a premium over and above the theoretical price of \$21.41, the simulations will have a zero mean. This is slightly higher than the 5% premium (i.e. half of the relative bid-ask spread of 10%) assumed for quarter-year vanilla options in the hedging portfolio.

We may conclude that, for this example, it is possible to select hedge weights that provide sufficient protection against jump risk while not incurring large transaction costs.

REMARK 1 (USING THE TSVD). If a relatively high cutoff is used within the TSVD when $\xi = 1$, the strategy may produce reasonable results, even though transaction costs are not taken into account when choosing the hedge weights. In this case, the TSVD solution procedure returns a vector of weights with a small norm—an ideal way to keep transaction costs down—while still providing adequate protection against jump risk. In general, any strategy which uses a regularization method for determining the portfolio weights will tend to keep transaction costs under control (e.g. the strategy in Cont et al. (2005)).

7.2. Varying the Rebalancing Frequency and Number of Options

Up to this point the hedge portfolio has consisted of five options and the underlying, and has been rebalanced every 0.025 years. We now vary the hedge portfolio composition—using three, five and seven options—and the frequency of rebalancing. When seven options are employed, those

Table 5 Relative P&L for different rebalancing frequencies and a varying number of hedging options in the hedge portfolio.

No. of Hedging Options	Rebalance Interval	Std.		
		Best ξ	Dev.	Mean
3	0.0125	0.001	0.0450	-0.0991
	0.025	0.001	0.0529	-0.0755
	0.05	0.0025	0.0641	-0.0715
5	0.0125	0.0025	0.0318	-0.0864
	0.025	0.0075	0.0384	-0.0843
	0.05	0.05	0.0453	-0.0934
7	0.0125	0.005	0.0363	-0.0995
	0.025	0.02	0.0424	-0.0964
	0.05	0.1	0.0501	-0.0981

The influence parameter from the discrete set (18) that yields the lowest standard deviation is termed the best ξ . The option being hedged is a one-year European straddle.

with strikes $K = [70, 80, 90, 100, 110, 120, 130]$ are used, while the three-option portfolio utilizes the middle strikes. The hedge is rebalanced a total of twenty, forty and eighty times over the one-year investment horizon. The hedging simulations are carried out under the same guidelines as before, and the results are presented in Table 5.

The “best” value of the influence parameter indicated in Table 5 comes from the simulation set that yields the lowest standard deviation. For a given number of hedging options, the best (i.e. smallest) standard deviation decreases as the rebalancing frequency increases, and the mean does not become considerably more negative. In general, as the rebalancing frequency increases, the best result is achieved by putting more weight on the transaction cost component of the objective.

In Appendix B we examine an idealized continuous trading environment and demonstrate that, by making the jump risk objective (10) and transaction cost objective (15) sufficiently small at each instant, the variance of the terminal hedging error may be made small. In relation to the discrete framework, the bound on the transaction cost objective should be $\mathcal{O}(\Delta t^2)$ as $\Delta t \rightarrow 0$ in order to ensure finite transaction costs (here, Δt is the length of the rebalancing interval). Within our objective function in (16), the appropriate tradeoff between the jump risk and transaction cost can be achieved with an influence parameter ξ that is $\mathcal{O}(\Delta t^2)$. In other words, as the rebalancing frequency increases, more and more weight should be put on the transaction cost component of the objective function. In practice, we clearly will not be able to simultaneously make both component objectives arbitrarily small, so at each rebalance time we simply attempt to make the objective in (16) as small as possible. The results of Table 5 are generally consistent with the theory. For example in the case of seven hedging options, which is our closest approximation to an idealized trading environment, the best value of the influence parameter is approximately $\mathcal{O}(\Delta t^2)$.

7.3. Using Calls and Puts with the Same Strike

For a given strike price, both calls and puts are typically available in the market, so limiting the hedge portfolio to holding either one or the other is not realistic. We therefore consider an augmented version of the example in Section 7.1 by doubling the number of available hedging options, such that there is now access to all European calls and puts of strike $K = [80, 90, 100, 110, 120]$. Note that all other settings, such as the constant relative bid-ask spread of 0.10 for options, remain the same. When only jump risk is considered ($\xi = 1$), the linear system resulting from the application of Lagrange multipliers to (16) is singular. This is due to put-call parity, as the redundancy inherent in this relationship manifests itself as a rank-deficient matrix. For ξ close to unity, we expect the matrix to be ill conditioned. As noted above, using a TSVD is a common way to deal with an ill-conditioned linear system. Nonetheless, the (usually low) range of influence parameters

that give the best hedging results tend to quell the ill conditioning. This is due to the fact the transaction cost component of the objective is not susceptible to degeneracy problems resulting from put-call parity.

For hedging with the ten options above, the influence parameter value $\xi = 0.0025$ yields a mean of -0.0581 and a standard deviation of 0.0228 , which are better results than can be achieved when only the five original options are used. In this case, put-call parity implies that anything achievable with the put, call and underlying can be accomplished with any two of these three instruments. Consequently, the hedger should include both puts and calls of similar strikes in the hedge portfolio.

7.4. A More Realistic Model of Bid-Ask Spreads

The constant relative bid-ask spread assumption is clearly deficient: out-of-the-money options tend to have higher relative bid-ask spreads than in-the-money options. Consider, for example, the 22Oct2005 option prices for Amazon.com, Inc. (AMZN) taken during trading on August 10, 2005; this data is presented in Table 8 of Appendix D. We use this Amazon.com option data to create a representative model of the relative bid-ask spread as a function of moneyness. The relative bid-ask spread for the puts and calls is found via

$$\text{Relative Bid-Ask Spread} = \frac{\text{Dollar Spread}}{\text{Midpoint Price}} = 2 \times \frac{\text{Ask} - \text{Bid}}{\text{Bid} + \text{Ask}}.$$

The discrete data is first smoothed using a simple moving average method, and the complete relative bid-ask spread curves for the calls and puts in Figure 6 are formed by linear interpolation and nearest neighbour extrapolation of this smoothed data. Note that the first four data points for the puts are discarded.

We have flat-topped the curves of Figure 6 in order to avoid unrealistically large values for the relative bid-ask spread. Options with large spreads will not be selected for the hedge portfolio (or will have very small weights), so as long as the transaction costs are included in the objective function, the precise form of the bid-ask spread for far out-of-the-money options should not be very important.

Up to this point, only options of strike $K = [70, 80, 90, 100, 110, 120, 130]$ have been used, regardless of the value of the underlying. More realistically, a wide range of options may be available for hedging. Of course, some of these will not contribute to a significant reduction of jump risk, and so will not be used in substantial amounts. Furthermore, the transaction costs associated with certain option positions may be prohibitive, and consequently these would be avoided. The objective in our optimization problem is designed to deal with these two facets of hedging. Consider the following scenario: in addition to using the relative bid-ask curves of Figure 6, access to a wider range of options is allowed, namely all puts and calls with strikes from \$10 to \$200 in increments of \$10 are available. As such, the linear system used to determine the hedge portfolio weights has dimension 42×42 . Table 6 contains representative results for different numbers of options in the hedging portfolio. Note that six hedging options corresponds to puts and calls of strike $[90, 100, 110]$, ten represents strikes of $[80, 90, 100, 110, 120]$, while fourteen represents $[70, 80, 90, 100, 110, 120, 130]$. We find that a portfolio with a large number of hedging instruments does not outperform a hedge with a smaller number, as long as the smaller hedge contains those short-term options that are best at replicating the target straddle position, i.e. calls and puts with strikes near \$100. The hedging results of Table 6 indicate our procedure can indeed successfully handle a more realistic model of bid-ask spreads.

7.5. An American Example

To demonstrate the applicability of our proposed technique to path-dependent options, we consider hedging an American put over its lifetime. The same \mathbb{P} and \mathbb{Q} parameters as in Table 1 are employed,

Figure 6 Relative bid-ask spread curves drawn from market data.

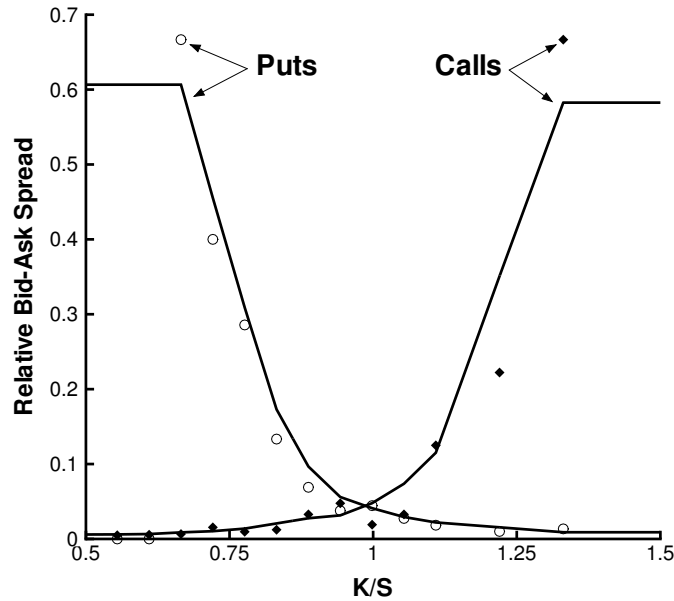


Table 6 Relative P&L for the European straddle hedging example, with the relative bid-ask curves of Figure 6.

No. of Hedging Options	Influence Parameter ξ	Mean	Std. Dev.	Percentiles			
				0.02%	0.2%	99.8%	99.98%
6	0.0001	-0.0593	0.0482	-0.2989	-0.2177	0.2095	0.5007
10	0.001	-0.0639	0.0230	-0.1536	-0.1254	0.0166	0.0661
14	0.0075	-0.0667	0.0206	-0.1257	-0.1153	-0.0130	-0.0016
40	0.02	-0.0770	0.0240	-0.1453	-0.1340	-0.0212	-0.0166

The weights are chosen by solving the optimization (16) for the given ξ .

except that a higher jump arrival rate of $\lambda^{\mathbb{P}} = 0.1$ is used in the simulations. The American prices are computed using the methods described in d'Halluin et al. (2004). The American put to be hedged has a strike of $K = \$100$, a half-year maturity, and is initially at-the-money. American calls, with an initial maturity of three months and strikes from $K = \$10$ to $K = \$200$ in increments of $\$10$, are available as hedging instruments. Quarter-year American puts, with strikes from $K = \$10$ to $K = \$100$ in increments of $\$10$, are also used. At $t = 0.25$ all hedging options are replaced. We must be mindful of the possibility of early exercise for the American puts (since we assume $q = 0$, it will never be optimal to exercise the calls before they expire). The early exercise region of every put is monitored at each tick mark of the simulated asset price path. If the price enters this region, the option should be exercised. If at any time it is deemed optimal for the target half-year put to be exercised, the hedging terminates and the portfolio is liquidated to cover the short position. Furthermore, if it is optimal for one of the shorter-term American hedging puts to be exercised, it is removed from the portfolio and the hedge is rebalanced.

A total of 250,000 simulations are carried out. The hedge is regularly rebalanced at intervals of 0.025 years, and is also rebalanced if a hedging put is removed due to early exercise. To incorporate transaction costs, the bid-ask spread model of Figure 6 is used. Similar to the previous example, using all thirty available options does not outperform a hedge that contains fewer instruments. Some representative results are presented in Table 7. Note the hedge compositions are the same

Table 7 Relative P&L for the American put hedging example, with the relative bid-ask curves of Figure 6.

No. of Hedging Options	Influence Parameter ξ	Mean	Std. Dev.	Percentiles			
				0.02%	0.2%	99.8%	99.98%
0 (delta hedge)	n.a.	0.0538	1.0436	-11.0525	-8.8367	0.7084	0.8265
5	0.0075	-0.0586	0.0245	-0.1733	-0.1378	0.0337	0.1501
8	0.03	-0.0605	0.0215	-0.1736	-0.1411	-0.0144	-0.0101
11	0.2	-0.0664	0.0217	-0.1740	-0.1415	-0.0173	-0.0133
30	0.05	-0.0776	0.0245	-0.1908	-0.1598	-0.0214	-0.0162

The weights are chosen by solving the optimization (16) for the given ξ .

as in Table 6, only now puts with strikes above \$100 are excluded due to early exercise provisions. The hedging results demonstrate that, for this American put, we can simultaneously protect our position from jumps without incurring prohibitive transaction costs.

REMARK 2 (OPTIMALITY OF THE HEDGING STRATEGY). Our hedging strategy is local in time, as it is only concerned with the instantaneous state of the overall hedged position. In general, it will not be globally optimal. As noted previously, though, solving the full stochastic control problem would be computationally infeasible. We expect that our hedging results can certainly be improved upon. However, even our (non-optimal) hedging results clearly demonstrate that the use of a dynamic hedge containing traded options is a viable technique for minimizing jump risk.

8. Conclusions

There is now overwhelming evidence that equities have jump risk. When the underlying follows a jump diffusion process, simple delta hedging is a very poor strategy. The only possible approach which can be used to mitigate jump risk is to include traded derivatives in the hedge portfolio. In this paper, we suggest a dynamic hedging strategy based on a portfolio consisting of the underlying and options. We solve an optimization problem at each hedge rebalance time to minimize a linear combination of a jump risk and transaction cost penalty function. This strategy has the advantage that it is easily applied to path-dependent options (e.g. American style). It is also easy to incorporate the most recently observed calibrated market parameters at each hedge rebalance time. We test this strategy by simulations in a synthetic market. We make the assumptions that:

- The underlying asset follows a Merton-type jump diffusion process. \mathbb{P} measure parameters are unobservable to the hedger.
- The midpoint prices of options are given by solving a jump diffusion PIDE with known \mathbb{Q} measure parameters.
- Relative bid-ask spreads are a known function of option moneyness.

Under these assumptions, simulations of our dynamic hedging strategy show the following:

- If the hedge portfolio is determined solely on the basis of minimizing jump risk (and ignoring transaction costs), the results are worse than simple delta hedging (which is itself quite poor). This is in accordance with conventional wisdom, which states that hedging with options is too expensive.

- On the other hand, if both jump risk and transaction costs are included in the objective function, our dynamic strategy is effective. In many cases only a small amount of buying and selling takes place while, at the same time, the overall position is protected against jumps. The standard deviation of the relative P&L is much reduced compared to simple delta hedging.

These results are very encouraging. Using a bid-ask spread model which captures the gross features of observed market prices (i.e. out-of-the-money options have larger relative spreads than near-the-money options) forces our strategy to reduce trading costs, while still minimizing jump risk. This indicates that if we are going to develop effective strategies for mitigating jump risk, it is necessary to include realistic market effects.

Appendix A: Derivation of Jump Risk

In this Appendix we derive a mathematical representation of jump risk. The derivation closely follows that provided in He et al. (2006). The hedge portfolio contains an amount B in cash, is long e units of the underlying asset S , and long N additional hedging instruments $\vec{I} = [I_1, I_2, \dots, I_N]$ (written on the underlying) with weights $\vec{\phi} = [\phi_1, \phi_2, \dots, \phi_N]$. When combined with a short position in the target option $-V$, the resulting overall hedged position has value

$$\Pi = -V + eS + \vec{\phi} \cdot \vec{I} + B,$$

where the explicit dependence on time t and asset price S has been dropped to ease notation. To represent changes in the components of Π due to a jump of size J , we use the notation $\Delta V = V(JS) - V(S)$, $\Delta S = S(J - 1)$ and $\Delta \vec{I} = \vec{I}(JS) - \vec{I}(S)$.

If a change in the short position $-V$ is always precisely neutralized by the hedge portfolio $(eS + \vec{\phi} \cdot \vec{I} + B)$, the hedge is considered perfect and Π will have zero variation over an instant dt . We must therefore consider the infinitesimal change of the overall hedged position value Π . Since we are concerned with the real-world evolution of this portfolio, the jump diffusion process of interest is governed by the objective measure \mathbb{P} . We have:

$$\begin{aligned} dS &= \eta^{\mathbb{P}} S dt + \sigma S dZ^{\mathbb{P}} + \Delta S d\pi^{\mathbb{P}} \\ dV &= \left[\frac{\partial V}{\partial t} + \frac{\sigma^2 S^2}{2} \frac{\partial^2 V}{\partial S^2} + \eta^{\mathbb{P}} S \frac{\partial V}{\partial S} \right] dt + \sigma S \frac{\partial V}{\partial S} dZ^{\mathbb{P}} + \Delta V d\pi^{\mathbb{P}} \\ d\vec{I} &= \left[\frac{\partial \vec{I}}{\partial t} + \frac{\sigma^2 S^2}{2} \frac{\partial^2 \vec{I}}{\partial S^2} + \eta^{\mathbb{P}} S \frac{\partial \vec{I}}{\partial S} \right] dt + \sigma S \frac{\partial \vec{I}}{\partial S} dZ^{\mathbb{P}} + \Delta \vec{I} d\pi^{\mathbb{P}} \\ dB &= rB dt, \end{aligned}$$

where $\eta^{\mathbb{P}} = \alpha^{\mathbb{P}} - \lambda^{\mathbb{P}} \kappa^{\mathbb{P}}$. The above implies that the instantaneous change in the value of the overall hedged position is

$$\begin{aligned} d\Pi &= -dV + e dS + \vec{\phi} \cdot d\vec{I} + dB \\ &= - \left[\frac{\partial V}{\partial t} + \frac{\sigma^2 S^2}{2} \frac{\partial^2 V}{\partial S^2} \right] dt + \vec{\phi} \cdot \left[\frac{\partial \vec{I}}{\partial t} + \frac{\sigma^2 S^2}{2} \frac{\partial^2 \vec{I}}{\partial S^2} \right] dt \\ &\quad + \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right] d\pi^{\mathbb{P}} + rB dt \\ &\quad + \eta^{\mathbb{P}} S \left[-\frac{\partial V}{\partial S} + e + \vec{\phi} \cdot \frac{\partial \vec{I}}{\partial S} \right] dt + \sigma S \left[-\frac{\partial V}{\partial S} + e + \vec{\phi} \cdot \frac{\partial \vec{I}}{\partial S} \right] dZ^{\mathbb{P}}, \end{aligned} \quad (19)$$

where e and $\vec{\phi}$ are regarded as constant over dt as they must be set at the beginning of this instant.

If the portfolio is delta neutral, then $\frac{\partial \Pi}{\partial S} = 0$, i.e.

$$-\frac{\partial V}{\partial S} + e + \vec{\phi} \cdot \frac{\partial \vec{I}}{\partial S} = 0. \quad (20)$$

Imposing delta neutrality within equation (19) eliminates the final two terms in the expression for $d\Pi$, including the one involving the Wiener process $dZ^{\mathbb{P}}$. The expression for $d\Pi$ consequently simplifies to

$$d\Pi = - \left[\frac{\partial V}{\partial t} + \frac{\sigma^2 S^2}{2} \frac{\partial^2 V}{\partial S^2} \right] dt + \vec{\phi} \cdot \left[\frac{\partial \vec{I}}{\partial t} + \frac{\sigma^2 S^2}{2} \frac{\partial^2 \vec{I}}{\partial S^2} \right] dt + rB dt + \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right] d\pi^{\mathbb{P}}, \quad (21)$$

indicating that $d\Pi$ is now a pure jump process with drift. Using an elementary rearrangement, the pricing PIDEs for the target and hedging options may be written as

$$\begin{aligned} \frac{\partial V}{\partial t} + \frac{\sigma^2 S^2}{2} \frac{\partial^2 V}{\partial S^2} &= rV + \{ \lambda^{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}}(\Delta S) - rS \} \frac{\partial V}{\partial S} - \lambda^{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}}(\Delta V) \\ \frac{\partial \vec{I}}{\partial t} + \frac{\sigma^2 S^2}{2} \frac{\partial^2 \vec{I}}{\partial S^2} &= r\vec{I} + \{ \lambda^{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}}(\Delta S) - rS \} \frac{\partial \vec{I}}{\partial S} - \lambda^{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}}(\Delta \vec{I}), \end{aligned} \quad (22)$$

where $\mathbb{E}^{\mathbb{Q}}(\Delta S) = \mathbb{E}^{\mathbb{Q}}(S[J-1]) = S\mathbb{E}^{\mathbb{Q}}(J-1) = S\kappa^{\mathbb{Q}}$. Substituting (22) into (21) yields

$$\begin{aligned} d\Pi &= - \left[rV + \{ \lambda^{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}}(\Delta S) - rS \} \frac{\partial V}{\partial S} - \lambda^{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}}(\Delta V) \right] dt \\ &\quad + \vec{\phi} \cdot \left[r\vec{I} + \{ \lambda^{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}}(\Delta S) - rS \} \frac{\partial \vec{I}}{\partial S} - \lambda^{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}}(\Delta \vec{I}) \right] dt + rB dt + [-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I})] d\pi^{\mathbb{P}} \\ &= r \left[-V + \vec{\phi} \cdot \vec{I} + S \left(\frac{\partial V}{\partial S} - \vec{\phi} \cdot \frac{\partial \vec{I}}{\partial S} \right) + B \right] dt \\ &\quad + \lambda^{\mathbb{Q}} \left[\mathbb{E}^{\mathbb{Q}}(\Delta V) - \vec{\phi} \cdot \mathbb{E}^{\mathbb{Q}}(\Delta \vec{I}) + \left(-\frac{\partial V}{\partial S} + \vec{\phi} \cdot \frac{\partial \vec{I}}{\partial S} \right) \mathbb{E}^{\mathbb{Q}}(\Delta S) \right] dt + [-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I})] d\pi^{\mathbb{P}}. \end{aligned}$$

Using the delta-neutral constraint (20) gives

$$\begin{aligned} d\Pi &= r \left[-V + eS + \vec{\phi} \cdot \vec{I} + B \right] dt + \lambda^{\mathbb{Q}} \left[\mathbb{E}^{\mathbb{Q}}(\Delta V) - \vec{\phi} \cdot \mathbb{E}^{\mathbb{Q}}(\Delta \vec{I}) - e\mathbb{E}^{\mathbb{Q}}(\Delta S) \right] dt \\ &\quad + [-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I})] d\pi^{\mathbb{P}} \\ &= r\Pi dt + \lambda^{\mathbb{Q}} dt \mathbb{E}^{\mathbb{Q}} \left[\Delta V - (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right] + d\pi^{\mathbb{P}} \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right]. \end{aligned} \quad (23)$$

The quantity

$$\lambda^{\mathbb{Q}} dt \mathbb{E}^{\mathbb{Q}} \left[\Delta V - (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right] + d\pi^{\mathbb{P}} \left[-\Delta V + (e\Delta S + \vec{\phi} \cdot \Delta \vec{I}) \right] \quad (24)$$

is termed the instantaneous jump risk.

Appendix B: Global Bound on the Hedging Error

In this appendix we consider a continuous-time treatment of our discrete hedging strategy in order to show that, in the limit, the variance of the terminal hedging error may be made arbitrarily small if our measure of jump risk and the transaction cost can be suitably bounded at each instant. A more intuitive approach would be to start with the discrete trading model and take the limit as $\Delta t \rightarrow 0$. However, this requires tedious algebraic manipulations: by making an assumption on the form of the transaction cost term in the continuous framework, this can be avoided. Also, to avoid complication, we shall assume the transaction cost for the underlying is zero. Hence, the total transaction cost can always be made zero by only adjusting the weight in the underlying to impose delta neutrality.

For future reference, let $\mathbb{E}_t[X_s]$ for $s \geq t$ denote the expected value of X_s conditioned on information known at time t , and define a proper weighting function in the following way:

Definition 1 (Proper Weighting Function) *A proper weighting function W with respect to g is such that*

$$\int_0^{\infty} f^2(J)g(J) dJ \leq \int_0^{\infty} f^2(J)W(J) dJ < \infty \quad (25)$$

for any function f satisfying $\int_0^{\infty} f^2(J)g(J) dJ < \infty$.

For a hedge portfolio consisting of N instruments, the total transaction cost of rebalancing at time t is

$$\Upsilon_t = \sum_{k=1}^N \Upsilon_t^k \geq 0, \quad (26)$$

where Υ_t^k is the transaction cost for the k^{th} instrument. We will construct the hedging strategy so that the transaction cost is proportional to dt ; that is,

$$\Upsilon_t = \vartheta_t dt, \quad (27)$$

with $\vartheta_t \geq 0$. This assumption is required in order to ensure finite transaction costs—we will indicate how this can be approximated in practice. Over an instant, the value of the overall hedged position will decrease by an amount $\Upsilon_t = \vartheta_t dt$ due to the transaction costs, such that the time evolution of the hedged portfolio

will be augmented by a negative drift, namely $-\vartheta_t dt$. The instantaneous change in the delta-neutral hedged position (23) may therefore be expressed as

$$d\Pi_t = \left(r\Pi_t + \lambda^{\mathbb{Q}} \mathbb{E}_t^{\mathbb{Q}} \left[-\Delta H_J(S_t, t) \right] - \vartheta_t \right) dt + \Delta H_J(S_t, t) d\pi^{\mathbb{P}}, \quad (28)$$

where

$$\Delta H_J(S_t, t) = -\left(V(JS_t, t) - V(S_t, t) \right) + eS_t(J-1) + \vec{\phi} \cdot \left(\vec{I}(JS_t, t) - \vec{I}(S_t, t) \right) \quad (29)$$

is the random jump component. Note that, as opposed to the derivation of jump risk in Appendix A, here we make explicit the dependence on time. This allows us to more easily consider the discounted overall hedged position $\tilde{\Pi}_t = \exp\{-rt\} \Pi_t$, such that

$$d\tilde{\Pi}_t = \exp\{-rt\} \left(-r\Pi_t dt + d\Pi_t \right). \quad (30)$$

Substituting (28) into (30) yields

$$d\tilde{\Pi}_t = \exp\{-rt\} \left(\left(\lambda^{\mathbb{Q}} \Theta_t^{\mathbb{Q}} - \vartheta_t \right) dt + \Delta H_J(S_t, t) d\pi^{\mathbb{P}} \right) \quad (31)$$

with

$$\Theta_t^{\mathbb{Q}} = \mathbb{E}_t^{\mathbb{Q}} \left[-\Delta H_J(S_t, t) \right]. \quad (32)$$

Our goal is to show that, by making the jump risk and transaction costs sufficiently small at each instant, the variance of the terminal hedging error Π_T can be made arbitrarily close to zero. To do this we will examine the expectation of $(\tilde{\Pi}_T)^2$, where $\tilde{\Pi}_t^2$ follows (by Itô's formula for semimartingales)

$$d\tilde{\Pi}_t^2 = 2 \exp\{-rt\} \left(\lambda^{\mathbb{Q}} \Theta_t^{\mathbb{Q}} - \vartheta_t \right) \tilde{\Pi}_t dt + \left(\tilde{\Pi}_t^2(JS_t) - \tilde{\Pi}_t^2(S_t) \right) d\pi^{\mathbb{P}}. \quad (33)$$

Here, $\tilde{\Pi}_t^2(JS_t) - \tilde{\Pi}_t^2(S_t)$ represents the change in the value of $\tilde{\Pi}^2$ assuming an asset price jump of size J occurs at time t . We will need the following result related to this and other quantities in order to establish a bound on the hedging error.

LEMMA 1 (Bounds on Expectations Involving $\tilde{\Pi}_t$). *Assume that, for time t , the following four conditions are met:*

(A1) *Jump risk is made small:* $\forall S_t > 0, \int_0^{\infty} \left[\Delta H_J(S_t, t) \right]^2 W(J) dJ < \epsilon;$

(A2) *Transaction cost is made small:* $\forall S_t > 0, \left(\vartheta_t(S_t) \right)^2 < \varrho^2 \epsilon, \text{ where } \varrho > 0;$

(A3) *The second moment of $\tilde{\Pi}_t$ exists:* $\mathbb{E}_0^{\mathbb{P}} \left[\tilde{\Pi}_t^2 \right] < \infty;$

(A4) *$W(J)$ is a proper weighting function with respect to both $g^{\mathbb{P}}(J)$ and $g^{\mathbb{Q}}(J)$.*

Then

$$\left| \mathbb{E}_0^{\mathbb{P}} \left[\tilde{\Pi}_t^2(JS_t) - \tilde{\Pi}_t^2(S_t) \right] \right| < \epsilon + 2\sqrt{\epsilon \mathbb{E}_0^{\mathbb{P}} \left[\tilde{\Pi}_t^2 \right]}, \quad (34)$$

$$\left| \mathbb{E}_0^{\mathbb{P}} \left[\Theta_t^{\mathbb{Q}} \tilde{\Pi}_t \right] \right| < \sqrt{\epsilon \mathbb{E}_0^{\mathbb{P}} \left[\tilde{\Pi}_t^2 \right]}, \quad (35)$$

and

$$\left| \mathbb{E}_0^{\mathbb{P}} \left[\vartheta_t \tilde{\Pi}_t \right] \right| < \varrho \sqrt{\epsilon \mathbb{E}_0^{\mathbb{P}} \left[\tilde{\Pi}_t^2 \right]}. \quad (36)$$

REMARK 3. Conditions (A1) and (A2) imply that, for time t , we have a well defined procedure for rebalancing the hedge portfolio (based on S_t and the existing weights) such that our measure of jump risk and the transaction cost can be made arbitrarily small.

REMARK 4. The expression $\mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2]$ represents the expected squared value of the discounted overall hedged position at time t . The expression $\mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2(JS_t) - \tilde{\Pi}_t^2(S_t)]$ represents the expected change in the value of $\tilde{\Pi}_t^2$ assuming a jump occurs at time t . Note that both of these expectations are taken at time zero, such that the only information known is S_0 .

Proof of Lemma 1. The value of the discounted overall hedged position at time t (after rebalancing) is

$$\tilde{\Pi}_t(S_t) = \exp\{-rt\} \left(-V(S_t, t) + eS_t + \vec{\phi} \cdot \vec{I}(S_t, t) + B(t) \right), \quad (37)$$

while after a jump at time t it has value

$$\tilde{\Pi}_t(JS_t) = \exp\{-rt\} \left(-V(JS_t, t) + eJS_t + \vec{\phi} \cdot \vec{I}(JS_t, t) + B(t) \right). \quad (38)$$

Solving for the bank account in (37) and substituting into (38) yields

$$\tilde{\Pi}_t(JS_t) = \tilde{\Pi}_t(S_t) + \exp\{-rt\} \Delta H_J(S_t, t),$$

where the definition of $\Delta H_J(S_t, t)$ in (29) is used. Therefore, we have

$$\begin{aligned} \left| \mathbb{E}_0^{\mathbb{P}} \left[\tilde{\Pi}_t^2(JS_t) - \tilde{\Pi}_t^2(S_t) \right] \right| &= \left| \mathbb{E}_0^{\mathbb{P}} \left[\left(\tilde{\Pi}_t(S_t) + \exp\{-rt\} \Delta H_J(S_t, t) \right)^2 - \tilde{\Pi}_t^2(S_t) \right] \right| \\ &= \left| \mathbb{E}_0^{\mathbb{P}} \left[\left(\exp\{-rt\} \Delta H_J(S_t, t) \right)^2 + 2 \exp\{-rt\} \tilde{\Pi}_t(S_t) \Delta H_J(S_t, t) \right] \right| \\ &\leq \mathbb{E}_0^{\mathbb{P}} \left[\left(\Delta H_J(S_t, t) \right)^2 \right] + 2 \left| \mathbb{E}_0^{\mathbb{P}} \left[\tilde{\Pi}_t(S_t) \Delta H_J(S_t, t) \right] \right|. \end{aligned} \quad (39)$$

Consider the first expectation on the right-hand side of (39), which only depends on the asset price at time t . By conditioning on S_t :

$$\begin{aligned} \mathbb{E}_0^{\mathbb{P}} \left[\left(\Delta H_J(S_t, t) \right)^2 \right] &= \int_0^\infty \left[\int_0^\infty \left(\Delta H_J(S_t, t) \right)^2 g^{\mathbb{P}}(J) dJ \right] p(S_t | S_0) dS_t \\ &\leq \int_0^\infty \left[\int_0^\infty \left(\Delta H_J(S_t, t) \right)^2 W(J) dJ \right] p(S_t | S_0) dS_t && \text{(by (A4))} \\ &< \int_0^\infty \epsilon p(S_t | S_0) dS_t && \text{(by (A1))} \\ &= \epsilon, \end{aligned} \quad (40)$$

where $p(S_t | S_0)$ is the transition density under \mathbb{P} . For the second expectation on the right-hand side of (39), since $\mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2(S_t)] < \infty$ by (A3) and $\mathbb{E}_0^{\mathbb{P}}[(\Delta H_J(S_t, t))^2] < \epsilon$ by (40), the Cauchy–Schwarz inequality gives

$$\begin{aligned} \left| \mathbb{E}_0^{\mathbb{P}} \left[\tilde{\Pi}_t(S_t) \Delta H_J(S_t, t) \right] \right| &\leq \sqrt{\mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2(S_t)] \mathbb{E}_0^{\mathbb{P}}[(\Delta H_J(S_t, t))^2]} \\ &< \sqrt{\epsilon \mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2(S_t)]}. \end{aligned} \quad (41)$$

Using the upper bounds (40) and (41) in (39) yields

$$\left| \mathbb{E}_0^{\mathbb{P}} \left[\tilde{\Pi}_t^2(JS_t) - \tilde{\Pi}_t^2(S_t) \right] \right| < \epsilon + 2\sqrt{\epsilon \mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2(S_t)]},$$

which is the first implication (34) of this Lemma.

For the second part of the Lemma, consider

$$\begin{aligned} \mathbb{E}_0^{\mathbb{P}} \left[\left(\Theta_t^{\mathbb{Q}} \right)^2 \right] &= \mathbb{E}_0^{\mathbb{P}} \left[\left(\mathbb{E}_t^{\mathbb{Q}} \left[-\Delta H_J(S_t, t) \right] \right)^2 \right] && \text{(by def. of } \Theta_t^{\mathbb{Q}} \text{ in (32))} \\ &\leq \mathbb{E}_0^{\mathbb{P}} \left[\mathbb{E}_t^{\mathbb{Q}} \left[\left(\Delta H_J(S_t, t) \right)^2 \right] \right] && \text{(since } \mathbb{E}[X]^2 \leq \mathbb{E}[X^2]) \\ &= \int_0^\infty \left[\int_0^\infty \left(\Delta H_J(S_t, t) \right)^2 g^{\mathbb{Q}}(J) dJ \right] p(S_t | S_0) dS_t \end{aligned}$$

$$\begin{aligned}
&\leq \int_0^\infty \left[\int_0^\infty (\Delta H_J(S_t, t))^2 W(J) dJ \right] p(S_t | S_0) dS_t && \text{(by (A4))} \\
&< \int_0^\infty \epsilon p(S_t | S_0) dS_t && \text{(by (A1))} \\
&= \epsilon. && (42)
\end{aligned}$$

As such, $\mathbb{E}_0^\mathbb{P}[(\Theta_t^\mathbb{Q})^2]$ exists and, since $\mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2] < \infty$ by (A3),

$$\begin{aligned}
\left| \mathbb{E}_0^\mathbb{P}[\Theta_t^\mathbb{Q} \tilde{\Pi}_t] \right| &\leq \sqrt{\mathbb{E}_0^\mathbb{P}[(\Theta_t^\mathbb{Q})^2]} \sqrt{\mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]} \\
&< \sqrt{\epsilon \mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]},
\end{aligned}$$

where the Cauchy–Schwarz inequality and the bound in (42) are employed.

Finally, condition (A2) guarantees that ϑ_t^2 is always bounded by $\varrho^2 \epsilon$, which means $\mathbb{E}_0^\mathbb{P}[(\vartheta_t)^2] < \varrho^2 \epsilon$. Similar to the above, an application of the Cauchy–Schwarz inequality gives the result (36). \square

We are now in a position to prove the main result of this appendix.

THEOREM 1 (Variance of the Hedging Error Can be Made Arbitrarily Small). *If, for all times t in the investment period $[0, T]$ the conditions (A1) – (A4) of Lemma 1 hold, then*

$$\mathbb{E}_0^\mathbb{P}[\Pi_T^2] < \epsilon \frac{\exp\{2rT\} (\lambda^\mathbb{P})^2}{4(\lambda^\mathbb{P} + \lambda^\mathbb{Q} + \varrho)^2} \left[W_{-1} \left(-\exp \left\{ -\frac{2(\lambda^\mathbb{P} + \lambda^\mathbb{Q} + \varrho)^2}{\lambda^\mathbb{P}} T - 1 \right\} \right) + 1 \right]^2, \quad (43)$$

where $W_{-1}(x)$ is the -1 branch of the Lambert W function (Corless et al. 1996).⁴ That is, $\mathbb{E}_0^\mathbb{P}[\Pi_T^2] \leq \mathcal{C}\epsilon$, where \mathcal{C} is a constant that only depends on $\lambda^\mathbb{P}$, $\lambda^\mathbb{Q}$, ϱ , r , and T .

Proof. The stochastic differential equation (33) may be used to derive an ordinary differential equation (ODE) relating the moments of $\tilde{\Pi}_t$, namely

$$\frac{d\mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]}{dt} = 2\exp\{-rt\} \lambda^\mathbb{Q} \mathbb{E}_0^\mathbb{P}[\Theta_t^\mathbb{Q} \tilde{\Pi}_t] - 2\exp\{-rt\} \mathbb{E}_0^\mathbb{P}[\vartheta_t \tilde{\Pi}_t] + \lambda^\mathbb{P} \mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2(JS_t) - \tilde{\Pi}_t^2(S_t)] \quad (44)$$

with initial condition $\mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_0^2] = 0$. Equation (44) yields

$$\frac{d\mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]}{dt} \leq 2\lambda^\mathbb{Q} \left| \mathbb{E}_0^\mathbb{P}[\Theta_t^\mathbb{Q} \tilde{\Pi}_t] \right| + 2 \left| \mathbb{E}_0^\mathbb{P}[\vartheta_t \tilde{\Pi}_t] \right| + \lambda^\mathbb{P} \left| \mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2(JS_t) - \tilde{\Pi}_t^2(S_t)] \right|,$$

which allows us to use the bounds established in Lemma 1 to set up the differential inequality

$$\frac{d\mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]}{dt} < 2\lambda^\mathbb{Q} \sqrt{\epsilon \mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]} + 2\varrho \sqrt{\epsilon \mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]} + \lambda^\mathbb{P} \left(\epsilon + 2\sqrt{\epsilon \mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]} \right);$$

the above may be written more succinctly as

$$\frac{d\mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]}{dt} < 2\sqrt{\epsilon}(\lambda^\mathbb{P} + \lambda^\mathbb{Q} + \varrho) \sqrt{\mathbb{E}_0^\mathbb{P}[\tilde{\Pi}_t^2]} + \epsilon \lambda^\mathbb{P}. \quad (45)$$

We define a *bounding function* $\beta(t)$ which satisfies the ODE

$$\frac{d\beta}{dt} = 2\sqrt{\epsilon}(\lambda^\mathbb{P} + \lambda^\mathbb{Q} + \varrho) \sqrt{\beta(t)} + \epsilon \lambda^\mathbb{P} \quad (46)$$

with initial condition $\beta(0) = 0$. This initial value problem has the exact solution

$$\beta(t) = \epsilon \frac{(\lambda^\mathbb{P})^2}{4(\lambda^\mathbb{P} + \lambda^\mathbb{Q} + \varrho)^2} \left[W_{-1} \left(-\exp \left\{ -\frac{2(\lambda^\mathbb{P} + \lambda^\mathbb{Q} + \varrho)^2}{\lambda^\mathbb{P}} t - 1 \right\} \right) + 1 \right]^2, \quad (47)$$

where $W_{-1}(x)$ is the -1 branch of the Lambert W function. Consider the relationship between $\beta(t)$ and $\mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2]$. Both are non-negative quantities with an initial value of zero. The differential inequality (45) for $\mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2]$ and the ODE (46) satisfied by $\beta(t)$ means

$$\frac{d\mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2]}{dt} < \frac{d\beta}{dt}.$$

Hence $\mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2] < \beta(t)$, which implies

$$\mathbb{E}_0^{\mathbb{P}}[\tilde{\Pi}_t^2] < \epsilon \frac{(\lambda^{\mathbb{P}})^2}{4(\lambda^{\mathbb{P}} + \lambda^{\mathbb{Q}} + \varrho)^2} \left[W_{-1} \left(-\exp \left\{ -\frac{2(\lambda^{\mathbb{P}} + \lambda^{\mathbb{Q}} + \varrho)^2}{\lambda^{\mathbb{P}}} t - 1 \right\} \right) + 1 \right]^2.$$

Using the fact that $\Pi_t = \exp\{rt\}\tilde{\Pi}_t$ and considering the above at $t = T$ yields the bound on $\mathbb{E}_0^{\mathbb{P}}[\Pi_T^2]$ in (43). \square

B.1. Discrete Rebalancing with Transaction Costs

With transaction costs present in the form of a bid-ask spread, the objective minimized at each rebalance time within our discrete hedging procedure is

$$\xi \left\{ \int_0^\infty [\Delta H_J(S_t, t)]^2 W(J) dJ \right\} + (1 - \xi) \sum_{k=1}^N (\Upsilon_t^k)^2, \quad (48)$$

where $\xi \in [0, 1]$ and $(\Upsilon_t^k)^2$ is the square of the transaction cost for rebalancing the k^{th} instrument. Note this is simply a more succinct form of the objective function introduced in (16). The following Lemma relates a bound on the objective (48) to bounds required to establish the continuous-time result in Theorem 1 for $\mathbb{E}^{\mathbb{P}}[\Pi_T^2]$, i.e. when $\Delta t \rightarrow 0$.

LEMMA 2 (Bounds Dictated by the Hedging Objective). *Assume that, for the current time t and asset price S_t , the condition*

$$\xi \left\{ \int_0^\infty [\Delta H_J(S_t, t)]^2 W(J) dJ \right\} + (1 - \xi) \sum_{k=1}^N (\Upsilon_t^k)^2 < \epsilon^* \Delta t^2 \quad (49)$$

holds for some $\xi \in (0, 1)$, where Δt is the rebalancing interval. In that case,

$$\int_0^\infty [\Delta H_J(S_t, t)]^2 W(J) dJ < \frac{\epsilon^* \Delta t^2}{\xi} \quad (50)$$

and

$$(\Upsilon_t)^2 = \left[\sum_{k=1}^N \Upsilon_t^k \right]^2 < \frac{\epsilon^* \Delta t^2}{1 - \xi} N^2. \quad (51)$$

Proof. The first result (50) is a simple consequence of the fact that both the jump risk and transaction cost component of the objective are positive. Similarly

$$\sum_{k=1}^N (\Upsilon_t^k)^2 < \frac{\epsilon^* \Delta t^2}{1 - \xi},$$

which implies $(\Upsilon_t^k)^2 < \frac{\epsilon^* \Delta t^2}{1 - \xi}$ for all $k = 1, \dots, N$. Now consider

$$\begin{aligned} \left[\sum_{k=1}^N \Upsilon_t^k \right]^2 &= \sum_{k=1}^N (\Upsilon_t^k)^2 + \sum_{k=1}^N \sum_{\substack{l=1 \\ l \neq k}}^N \Upsilon_t^k \Upsilon_t^l \\ &\leq \sum_{k=1}^N (\Upsilon_t^k)^2 + \sum_{k=1}^N \sum_{\substack{l=1 \\ l \neq k}}^N \frac{(\Upsilon_t^k)^2 + (\Upsilon_t^l)^2}{2} \\ &< \sum_{k=1}^N \sum_{l=1}^N \frac{\epsilon^* \Delta t^2}{1 - \xi} \\ &= \frac{\epsilon^* \Delta t^2}{1 - \xi} N^2, \end{aligned} \quad (52)$$

which is the desired result in (51).

REMARK 5 (LINK BETWEEN THE DISCRETE AND CONTINUOUS HEDGING STRATEGIES). The results (50) and (51) of Lemma 2 demonstrate that by choosing $\xi = C\Delta t^2$ within the discrete framework, where C is a positive constant, we obtain

$$\int_0^\infty \left[\Delta H_J(S_t, t) \right]^2 W(J) dJ < \frac{\epsilon^*}{C} \quad (53)$$

and

$$\left(\Upsilon_t \right)^2 = \left[\sum_{k=1}^N \Upsilon_t^k \right]^2 < \epsilon^* \Delta t^2 N^2 + \mathcal{O}(\Delta t^4). \quad (54)$$

Define $\epsilon = \frac{\epsilon^*}{C}$ and $\varrho = N\sqrt{C}$. Then, (53) becomes

$$\int_0^\infty \left[\Delta H_J(S_t, t) \right]^2 W(J) dJ < \epsilon \quad (55)$$

and (54) is

$$\left(\Upsilon_t \right)^2 < \varrho^2 \epsilon \Delta t^2 + \mathcal{O}(\Delta t^4).$$

Let $\Upsilon_t = \vartheta_t \Delta t$, so that

$$\vartheta_t^2 < \varrho^2 \epsilon + \mathcal{O}(\Delta t^2). \quad (56)$$

Assume that both (55) and (56) can be satisfied at any time and for any S_t , where N represents the maximum number of hedging instruments required to do this. In the limit as $\Delta t \rightarrow 0$, conditions (A1) and (A2) of Theorem 1 are satisfied from (55) and (56), respectively, and we obtain the bound (43) on $\mathbb{E}^{\mathbb{P}}[\Pi_T^2]$ (assuming, of course, that (A3) and (A4) in Theorem 1 also hold).

REMARK 6 (MINIMIZING THE HEDGING OBJECTIVE IN PRACTICE). It is always possible to satisfy the transaction cost objective (A2) by trading only in the underlying, although in this case the jump risk is expected to be substantial. Indeed, we may not be able to make both the jump risk and transaction cost small in the manner required by (49). As a practical matter, at each rebalance time we simply attempt to make the objective (48) as small as possible for a given ξ .

REMARK 7 (CHOICE OF ξ). If we fix ξ and minimize the objective (48), the above analysis implies that the best choice of ξ should be $\mathcal{O}(\Delta t^2)$.

Appendix C: \mathbb{P} and \mathbb{Q} Linkage

This Appendix provides the expressions linking the \mathbb{P} and \mathbb{Q} measure parameters. As our intention here is to just obtain an approximate transformation from the pricing measure \mathbb{Q} to the real-world measure \mathbb{P} , we follow Naik and Lee (1990) and Bates (1991) and assume the simple case of power utility. Letting the coefficient of relative risk aversion be $1 - \beta$, we have the following relations:

$$\begin{aligned} \sigma^{\mathbb{P}} &= \sigma^{\mathbb{Q}} \\ \gamma^{\mathbb{P}} &= \gamma^{\mathbb{Q}} \\ \mu^{\mathbb{P}} &= \mu^{\mathbb{Q}} + (1 - \beta)(\gamma^{\mathbb{Q}})^2 \\ \lambda^{\mathbb{P}} &= \lambda^{\mathbb{Q}} \exp \left\{ (1 - \beta) \left(\mu^{\mathbb{Q}} + \frac{1}{2}(1 - \beta)(\gamma^{\mathbb{Q}})^2 \right) \right\} \\ \alpha^{\mathbb{P}} &= r + (1 - \beta)\sigma^2 + (\kappa^{\mathbb{P}}\lambda^{\mathbb{P}} - \kappa^{\mathbb{Q}}\lambda^{\mathbb{Q}}). \end{aligned} \quad (57)$$

The above expressions were used to generate the \mathbb{P} measure parameters provided in Table 1, with $1 - \beta = 2$ and $r = .05$. We also assume that the dividend yield $q = 0$. The \mathbb{Q} measure parameters are similar to those reported in Andersen and Andreasen (2000), which were obtained by calibrating to observed prices of S&P 500 index options.

Appendix D: Amazon.com, Inc. (AMZN) Option Price Data

The option price data in Table 8 is used to generate the relative bid-ask spread curves of Figure 6.

Table 8 Option data used to generate the relative bid-ask spread curves of Figure 6.

Strike	Moneyness	Calls			Puts		
		Bid	Ask	Relative Spread	Bid	Ask	Relative Spread
\$20.00	0.4437	\$25.10	\$25.30	0.0079	\$ 0.05	\$ 0.05	0.0000
\$22.50	0.4991	\$22.70	\$22.80	0.0044	\$ 0.05	\$ 0.05	0.0000
\$25.00	0.5546	\$20.20	\$20.30	0.0049	\$ 0.05	\$ 0.05	0.0000
\$27.50	0.6100	\$17.70	\$17.80	0.0056	\$ 0.05	\$ 0.05	0.0000
\$30.00	0.6655	\$15.30	\$15.40	0.0065	\$ 0.05	\$ 0.10	0.6667
\$32.50	0.7209	\$12.80	\$13.00	0.0155	\$ 0.10	\$ 0.15	0.4000
\$35.00	0.7764	\$10.40	\$10.50	0.0096	\$ 0.15	\$ 0.20	0.2857
\$37.50	0.8319	\$ 8.10	\$ 8.20	0.0123	\$ 0.35	\$ 0.40	0.1333
\$40.00	0.8873	\$ 6.00	\$ 6.20	0.0328	\$ 0.70	\$ 0.75	0.0690
\$42.50	0.9428	\$ 4.10	\$ 4.30	0.0476	\$ 1.30	\$ 1.35	0.0377
\$45.00	0.9982	\$ 2.60	\$ 2.65	0.0190	\$ 2.20	\$ 2.30	0.0444
\$47.50	1.0537	\$ 1.50	\$ 1.55	0.0328	\$ 3.60	\$ 3.70	0.0274
\$50.00	1.1091	\$ 0.75	\$ 0.85	0.1250	\$ 5.40	\$ 5.50	0.0183
\$55.00	1.2201	\$ 0.20	\$ 0.25	0.2222	\$ 9.90	\$10.00	0.0101
\$60.00	1.3310	\$ 0.05	\$ 0.10	0.6667	\$14.80	\$15.00	0.0134

Option price data for 22Oct2005 puts and calls on Amazon.com, Inc. (AMZN), taken during trading on August 10, 2005. The spot value of the underlying is \$45.08.

Endnotes

¹The triangular tails ensure the weighting function is continuous. This results in better numerical behaviour within the FFT procedure we use to compute the integrals necessary for solving the optimization.

²It would also include digital options, but such contracts are not liquidly traded on financial exchanges and so they will not be considered here.

³Note that this is the theoretical value, unadjusted for a spread. Doing this allows us to calculate the appropriate premium to be charged by the financial institution for the cost of following our hedging strategy.

⁴The Lambert W function is defined as the inverse of $y = x \exp\{x\}$. For any $y \in (-\exp\{-1\}, 0)$, there are two real values x^* such that $y = x^* \exp\{x^*\}$, with one root $x_1^* \in (-1, 0)$ and the other root $x_2^* \in (-\infty, -1)$. The branch W_{-1} corresponds to those x^* that are less than or equal to -1. Therefore, $W_{-1}(-\exp\{-1\}) = -1$ and $W_{-1}(0) = -\infty$.

Acknowledgments

This work was supported by the Natural Sciences and Engineering Research Council of Canada and the Social Sciences and Humanities Research Council of Canada.

References

- Andersen, L. and J. Andreasen (2000). Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing. *Review of Derivatives Research* 4, 231–262.
- Asmussen, S. and J. Rosiński (2001). Approximations of small jumps of Lévy processes with a view to simulation. *Journal of Applied Probability* 38, 482–493.
- Barles, G. (1997). Convergence of numerical schemes for degenerate parabolic equations arising in finance. In L. C. G. Rogers and D. Talay, eds., *Numerical Methods in Finance*, pp. 1–21. Cambridge University Press, Cambridge.
- Bates, D. S. (1988). Pricing options under jump-diffusion processes. Working paper 37-88, Rodney L. White Center for Financial Research, The Wharton School, University of Pennsylvania.
- Bates, D. S. (1991). The crash of '87: Was it expected? The evidence from options markets. *Journal of Finance* 46, 1009–1044.

- Carr, P., K. Ellis, and V. Gupta (1998). Static hedging of exotic options. *Journal of Finance* 53, 1165–1190.
- Carr, P. and L. Wu (2004). Static hedging of standard options. Working paper, Courant Institute, New York University.
- Cont, R. and P. Tankov (2004). *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, Boca Raton.
- Cont, R., P. Tankov, and E. Voltchkova (2005). Hedging with options in models with jumps. *Stochastic Analysis and Applications*, forthcoming.
- Corless, R., G. Gonnet, D. Hare, D. Jeffrey, and D. Knuth (1996). On the Lambert W function. *Advances in Computational Mathematics* 5, 329–359.
- Davis, M., V. Panas, and T. Zariphopoulou (1993). European option pricing with transaction costs. *SIAM Journal of Optimization and Control* 31(2), 470–493.
- Derman, E., D. Ergener, and I. Kani (1995). Static options replication. *Journal of Derivatives* 2(4), 78–95.
- d’Halluin, Y., P. Forsyth, and K. Vetzal (2005). Robust numerical methods for contingent claims under jump diffusion processes. *IMA Journal of Numerical Analysis* 25, 65–92.
- d’Halluin, Y., P. A. Forsyth, and G. Labahn (2004). A penalty method for American options with jump diffusion processes. *Numerische Mathematik* 97, 321–352.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Grandits, P. and W. Schachinger (2001). Leland’s approach to option pricing: The evolution of a discontinuity. *Mathematical Finance* 11(3), 347–355.
- Hansen, P. (1987). The truncated SVD as a method for regularization. *BIT* 27, 534–553.
- He, C., J. Kennedy, T. Coleman, P. Forsyth, Y. Li, and K. Vetzal (2006). Calibration and hedging under jump diffusion. *Review of Derivatives Research* 9(1), 1–35.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 327–343.
- Hoggard, T., E. Whalley, and P. Wilmott (1994). Hedging option portfolios in the presence of transaction costs. *Advances in Futures and Options Research* 7, 21–35.
- Kabanov, Y. and M. Safarian (1997). On Leland’s strategy of option pricing with transactions costs. *Finance and Stochastics* 1, 239–250.
- Keppo, J. and S. Peura (1999). Optimal portfolio hedging with nonlinear derivatives and transaction costs. *Computational Economics* 13(2), 117–145.
- Leland, H. E. (1985). Option pricing and replication with transaction costs. *Journal of Finance* 40(5), 1283–1301.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.
- Mocioalca, O. (2003). Jump diffusion option with transaction costs. Working paper, Kent State University.
- Naik, V. and M. Lee (1990). General equilibrium pricing of options on the market portfolio with discontinuous returns. *Review of Financial Studies* 3, 493–521.
- Press, W., B. Flannery, S. Teukolsky, and W. Vetterling (1993). *Numerical Recipes in C*. Cambridge University Press, second edition.