Calibration and Hedging under Jump Diffusion

C. He [*] J.S. Kennedy [†] T. Coleman [‡] P.A. Forsyth [§] Y. Li [¶] K. Vetzal [‖]

Abstract

A jump diffusion model coupled with a local volatility function has been suggested in [1]. This model appears capable of capturing the observed market cross-sectional implied volatilities, without being unduly complex. By generating a discrete set of American option prices assuming a jump diffusion with known parameters (a synthetic market), we investigate two crucial challenges intrinsic to this type of model: calibration of model parameters and hedging jump risk. Our investigation suggests that, while it seems more difficult to estimate the model parameters that govern the jump size distribution, the local volatility function is relatively easier to estimate when an appropriate regularization (e.g. splines) is used to avoid over-fitting. In addition it appears that, in spite of the ill-posedness of the estimation problem, a jump diffusion model with a local volatility function can be estimated to price and hedge portfolios of liquid standard options sufficiently accurately. Two different hedging strategies are investigated: a semi-static approach which uses a portfolio of the underlying and traded shorter maturity options to replicate the option value of a longer maturity, and a dynamic technique which involves frequent trading of options and the underlying. Simulation experiments in the synthetic market suggest that both of these methods can be used to dramatically reduce the standard deviation of the hedging portfolio profit & loss distribution in the presence of jumps.

Keywords: Jump diffusion, calibration, static hedging, dynamic hedging

(This version: March 17, 2005)

1 Introduction

Limitations of the constant volatility Black-Scholes model are well known. Various approaches have been suggested to account for the volatility smile, such as a deterministic local volatility, stochastic volatility, and jump diffusions. As outlined in [19], there is much empirical evidence to suggest that Brownian motion should be augmented by discontinuous jump processes. Cont and Tankov [12] provide an extensive discussion of the arguments for including jump processes in modeling stock returns. Andersen and Andreason [1] propose a jump diffusion model coupled with a deterministic local volatility, and report that this model shows some promise in terms of capturing observed implied volatility behavior.

However, there are two crucial challenges when the underlying asset follows a jump process. Model parameter estimation using market option prices is an ill-posed problem. In addition, it is impossible to perfectly hedge contingent claims (when the underlying asset follows a jump diffusion with a continuum of possible jump sizes) using a finite number of instruments, even in infinitesimal time.

A good pricing method is essential for any hedging strategy, and such methods require calibration of the model parameters to market prices. Unfortunately, even for simple jump diffusion models, the calibration problem is ill-posed [13]. Hence these calibration errors may have a serious impact on hedging performance.

[*]J.P. Morgan Securities Inc., 270 Park Ave, Floor 6, New York, NY, 10017-2070, email: changhong.he@jpmchase.com

[†]Department of Applied Mathematics, University of Waterloo, Waterloo ON, Canada N2L 3G1, email: jskennedy@uwaterloo.ca

[‡]Department of Computer Science and Center for Applied Mathematics, Cornell University, Ithaca, NY, USA, 14851, email: coleman@cs.cornell.edu

[§]School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, email: paforsyt@uwaterloo.ca

[¶]Department of Computer Science,Cornell University, Ithaca, NY, USA, 14851, email: yuying@cs.cornell.edu

[‖]Center for Advanced Studies in Finance, University of Waterloo, Waterloo ON, Canada N2L 3G1, email: kvetzal@watarts.uwaterloo.ca

There have been comparatively few studies of hedging strategies in the case of jump diffusions. Bates [4] suggests using a dynamic self-financing portfolio consisting of the underlying asset and additional options to minimize the jump risk. This idea has also been suggested in [1]. It is not clear how many additional options are required in the hedging portfolio in order to achieve a satisfactory risk reduction. In addition, if the portfolio is rebalanced frequently, bid-ask spreads for option prices may result in excessive trading costs.

Another approach is semi-static hedging [7]. In [7], the starting point is the exact relationship between the primary option value in terms of the risk-adjusted density function and the primary option payoff. By an algebraic manipulation, this relationship is transformed into a weighted integration over a continuum of short term options (at all strikes). The weights of the spanning options are given by the gamma of the primary option at the expiry date of the short term options. In [7], the integral is approximated based on a quadrature rule using a finite number of integration nodes (and hence a finite number of hedging options). The hedge portfolio may need to be rolled over repeatedly in order to hedge for a desired time horizon; hence the term semi-static hedging.

Of course, a model is necessary in order to evaluate the gamma of the primary option at some future time. In practice, since the hedger does not know precisely the market dynamics, a fitted model is used to compute the portfolio weights. This is analogous to the determination of the model parameters by calibration. In [7], the spanning relation is based on the assumption that short maturity options of a continuum of strikes are available and the market is effectively complete; the objective $\mathbb{P}$ measure density is thus not needed.

In this paper, we investigate the specific characteristics of the ill-posedness of the calibration problem for a jump diffusion model with a local volatility function. In addition, we evaluate the impact of calibration error on pricing and hedging. Our investigation suggests that, in spite of the fact that it is difficult to estimate model parameters for the jump size distribution, a local volatility function can be estimated relatively accurately, when a suitable regularization technique is used for the volatility function to avoid over-fitting.

We evaluate the effectiveness of hedging strategies in the presence of calibration errors in simulation studies. We compare both the dynamic and semi-static hedging methods for reduction of the jump risk. In order to carry out our studies, we assume the existence of a synthetic market, with known jump diffusion parameters under the objective $\mathbb{P}$ measure and risk-adjusted $\mathbb{Q}$ measure respectively. Calibration effects are then studied by assuming that a practitioner does not know the synthetic market $\mathbb{Q}$ measure parameters, and attempts to discover these parameters on the basis of a finite number of observed American option prices. Given the estimated $\mathbb{Q}$ measure parameters (which may be in error), the practitioner then attempts to construct a hedge for a short position in a contingent claim, using either a semi-static or a dynamic hedging strategy. Since we explicitly note that the hedging strategy is not exact, both of these strategies formally require an estimate of the $\mathbb{P}$ measure parameters. However, since the $\mathbb{P}$ measure parameters appear only as a weighting function in a minimization problem, the precise form of this weighting function is not crucial, at least when several near the money options are used as hedging instruments. We determine the effect of errors in these estimates (for both $\mathbb{P}$ and $\mathbb{Q}$ measure parameters) on the hedging strategy. In order to compare these different hedge strategies precisely, all the simulations are carried out in a synthetic model with (true) known parameters. We focus solely on the effectiveness of the hedging strategies, especially in the presence of calibration errors. We ignore transaction costs in the following, leaving this topic for future work.

This paper has two major thrusts. The first thrust concerns model calibration. Our objective in this work is to further understand the characteristics of the ill-posedness of the calibration problem for a jump model with a local volatility function and devise appropriate techniques to alleviate the estimation difficulties. One major difference of our investigation from [13] is that we consider the calibration problem for a jump diffusion model with a local volatility function rather than a non-parametric Levy density.

The second thrust concerns hedging under a jump diffusion model. We investigate two approaches: dynamic and semi-static strategies. The dynamic strategy implemented in this study is essentially that proposed in [4], though approached from a different point of view. Although similar in philosophy to the method described in [7], our semi-static hedging assumes that the market is incomplete with only a limited set of liquid options as hedging instruments and we formulate an appropriate risk minimization problem to compute optimal hedging positions. Moreover, we illustrate that lack of the information on the objective measure $\mathbb{P}$ density does not pose any serious difficulty in hedging if a sufficient number of liquid options are available.

In section 2, we focus on the jump model calibration problem. The appropriate mathematical formulation for dynamic hedging and semi-static hedging under jump risk are presented in §3. Computational results for

hedge effectiveness evaluation under calibration errors are presented in §4. Finally, we include concluding remarks in §5.

2 Calibration

Before evaluating hedging strategies, we first investigate characteristics of the ill-posedness when calibrating a jump diffusion model to a finite number of American price observations. Specifically, we analyze in detail two major components of ill-posedness: insufficient information regarding the tails of the price distribution and finite market option price observations in determining a volatility function.

In §2.1, we first formulate the calibration problem of a general jump diffusion model with a local volatility function $\sigma(S,t)$, using liquid standard American option prices. We investigate the characteristics of the ill-posedness of the jump model calibration problem with a constant volatility and analyze its impact on pricing and hedging in §2.2. We discuss estimation of both the jump parameters and a local volatility function in a more general jump diffusion model in §2.3.

2.1 A Jump Model Estimation Problem

Under a jump diffusion model with a deterministic local volatility function, the risk-adjusted evolution of the underlying asset price $S(t)$ is governed by

$$\frac{dS(t)}{S(t^-)} = \big(r - q - \kappa^{\mathbb{Q}}\lambda^{\mathbb{Q}}\big)dt + \sigma\big(S(t^-),t\big)dZ_t^{\mathbb{Q}} + \big(J-1\big)d\pi_t^{\mathbb{Q}} \tag{2.1}$$

where t^- denotes the instant immediately before time t, r is the risk free rate, q is the dividend yield, and $\sigma(S,t)$ is a deterministic local volatility function. The superscript $\mathbb{Q}$ denotes the pricing measure. In addition, $\pi_t^{\mathbb{Q}}$ is a Poisson counting process, $\lambda^{\mathbb{Q}} > 0$ is the jump intensity, and J is a random variable representing the jump amplitude with $\kappa^{\mathbb{Q}} = \mathbb{E}^{\mathbb{Q}}(J-1)$. For simplicity, $\log J$ is assumed to be normally distributed with constant mean $\mu^{\mathbb{Q}}$ and standard deviation $\gamma^{\mathbb{Q}}$ (thus $\mathbb{E}^{\mathbb{Q}}(J) = e^{\mu^{\mathbb{Q}}+\frac{1}{2}(\gamma^{\mathbb{Q}})^2}$). More specifically, we assume that the probability of a jump size in $[J, J+dJ]$ is given by $g^{\mathbb{Q}}(J)dJ$, where

$$g^{\mathbb{Q}}(J) = \frac{e^{\left(-\frac{\left(\log(J)-\mu^{\mathbb{Q}}\right)^2}{2(\gamma^{\mathbb{Q}})^2}\right)}}{\sqrt{2\pi}\gamma^{\mathbb{Q}}J} \ . \tag{2.2}$$

In the following, we refer to the risk-adjusted process (2.1), with a constant local volatility σ and a lognormal jump density (2.2), as Merton's jump diffusion model. Note that we omit the superscript $\mathbb{Q}$ from σ since the local volatility is the same under the objective and risk-adjusted measures in an equilibrium economy [4, 22].

The value of a European option under the process (2.1) is given by [1, 23]

$$V_\tau = \frac{\sigma^2 S^2}{2}V_{SS} + \big(r - q - \kappa^{\mathbb{Q}}\lambda^{\mathbb{Q}}\big)SV_S - rV + \lambda^{\mathbb{Q}}\left(\int_0^\infty V(S\eta,\tau)g^{\mathbb{Q}}(\eta)d\eta - V(S,\tau)\right) \ , \tag{2.3}$$

where $\tau = T - t$, and T is the expiry date of the contract. Defining

$$\mathcal{L}V \equiv V_\tau - \left(\frac{\sigma^2 S^2}{2}V_{SS} + \big(r - q - \kappa^{\mathbb{Q}}\lambda^{\mathbb{Q}}\big)SV_S - \big(r+\lambda^{\mathbb{Q}}\big)V + \lambda^{\mathbb{Q}}\int_0^\infty V(S\eta,\tau)g^{\mathbb{Q}}(\eta)d\eta\right) \tag{2.4}$$

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

$$\min(\mathcal{L}V, V - V_e) = 0 \ . \tag{2.5}$$

Assume that the market prices of current, standard American options $\{\bar{V}_j\}_{j=1}^m$ are given, where $\bar{V}_j = \bar{V}(S_0, 0; K_j, T_j)$ denotes the time $t=0$ standard option price with strike K_j and maturity T_j. To accurately price options, one typically first determines a risk-adjusted model from the current market information $\{\bar{V}_j\}_{j=1}^m$.

Let $\mathcal{H}$ denote the space of measurable functions in the region $[0, +\infty) \times [0, T_{\max}]$; here $T_{\max}$ denotes the longest option maturity of interest. Throughout the paper, $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$ and σ will represent the true parameters of the risk-adjusted jump process that governs the underlying dynamics, with $y = (\lambda^{\mathbb{Q}'}, \mu^{\mathbb{Q}'}, \gamma^{\mathbb{Q}'})$ and the local volatility σ' denoting their corresponding estimates obtained from a calibration process. The model option value with strike K and maturity T under the parameters $y = (\lambda^{\mathbb{Q}'}, \mu^{\mathbb{Q}'}, \gamma^{\mathbb{Q}'})$ and the local volatility σ' is denoted by $V(S, t; K, T, y, \sigma')$. Given a set of market option prices $\{\bar{V}_j\}_{j=1}^m$, the estimation problem for a jump model (2.1) can be formulated as the variational least squares problem

$$\min_{y \in \mathbb{R}^3, \sigma'(S,t) \in \mathcal{H}} \quad \left(\sum_{j=1}^{m} \left(V(S_0, 0; K_j, T_j, y, \sigma') - \bar{V}_j \right)^2 \right). \tag{2.6}$$

We assume that the dividend yield q is known, though this quantity could also be estimated by the calibration.

Note that, when calibrating the jump diffusion model (2.5), both the local volatility function $\sigma'(S, t)$ and the jump parameters y need to be determined from a finite set of current market prices $\{\bar{V}_j\}, j = 1, 2, \cdots, m$. This is clearly an ill-posed inverse problem. In [13], Cont et al. illustrate that the calibration problem for a Lévy process model using European option prices is ill-posed. Consequently, there can be many model parameters that yield sufficient agreement with market prices. In addition, even with fixed jump parameters y, a finite set of market option prices is insufficient to determine a unique local volatility function.

2.2 Calibrating Merton's Jump Diffusion Model

For Merton's jump diffusion model, the local volatility is simply a constant. In this section we consider calibration of Merton's model from standard American option prices, and analyze calibration difficulties and their potential effects on option pricing and hedging.

The calibration problem for Merton's jump model can be simplified as

$$\begin{aligned} \min_{y \in \mathbb{R}^3, \sigma' \in \mathbb{R}} \quad & \left(\sum_{j=1}^{m} \left(V(S_0, 0; K_j, T_j, y, \sigma') - \bar{V}_j \right)^2 \right) \\ \text{subject to} \quad & l_\sigma \le \sigma' \le u_\sigma, \; l_y \le y \le u_y \; . \end{aligned} \tag{2.7}$$

In this formulation, *a priori* information for the model parameters in the form of simple bounds can be included.

The calibration problem (2.7) of Merton's model from standard liquid option prices is ill-posed because these prices are relatively insensitive to the tails of the price distribution. We illustrate this with an example. Suppose that the current stock price is $S_0 = 100$, the interest rate is $r = 0.05$, and the dividend rate is $q = 0.02$. Assume that the risk-adjusted price dynamics are governed by the constant volatility jump model (2.1), with the jump density given by (2.2) and the synthetic market parameters

$$\lambda^{\mathbb{Q}} = 0.1, \; \mu^{\mathbb{Q}} = -0.92, \; \gamma^{\mathbb{Q}} = 0.425, \quad \text{and} \quad \sigma = 0.2 \; . \tag{2.8}$$

Andersen and Andreasen [1] calibrate Merton's model to a set of implied volatilities from the S&P 500 index, and determine the best-fit pricing parameters to be $[\lambda^{\mathbb{Q}}_{AA} = 0.089, \mu^{\mathbb{Q}}_{AA} = -0.8898, \gamma^{\mathbb{Q}}_{AA} = 0.4505, \sigma_{AA} = 0.1765]$. Clearly, our choice of synthetic market parameters in (2.8) are close to these values. A set of American option prices $\{\bar{V}_j\}_{j=1}^{30}$ are computed from the assumed jump model – these contracts are currently liquid at-the-money and out-of-the-money options. Their maturities are 1 month, half a year, and one year respectively with put option strikes of $[80, 85, 90, 95, 100]$ and call option strikes of $[100, 105, 110, 115, 120]$. Problem (2.7) is a nonlinear minimization problem with simple bounds, and we compute the parameters using a trust region method [9].

Table 2.1 displays the estimated model parameters and their corresponding calibration errors using various starting points. The results in Table 2.1 show that there are significantly different parameter estimations which yield an acceptable match of model prices to the synthetic market option values. For these results, the jump intensity $\lambda^{\mathbb{Q}'}$ and the volatility σ' are close to the true values while the estimates for the jump size parameters, $\mu^{\mathbb{Q}'}$ and $\gamma^{\mathbb{Q}'}$, are considerably different from the true values, with up to 15% error in estimation of the standard deviation $\gamma^{\mathbb{Q}}$. This suggests that the constant volatility jump diffusion calibration problem, using the liquid American option prices, is ill-posed.

starting point $(\lambda_0^{\mathbb{Q}'}, \mu_0^{\mathbb{Q}'}, \gamma_0^{\mathbb{Q}'}, \sigma_0')$	$\lambda^{\mathbb{Q}'}$	$\mu^{\mathbb{Q}'}$	$\gamma^{\mathbb{Q}'}$	σ'	$\frac{1}{2}\|V - \bar{V}\|^2$
$(0.4, 0.4, 0.4, 0.4)$	0.1077	-0.8639	0.4906	0.1991	9.0730e-05
$(0.1, 0.8, 0.1, 0.1)$	0.1008	-0.9144	0.4367	0.1999	1.9542e-06
$(0.2, -0.8, 0.2, 0.2)$	0.1039	-0.8906	0.4584	0.1996	1.3985e-05

$$\lambda^{\mathbb{Q}} = 0.1, \quad \mu^{\mathbb{Q}} = -0.92, \quad \gamma^{\mathbb{Q}} = 0.425, \quad \sigma = 0.2$$

Table 2.1: Estimate of the volatility and the jump parameters from 30 standard American puts and calls, Merton jump diffusion model. Bounds setting: $\lambda^{\mathbb{Q}'} \in [0, 1]$, $\mu^{\mathbb{Q}'} \in [-2, 2]$, $\gamma^{\mathbb{Q}'} \in [0, 1]$, $\sigma' \in [0, 1]$.

To further illustrate, Figure 2.1 plots the calibration error function $\frac{1}{2}\|V - \bar{V}\|^2$ and its contours in $(\gamma^{\mathbb{Q}'}, \mu^{\mathbb{Q}'})$-space with fixed $\lambda^{\mathbb{Q}}$ and σ. These graphs clearly indicate that there is a large, nearly flat region for the calibration error function; indeed there are many parameter sets for which the constant volatility jump model prices match the synthetic market prices within a couple of cents. This effect was also noted in [13]. Figure 2.2 plots the calibration error function $\frac{1}{2}\|V - \bar{V}\|^2$ and its contours in $(\lambda^{\mathbb{Q}'}, \sigma')$-space with fixed parameters $\mu^{\mathbb{Q}}$ and $\gamma^{\mathbb{Q}}$ for the jump amplitude. It can be observed that the intensity $\lambda^{\mathbb{Q}'}$ and volatility σ' are uniquely defined under this assumption, since the objective function has a clearly defined minimum. This suggests that the calibration problem (2.7) is well-posed if the parameters specifying the distribution of the jump amplitude are given. It may be possible to estimate the standard deviation parameter $\gamma^{\mathbb{Q}}$, based on the historical price data, in the equilibrium economy described in [4, 22] since $\gamma^{\mathbb{Q}} = \gamma^{\mathbb{P}}$ (where $\mathbb{P}$ is the objective measure). However the mean of $\log(J)$, $\mu^{\mathbb{Q}}$, under the risk-adjusted measure cannot be determined based on historical price data since $\mu^{\mathbb{Q}} \neq \mu^{\mathbb{P}}$.

How do errors in the estimated jump model parameters affect option pricing and hedging? Today's price of a simple European option depends only on the risk-adjusted transition density $p(S_0; S_T)$, i.e. the transition density from today's stock price to the stock price at T. However, as pointed out by various authors, e.g. [3], path-dependent option values, such as barrier and forward starting options, also depend on risk-adjusted transition densities from stock prices at one future time to those at another future time – the conditional transition densities. Indeed, one should be aware of this when calibration is carried out using vanillas only, as the procedure will most likely fail to capture the "fine-grain structure" of the underlying stochastic process that is involved in pricing exotics [25]. Independent of a model, vanilla option prices only provide information about $p(S_0; S_T)$ (for various T). They do not tell us about $p(S_{t_1}; S_{t_2})$, for $0 < t_1 < t_2 \leq T$. However, due to the structure imposed by Merton's jump diffusion model, all of the $p(S_{t_1}; S_{t_2})$ conditional transition densities are of the same form as $p(S_0; S_T)$. Hence, analysis of $p(S_0; S_T)$ offers insight into the ill-posedness of the calibration. We compare the transitional probability densities of the estimated model parameter sets with that of the assumed model parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$ and σ. We note that the probability density function under Merton's jump model can be computed analytically, see, e.g. [20]. In Figure 2.3, the probability density functions of $X_t = \log(\frac{S_t}{S_0})$ with $t = 1$ are plotted for varying $\gamma^{\mathbb{Q}}$, with other model parameters fixed at their assumed values. Figure 2.4 presents the probability density functions of $\log(\frac{S_t}{S_0})$ for the true model (i.e. in the synthetic market) as well as those from the calibrated models using 30 American option prices. It can be observed that the probability densities of the calibrated models are all very close to that of the true model.

Figures 2.3 and 2.4 suggest that the calibration error function has a nearly flat region surrounding the solution, due to the fact that the probability density function of the constant volatility jump diffusion model is relatively insensitive to the jump parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$. Specifically, the jump parameters y mainly affect the tail of the probability density function which has minimal effect in pricing liquid standard options, whose values are relatively insensitive to the tail of the price distribution. Moreover, we note the following:

1. It is computationally difficult and expensive to obtain accurate jump parameter estimates $(\lambda^{\mathbb{Q}'}, \mu^{\mathbb{Q}'}, \gamma^{\mathbb{Q}'})$ since the calibration error function is nearly flat in a large region surrounding the solution; this implies that the Hessian matrix of the objective function is almost singular in the nearly flat region. Indeed, our computational experience suggests that it can take hundreds of iterations of the trust region method to achieve an estimation of the minimizer with high accuracy.

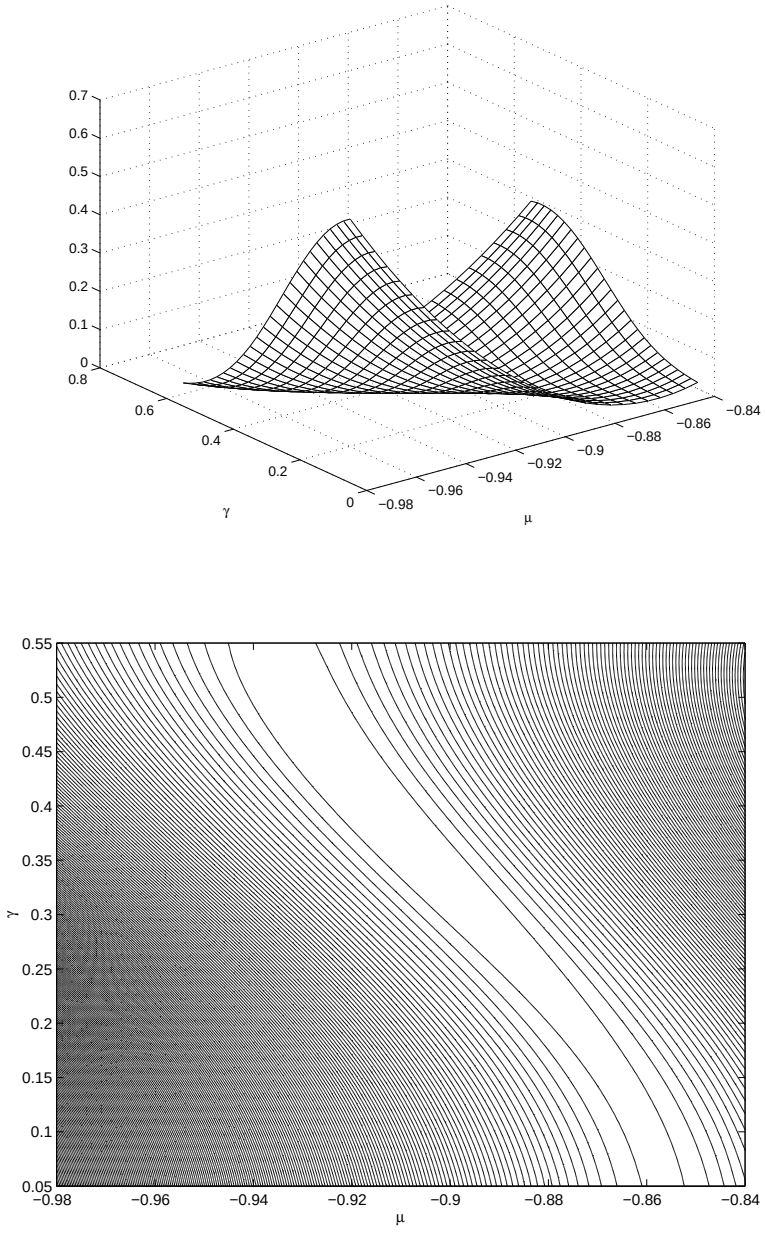

Figure 2.1: Calibration error and associated contours, with $\lambda^{\mathbb{Q}} = 0.1$ and $\sigma = 0.20$ fixed, for Merton's constant volatility jump diffusion model. In this case, for fixed $\lambda^{\mathbb{Q}}$ and σ, the objective function has a very flat region, suggesting that the minimization problem is ill-posed.

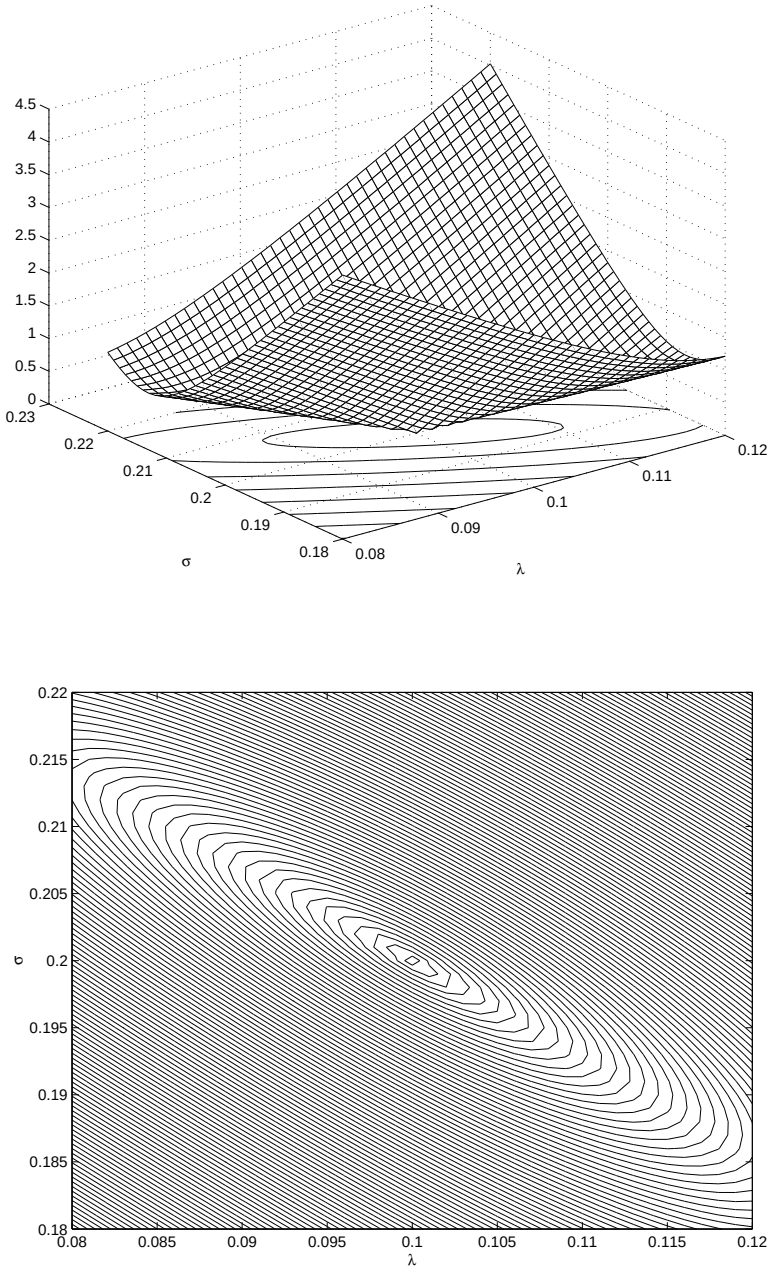

Figure 2.2: Calibration error and associated contours, with $\mu^{\mathbb{Q}} = -0.92$ and $\gamma^{\mathbb{Q}} = 0.425$ fixed, for Merton's constant volatility jump diffusion model. Note that the objective function has a well defined minimum, suggesting that if $\mu^{\mathbb{Q}}$ and $\gamma^{\mathbb{Q}}$ are known, the minimization problem is well-posed.

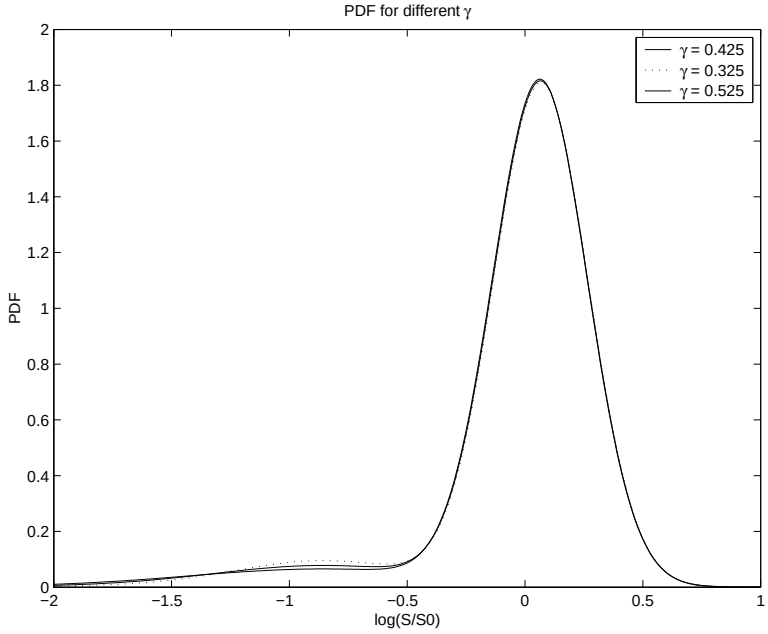

Figure 2.3: Transition probability densities for a 1-year time horizon, with assorted values of γ and other model parameters fixed ($\mu^{\mathbb{Q}} = -0.92$, $\lambda^{\mathbb{Q}} = 0.1$, $\sigma = 0.20$). Merton constant volatility jump diffusion model, lognormally distributed jump amplitude.

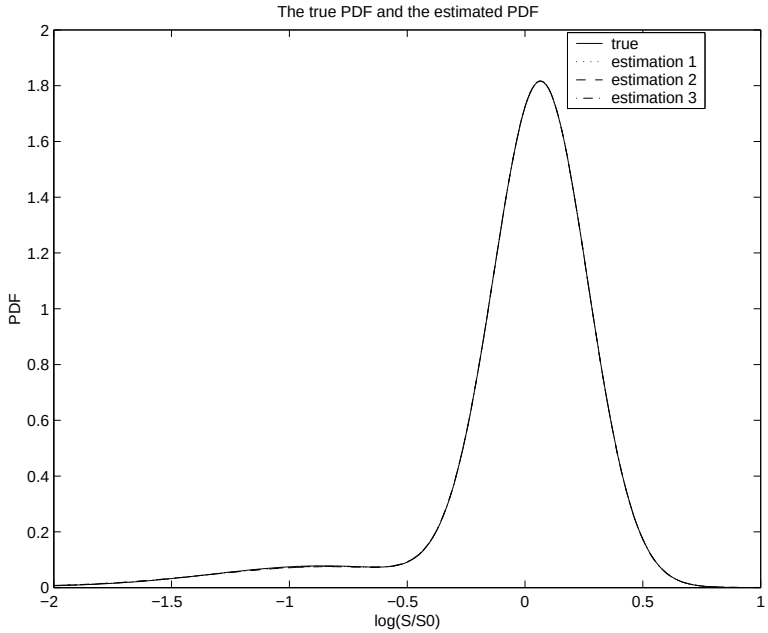

Figure 2.4: Transition probability density functions for a 1-year time horizon, with the synthetic market parameters (2.8) and the estimated parameters from 30 American puts and calls. The different estimations are obtained using various starting points in the calibration algorithm (see Table 2.1). Merton constant volatility jump diffusion model, lognormally distributed jump amplitude.

starting point $(\lambda_0^{\mathbb{Q}'}, \mu_0^{\mathbb{Q}'}, \gamma_0^{\mathbb{Q}'}, \sigma_0')$	$\lambda^{\mathbb{Q}'}$	$\mu^{\mathbb{Q}'}$	$\gamma^{\mathbb{Q}'}$	σ'	$\frac{1}{2}\|V - \bar{V}\|^2$
$(0.4, 0.4, 0.4, 0.4)$	0.1017	-0.9078	0.4424	0.1998	2.2374e-06
$(0.1, 0.8, 0.1, 0.1)$	0.1030	-0.8988	0.4615	0.1997	1.8937e-05
$(0.2, -0.8, 0.2, 0.2)$	0.1029	-0.8985	0.4518	0.1997	7.652e-06

$$\lambda^{\mathbb{Q}} = 0.1, \quad \mu^{\mathbb{Q}} = -0.92, \quad \gamma^{\mathbb{Q}} = 0.425, \quad \sigma = 0.2$$

Table 2.2: Estimate of the volatility and the jump parameters from 30 American puts and calls and 3 forward start options, Merton constant volatility jump diffusion model. Bounds setting: $\lambda^{\mathbb{Q}'} \in [0, 1]$, $\mu^{\mathbb{Q}'} \in [-2, 2]$, $\gamma^{\mathbb{Q}'} \in [0, 1]$, $\sigma' \in [0, 1]$.

2. When American option prices are replaced by European option prices in the calibration problem, we have observed similar phenomena regarding jump model parameter estimation accuracy. This suggests that the ill-posedness of the jump parameter estimation does not arise from the use of American-type options rather than European contracts.

3. For the purpose of pricing and hedging options whose values are relatively insensitive to the tails of the price distribution, it is not necessary to obtain estimates of the jump parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$ with high accuracy; we illustrate this computationally in greater detail in sections 3 and 4.

There are different approaches that one might consider if better accuracy in jump parameter estimation is desired. One possible idea is to add additional market price information, e.g. exotic option prices. To investigate this approach, three one-year maturity forward start put options, which begin at 1 month, 6 months and 9 months from the current time respectively, were added to the given set of American option prices. The calibration results are presented in Table 2.2. It can be observed that, similar to Table 2.1, there are significantly different jump models that produce a price agreement within cents for the given 33 options; the calibration problem remains ill-posed. The probability density functions for $X_t = \log(\frac{S_t}{S_0})$ in the synthetic market, as well as those from the calibrated models for 30 American puts and calls and three forward start prices, are similar to Figure 2.4.

The inclusion of additional forward start options does not help to rectify the ill-posedness of the jump model calibration problem, due to the fact that their values are still not sufficiently sensitive to the tail of the underlying price distribution. If accurate, deeply out-of-the-money option prices are available, e.g. American put options (or digital puts) with strikes in the range $[30, 60]$ for the example considered, preliminary tests indicate that the calibration problem becomes better posed and accurate jump parameter estimates can be more easily obtained. This suggests that, to ensure accurate jump parameter estimation, option instruments whose values are more sensitive to the corresponding tail of the underlying price distribution are necessary. Equity default swaps (EDS) are an example of the type of options which contain path-dependent information far from at-the-money. EDSs [30] are essentially one touch digital puts with strike set at 30% of the spot price. If a liquid market develops in these instruments, then we can expect the calibration problem to be better posed.

To overcome the ill-posedness in the jump model calibration problem, an alternative approach is to provide additional information about the model parameters directly. For example, one may assume that a prior measure exists and seek a set of model parameters that are closest to the prior in some sense, but which also yield model prices which sufficiently match the market prices. Cont et al. [13] gauge the closeness to the prior using the relative entropy, assuming that the return of the underlying price follows a Lévy process. Challenges with this approach include determination of an appropriate prior, and the balance between closeness to the prior and calibration accuracy of the available prices. The fact that the desired model parameters are risk-adjusted, rather than the real-world jump parameters, can potentially make the task of obtaining a reasonable prior challenging.

Interestingly, it seems that the most inaccurate parameter estimation in Merton's model with lognormal jumps (at least in our computational experience) is the standard deviation $\gamma^{\mathbb{Q}}$ of $\log(J)$. Thus, it is practically feasible to impose some reasonable range restriction on this parameter based on historical information, since $\gamma^{\mathbb{P}} = \gamma^{\mathbb{Q}}$. This can potentially ease some of the difficulties in the jump model calibration problem.

Test	$\|\bar{V} - V^*\|_\infty$	$\|\frac{\bar{V}-V^*}{V^*}\|_\infty$
1	0.1901	4.69%
2	0.2422	6.00%
3	0.2911	4.72%
4	0.1879	4.66%

Table 2.3: Noise level for each test

Unfortunately, additional complications exist for calibration. Market prices available for calibration are typically noisy. This means that it is not reasonable to insist that model prices and noisy market prices match with a high accuracy. This makes the calibration problem more difficult since less price information can be used to determine the model parameters. To analyze the potential effect of noisy prices on model calibration, we first modify the objective function (2.7) to account for the fact that magnitude of the noise typically depends on the moneyness of the option; we consider instead the calibration problem

$$\begin{aligned} \min_{y\in\mathbb{R}^3,\sigma'\in\mathbb{R}} \quad & \left(\sum_{j=1}^{m} w_j \big(V(S_0,0;K_j,T_j,y,\sigma') - \bar{V}_j\big)^2\right) \\ \text{subject to} \quad & l_\sigma \le \sigma' \le u_\sigma,\ l_y \le y \le u_y \end{aligned} \tag{2.9}$$

where $w_j \ge 0$ allows different degrees of trust to be placed on different prices.

We analyze the calibration problem using noisy prices by considering a similar example as before. Let V^* denote the exact synthetic market option prices under the assumed model with parameters in (2.8). We assume that the market prices $\bar{V}$ are noisy; specifically $\bar{V}$ is set as

$$\bar{V}_j \quad = \quad V_j^* + \left(0.01 + \left(\frac{K}{S_0} - 1\right)^2\right) \cdot \phi \cdot V_j^* \tag{2.10}$$

where ϕ is drawn from a standard normal distribution. Hence, at-the-money options have relatively smaller noise errors. The maximum percentage price errors in the test samples are given in Table 2.3. In the calibration problem (2.9), we set the weights as

$$w_j \quad = \quad e^{-20\left|\frac{K_j}{S_0}-1\right|}\ . \tag{2.11}$$

Roughly, the weights are $1, 0.1, 0.01$ when $|\frac{K_j}{S_0} - 1| = 0, 0.10, 0.20$ respectively.

Table 2.4 presents the estimates from the noisy prices $\bar{V}$ under the constraint $\gamma^{\mathbb{Q}'} \in [0.3, 0.55]$; other parameters have much wider ranges and their bounds are the same as those in Table 2.1. It can be observed that the estimation of the constant volatility consistently has relative high accuracy, even when calibrating from noisy prices. However, different sets of noisy prices lead to significantly different estimations of jump parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$, particularly for the average log jump size $\mu^{\mathbb{Q}}$ (recall that $\gamma^{\mathbb{Q}'}$ is restricted to the interval $[0.3, 0.55]$). Figure 2.5 compares the (one-year transitional) probability density functions corresponding to the estimated parameter sets. It can be observed that the estimated parameter sets all lead to similar probability density functions, with errors in the jump parameters predominantly affecting the left tails.

2.3 Calibrating the Local Volatility Function

In §2.2, we have assumed that the volatility is a constant and focused on analyzing the ill-posedness of the calibration problem with respect to the jump parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$ and the constant volatility σ. Unfortunately, a simple constant volatility jump model (2.1) is generally not able to calibrate the market option prices of different maturities and strikes. In this section, we consider the calibration problem of a more complex jump model (2.1) for which the local volatility is a function of the underlying price S and time

Test	$(\lambda^{\mathbb{Q}'}, \mu^{\mathbb{Q}'}, \gamma^{\mathbb{Q}'}, \sigma')$	$\frac{1}{2}\|V - V^*\|_2^2$	$\frac{1}{2}\|V - \bar{V}\|_2^2$	$\|V - V^*\|_\infty$	$\|V - \bar{V}\|_\infty$
1	$(0.0600, -1.7721, 0.4074, 0.2028)$	0.0302	0.0507	0.1904	0.2074
2	$(0.1538, -0.6069, 0.3001, 0.1959)$	0.0395	0.1046	0.2139	0.2326
3	$(0.0650, -1.6068, 0.4245, 0.2011)$	0.0426	0.1900	0.2123	0.3374
4	$(0.0929, -1.0156, 0.4380, 0.2008)$	0.0178	0.0542	0.0844	0.1592

$$\lambda^{\mathbb{Q}} = 0.1, \quad \mu^{\mathbb{Q}} = -0.92, \quad \gamma^{\mathbb{Q}} = 0.425, \quad \sigma = 0.2$$

Table 2.4: Estimation of constant volatility Merton jump diffusion model parameters from noisy prices, with the restriction $\gamma^{\mathbb{Q}'} \in [0.3, 0.55]$.

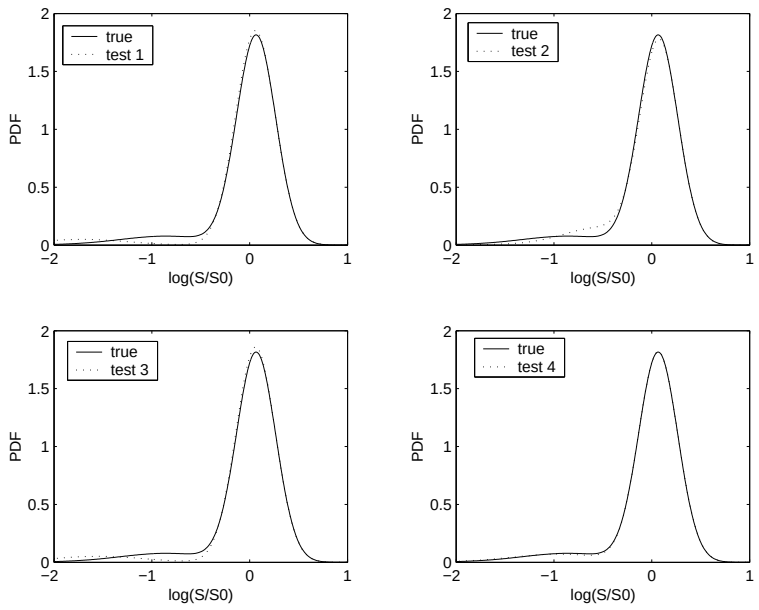

Figure 2.5: Comparison of the synthetic (true) PDF, and the PDFs estimated using the calibration results in Table 2.4. Compared to the true PDF, the most pronounced effect of the noisy data can be seen in the left tail of the estimated PDF.

t, as has been suggested in [1]. Similar to the calibration of a local volatility function in a generalized Black-Scholes model [10], we represent a local volatility function as a spline. In particular, we illustrate that with proper use of splines, the ill-posedness of the calibration problem with the jump parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$ does not seem to significantly affect the local volatility function calibration.

Assuming that the risk-adjusted price dynamics are described by a jump diffusion model (2.1) with a local volatility function, it has been illustrated in [10, 8] that it is important to estimate the local volatility function sufficiently accurately for the purposes of pricing and hedging. Given a finite set of current option prices, the jump model calibration problem (2.6) is ill-posed even when the jump parameters $y = (\lambda^{\mathbb{Q}'}, \mu^{\mathbb{Q}'}, \gamma^{\mathbb{Q}'})$ are fixed, and the impact of this ill-posedness can be significant since the option prices are more sensitive to volatilities. Thus it is important to regularize the ill-posedness of the estimation problem with respect to the local volatility function estimation and compute a parsimonious model for accurate pricing and hedging.

Smoothness has long been used [26, 27, 28] as a regularization condition for a function approximation problem with limited observation data. In addition, smoothness of the local volatility function can be important in computational option valuation schemes. In [21] it is proposed to use smoothness as a regularization condition to approximate the local volatility function in a generalized Black-Scholes model when the market European option prices are given. For the regularized optimization problem proposed in [21], the change of the first order derivative of a local volatility function is minimized depending on the regularization parameter (for which determining a suitable value may not be easy). In addition, computational implementation of this method requires solving a large-scale discretized optimization problem; see [10] for a more detailed discussion on these issues.

We use a 2-dimensional spline functional to directly approximate a local volatility function. Let the number of spline knots $\ell \leq m$ be given. We choose a set of fixed spline knots at positions $\{(\bar{S}_j, \bar{t}_j)\}_{j=1}^{\ell}$ in the region $[0, \infty) \times [0, T_{\max}]$. Given $\{(\bar{S}_j, \bar{t}_j)\}_{j=1}^{\ell}$ spline knots, an interpolating cubic spline $\varsigma(S, t; \bar{\sigma})$ with a fixed end condition is uniquely defined under the condition $\varsigma(\bar{S}_j, \bar{t}_j) = \bar{\sigma}_j, j = 1, \cdots, \ell$, with $\bar{\sigma}_j \stackrel{\text{def}}{=} \sigma'(\bar{S}_j, \bar{t}_j)$ corresponding to the local volatility value at the knot. We then determine the local volatility values $\bar{\sigma}_j$ (hence the spline) by calibrating the market observable option prices. The unknowns in this problem are the volatility values $\{\bar{\sigma}_j\}$ at the given knots $\{(\bar{S}_j, \bar{t}_j)\}$. If $\bar{\sigma}$ is an ℓ-vector, $\bar{\sigma} = (\bar{\sigma}_1, \cdots, \bar{\sigma}_\ell)$, then we denote the corresponding interpolating spline with the specified end condition as $\varsigma(S, t; \bar{\sigma})$.

Let $V_j^0(y, \bar{\sigma})$ denote the current model option price for a given (spline representation) model specification (y, ς), where $\varsigma = \varsigma(S, t; \bar{\sigma})$, i.e.,

$$V_j^0(y, \bar{\sigma}) \stackrel{\text{def}}{=} V\big(S_0, 0; K_j, T_j, y, \varsigma(S, t; \bar{\sigma})\big), \quad j = 1, \cdots, m \ .$$

To incorporate additional a priori information, lower and upper bounds (l_y, l_σ) and (u_y, u_σ) can be imposed on the jump parameters and local volatilities at the knots. Thus, we consider the *inverse spline estimation problem for a jump model (2.1)*: given ℓ spline knots $(\bar{S}_1, \bar{t}_1), \cdots, (\bar{S}_\ell, \bar{t}_\ell)$, solve for the ℓ-vector $\bar{\sigma}$

$$\begin{aligned} \min_{y \in \mathbb{R}^3, \bar{\sigma} \in \mathbb{R}^\ell} \quad & \left(f(y, \bar{\sigma}) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{j=1}^{m} w_j \left(V_j^0(y, \bar{\sigma}) - \bar{V}_j \right)^2 \right) \\ \text{subject to} \quad & l_\sigma \leq \bar{\sigma} \leq u_\sigma, \ l_y \leq y \leq u_y \end{aligned} \tag{2.12}$$

where the positive constants $\{w_j\}_{j=1}^m$ are weights, accounting for different accuracies of $\bar{V}_j$ or computed V_j. The determination of an approximation in the l_1 or l_∞ norm may be a valuable alternative, although the problem becomes even more computationally difficult.

The inverse spline estimation problem (2.12) is a minimization with respect to the jump parameters y and local volatility $\bar{\sigma}$ at the spline knots. The computed volatility function depends on the number of knots ℓ and their location, $\{(\bar{S}_j, \bar{t}_j)\}_{j=1}^{\ell}$. The choice of the number of knots and their placement in spline approximation is generally a complicated issue [18, 28]; we simply choose the minimum number of knots placed around $(S_0, 0)$ which leads to sufficiently accurate calibration of the market prices.

Compared to the calibration problem for a generalized Black-Scholes model with a deterministic volatility function in [10], the calibration problem under the jump diffusion model (2.1) is, in terms of computation, substantially more expensive. In addition it is assumed here that the market prices are those of American-style contracts, which is typically the case for equity options. The price of an American option is given by (2.5) while in the European case, $V(S, t; K, T)$ (notational dependence on y, σ is omitted here) satisfies the

forward PIDE

$$-V_T + (q - r + \kappa^{\mathbb{Q}}\lambda^{\mathbb{Q}})KV_K + \frac{1}{2}\sigma^2 K^2 V_{KK} + (1+\kappa^{\mathbb{Q}})\lambda^{\mathbb{Q}}\left(\int_0^\infty V(S,t;K/z,T)g'(z)dz - V(S,t)\right) = qV \qquad (2.13)$$

where $g'(z) = \frac{z}{1+\kappa^{\mathbb{Q}}}g^{\mathbb{Q}}(z)$.

The model estimation problem (2.12) is typically solved by an iterative method which requires at least function and gradient (Jacobian of $[V_1, \cdots, V_j]$) evaluations at each iteration. For European options, the forward equation (2.13) above can be solved to compute option prices with different strikes and maturities simultaneously. Unfortunately, for American options, each contract needs to be valued by appropriately solving a backward partial differential complementarity problem [11, 16, 29]. This becomes computationally expensive, with the Jacobian matrix computation being the most costly calculation. Note that, using finite differences, the Jacobian matrix computation requires an additional ℓm option valuations, where ℓ is the total number of spline knots and m is the total number of option prices given.

The calibration problem (2.7) for Merton's jump model is, on the other hand, far less costly since there are only four variables, and the most expensive calculation is the Jacobian evaluation which is proportional to the number of variables. By obtaining a reasonably good estimation of model parameters from the calibration problem (2.7) and using it as a starting point for the more general model (2.12), computational cost can be significantly reduced. Thus we consider a 2-stage calibration method using splines for estimating the jump parameters y and $\bar{\sigma}$.

Stage 1: Estimating a constant volatility jump model $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}}, \sigma)$. In this first stage, we assume that the volatility σ is constant (we can regard this as choosing $\ell = 4$ knots with the condition that the gradient equals zero at each knot). We compute a first approximation to model parameters y^0 and $\bar{\sigma}^0$ by solving (2.7)

$$\min_{y\in\mathbb{R}^3, \sigma'\in\mathbb{R}} \quad \left(\sum_{j=1}^m w_j\big(V(S_0,0;K_j,T_j,y,\sigma') - \bar{V}_j\big)^2\right)$$
$$\text{subject to} \qquad l_\sigma \le \sigma' \le u_\sigma,\ l_y \le y \le u_y$$

with some assumed bounds $(l_y, l_\sigma) = (l_y^0, l_\sigma^0)$ and $(u_y, u_\sigma) = (u_y^0, u_\sigma^0)$.

If the market price dynamics are mostly a jump component plus a small non-constant volatility, estimation under the constant volatility assumption is a reasonable first step in determining model parameters, as noted in [1]. Recall that this first approximation also makes sense from the point of view of computational efficiency.

Stage 2: Estimating a local volatility function. In the second stage, the spline calibration problem (2.12) is solved with the estimation from the first stage calibration problem used as the starting point.

To illustrate, we assume that the underlying price is governed by a constant elasticity of variance (CEV) model described in [14]. Specifically, we let the risk-adjusted price dynamics of the synthetic market satisfy the jump model (2.1) with $\lambda^{\mathbb{Q}} = 0.1$, $\mu^{\mathbb{Q}} = -0.92$, $\gamma^{\mathbb{Q}} = 0.425$ and the local volatility function given by

$$\sigma(S) = \frac{25}{S}\ ; \qquad S > 0$$

and we let the origin be an absorbing barrier to impose limited liability. Note that the calibration is performed without any knowledge of the parametric form of the local volatility function. The initial underlying price is assumed to be 100 and other parameters are the same as in §2.2. A volatility function is represented by a cubic spline with knots placed at $[56, 96, 136, 176, 216] \times [0.5, 1]$.

Table 2.5 presents the estimation of the parameters obtained from each stage. It can be observed that, similar to Merton's jump model calibration in §2.2, different starting points produce significantly different $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$ estimation at the end of the second stage.

Figure 2.6 compares the estimated local volatility function with the true $\sigma(S)$. This figure illustrates that, in spite of the difficulty in determining an accurate estimation of the jump parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$, a fairly accurate estimation of the local volatility function is obtained from a limited set of American option prices with appropriate smooth regularization using splines. In addition, we note that different starting points yield

	starting point $(\lambda_0^{\mathbb{Q}'}, \mu_0^{\mathbb{Q}'}, \gamma_0^{\mathbb{Q}'}, \sigma_0')$	$\lambda_p^{\mathbb{Q}'}$	$\mu_p^{\mathbb{Q}'}$	$\gamma_p^{\mathbb{Q}'}$	σ_p'	$\frac{1}{2}\|V - \bar{V}\|^2$
1^{st} stage	$(0.4, 0.4, 0.4, 0.4)$	0.1621	-0.7713	0.3056	0.2201	
2^{nd} stage		0.0955	-0.9761	0.4535		4.6104e-05
1^{st} stage	$(0.1, 0.8, 0.1, 0.1)$	0.1620	-0.7717	0.3071	0.2201	
2^{nd} stage		0.0983	-0.9528	0.5043		2.8060e-05
1^{st} stage	$(0.2, -0.8, 0.2, 0.2)$	0.1618	-0.7724	0.3050	0.2201	
2^{nd} stage		0.0931	-1.0023	0.4488		4.2284e-05

Table 2.5: Estimation of the volatility and jump parameters from 30 American options. Bounds setting: $\lambda^{\mathbb{Q}'} \in [0, 1]$, $\mu^{\mathbb{Q}'} \in [-2, 2]$, $\gamma^{\mathbb{Q}'} \in [0, 1]$, $\sigma' \in [0, 1]$. Note that σ_p' is the constant volatility predicted in the first stage of the algorithm.

very similar volatility function estimations. While the ill-posedness of the constant volatility jump diffusion model calibration problem with respect to the jump parameters y can be improved with addition of options far from at-the-money, adding option prices with more strikes and maturities will also further improve the accuracy of estimating the local volatility function.

To further illustrate, we calibrate a jump diffusion model with a volatility function from market American option price data for Brocade Communications Systems Inc. (Nasdaq-NM:BRCD) on April 21,2004. The spot price is $S_0 = 6.19$, the dividend rate is zero, and the zero rate is time-dependent; the employed zero rates, which are the yield to maturity on various zero-coupon government bonds, are listed in Table A.1 in the Appendix. Out-of-the-money calls and puts are used in our calibration problem because they are more liquid. The strike prices range from 2.5 to 20 and maturities range from $\frac{1}{12}$ to 2 years. Table A.2 in the Appendix tabulates option prices together with associated strikes and maturities. The corresponding implied volatilities are graphed in Figure 2.7.

Using $(\lambda_0^{\mathbb{Q}'}, \mu_0^{\mathbb{Q}'}, \gamma_0^{\mathbb{Q}'}, \sigma_0') = (0.4, 0.4, 0.4, 0.4)$ as the starting point, the calibrated parameters from the first stage are $\sigma_p' = 0.5299$ and $(\lambda_p^{\mathbb{Q}'}, \mu_p^{\mathbb{Q}'}, \gamma_p^{\mathbb{Q}'}) = (0.0098, -1.6269, 1.9761)$. In the second stage, the spline knots are placed on the rectangular grid

$$[0.1S_0; 1.55S_0; 3S_0] \times [0; 1; 2] \ ,$$

and the jump parameters estimated after the second stage are

$$(\lambda^{\mathbb{Q}'}, \mu^{\mathbb{Q}'}, \gamma^{\mathbb{Q}'}) = (0.0323, -0.1287, 0.9838) \ .$$

The calibrated local volatility function is graphed in Figure 2.8. The optimal objective function value is 0.0029 and the maximum the calibration error is 0.0357; the largest calibration error occurred at pricing an option with the shortest maturity.

We note that the (risk-adjusted) jump intensity obtained from the calibration is relatively small for this particular data set; here the local volatility function is the main contributor to the total volatility and the volatility skewness. In addition, a spline representation of the local volatility has led to fairly accurate calibration to the liquid option prices and the shape of the local volatility function resembles the typically observed volatility smiles.

Before moving to the study of hedging strategies, we summarize below the main findings from our investigations of calibrating a jump diffusion model from a finite set of option prices:

1. Firstly, even calibrating a simple Merton's jump model from liquid standard option prices is difficult, due to the fact that the calibration error has a large nearly flat region surrounding the optimal solution. However, the volatility estimation is consistently and relatively accurate. In addition, accurate model parameters can be eventually estimated after extensive optimization iterations, assuming that there is no error in price and model assumption.

2. Relatively accurate estimation of the transitional density function (in the case of a constant volatility) can be obtained even when the fitted jump parameters $(\lambda^{\mathbb{Q}'}, \mu^{\mathbb{Q}'}, \gamma^{\mathbb{Q}'})$ contain large relative errors. Adding standard options that are deeply out-of-the-money or deeply in-the-money should make the calibration problem easier.

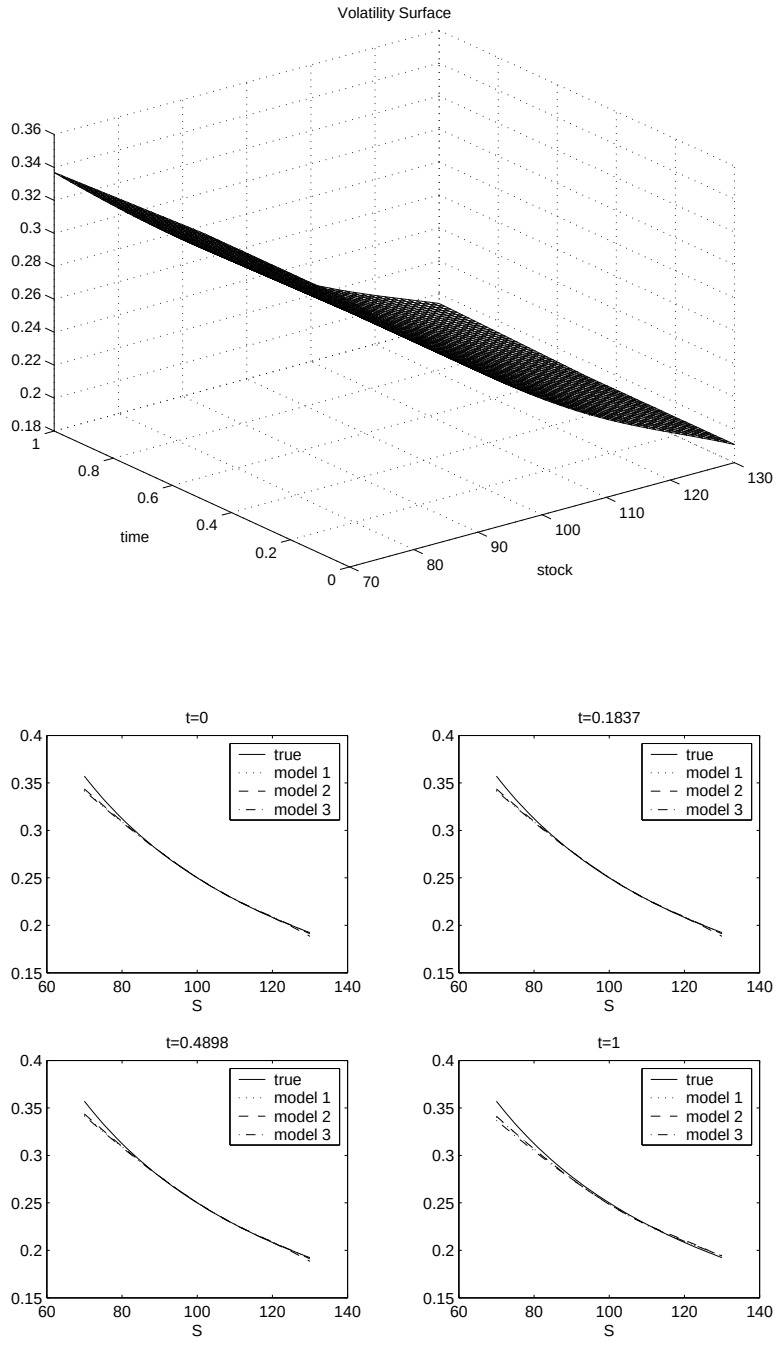

Figure 2.6: Estimation of the local volatility function from 30 American options. Models $1-3$ refer to the local volatility function obtained using different starting values in the calibration algorithm, as in Table 2.5.

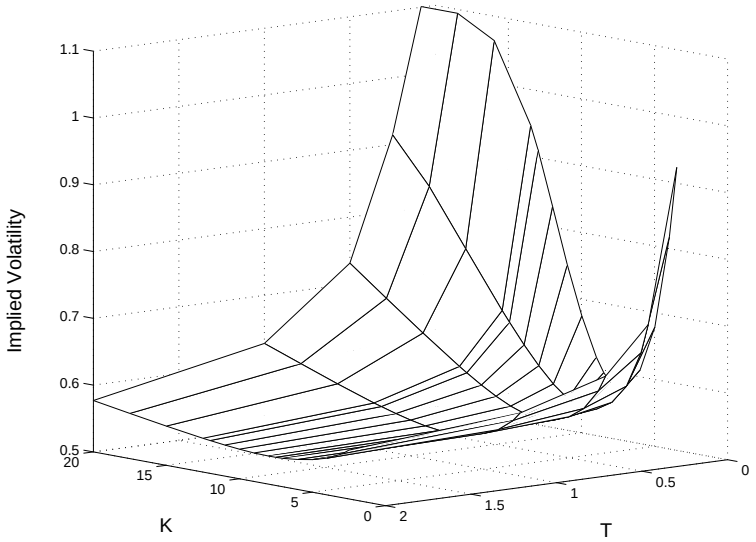

Figure 2.7: Implied volatilities, Nasdaq-NM:BRCD on April 21,2004.

3. When the available market prices are noisy, it is even more difficult to calibrate the jump parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$ accurately. However the estimation of volatility achieves relatively high accuracy, and the calibrated model yields relatively accurate pricing for liquid American options whose values are relatively insensitive to the tail of the price distribution.

4. For a more general jump diffusion model (2.1), a relatively accurate estimate to the local volatility function can be obtained (in spite of the difficulty in the estimation of the jump parameters) when a suitable regularization technique such as a spline is used to represent the unknown local volatility function.

3 Hedging Strategies

In the following, we investigate the effect of jump model calibration error on hedging and evaluate the performance of two approaches to the hedging of jump risk. Both techniques use the underlying asset and additional options to form a hedge portfolio. In the first approach, which we term dynamic hedging, the weights of the hedge portfolio are selected so as to minimize jump risk and impose delta neutrality (thus eliminating the diffusion risk) over the next infinitesimal interval. This dynamic hedging technique requires frequent buying/selling of options, and hence may incur significant transaction costs. However, knowledge of the objective measure is not required. A second approach, which we term semi-static hedging, involves selecting a hedge portfolio that attempts to replicate the value of the primary option at some future time. Semi-static hedging reduces the transaction costs, but requires estimates of the $\mathbb{P}$-measure probability density of the underlying process. For simplicity and without loss of generality, we assume that the dividend yield is zero ($q = 0$) for the subsequent hedging analysis.

3.1 Dynamic Hedging

Due to the incomplete nature of a market possessing an infinite number of possible jump sizes [1], the dynamic hedging of a contingent claim under the jump diffusion process is a far greater challenge than it otherwise would be in the (complete) Black-Scholes universe. The only source of randomness in the Black-Scholes model is the Wiener process that drives the geometric Brownian motion – the intrinsic diffusion risk may be eliminated at each instant by enforcing delta neutrality. When implemented within the Black-Scholes framework, a dynamic delta-neutral hedging strategy with discrete rebalancing tends to yield small hedging

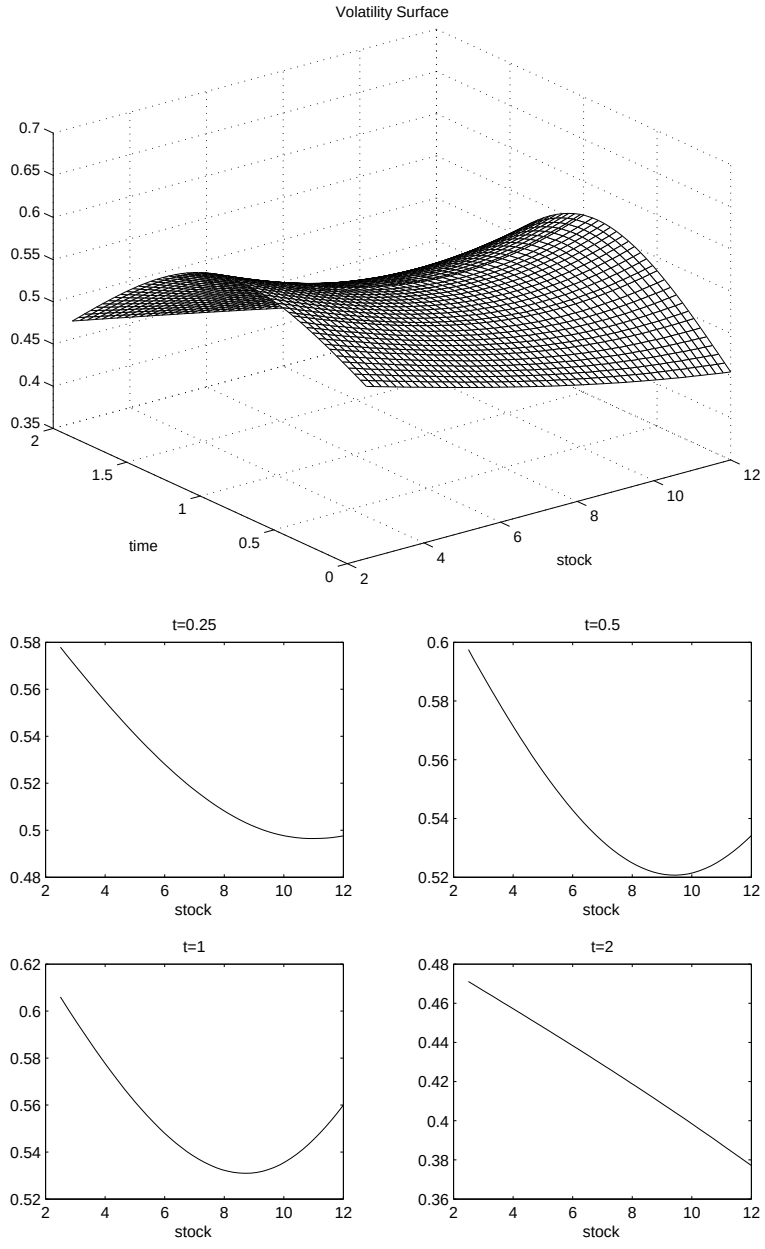

Figure 2.8: The calibrated local volatility function from Nasdaq-NM:BRCD

errors. [1] However, the risk inherent in the jump diffusion model (with a continuum of possible jump sizes) can only be removed completely by utilizing an infinite number of hedging instruments. The diffusion risk may again be eliminated by imposing delta neutrality, only now the compound Poisson process that governs the arrival and magnitude of possible jumps precludes the removal of the associated jump risk. As such, no perfect hedge exists in a market with a continuum of jump amplitudes, even in the theoretical limit of continuous rebalancing.

The unfortunate realization that jump risk cannot be totally removed is further aggravated by the fact that this risk may indeed be substantial. Consider a bank that is short an option with a convex payoff (e.g. a call) and chooses to construct a continuously rebalanced delta-neutral hedge using the underlying. Consequently, the institution is only afforded protection against the diffusion risk. If a jump does occur, the convexity of the option essentially guarantees the bank's position will lose money regardless of the size or direction of the movement [23]. Even if a jump event is deemed highly unlikely, as is often assumed the case, the potentially disastrous aftermath is reason enough to force a competent option writer to hedge the jump risk. In what follows, we derive an expression for the instantaneous jump risk and develop a measure of the overall exposure to this risk. Within our dynamic hedging strategy, the representation of jump risk exposure is minimized at each rebalance time using an appropriate choice of hedge portfolio weights. Furthermore any desired linear equality constraints, such as delta neutrality, may also be imposed.

3.1.1 Jump Risk

The concept of jump risk has been alluded to in the above discussion. In this section we shall derive a mathematical representation of this type of risk under the jump diffusion model. The hedge portfolio contains an amount B in cash, is long e units of the underlying asset S, and long N additional hedging instruments $\vec{I} = [I_1, I_2, ..., I_N]^T$ (written on the underlying) with weights $\vec{\phi} = [\phi_1, \phi_2, ..., \phi_N]^T$. For notational simplicity, we assume that $\vec{I}$ denotes all the possible hedging instruments for the entire hedging horizon with the understanding that the holdings of hedging instruments that are not traded at the current time are explicitly set to zero. When combined with the short position in the primary option $-V$, the resulting hedged portfolio has value

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

where, to ease notation, the explicit dependence on time t and/or asset price S has been dropped. The Δ-notation that represents the change due to a jump of size J,

$$\begin{aligned} \Delta V &= V(JS) - V(S) \\ \Delta S &= S(J-1) \\ \Delta\vec{I} &= \vec{I}(JS) - \vec{I}(S)\ , \end{aligned}$$

is used in the following.

If a change in the short position $-V$ is always precisely neutralized by the hedge portfolio $(eS + \vec{\phi}\cdot\vec{I} + B)$ created by the option writer, the hedge is considered perfect and Π will have zero variation over an instant dt. It is therefore necessary to consider the infinitesimal change of the hedged portfolio value Π in order to explore its risk characteristics. Since we are concerned with the real-world evolution of this portfolio, the underlying jump diffusion process of interest is the one governed by the objective measure $\mathbb{P}$. Using standard stochastic calculus techniques, the differentials of the instruments in the hedged portfolio are found to be

$$\begin{aligned} dS &= \xi^{\mathbb{P}} S\ dt \;+\; \sigma S\ dZ^{\mathbb{P}} \;+\; \Delta S\ d\pi^{\mathbb{P}} \\ dV &= \left[V_t + \frac{\sigma^2 S^2}{2} V_{SS} + \xi^{\mathbb{P}} S V_S\right] dt \;+\; \sigma S V_S\ dZ^{\mathbb{P}} \;+\; \Delta V\ d\pi^{\mathbb{P}} \\ d\vec{I} &= \left[\vec{I}_t + \frac{\sigma^2 S^2}{2}\vec{I}_{SS} + \xi^{\mathbb{P}} S \vec{I}_S\right] dt \;+\; \sigma S\vec{I}_S\ dZ^{\mathbb{P}} \;+\; \Delta\vec{I}\ d\pi^{\mathbb{P}} \\ dB &= rB\ dt \end{aligned}$$

[1] Of course even for pure diffusion cases, dynamic hedging strategies applied to barrier options become problematic if the spot approaches the barrier, since the gamma becomes infinite at this point. This has led some researchers to suggest static hedging for barrier options [15, 6, 2].

where $\xi^{\mathbb{P}} = \alpha^{\mathbb{P}} - \lambda^{\mathbb{P}}\kappa^{\mathbb{P}}$. Using the above, the instantaneous change in the value of the hedged portfolio is

$$\begin{aligned} d\Pi &= -dV \;+\; e\,dS \;+\; \vec{\phi}\cdot d\vec{I} \;+\; dB \\ &= -\Big[V_t + \tfrac{\sigma^2 S^2}{2} V_{SS}\Big]\,dt \;+\; \vec{\phi}\cdot\Big[\vec{I}_t + \tfrac{\sigma^2 S^2}{2}\vec{I}_{SS}\Big]\,dt \\ &\quad +\; \Big[-\Delta V + \vec{\phi}\cdot\Delta\vec{I} + e\Delta S\Big]\,d\pi^{\mathbb{P}} \;+\; rB\,dt \\ &\quad +\; \xi^{\mathbb{P}} S\Big[-V_S + \vec{\phi}\cdot\vec{I}_S + e\Big]\,dt \;+\; \sigma S\Big[-V_S + \vec{\phi}\cdot\vec{I}_S + e\Big]\,dZ^{\mathbb{P}} \end{aligned} \tag{3.1}$$

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

Delta neutrality compels the overall delta $\frac{\partial \Pi}{\partial S}$ of the hedged portfolio to be zero, and is represented mathematically by the linear equation

$$-V_S \;+\; \vec{\phi}\cdot\vec{I}_S \;+\; e = 0\ . \tag{3.2}$$

Imposing delta neutrality within equation (3.1) eliminates the final two terms in the expression for $d\Pi$, including the one involving the increment of the Wiener process $dZ^{\mathbb{P}}$; a delta-neutral portfolio has no instantaneous diffusion risk. The expression for $d\Pi$ consequently simplifies to

$$\begin{aligned} d\Pi &= -\Big[V_t + \tfrac{\sigma^2 S^2}{2} V_{SS}\Big]\,dt \;+\; \vec{\phi}\cdot\Big[\vec{I}_t + \tfrac{\sigma^2 S^2}{2}\vec{I}_{SS}\Big]\,dt \;+\; rB\,dt \\ &\quad +\; \Big[-\Delta V + \vec{\phi}\cdot\Delta\vec{I} + e\Delta S\Big]\,d\pi^{\mathbb{P}}\ , \end{aligned} \tag{3.3}$$

indicating that $d\Pi$ is now a pure jump process with drift. We make the assumption that at state (S,t), the option V and none of the hedging instruments in $\vec{I}$ are exercised early. If the short position in V is exercised, then the hedge portfolio is liquidated to cover the short position. If any of the hedging instruments are exercised at (S,t), then they are replaced by another instrument. Using an elementary rearrangement, the pricing PIDEs [1, 23] for V and $\vec{I}$ in the continuation region of the form (2.3) may be written as

$$\begin{aligned} V_t \;+\; \tfrac{\sigma^2 S^2}{2} V_{SS} &= rV \;+\; \big\{\lambda^{\mathbb{Q}}\mathbb{E}^{\mathbb{Q}}(\Delta S) - rS\big\}V_S \;-\; \lambda^{\mathbb{Q}}\mathbb{E}^{\mathbb{Q}}(\Delta V) \\ \vec{I}_t \;+\; \tfrac{\sigma^2 S^2}{2}\vec{I}_{SS} &= r\vec{I} \;+\; \big\{\lambda^{\mathbb{Q}}\mathbb{E}^{\mathbb{Q}}(\Delta S) - rS\big\}\vec{I}_S \;-\; \lambda^{\mathbb{Q}}\mathbb{E}^{\mathbb{Q}}(\Delta\vec{I}) \end{aligned} \tag{3.4}$$

where

$$S\kappa^{\mathbb{Q}} = S\mathbb{E}^{\mathbb{Q}}(J-1) = \mathbb{E}^{\mathbb{Q}}\big(S[J-1]\big) = \mathbb{E}^{\mathbb{Q}}(\Delta S)$$

is involved in the rearrangement. Note the $\mathbb{Q}$ superscript signifies that the corresponding quantities are risk-adjusted, e.g. $\lambda^{\mathbb{Q}} \neq \lambda^{\mathbb{P}}$ in general. Substituting (3.4) into (3.3) yields

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

Employing the delta-neutral constraint (3.2) in the above gives

$$\begin{aligned} d\Pi &= r\Big[-V+\vec{\phi}\cdot\vec{I}+eS+B\Big]\,dt \\ &\quad + \lambda^{\mathbb{Q}}\Big[\mathbb{E}^{\mathbb{Q}}(\Delta V)-\vec{\phi}\cdot\mathbb{E}^{\mathbb{Q}}(\Delta\vec{I})-e\mathbb{E}^{\mathbb{Q}}(\Delta S)\Big]\,dt \\ &\quad + \Big[-\Delta V+\vec{\phi}\cdot\Delta\vec{I}+e\Delta S\Big]\,d\pi^{\mathbb{P}} \\ &= r\Pi\,dt \\ &\quad + \lambda^{\mathbb{Q}}\,dt\,\mathbb{E}^{\mathbb{Q}}\Big[\Delta V-\big(\vec{\phi}\cdot\Delta\vec{I}+e\Delta S\big)\Big] \\ &\quad + d\pi^{\mathbb{P}}\,\Big[-\Delta V+\big(\vec{\phi}\cdot\Delta\vec{I}+e\Delta S\big)\Big]\,. \end{aligned} \tag{3.5}$$

Therefore equation (3.5) indicates that the hedged portfolio grows at the riskfree rate (as usual, if the portfolio is delta-neutral), but has additional terms due to the jump component:

$$\underbrace{\lambda^{\mathbb{Q}}\,dt\,\mathbb{E}^{\mathbb{Q}}\Big[\Delta V-\big(\vec{\phi}\cdot\Delta\vec{I}+e\Delta S\big)\Big]\;+\;d\pi^{\mathbb{P}}\,\Big[-\Delta V+\big(\vec{\phi}\cdot\Delta\vec{I}+e\Delta S\big)\Big]}_{\text{Instantaneous Jump Risk}}\,. \tag{3.6}$$

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 and the size of the resulting jump. Note that if the jump processes under $\mathbb{P}$ and $\mathbb{Q}$ are the same, the real-world expected value of the instantaneous jump risk is zero.

3.1.2 Minimizing Jump Risk

The size of the stochastic element of the jump risk in (3.6), assuming a jump event occurs ($d\pi^{\mathbb{P}}=1$), is treated as a random variable ΔH dependent on the jump amplitude J:

$$\Delta H(J) \;=\; -\Delta V\;+\;\vec{\phi}\cdot\Delta\vec{I}\;+\;e\Delta S\,.$$

This function simply represents the change in the hedged portfolio value due to a jump of amplitude J. We consider only the stochastic component of the jump risk since the deterministic constituent becomes small when $\Delta H(J)$ is minimized (in a manner that will be made clear below). Alternatively, the deterministic component can be set to zero by prescribing a suitable linear constraint, namely $\mathbb{E}^{\mathbb{Q}}\big[\Delta H\big]=0$.

Consider the writer of an option who, for the current time t and asset value S_t, wants to initiate a hedge portfolio so that her position is insulated from diffusion and jump risk over the next instant dt. If the jump amplitudes are drawn from a finite set of size M, the jump risk may be totally eliminated by introducing M hedging instruments into $\vec{I}$. The linear system

$$\begin{aligned} -\Big[V(J_iS_t)-V(S_t)\Big]\;+\;\vec{\phi}\cdot\Big[\vec{I}(J_iS_t)-\vec{I}(S_t)\Big]\;+\;eS_t\Big[J_i-1\Big] &= 0 \qquad i=1,...,M \\ -\frac{\partial V}{\partial S}\Big|_{S_t}\;+\;\vec{\phi}\cdot\frac{\partial\vec{I}}{\partial S}\Big|_{S_t}\;+\;e &= 0 \end{aligned} \tag{3.7}$$

ensures that for the current time t and asset price S_t, the diffusion risk is removed and the hedged portfolio is invariant to jumps of size $J^*=[J_1,J_2,...,J_M]$. In the more likely case where the jump amplitudes are drawn from an infinite set, this idea can be used as the basis for a simple hedging strategy. By eliminating the jump risk for a finite number of suitably chosen jump amplitudes, it is hoped there will be very little hedging error in the event that an actual jump occurs between those values in J^*. To ensure the linear system (3.7) has a solution, some care should be taken when selecting the hedging instruments for $\vec{I}$. For example if $\vec{I}$ contains a European call and put of the same maturity and strike, the resulting matrix will be singular due to put-call parity.

It is most often assumed the jump sizes are drawn from a continuum, which implies an infinite number of instruments are needed to eliminate the jump risk. Therefore the goal is to reduce the jump risk in some

optimal sense using a finite and practical number of instruments. This is done by minimizing the integral of $\Delta H(J)^2$ times a positive weighting function $W(J)$

$$\min_{e,\vec{\phi}} \int_0^\infty \left[-\Delta V + \{\vec{\phi}\cdot\Delta\vec{I} + e\Delta S\} \right]^2 W(J) dJ \tag{3.8}$$

where $W(J)$ has the properties of a probability density, namely

$$\begin{aligned} W(J) &\geq 0 \,, \\ \int_0^\infty W(J)\, dJ &= 1 \,. \end{aligned}$$

In essence the hedge portfolio weights e and $\vec{\phi}$ are chosen to minimize the integral in (3.8), which may be regarded as a metric for the overall exposure to jump risk.

At time t the asset price and all option values are known – the only user controlled input for the optimization problem (3.8) is the weighting function. The selection of $W(J) = g^{\mathbb{P}}(J)$ [4, 1], the distribution of jump amplitudes under the objective measure, seems a logical choice. In this case the optimization problem is similar to a local variance minimization. To see this, note the instantaneous change in the delta-neutral hedged portfolio (3.5) may be written as

$$d\Pi = a\ dt + b\ d\pi^{\mathbb{P}} \quad \text{with} \quad b = -\Delta V + \{\vec{\phi}\cdot\Delta\vec{I} + e\Delta S\} \,.$$

Neglecting terms of order $O(dt^2)$ and using the fact that $\mathbb{E}^{\mathbb{P}}\left[(d\pi^{\mathbb{P}})^n\right] = \lambda^{\mathbb{P}} dt$ (for all n) gives

$$\begin{aligned} \text{Var}[d\Pi] \approx \lambda^{\mathbb{P}}\ dt\ \mathbb{E}^{\mathbb{P}}\left[b^2\right] &= \lambda^{\mathbb{P}}\ dt\ \mathbb{E}^{\mathbb{P}}\left[\left(-\Delta V + \{\vec{\phi}\cdot\Delta\vec{I} + e\Delta S\} \right)^2 \right] \\ &= \lambda^{\mathbb{P}}\ dt \int_0^\infty \left[-\Delta V + \{\vec{\phi}\cdot\Delta\vec{I} + e\Delta S\} \right]^2 g^{\mathbb{P}}(J) dJ \,. \end{aligned}$$

Consequently, in order to locally minimize the variance, the optimization problem (3.8) is solved with $W(J) = g^{\mathbb{P}}(J)$. When supplemented with the linear constraint $\mathbb{E}^{\mathbb{P}}[\Delta H] = 0$ that hedges away the mean portfolio jump, the resulting hedging strategy is that suggested in [1].

Unfortunately, the density $g^{\mathbb{P}}(J)$ is usually not explicitly known – recall it is the risk-adjusted distribution $g^{\mathbb{Q}}(J)$ that is required for pricing, and that is implicitly determined by calibration to market prices. One may hazard a guess for the real-world distribution. For example, in Merton's constant volatility jump diffusion model the selection of an appropriate coefficient of relative risk aversion $1-\beta$ could be used via the relations

$$\begin{aligned} \gamma^{\mathbb{P}} &= \gamma^{\mathbb{Q}} \\ \mu^{\mathbb{P}} &= \mu^{\mathbb{Q}} + (1-\beta)(\gamma^{\mathbb{Q}})^2 \end{aligned} \tag{3.9}$$

to link the known risk-adjusted lognormal distribution $g^{\mathbb{Q}}(J)$ with the unknown density $g^{\mathbb{P}}(J)$. On the other hand a total lack of information may be encapsulated using a uniform distribution as the weighting function, with the density set as non-zero for the range of jump amplitudes deemed plausible. The main point is that knowledge of the objective measure $\mathbb{P}$ is not a prerequisite for the dynamic hedging strategy.

The optimization problem (3.8) may be expressed as an expectation

$$\min_{e,\vec{\phi}} \underbrace{\mathbb{E}\left[\frac{1}{2} \left[\Delta V - \{\vec{\phi}\cdot\Delta\vec{I} + e\Delta S\} \right]^2 \right]}_{F(e,\vec{\phi})} \tag{3.10}$$

where the expectation is taken with respect to $W(J)$. This is clearly acceptable given that the weighting function has the properties of a probability density. The objective function is quadratic in the unknowns e and $\vec{\phi}$; as such, optimality requires

$$\begin{aligned} \tfrac{\partial F}{\partial e} &= \mathbb{E}\left[\left(\Delta V - \{\vec{\phi}\cdot\Delta\vec{I} + e\Delta S\} \right) \left(-\Delta S \right) \right] &= 0 \\ \tfrac{\partial F}{\partial \phi_k} &= \mathbb{E}\left[\left(\Delta V - \{\vec{\phi}\cdot\Delta\vec{I} + e\Delta S\} \right) \left(-\Delta I^k \right) \right] &= 0 \qquad \text{for } k = 1, ..., N \end{aligned} \tag{3.11}$$

where N is the number of hedging instruments in $\vec{I}$. Furthermore any desired linear constraints, such as delta neutrality, may be included via Lagrange multipliers. The entries of the linear system involving expectations (e.g. $\mathbb{E}[\Delta V \Delta S]$) are calculated using strictly numerical techniques. The expectations are first transformed to correlation integrals and the Fast Fourier Transform is used to compute these integrals in an efficient manner [17]. This allows a large range of option types and weighting functions to be handled. In order to accommodate the possibility of an ill-conditioned linear system in (3.11), a least squares solution via a Singular Value Decomposition (SVD) is computed. We use the standard approach of setting small singular values to zero after the SVD is completed, and then using this modified decomposition to compute the portfolio weights.

This strategy is essentially that proposed in [4], though approached from a different point of view. In [4] the author considers hedging the primary option V by closely replicating ΔV with $\vec{\phi} \cdot \Delta \vec{I} + e\Delta S$. Assuming a jump will occur over the next instant, the goal is to mimic ΔV for all possible jump sizes. By using the jump probability density associated with $\mathbb{P}$ or $\mathbb{Q}$, a minimization in the form of (3.8) is obtained. Delta neutrality is also imposed. The replication approach is easily linked to our concept of jump risk. For a portfolio that does a good job of replication, the quantity $-\Delta V + \vec{\phi} \cdot \Delta \vec{I} + e\Delta S$ should be small for the appropriate range of jump amplitudes. Therefore the deterministic and stochastic components of the jump risk (3.6) should also be small.

To recapitulate: for a certain time t and corresponding asset value S, the overall jump risk in the sense of integral (3.8) may be minimized. This optimization involves solving a linear system, within which any desired linear equality constraints may also be incorporated. Consider an option writer who has just sold a one-year at-the-money American straddle with strike $K = \$100$, and who consequently wants to form a hedge portfolio at time zero that imposes delta neutrality and reduces her jump risk. As hedging instruments $\vec{I}$, the option writer has access to a range of 90-day European call options with strikes in increments of \$10. The financial parameters are as in Table 2.1 and the weighting function is a uniform-like distribution, as depicted in Figure 3.1. This weighting function is very much like a uniform density, except for the ramping tails – the imposition of continuity results in better numerical behavior. Presented in Figure 3.2 is the jump risk profile $\Delta H(J)$ for the optimized hedge portfolios resulting from various hedging strategies. These curves represent the change in the hedged portfolio value Π as a function of jump size.

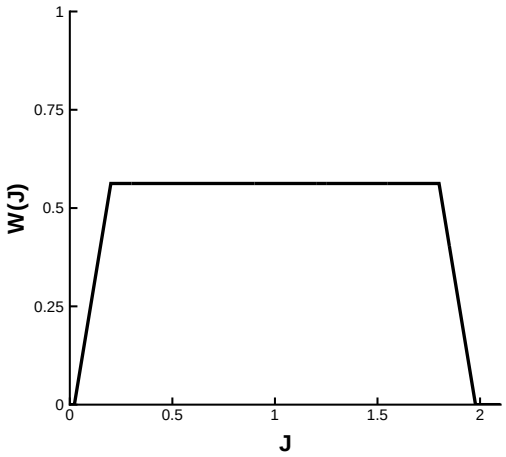

Figure 3.1: Uniform-like weighting function $W(J)$ (as in equation (3.8)) used in some of the simulations. Numerical experiments showed that use of the ramp tails resulted in better numerical behavior.

There is a large potential loss to a portfolio that ignores jumps, which is evident in the *Delta Hedge Only* curve of Figure 3.2. This is due to the fact that a delta-neutral portfolio short a convex option and long the underlying asset will always suffer a loss when a jump occurs. The $J-$intercepts where the curves cross the zero line are fortuitous points as there is no change in the hedged portfolio for jumps corresponding to these amplitudes – this is exactly what is desired. As more instruments are added to the hedge portfolio, the

curves more closely hug the zero line and the number of $J-$intercepts increases. In fact for the case of ten hedging instruments, the curve remains so close to zero that its structure is not clearly seen; an enlargement is provided in the right plot of Figure 3.2. In this instance the maximum change in portfolio value is less than 15 cents for the range of jump amplitudes $J \in [0, 1.5]$, which is a most favorable result given that the straddle costs \$21.74. Both [1] and [4] contain similar graphs. Overall, it is clear from these examples that the outlined procedure may indeed be utilized to reduce the jump risk of an option seller.

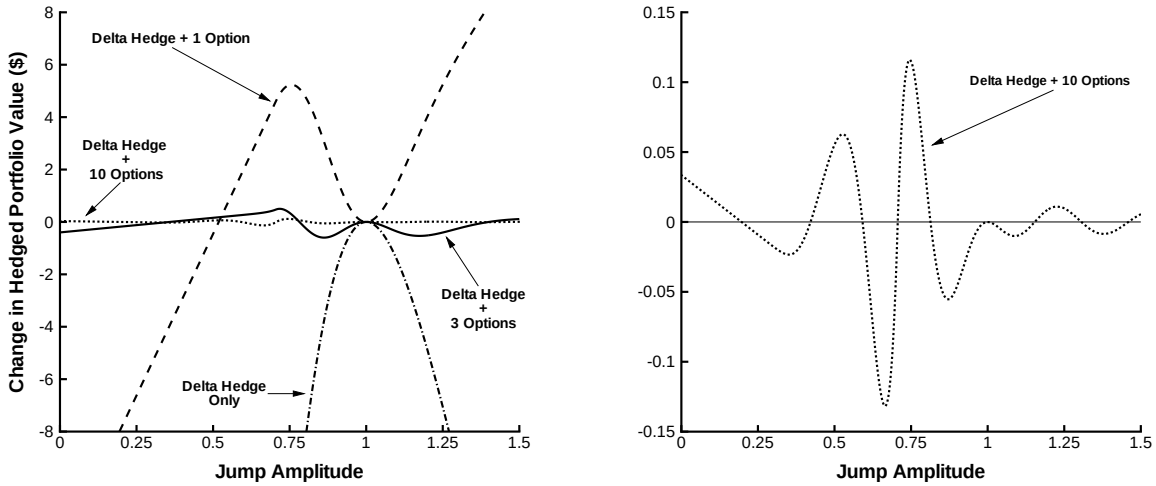

Figure 3.2: Dollar change in portfolio value resulting from a jump. The plot on the left represents the performance of several hedging strategies, while the right plot is an enlargement of the curve when ten options are used in the hedge. The option being hedged is a 1-year American straddle.

3.1.3 Implementing the Jump Risk Minimization Dynamic Hedging Strategy

In his seminal paper on the jump diffusion model [23], Merton considered the characteristics of a dynamic delta-neutral hedge that uses the underlying. Since this hedging strategy removes the diffusion component from the hedged portfolio dynamics, the return on the portfolio is a pure jump process with drift. Consequently most of the time the change in the portfolio is small. However when a jump occurs, the variation can be quite large. With Merton's assumption of diversifiable jump risk, the jump processes under $\mathbb{P}$ and $\mathbb{Q}$ are the same [22]. As such, the real-world expected value of the instantaneous jump risk (3.6) is zero. With the simple delta-neutral hedge, the many small changes and the infrequent large changes will, on average, balance each other over a long time period. In this case the hedge accounts for the average effect of jumps [12], though is obviously not perfect due to the random arrival of the jumps. However, in general, jump risk will not be diversifiable, meaning the jump process under $\mathbb{P}$ and $\mathbb{Q}$ will not be the same.

A delta-neutral hedge based on option values found by ignoring jumps (i.e. using the Black-Scholes PDE) is studied in [24]. The authors initially assume the volatility used in the pricing PDE is σ, corresponding to the diffusion component of the underlying stochastic process. If no jump event occurs over the investment horizon, then the continuous hedge is perfect as the market behaves as if it exists in the Black-Scholes universe. However a jump will lead to failure of the hedge as a cash inflow/outflow is needed to maintain replication. The deficiencies remain when a higher value of volatility, that captures some of the added variability due to jumps ($\sqrt{\sigma^2 + \lambda\gamma^2}$), is used in the pricing PDE.

Delta hedging alone is clearly not sufficient, so it is now desired to incorporate the jump risk minimization procedure into an overall hedging strategy. Ideally the hedge portfolio would be continuously updated to maintain delta neutrality and a minimal exposure to jump risk at each instant. However, as with the simple delta-neutral strategy in the Black-Scholes model, in practice the hedge can only be rebalanced at discrete times. As such, the first decision concerns the length of time between rebalancing over the hedging time horizon T.

At time zero the hedge portfolio weights $e(0)$ and $\vec{\phi}(0)$ are chosen such that the jump risk is minimized and any desired linear constraints are satisfied. The bank account is then endowed with an amount of money that gives rise to a initial hedged portfolio value of zero; therefore

$$B(0) = V(S_0, 0) - e(0)S_0 - \vec{\phi}(0) \cdot \vec{I}(S_0, 0) \ .$$

At each rebalance time t_i the hedge portfolio weights are recalculated. The long position in the underlying asset is updated by purchasing $e(t_i) - e(t_{i-1})$ shares, where $e(t_i)$ is the new computed weight and t_{i-1} denotes the time of the last rebalancing. In addition, the long positions in the hedging instruments are updated in a similar fashion by purchasing $\vec{\phi}(t_i) - \vec{\phi}(t_{i-1})$ units. These transactions must be financed by the cash account, which after the rebalancing contains

$$B(t_i) = e^{r(t_i - t_{i-1})} B(t_{i-1}) \ - \ \Big[e(t_i) - e(t_{i-1})\Big] S_{t_i} \ - \ \Big[\vec{\phi}(t_i) - \vec{\phi}(t_{i-1})\Big] \cdot \vec{I}(S_{t_i}, t_i)$$

and the hedged portfolio, an instant after rebalancing, will have value

$$\Pi(t_i) = -V(S_{t_i}, t_i) \ + \ e(t_i) S_{t_i} \ + \ \vec{\phi}(t_i) \cdot \vec{I}(S_{t_i}, t_i) \ + \ B(t_i) \ .$$

Hedging instruments may be added to or removed from the hedge portfolio at a rebalance time t_i. Furthermore the hedger may want to only adjust the position in the underlying, keeping constant the weights in the other hedging options. While this will impose delta neutrality at the rebalance time, the protection against jumps will be altered in a non-optimal way – this will be addressed in the sequel.

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

$$\text{Relative P\&L} \quad = \quad \frac{e^{-rT^*} \Pi(T^*)}{V(S_0, 0)} \ . \tag{3.12}$$

3.2 Semi-Static Hedging

As an alternative to the dynamic hedging approach described in §3.1, a semi-static hedging strategy can be implemented by minimizing directly the difference between the values of the primary option and the hedge portfolio at a specified future time. The hedge portfolio may need to be rolled over repeatedly before reaching the end of the hedging time horizon. This strategy is based on the semi-static hedging idea proposed by Carr [5], which utilizes the principle that a complex option payoff can be replicated by the payoffs of standard options with a continuum of strikes. Since the determined hedging strategy replicates the option value at a specified time, rebalancing is typically infrequent. In practice a semi-static hedge does not maintain instantaneous delta-neutrality, since it is impossible to use options with a continuum of strikes as hedging instruments.

In this section, we briefly review the main idea of semi-static hedging and discuss a computational implementation. In §4.1.2, we illustrate the accuracy of the estimated jump models in semi-static hedging.

Assume that we have an existing primary option with a T-year time to maturity, whose time t value is denoted by $V(S_t, t)$; here we explicitly emphasize the dependence on S_t. Assume that the hedge portfolio is rebalanced at a set of discrete times

$$0 = t_0 < t_1 < \cdots < t_M = T$$

where $t_k = k\delta t$.

It can be shown that the value of the primary option may be computed based on the option replication principle

$$\begin{aligned} V(S_T, T) \quad = \quad & \Big[V(\hat{K}, T) \ - \ V_K(\hat{K}, T)\hat{K}\Big] \ + \ V_K(\hat{K}, T) S_T \\ & + \ \int_0^{\hat{K}} V_{KK}(K, T)(K - S_T)^+ dK \ + \ \int_{\hat{K}}^{\infty} V_{KK}(K, T)(S_T - K)^+ dK \ , \end{aligned} \tag{3.13}$$

where V_K and V_{KK} denote the first and second derivatives with respect to the first argument in V, respectively, and the reference $\hat{K}$ can be set to $S_{t_{M-1}}$. The replication equation states that the primary option value $V(S_T, T)$ at time T can be replicated by a bond (the first term), the underlying (the second term), and a portfolio of out-of-the-money options. Thus, under the no-arbitrage assumption, the value $V(S_{t_{M-1}}, t_{M-1})$ of the primary option at time t_{M-1} is the value of the replicating portfolio at time t_{M-1}.

Stepping back to time $t_{M-1}, \cdots, t_0$, the no-arbitrage value at time zero for the primary option can be computed if the future prices of out-of-the-money standard options of all strikes are known, and the option can be semi-statically hedged by standard options with a continuum of strike prices and maturity of $t_{i+1} - t_i$.

Unfortunately, in practice, liquid standard options are restricted to a few strike prices. To account for the limited strike availability of the liquid standard options in the market, instead of the integration equation, we compute an optimal replication in a weighted least squares sense. Ideally, the weights are the transitional probability density function for the underlying price under the objective probability measure $\mathbb{P}$. When this is not available, an approximation (e.g. the transitional density under $\mathbb{Q}$) can be used. In addition, we compute an optimal hedging strategy which satisfies the self-financing constraint. More precisely, the self-financed semi-static hedging strategy is computed as follows.

Assume that the hedging instruments at time t_k have values denoted by $\vec{I}(t_k)$, with tacit dependence on the underlying asset value S_{t_k}. Let $\vec{\phi}(t_k)$ represent the positions taken in the hedging instruments $\vec{I}$ at time t_k, $e(t_k)$ be the position in the underlying asset, and $B(t_k)$ denote the amount of cash in the money market account at this time. Recall that, for notational simplicity, we assume that $\vec{I}$ denotes all the hedging instruments for the entire hedging horizon with the understanding that the holdings of hedging instruments that are not in the hedge portfolio at any rebalancing time are explicitly set to zero. Initially, $V(0) = B(0) + e(0)S(0) + \vec{\phi}(0) \cdot \vec{I}(0)$. For $k \geq 1$, $B(t_k)$ satisfies the self-financing equation

$$\underbrace{\vec{\phi}(t_k) \cdot \vec{I}(t_k) + e(t_k)S(t_k) + B(t_k)}_{\text{Instant After Rebalancing}} = \underbrace{\vec{\phi}(t_{k-1}) \cdot \vec{I}(t_k) + e(t_{k-1})S(t_k) + B(t_{k-1})e^{r\delta t}}_{\text{Instant Before Rebalancing}} \ . \tag{3.14}$$

Assuming a short position in the primary option, the values of the hedged portfolio at time t_k and t_{k+1} an instant before rebalancing are

$$\Pi(t_k) \ = \ -V(t_k) \ + \ \vec{\phi}(t_{k-1}) \cdot \vec{I}(t_k) \ + e(t_{k-1})S(t_k) \ + \ B(t_{k-1})e^{r\delta t}, \tag{3.15}$$

$$\Pi(t_{k+1}) \ = \ -V(t_{k+1}) \ + \ \vec{\phi}(t_k) \cdot \vec{I}(t_{k+1}) \ + \ e(t_k)S(t_{k+1}) \ + \ B(t_k)e^{r\delta t} \ . \tag{3.16}$$

Substituting (3.14) and (3.15) into (3.16) yields

$$\begin{aligned} \Pi(t_{k+1}) \ &= \ -V(t_{k+1}) \ + \ \vec{\phi}(t_k) \cdot \vec{I}(t_{k+1}) \ + \ e(t_k)S(t_{k+1}) \\ &\quad + \ \big(\vec{\phi}(t_{k-1}) \cdot \vec{I}(t_k) \ + \ e(t_{k-1})S(t_k) \ + \ B(t_{k-1})e^{r\delta t} \ - \ \vec{\phi}(t_k) \cdot \vec{I}(t_k) \ - \ e(t_k)S(t_k)\big)e^{r\delta t} \\ &= \ -\big(V(t_{k+1}) - V(t_k)\big) \ + \ \vec{\phi}(t_k) \cdot \big(\vec{I}(t_{k+1}) - \vec{I}(t_k)\big) + e(t_k)\big(S(t_{k+1}) - S(t_k)\big) \\ &\quad + \ \big(V(t_k) - \vec{\phi}(t_k) \cdot \vec{I}(t_k) - e(t_k)S(t_k)\big)\big(e^{r\delta t} - 1\big) \ + \ \Pi(t_k)e^{r\delta t} \ . \end{aligned} \tag{3.17}$$

Thus the optimal holding $\{e(t_k), \vec{\phi}(t_k)\}$ is chosen to minimize

$$\min_{\{e, \vec{\phi}\}} \mathbb{E}_k^{\mathbb{P}}\Big[\hat{F}(e, \vec{\phi})^2\Big] \tag{3.18}$$

where

$$\begin{aligned} \hat{F}(e, \vec{\phi}) \ &= \ \big(V(t_{k+1}) - V(t_k)\big) \ - \ \vec{\phi}(t_k) \cdot \big(\vec{I}(t_{k+1}) - \vec{I}(t_k)\big) - e(t_k)\big(S(t_{k+1}) - S(t_k)\big) \\ &\quad - \ \big(V(t_k) - \vec{\phi}(t_k) \cdot \vec{I}(t_k) - e(t_k)S(t_k)\big)\big(e^{r\delta t} - 1\big) \end{aligned} \tag{3.19}$$

and $\mathbb{E}_k^{\mathbb{P}}$ denotes the conditional expectation under the probability measure $\mathbb{P}$, or a suitable approximation. Note that effectively the $\mathbb{E}_k^{\mathbb{P}}$ enters into the optimization problem as a weighting function in the integral of equation (3.18). Hence, as in the dynamic hedging strategy, if the objective function $(\hat{F}(e, \vec{\phi}))^2$ is small over

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

Table 4.1: The pricing $\mathbb{Q}$-measure and corresponding historical $\mathbb{P}$-measure that characterize the synthetic market (Merton's constant volatility jump diffusion model). The coefficient of relative risk aversion is $1 - \beta = 2$.

a range of S values, then the precise estimation of this weighting function (the $\mathbb{P}$-measure density) should not be too crucial. We note that under Merton's constant volatility jump diffusion model, the optimal hedging positions can be computed via numerical integration using an explicit formula for the transitional density function. This is used in our implementation of semi-static hedging. Alternatively, an approximation to (3.18) based on Monte Carlo simulations can be solved to obtain the hedge positions. Once the optimal hedge positions are determined, the hedging process and the corresponding relative P&L are computed in the same fashion as described in §3.1.3. We note, however, that the semi-static and dynamic hedging strategies have completely different rebalancing times.

4 Hedging Evaluation Using Simulation

4.1 A European Example

To illustrate the various hedging strategies under consideration, we take as the primary option a 1-year European straddle with strike \$100. As the combination of a call and put, the payoff is convex and a short position will be highly sensitive to jumps of any direction. The hedging horizon is half a year, and 3-month call & put options with strikes in intervals of \$5 are used as hedging instruments – these options exist for $t \in [0.0, 0.25]$ and $t \in [0.25, 0.5]$. When the hedge portfolio is rebalanced, the choice of options is based on liquidity considerations as the active instruments should ideally have strikes in $\pm 10\%$ increments of the underlying's current value. For example if five options are to be held in the hedge at inception, those with strikes closest to $[0.8S_0, 0.9S_0, 1.0S_0, 1.1S_0, 1.2S_0]$ should be used.

In order to establish an objective measure under which to carry out the simulations, the unknown $\mathbb{P}$-measure parameters are linked to the risk-adjusted measure in (2.8) via the utility-based equilibrium relations [22]

$$\sigma^{\mathbb{P}} = \sigma^{\mathbb{Q}}, \quad \gamma^{\mathbb{P}} = \gamma^{\mathbb{Q}}, \quad \mu^{\mathbb{P}} = \mu^{\mathbb{Q}} + (1-\beta)(\gamma^{\mathbb{Q}})^2, \quad \lambda^{\mathbb{P}} = \lambda^{\mathbb{Q}} e^{\left(1-\beta\right)\left(\mu^{\mathbb{Q}} + \frac{1}{2}(1-\beta)(\gamma^{\mathbb{Q}})^2\right)}$$

where $1 - \beta$ is the coefficient of relative risk aversion. With $1 - \beta = 2$, the parameter set of Table 4.1 is obtained. This market will experience a jump event, on average, once every 43.9 years with a mean jump size of -37.4%. The equilibrium expected return of the asset is given by

$$\alpha^{\mathbb{P}} = r + (1-\beta)\sigma^2 + (\kappa^{\mathbb{P}}\lambda^{\mathbb{P}} - \kappa^{\mathbb{Q}}\lambda^{\mathbb{Q}}) \ .$$

With a risk-free rate of $r = 0.05$, the drift is found to be $\alpha^{\mathbb{P}} = 0.1779$; for $S_0 = 100$ the expected value of the asset in half a year is \$109.30.

In addition to assessing the efficacy of the hedging strategies, the simulations also investigate the consequences of using an estimated pricing measure. Recall from §2 that many different pricing measure parameter sets produce close agreement to the values of standard options generated by the $\mathbb{Q}$-measure of Table 4.1. This $\mathbb{Q}$-measure in Table 4.1 is termed the true pricing measure and may be thought of as the pricing rule responsible for all option values observed in the market. For simplicity in the following, the $\mathbb{Q}$ volatility is taken to be a constant, since we are focusing on hedging strategies.

In the experiments to come, the estimated pricing measure $\mathbb{Q}'$ is set to

$$(\lambda^{\mathbb{Q}'}, \mu^{\mathbb{Q}'}, \gamma^{\mathbb{Q}'}, \sigma') = (0.1077, -0.8639, 0.4906, 0.1991) \ ,$$

which is the calibration result from the first row of Table 2.1. The associated objective measure $\mathbb{P}'$ found using $1 - \beta = 2$ is

$$(\lambda^{\mathbb{P}'}, \mu^{\mathbb{P}'}, \gamma^{\mathbb{P}'}, \sigma') = (0.0310, -0.3825, 0.4906, 0.1991) \ .$$

Simulation Set #	Measure for Calculating Option Prices	Measure for Constructing the Hedge Portfolio
1	$\mathbb{Q}$	$\mathbb{P}$
2	$\mathbb{Q}'$	$\mathbb{P}$
3	$\mathbb{Q}'$	$\mathbb{P}'$
4	$\mathbb{Q}$	$\mathbb{Q}$
5	$\mathbb{Q}$	Lognormal ($\mu = 0.5, \gamma = 0.1$)
For all simulations, the true objective measure $\mathbb{P}$ is used to generate asset price paths. Furthermore, all observable market prices come from the true pricing measure $\mathbb{Q}$.		

Table 4.2: Summary of the simulation experiments.

A common suite of five simulation sets is carried out using both dynamic and semi-static hedging. In all cases, the underlying evolves according to a synthetic $\mathbb{P}$-measure and the market prices of options are determined by a synthetic $\mathbb{Q}$-measure – these two measures are given in Table 4.1. The first set of simulations may be regarded as the baseline case as they only involve the true model parameters of Table 4.1. The $\mathbb{Q}$-measure is used to compute all necessary option quantities (i.e. prices, deltas), while the objective probability measure $\mathbb{P}$ is used to construct the hedge portfolio. In the case of dynamic hedging, this means the weighting function is set to the lognormal density suggested by the objective measure, $W(J) = \text{logn}(-0.5588, 0.425)$. For semi-static hedging, the conditional expectation in (3.18) is taken with respect to the transitional probability density under $\mathbb{P}$.

The second set of simulations is almost identical to the first, only now the estimated pricing measure $\mathbb{Q}'$ is used to value the options. For example, with dynamic hedging the option prices found from the $\mathbb{Q}'$-measure are used in calculating the expectation integrals that constitute the coefficients of the linear system (3.11). The estimated pricing measure is also used to establish the delta-neutral constraint. For semi-static hedging, the option values involved in (3.18) are generated from the estimated pricing measure. It is important to note that it is the true option values under $\mathbb{Q}$ that are used when rebalancing the hedge portfolio, as these are the prices observed in the market. If $\mathbb{Q}$ and $\mathbb{Q}'$ give similar prices over the time range being considered, the error introduced by using $\mathbb{Q}'$ is expected to be small.

The third set of simulations gauges the effect of utilizing the estimated pricing measure in a potentially unintelligent way. As with simulation set #2, $\mathbb{Q}'$ is used to price the options. However we assume the coefficient of relative risk aversion is known, and is used in combination with $\mathbb{Q}'$ to find an estimate for the objective measure $\mathbb{P}'$. In turn this estimated objective measure is used to establish the weighting function (for dynamic hedging) or calculate the expectation in (3.18) (semi-static hedging).

The intention of the final two simulation sets (#4 & #5) is to explore further the choice of measure when constructing the hedge portfolio. For simulation set #4, the true pricing measure $\mathbb{Q}$ is used to both calculate the option values and construct the hedge portfolio. For simulation set #5 a highly unreasonable parameter set is used as an estimate for the distribution of jump amplitudes, which is subsequently employed when constructing the hedge portfolio. The simulation experiments are summarized in Table 4.2.

4.1.1 Dynamic Hedging Experiments

Only call options are used as hedging instruments in the dynamic hedge, as put-call parity allows any desired position in a put to be replicated by positions in a call and the underlying. This is acceptable here since all options are European – in a forthcoming section involving American options, out-of-the-money calls and puts are both used as an exact put-call parity relationship no longer holds. All required option values, option deltas and correlation integrals required to form the coefficients of the Lagrange multiplier linear system are precomputed before any simulations are carried out. For M hedging instruments that may possibly become active during the course of a simulation, a total of $\frac{1}{2}[M^2 + 7M + 8]$ computational grids must be determined. The time spacing of these grids is such that each rebalance time is represented exactly. If a required quantity for one of these times does not correspond to a mesh node, a cubic Lagrange interpolating polynomial is employed.

The hedge portfolio contains five hedging options and the underlying. It is completely rebalanced once

every 0.0125 years, a total of 40 times over the half-year hedging horizon. A complete rebalancing refers to the situation where all hedging weights are updated according to the jump risk minimization procedure. Between each complete rebalancing the position in the underlying alone is adjusted, at three equally spaced times, in order to impose delta neutrality. For each set described in Table 4.2, there are 500000 simulations carried out – with a mean arrival rate of $\lambda^{\mathbb{P}} = 0.0228$, there are expected to be approximately 5698 jump events over each simulation set. Note that at times near the midpoint of the investment horizon $t = 0.25$, the hedge still employs options that expire at this time (as explained above, the last possible reshuffling of these instruments occurs at $t = 0.2375$). This is not the optimal strategy, since the hedging options lose curvature as the end of their lifetime is approached. However protection against the jump risk should not be adversely affected. In the future, more sophisticated hedging regimes will be considered.

The distributions of relative profit and loss are presented in Figure 4.1, with the associated means and standard deviations in Table 4.3. The outliers of these distributions are very important when considering the hedging error, as it is precisely these potentially catastrophic events one tries to prevent. Various outliers (in the form of the 0.02%, 0.2%, 99.8% & 99.98% percentiles) are given in Table 4.4. The data Tables and Figures contain results for two simulation experiments particular to dynamic hedging, namely (i) a simple delta-neutral hedge that employs the underlying alone and (ii) use of a uniform-like weighting function $W(J)$. This weighting function is basically a uniform density between $J = 0.2$ and $J = 1.8$, with ramp tails that ensure continuity. A plot of this density is contained in Figure 3.1.

The first row of Tables 4.3 and 4.4 give the summary statistics for the case where only delta hedging is used, with the corresponding distribution plot of relative profit and loss contained in Figure 4.1. It appears the seller of the option is almost always making money, which they indeed are; in fact the 99% VAR is approximately zero. This corresponds primarily to those simulations where no jump event occurs. The mean of the distribution is positive, which agrees with theory since $\mathbb{Q}$ is more pessimistic than $\mathbb{P}$. However, the positive bias of the mass is offset by the losses that result when a jump event takes place. These outliers may be quite large – for approximately 20% of the jump events, the option seller will lose almost three times the value of the initial option premium. As was evident in Figure 3.2, a jump of any size will lead to a loss for the hedge portfolio, a loss which is often quite substantial. A cynical observer might consider this strategy to be an excellent technique for a hedge fund. Most of the time, this strategy makes money, and consequently large bonuses are paid to fund managers. Occasionally, large losses occur, but this affects only the capital of the investors.

Simulation sets #1 through #4 demonstrate the dynamic hedging strategy of jump risk minimization does work for this straddle. The means are all zero, with standard deviations that are fairly small. Furthermore the distributions of relative profit and loss have the desired shape, namely a symmetric, mean-zero density with limiting form the Dirac δ-functional. More importantly there are no significantly negative outliers, as was the case with delta hedging.

The use of the estimated pricing measure has little effect, as the results for simulation sets #1 and #2 are almost identical. This is expected since both $\mathbb{Q}$ and $\mathbb{Q}'$ yield similar option prices. As for the choice of the weighting function, there appears to be a certain amount of leeway. Using $W(J)$ motivated by $\mathbb{P}$ and $\mathbb{P}'$ give similar results, though there is a slight degradation when the weighting function suggested by $\mathbb{Q}$ is employed. The salient conclusion from simulation #5, whose summary statistics are poor, is that the use of a totally inappropriate weighting function will not yield good results. On the other hand the results from using a uniform-like weighting function are heartening – this $W(J)$ encapsulates a lack of knowledge concerning the distribution of the jump amplitudes, but nonetheless yields the best results.

So far we have considered the trading regime of 40 complete rebalances, with 3 adjustments of the underlying interspersed between each complete rebalance. How do these adjustments that re-impose delta neutrality influence the results? In the case of uniform-like weighting, two plots of relative profit and loss are presented in Figure 4.2. The solid curve represents the trading regime previously mentioned – that is, 40 complete rebalances with 3 adjustments of the underlying between each rebalance. The dashed curve corresponds to the case where only the 40 complete rebalances are carried out. By incorporating frequent updating of the underlying that imposes delta neutrality, we expect a decrease in the diffusion risk. In fact, the standard deviation drops from 0.0359 to 0.0229. By adjusting only the underlying between complete rebalances, the protection against jump risk is no longer optimal. However this does not seem to have a large effect in this case, perhaps due to the infrequent arrival of jumps.

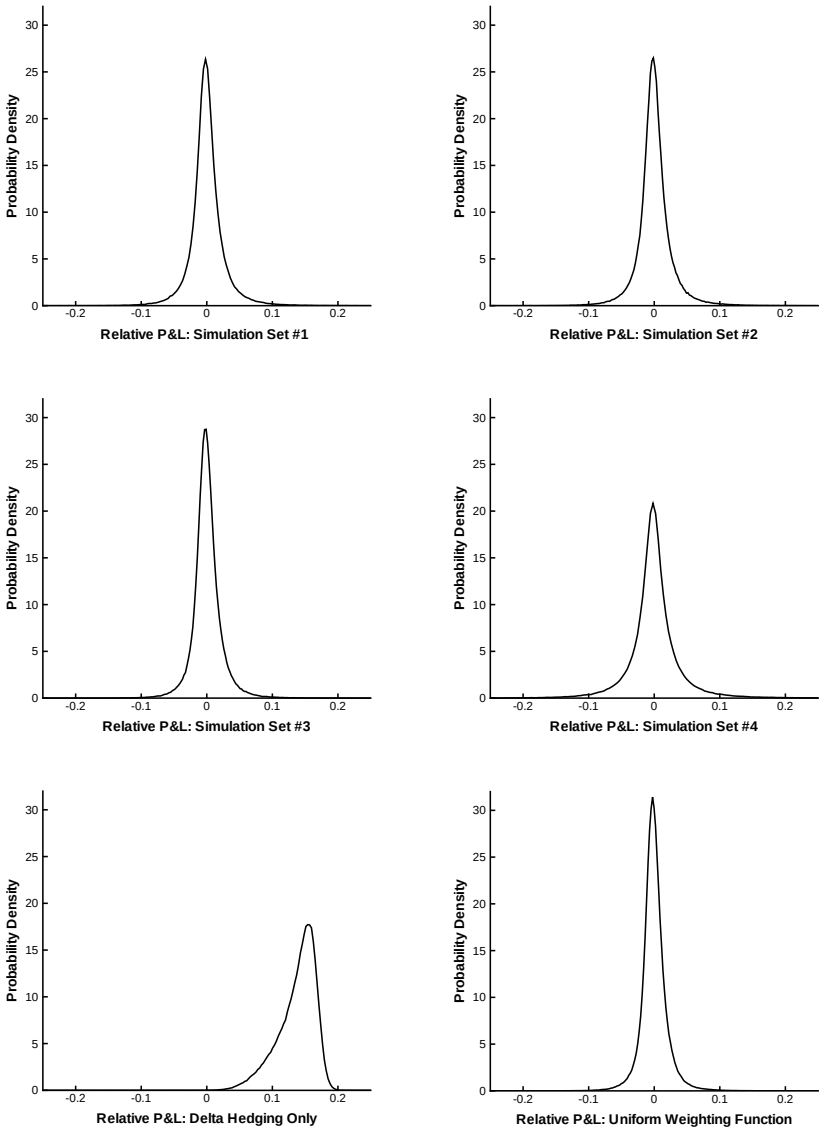

Figure 4.1: Dynamic Hedging: distributions of relative profit and loss for the European example.

Simulation Set	Mean	Standard Deviation
Delta Hedging Only	0.1172	0.2377
1	0.0003	0.0311
2	0.0003	0.0299
3	0.0004	0.0249
4	-0.0001	0.0464
5	0.0195	0.2191
Uniform Weighting	-0.0004	0.0229

Table 4.3: Dynamic Hedging: mean and standard deviation of the relative profit and loss for the European example.

Simulation Set	$\frac{100}{500000}$ Percentile	$\frac{1000}{500000}$ Percentile	$\frac{499000}{500000}$ Percentile	$\frac{499900}{500000}$ Percentile
Delta Hedging Only	-4.6026	-3.0158	0.1907	0.2201
1	-0.3403	-0.1166	0.1494	0.3947
2	-0.3640	-0.1173	0.1447	0.3517
3	-0.2860	-0.0964	0.1180	0.2915
4	-0.5027	-0.1897	0.2321	0.4987
5	-2.7946	-1.6948	0.4301	3.9394
Uniform Weighting	-0.2413	-0.0879	0.1061	0.2821

Table 4.4: Dynamic Hedging: outliers of the relative profit and loss for the European example. For the $\frac{\alpha}{500000}$ percentile, approximately α relative profit and loss values are less than this quantity.

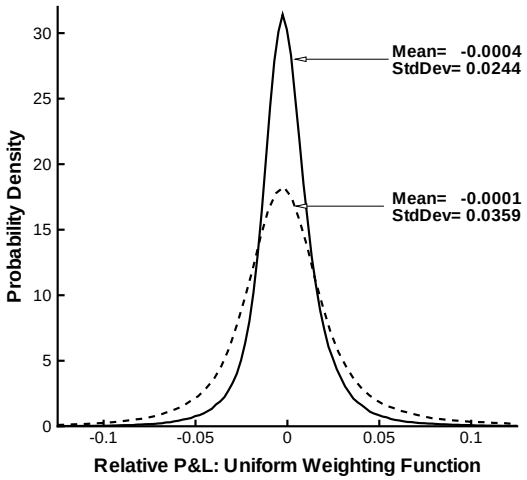

Figure 4.2: Dynamic Hedging: distributions of relative profit and loss for the European example when a uniform-like weighting is used. The solid curve corresponds to the case of 40 complete rebalances, with 3 adjustments of the underlying between each complete rebalance. The dashed curve represents the case when only the 40 complete rebalances are carried out.

Simulation Set	Mean	Standard Deviation
1	0.0008	0.0116
2	0.0007	0.0116
3	0.0009	0.0116
4	-0.0001	0.0121
5	0.0019	0.0179

Table 4.5: Semi-static Hedging: mean and standard deviations of the relative profit and loss for each simulation set. Primary option: European straddle.

Simulation set	$\frac{100}{500000}$ Percentile	$\frac{1000}{500000}$ Percentile	$\frac{499000}{500000}$ Percentile	$\frac{499900}{500000}$ Percentile
1	-0.1235	-0.0463	0.0529	0.1889
2	-0.1238	-0.0463	0.0528	0.1881
3	-0.1204	-0.0463	0.0526	0.1841
4	-0.1169	-0.0543	0.0579	0.1996
5	-0.4525	-0.0583	0.0898	0.3404

Table 4.6: Semi-static hedging: outliers of the relative profit and loss in each simulation. Primary option: European straddle.

4.1.2 Semi-static Hedging Experiments

In this section, we carry out a study of semi-static hedging, as discussed in §3.2. We consider the same scenarios as used in the study of dynamic hedging in §4.1.1 As in the case of dynamic hedging, only call options and the underlying are used as hedging instruments. At each rebalance time, call options with 3 months time to maturity and strikes closest to $[0.8, 0.9, 1.0, 1.1, 1.2]S_t$ are considered as the most liquid options and thus used as hedging instruments. The hedge portfolio is rebalanced once at $t_1 = 0.25$. Specifically, we compute a hedge portfolio at time zero that minimizes the quadratic hedging error (3.18) at time t_1, and hold this portfolio until t_1. The hedge portfolio is rebalanced at t_1 in a self-financing fashion and the optimal portfolio for the next time period $[t_1, T]$ is computed from (3.18) to minimize the quadratic hedging error at $T = 0.5$, again using the underlying and European options with 3 months to maturity and strikes in intervals of \$5 and close to $[0.8, 0.9, 1.0, 1.1, 1.2]S_{t_1}$. The five simulation experiments described in Table 4.2 are performed, with the same 500000 underlying paths used for each simulation set.

Table 4.5 presents the means and standard deviations of the relative profit and loss for each simulation set, while Table 4.6 shows the outlier information. The results in simulation 1 and simulation 2 are very close, due to the fact the estimated measure $\mathbb{Q}'$ leads to fairly accurate pricing of liquid standard options. In addition, the results of simulation 3 are quite similar to those in the first two experiments. This suggests that the estimated probability measures are adequate for hedging as well as pricing standard liquid options. In simulation 4 we assume that the objective measure $\mathbb{P}$ is not available, and the risk-adjusted measure is instead used in the hedging computation. The results of simulation 4 are close to those using the objective measure $\mathbb{P}$. In simulation 5, the probability measure used to compute the hedging strategy is significantly different from the objective measure $\mathbb{P}$. It is interesting to observe that the hedging error experiences only a slight increase, in spite of the very poor choice for the transition PDF. This is an encouraging result, since it suggests that this hedging strategy is robust to misspecification of the objective PDF. This observation is consistent with the results obtained in the study of semi-static hedging in [7].

In Figure 4.3, the estimated densities of relative profit and loss for the first four simulation sets are graphed. It can be observed that the densities of profit and loss for these experiments are very close. This suggests that, even though some of the model parameters have relatively large errors, the estimated jump models lead to acceptably accurate pricing and hedging for the straddle.

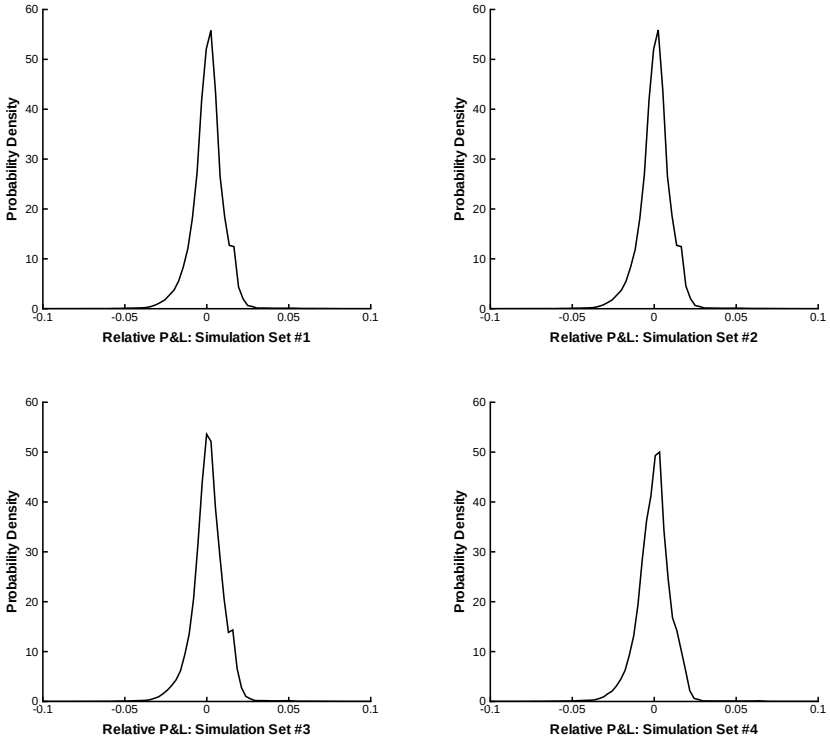

Figure 4.3: Semi-static hedging: distributions of relative profit and loss for the European straddle example.

4.2 Dynamic Hedging of an American Option

In this section we demonstrate that the jump risk minimization procedure may be used to dynamically hedge a contract with American-style exercise rights. The primary option to be hedged is an American put with a half-year maturity. The hedging horizon is a quarter-year, and 3-month out-of-the-money American puts and calls are used as the hedging instruments. Similar hedge portfolio selection rules apply as in the European case, only now out-of-the-money options are (ideally) used, as well as the underlying asset. For instance if five options are to be held in the hedge portfolio at inception, those puts with strikes closest to $[0.8S_0, 0.9S_0, 1.0S_0]$ and calls with strikes nearest $[1.1S_0, 1.2S_0]$ are used. Note that an American hedging put may have to be exercised before the end of the quarter-year hedging horizon. Essentially this means that the early exercise regions of all active American hedging options are monitored, and if the asset price path enters such a region, it is assumed both long and short positions in the corresponding put are exercised optimally. In this case, the hedging option is suitably replaced with a put of lower strike. If the half-year put being hedged is exercised early, the hedge portfolio is immediately liquidated to cover the payoff.

The same financial parameters as in the European case are utilized, with the only exception being that the mean arrival rate of jumps in the market is set about four times higher, to $\lambda^{\mathbb{P}} = 0.1$. This will provide a moderately higher stress on the hedging system than before, as over the course of 500000 simulations approximately 12500 jump events are expected to occur. A uniform-like weighting function (in equation (3.8)) is used, and the number of hedging instruments and rebalancing frequency is varied. Note that only complete rebalances are carried out – there are no intermediate adjustments made to the underlying in order to re-impose delta neutrality. Table 4.7 contains the summary statistics while the distributions of relative profit and loss are presented in Figure 4.4 for the case of 20 and 5 rebalances.

The results are consistent with the European case. When a simple delta-neutral hedge is employed, the seller of the put makes money both most of the time and on average. However this strategy may again result in catastrophic losses when a jump occurs, as the put has a convex payoff. The use of other instruments within our dynamic hedge mitigates this jump risk, as is evident with the $\frac{100}{500000}$ and $\frac{1000}{500000}$ percentiles.

Summary Statistics for the American Example							
Number of Rebalances	Hedging Options	Mean	Standard Deviation	$\frac{100}{500000}$ Ptile	$\frac{1000}{500000}$ Ptile	$\frac{499900}{500000}$ Ptile	$\frac{499990}{500000}$ Ptile
20	0	0.0427	0.6758	-9.5941	-7.0525	0.2987	0.3224
	3	-0.0011	0.0757	-0.8719	-0.5114	0.1145	0.2763
	5	0.0011	0.0193	-0.1367	-0.0738	0.0854	0.2067
10	0	0.0414	0.6789	-9.6190	-7.0703	0.3408	0.3609
	3	-0.0008	0.0959	-1.0719	-0.6372	0.3626	0.5695
	5	0.0010	0.0244	-0.1748	-0.0882	0.1145	0.2736
5	0	0.0392	0.6799	-9.5877	-6.9625	0.3754	0.3860
	3	-0.0021	0.1196	-1.2362	-0.7441	0.4792	0.7353
	5	0.0011	0.0311	-0.2577	-0.1288	0.1397	0.3695
3	0	0.0346	0.6900	-9.4573	-6.8169	0.3887	0.3922
	3	-0.0024	0.1238	-1.1196	-0.6945	0.5578	0.7862
	5	0.0008	0.0354	-0.3543	-0.1652	0.1509	0.3701

Table 4.7: Dynamic Hedging: summary statistics for the American example. The hedging horizon is 0.25 years.

The use of more options in the hedge portfolio leads to better performance, which is not surprising. Note that the standard deviation and negative outliers for the case of zero hedging options (delta hedging only) are essentially the same, regardless of the rebalancing frequency. This should reinforce the idea that the ignored jump risk inherent to the simple delta-neutral hedge will manifest itself regardless of the frequency of rebalancing.

It is interesting to observe from Table 4.7 that using the hedging strategy based on minimization of the infinitesmal variance (the dynamic strategy) does a remarkable job of reducing the standard deviation of the $P\&L$ even if the the rebalancing interval is monthly. In this case, the distinction between dynamic and semi-static becomes blurred.

5 Conclusion

Empirical evidence indicates that a jump diffusion model coupled with a local volatility function is able to capture the observed market cross-sectional implied volatilities [1]. In this paper, we investigate two crucial challenges associated with such a jump model, namely model calibration and hedging of jump risk. Previous authors, e.g. [13], have shown that the calibration problem of a basic Merton constant volatility jump diffusion model using near the money vanilla puts and calls is ill-posed. There is a large range of parameter values which match the observed market prices to within a reasonable error tolerance.

When a jump model is coupled with a local volatility function, the estimation problem becomes computationally difficult. In addition to the fact that typical liquid standard options lack information regarding tails of the price distribution, there are insufficient option prices to completely determine a local volatility function. By generating synthetic market prices where the true market follows a jump diffusion process with known parameters, we illustrate that, while it is difficult to obtain accurate parameter estimations for the jump size distribution, it is relatively easier to obtain the local volatility. In the context of a general jump model with a local volatility function, it is crucial to employ suitable regularization technique, e.g. splines, to overcome the difficulties associated with sparse market data.

An additional challenge with a jump model is how to effectively hedge away the diffusion and jump risk, particularly in the presence of calibration error. Moreover, only estimated risk-adjusted price dynamics, rather than the real-world price dynamics, are available. As is well known, if the price process is a Merton constant volatility jump diffusion, then the market is incomplete and it is not possible to perfectly hedge a contingent claim using a hedging portfolio with a finite number of instruments. Assuming a synthetic market with given $\mathbb{Q}$ and $\mathbb{P}$ measure parameters, we have investigated two strategies for hedging options under a jump diffusion. Monte Carlo experiments were undertaken to determine the distribution of the $P\&L$ of

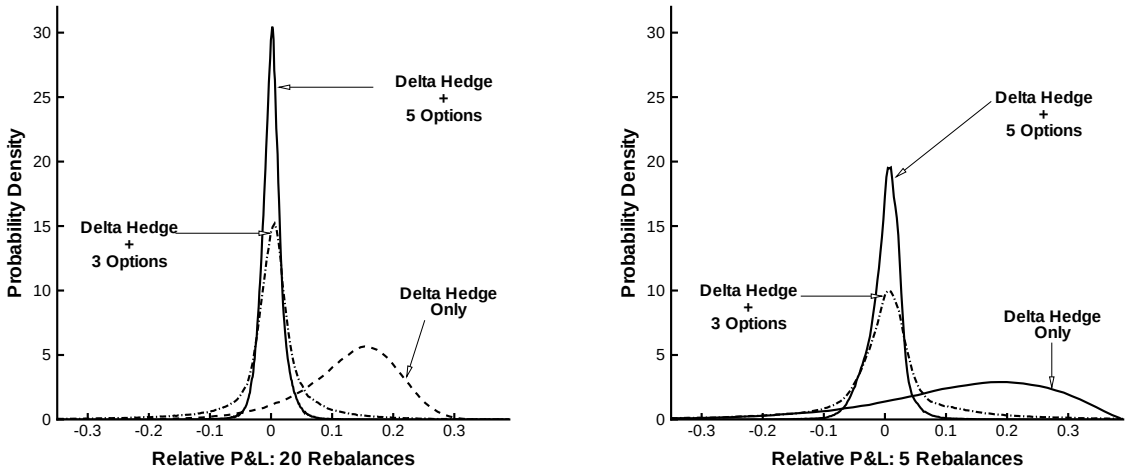

Figure 4.4: Dynamic Hedging: distributions of relative profit and loss for the American example. The hedging horizon is 0.25 years.

these strategies. Note that the experiments included an investigation of estimation errors in the parameters, assuming that a hedger would, of course, not use the exact synthetic market $\mathbb{Q}$ and $\mathbb{P}$ parameters.

The first hedging strategy tested used a dynamic portfolio of options and the underlying asset to hedge the contingent claim. Delta neutrality was imposed, and the jump risk was minimized at each rebalancing time. A portfolio consisting of the underlying and five options is effective in reducing the hedging portfolio standard deviation. This strategy requires frequent rebalancing of the portfolio.

The second hedging strategy is based on a semi-static method. However, in contrast to [7], we assume that the market is incomplete with a limited set of options available as hedging instruments. Hence a quadratic risk minimization formulation is used to construct a portfolio of options and the underlying to optimally replicate the primary option value at infrequent rebalancing times. Again, a portfolio of five options and the underlying asset does an excellent job of reducing the hedging portfolio standard deviation. For the European example considered, this strategy required only an initial purchase of hedging options and one rebalancing during the half-year hedging horizon.

The dynamic strategy has the advantage that estimation of the PDF of the objective price process appears to be less important. However, frequent rebalancing of the hedge portfolio may lead to excessive transaction costs. On the other hand, there is no difficulty in hedging American-style contingent claims.

The semi-static strategy appears to be at a disadvantage compared to the dynamic strategy, since this technique requires knowledge of the objective measure PDF of the price process. However, the numerical tests indicated that this strategy was very successful, even if a very poor estimate of the objective measure PDF was used. It appears that the PDF is dominated by the diffusion volatility, which is comparatively easy to estimate. The jump parameters mainly affect the tails of the distribution. Since the PDF appears as a weighting function in the semi-static objective function, it seems that the precise form of the weighting function is not important, at least when a sufficient number of hedging options are available (i.e., the market is nearly complete).

The results of this study are encouraging:

- Even though estimation of $\mathbb{Q}$ measure parameters is ill-posed in a Merton constant volatility jump diffusion market, the performance of hedging strategies is not largely affected by this poor estimation.

- Both dynamic and semi-static hedging strategies performed very well with hedging portfolios containing five options. The semi-static approach was robust to poor estimation of the $\mathbb{P}$ measure pricing PDF.

Consequently, with the addition of a small number of options to the hedging portfolio, both dynamic and semi-static methods can dramatically reduce jump risk. With this in mind, the choice of which of these two

techniques is more useful in practice should be based on other considerations, such as transaction costs and and effectiveness of hedging path dependent options.

In future work, we will consider the effect of transaction costs on the dynamic hedging strategy. It is, of course, possible to include a penalty term in our objective function so that buying and selling of options at each rebalance date is minimized. For non-path dependent options, the semi-static approach will not incur much in the way of transaction costs, so that this method will be used to benchmark the dynamic hedging technique. In the case of path-dependent options, with discrete monitoring dates, the semi-static approach must be rebalanced at each monitoring date, as described in [7]. Hence, both hedging methods will incur significant transaction costs in this case.

Due to its low cost, the semi-static approach to hedging jump risk looks very promising. However, this method appears to require frequent rebalancing for path-dependent options.

Finally, we remark that hedging exotic options which are sensitive to the tails of the jump size distribution may require calibration with similar exotics, in order to ensure that the tail information is captured in the calibrated pricing parameters.

References

[1] L. Andersen and J. Andreasen. Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing. *Review of Derivative Research*, 4:231–262, 2000.

[2] L. Andersen, J. Andreasen, and D. Eliezer. Static replication of barrier options: Some general results. *Journal of Computational Finance*, 5(4):1–25, 2002.

[3] G. Baker, R. Beneder, and A. Zilber. FX barriers with smile dynamics. Working paper, 2004.

[4] D. S. Bates. Pricing options on jump-diffusion processes. Working paper 37-88, University of Pennsylvania, Rodney L. White Center, 1988.

[5] P. Carr. Semi-static hedging. Technical Report Feburary, Courant Institute, NYU, 2002.

[6] P. Carr, K. Ellis, and V. Gupta. Static hedging of exotic options. *Journal of Finance*, 53(3):1165–91, 1998.

[7] P. Carr and L. Wu. Static hedging of standard options. Technical Report, Courant Institute, NYU, 2004.

[8] T. F. Coleman, Y. Kim, Y. Li, and A. Verma. Dynamic hedging with a deterministic volatility function model. *The Journal of Risk*, 4(1):63–90, 2001.

[9] T. F. Coleman and Y. Li. An interior, trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on Optimization*, 6(2):418–445, 1996.

[10] T. F. Coleman, Y. Li, and A. Verma. Reconstructing the unknown local volatility function. *The Journal of Computational Finance*, 2(3):77–102, 1999.

[11] T. F. Coleman, Y. Li, and A. Verma. A Newton method for American option pricing. *The Journal of Computational Finance*, 5:51–78, 2002.

[12] R. Cont and P. Tankov. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, Boca Raton, 2004.

[13] R. Cont and P. Tankov. Non-parametric calibration of jump-diffusion option pricing models. *Journal of Computational Finance*, 7(2):1–49, 2004.

[14] J. C. Cox and S. A. Ross. The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3:145–166, 1976.

[15] E. Derman, D. Ergener, and I. Kani. Static options replication. *Journal of Derivatives*, 2(4):78–95, 1995.

[16] Y. d'Halluin, P.A. Forsyth, and G. Labahn. A penalty method for American options with jump diffusion processes. *Numerische Mathematik*, 97:321–352, 2004.

[17] Y. d'Halluin, P.A. Forsyth, and K.R. Vetzal. Robust numerical methods for contingent claims under jump diffusion processes. *IMA Journal on Numerical Analysis*, 25:65–92, 2005.

[18] P. Dierckx. *Curve and Surface Fitting with Splines*. Oxford Science Publications, 1993.

[19] B. Eraker, M. Johannes, and N. Polson. The impact of jumps in volatility and returns. *Journal of Finance*, 58:1269–1300, 2003.

[20] G. Labahn. Closed form PDF for Merton's jump diffusion model. Technical report, School of Computer Science, Unviersity of Waterloo, Waterloo, Ont., Canada N2L 3G1, 2003.

[21] R. Lagnado and S. Osher. Reconciling differences. *Risk*, 10:79–83, 1997.

[22] A. Lewis. Fear of jumps. *Wilmott magazine*, December 2002.

[23] R. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3:124–144, 1976.

[24] V. Naik and M. Lee. General equilibrium pricing of options on the market portfolio with discontinuous returns. *The Review of Financial Studies*, 3(4):493–521, 1990.

[25] W. Schoutens, E. Simons, and J. Tistaert. A perfect calibration! Now what? *Wilmott magazine*, March 2004.

[26] A. N. Tikhonov and V. Y. Arsenin. *Solutions of Ill-posed Problems*. W. H. Winston, Washington, D.C., 1977.

[27] V. N. Vapnik. *Estimation of Dependences Based on Empirical Data*. Springer-Verlag, Berlin, 1982.

[28] G. Wahba. *Splines Models for Observational Data*. Series in Applied Mathematics, Vol. 56, SIAM, Philadelphia, 1990.

[29] P. Wilmott, S. Howison, and J. Dewynne. *The Mathematics of Financial Derivatives*. Cambridge, 1995.

[30] R. Wolcott. Two of a kind? *Risk*, 17(3):24–26, 2004.

A Appendix: Market Data

Market data details for Brocade Communications Systems Inc. on April 21 2004, are presented in the subsequent Tables. This data was used as test of the calibration algorithm in §2.3.

Table A.1: The zero rates – yield to maturity of various zero-coupon government bonds (%) used in the calibration of option prices obtained from Nasdaq-NM:BRCD on April 21,2004.

1-month	1.09
3-month	1.17
6-month	1.28
1-year	1.64
2-year	2.43

Table A.2: Market data: American option prices on Nasdaq-NM:BRCD on April 21, 2004. Spot price \$6.19. Call prices reported when strike greater than spot. Put prices reported when strike less than spot.

Strikes(K)	Maturities(T)	OTM Call/Put prices (\$)	Implied Volatility (%)
5	0.083333	0.06	61.76
6	0.083333	0.32	58.80
7	0.083333	0.15	58.9
7.5	0.083333	0.08	60.32
8	0.083333	0.04	62.43
9	0.083333	0.02	68.21
4	0.25	0.05	60.24
5	0.25	0.21	57.47
6	0.25	0.58	56.14
7	0.25	0.40	55.82
7.5	0.25	0.28	56.18
8	0.25	0.20	56.78
9	0.25	0.11	58.58
10	0.25	0.06	60.92
11	0.25	0.04	63.63
12	0.25	0.03	66.61
12.5	0.25	0.02	68.17
15	0.25	0.01	76.42
17.5	0.25	0.01	84.75
3	0.5	0.03	60.29
4	0.5	0.14	57.79
5	0.5	0.40	56.20
6	0.5	0.83	55.40
7	0.5	0.67	54.98
7.5	0.5	0.54	55.10
8	0.5	0.43	55.34
9	0.5	0.28	56.13
10	0.5	0.19	57.20
11	0.5	0.13	58.48
12	0.5	0.09	59.90
12.5	0.5	0.08	60.65
15	0.5	0.04	64.65
17.5	0.5	0.03	68.85

Continued on next page

Strikes(K)	Maturities(T)	OTM Call/Put prices ($)	Implied Volatility (%)
20	0.5	0.02	73.11
25	0.5	0.02	81.42
30	0.5	0.02	88.15
40	0.5	0.01	96.91
2.5	1	0.05	57.41
3	1	0.11	57.47
4	1	0.33	56.44
5	1	0.68	55.55
6	1	1.17	55.12
7	1	1.08	54.52
7.5	1	0.94	54.52
8	1	0.82	54.59
9	1	0.62	54.89
10	1	0.48	55.35
11	1	0.37	55.92
12	1	0.30	56.56
12.5	1	0.27	56.91
15	1	0.17	58.78
17.5	1	0.12	60.77
20	1	0.09	62.81
25	1	0.06	66.85
30	1	0.04	70.74
40	1	0.03	77.98
45	1	0.03	81.23
2.5	2	0.16	56.26
3	2	0.28	56.25
4	2	0.61	55.92
5	2	1.05	55.55
6	2	1.59	55.51
7	2	1.69	54.26
7.5	2	1.55	54.21
8	2	1.42	54.20
9	2	1.20	54.27
10	2	1.02	54.43
11	2	0.88	54.65
12	2	0.76	54.91
12.5	2	0.71	55.06
15	2	0.53	55.87
17.5	2	0.40	56.76
20	2	0.32	57.70
25	2	0.22	59.57
30	2	0.17	61.41
40	2	0.12	64.85
45	2	0.10	66.45