Calibration and Hedging under Jump Diffusion

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

April 28, 2005

Abstract

A jump diffusion model coupled with a local volatility function has been suggested by Andersen and Andreasen (2000). This model is attractive in that it shows promise in terms of being able to capture 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 (i.e. in a synthetic market), we investigate two crucial challenges intrinsic to this type of model: calibration of parameters and hedging of jump risk. Our investigation suggests that it can be difficult to estimate the model parameters that govern the jump size distribution. However, the local volatility function is easier to estimate when an appropriate regularization (e.g. splines) is used to avoid over-fitting. In general, even though the estimation problem is ill-posed, it appears that combining jump diffusion with a local volatility function produces a model which can be calibrated with sufficient accuracy to prices of liquid vanilla options. With regard to hedging jump risk, two different hedging strategies are explored: a semi-static approach which uses a portfolio of the underlying and traded short maturity options to hedge a long maturity option, 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 sharply reduce the standard deviation of the hedging portfolio relative profit and loss distribution.

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

1 Introduction

The inability of the constant volatility Black-Scholes option valuation model to adequately capture features of empirically observed option prices is well known. In particular, the volatility smile for equity options shows a consistent pattern in which implied volatilities are lower for options with higher strike prices, especially for options with short maturities. This stands in obvious contradiction to the basic Black-Scholes model. Various approaches have been suggested to account for the volatility smile. One possibility is to stay in the Black-Scholes setting of single factor diffusion models, but to allow volatility to be state dependent and time dependent (see, e.g. Dupire, 1994; Rubinstein, 1994). However, while such deterministic local volatility functions do an excellent job of matching observed prices of vanilla options, there is an obvious danger of overfitting. Indeed, empirical evidence provided by Dumas et al. (1998) shows that this type of approach performs quite poorly in a predictive sense, and it seems to give unreliable estimates of hedging parameters.

Given the inadequacy of univariate diffusions, two obvious alternatives are stochastic volatility models (e.g. Hull and White, 1987; Heston, 1993) and jump diffusion models (e.g. Merton, 1976; Naik and Lee, 1990; Bates, 1991). Available evidence indicates that stochastic volatility models are a significant improvement, but

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

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

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

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

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

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

they require implausible levels of correlation between volatility and returns and excessively high "volatility of volatility" to match observed prices of index options (Bakshi et al., 1997; Bates, 2000). Given this, much of the more recent research has focussed on augmenting stochastic volatility models with jumps, possibly in both returns and volatility. Closed form solutions for vanilla European option values in this general context have been derived by Duffie et al. (2000). Empirical investigations have been conducted by Eraker et al. (2003) for the case of index options and Bakshi and Cao (2002) for individual stocks. Overall, these studies are supportive of including both jumps in volatility and jumps in returns in valuation models for vanilla equity options.

However, such models are rather difficult to use for option contracts which do not have analytic solutions. Even in the relatively simple context of vanilla American options, it is worth noting that the study of Bakshi and Cao (2002) uses European option calculations on a restricted set of options which are unlikely to be exercised early (i.e. out of the money options with short maturities). A simpler alternative was proposed by Andersen and Andreasen (1999, 2000). The idea was to extend the analysis of Dupire (1994) to the case of jumps, resulting in a model which combines jumps with a deterministic volatility surface. While such an approach has not been subjected to extensive empirical tests (as far as we are aware), Andersen and Andreasen (2000) report preliminary evidence indicating that this approach shows some promise in terms of capturing observed behavior of implied volatilities. It also has the virtue of being much easier to use for contracts where no analytic solution is available. Finally, it is worth noting that the deterministic local volatility function approach and its extension to jumps by Andersen and Andreasen (2000) has a significant advantage in terms of calibration to European option prices. Rather than having to solve a collection of backward equations, one for each option of different strike and maturity, a single one-dimensional forward equation can be solved for all strikes and maturities, for a given initial value of the underlying asset.

There are two crucial challenges when the underlying asset follows a jump process. First, the calibration problem is ill-posed, even for simple jump diffusion models (Cont and Tankov, 2004b), and resulting calibration errors may have serious effects on hedging performance and valuation of exotic contracts. Second, it is not possible to perfectly hedge a contingent claim with a finite number of hedging instruments when the underlying asset follows a jump diffusion with a continuum of possible jump sizes, even in infinitesimal time (see, e.g. Naik and Lee, 1990).

The two major thrusts of this paper address these challenges. First, we focus in detail on the calibration problem. In particular, we investigate the specific characteristics of the ill-posedness of the calibration problem for a jump diffusion model with a local volatility function.[1] Our analysis suggests that it is difficult to estimate model parameters for the jump size distribution, but nevertheless a local volatility function can be estimated with a reasonable level of accuracy provided that a suitable regularization technique is used to avoid overfitting.

The second thrust of this paper explores issues related to hedging. There have been comparatively few studies of hedging strategies in the case of jump diffusions. Bates (1988) 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 Andersen and Andreasen (2000). However, 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, relatively wide bid-ask spreads for option prices may result in high trading costs. An alternative approach is semi-static hedging. In Carr and Wu (2004), the starting point is the exact relationship between value of the primary option (i.e. the option being hedged) in terms of its payoff and the risk-adjusted density function. By an algebraic manipulation, this relationship is transformed into a weighted integration over a continuum (at all strikes) of short term options used for hedging the primary option. The weights of the spanning options are given by the gamma of the primary option at the expiry date of the short term options. Carr and Wu approximate the integral using a quadrature rule with 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; thus the term "semi-static" hedging. Of course, a model is necessary in order to evaluate the gamma of the primary option at some

[1] Note that this differs from the approach used by Cont and Tankov (2004b), who calibrate with a non-parametric Lévy measure rather than a local volatility function.

future time. In practice, since the hedger lacks precise knowledge of market dynamics, a fitted model is used to compute the portfolio weights. This is analogous to determining model parameters by calibration. Carr and Wu base their spanning relation on the assumption that short maturity options of a continuum of strikes are available and the market is effectively complete. An implication of this is that the objective $\mathbb{P}$ measure density is not needed. In the following, we will investigate both dynamic and semi-static hedging strategies. The dynamic strategy is essentially that proposed in Bates (1988), though approached from a different point of view. The semi-static hedging strategy is similar in philosophy to the method described in Carr and Wu (2004). However, 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.

Our overall strategy for exploring calibration and hedging issues in the presence of jumps is as follows. We assume the existence of a synthetic market, with known jump diffusion parameters under both the objective $\mathbb{P}$ measure and the risk-adjusted $\mathbb{Q}$ measure. Calibration effects are then examined by assuming that a practitioner does not know the synthetic market $\mathbb{Q}$ measure parameters, but tries to estimate these from a finite number of observed American option prices. The use of American rather than European options is a distinguishing feature of our work as compared to other papers in the literature. Most traded options are in fact American-style. Calibration to American option prices is a much more intensive problem from a computational perspective because in general a single forward equation cannot be used to handle all strikes and maturities.[2] Moreover, early exercise has important ramifications for hedging, whether the hedger has a long or short position. Given the estimated $\mathbb{Q}$ measure parameters (which clearly may contain errors), 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 hedging strategies precisely, all of the simulations are carried out in a synthetic model where we know the true parameters. We focus solely on the effectiveness of the hedging strategies, especially in the presence of calibration errors. We ignore transaction costs, leaving this topic for future work.

The outline of the paper is as follows. In Section 2, we focus on the jump model calibration problem. The mathematical formulations for dynamic hedging and semi-static hedging under jump risk are presented in Section 3. Computational results for hedge effectiveness evaluation under calibration errors are provided in Section 4, and Section 5 contains concluding remarks.

2 Calibration

This section investigates the characteristics of the ill-posedness of the calibration problem when a jump diffusion model is fit to a finite number of vanilla American option prices. In particular, we analyze in detail two major components which contribute to the problem being ill-posed: insufficient information regarding the tails of the price distribution and finite market option price observations in determining a volatility function. In Section 2.1, we first formulate the calibration problem for a general jump diffusion model with a local volatility function $\sigma(S,t)$. We investigate the characteristics of the jump model calibration problem with constant volatility and analyze its impact on pricing and hedging in Section 2.2. We discuss estimation of both the jump parameters and a local volatility function in a more general jump diffusion model in Section 2.3.

[2]Note that Carr and Hirsa (2003) derive a forward equation for American options in the special case where the logarithm of the underlying asset price follows a Lévy process but the volatility function and drift terms are independent of time and the level of the underlying. Of course, this rules out a local volatility function.

2.1 A Jump Model Estimation Problem

In 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 $\gamma^{\mathbb{Q}}$. We will refer to the risk-adjusted process (2.1), with a constant local volatility σ and a lognormal jump density 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.

Following standard arguments (e.g. Merton, 1976; Andersen and Andreasen, 2000), the value of a European option under the process (2.1) is given by

$$V_\tau = \frac{\sigma^2 S^2}{2}V_{SS} + \big(r - q - \kappa^{\mathbb{Q}}\lambda^{\mathbb{Q}}\big)SV_S - rV + \lambda^{\mathbb{Q}}\left(\int_0^\infty V(S\eta,\tau)g^{\mathbb{Q}}(\eta)\,d\eta - V(S,\tau)\right), \tag{2.2}$$

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

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

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

$$\min(\mathcal{L}V, V - V_e) = 0. \tag{2.4}$$

Assume that the current market prices of vanilla 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 $t = 0$ vanilla option price with strike K_j and maturity T_j. To accurately value over-the-counter options, one typically first determines a risk-adjusted model from the currently available 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 this 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

$$\min_{y\in\mathbb{R}^3, \sigma'(S,t)\in\mathcal{H}} \quad \left(\sum_{j=1}^m \big(V(S_0,0;K_j,T_j,y,\sigma') - \bar{V}_j\big)^2\right). \tag{2.5}$$

Note that we assume that the dividend yield q is known, though this quantity could also be estimated in the calibration.

When calibrating the jump diffusion model (2.4), 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 $\{\bar{V}_j\}, j = 1, 2, \ldots, m$. This is clearly an ill-posed inverse problem. 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.

Starting Points				Final Estimates				Calibration Error
$\lambda_0^{\mathbb{Q}'}$	$\mu_0^{\mathbb{Q}'}$	$\gamma_0^{\mathbb{Q}'}$	σ_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

TABLE 2.1: *Parameter estimation for Merton's jump diffusion model using 30 vanilla American puts and calls in a synthetic market. Bounds setting:* $\lambda^{\mathbb{Q}'} \in [0,1]$, $\mu^{\mathbb{Q}'} \in [-2,2]$, $\gamma^{\mathbb{Q}'} \in [0,1]$, $\sigma' \in [0,1]$. *True values:* $\lambda^{\mathbb{Q}} = 0.1$, $\mu^{\mathbb{Q}} = -0.92$, $\gamma^{\mathbb{Q}} = 0.425$, $\sigma = 0.2$.

2.2 Calibrating Merton's Jump Diffusion Model

In Merton's jump diffusion model, the local volatility is simply a constant. We now consider calibration of Merton's model from standard American option prices, and analyze calibration difficulties and their potential effects on option pricing and hedging. In this case, the calibration problem can be simplified to

$$\min_{y\in\mathbb{R}^3,\sigma'\in\mathbb{R}} \left(\sum_{j=1}^{m} \left(V(S_0;0;K_j,T_j,y,\sigma') - \bar{V}_j \right)^2 \right) \quad \text{subject to} \quad l_\sigma \le \sigma' \le u_\sigma,\ l_y \le y \le u_y. \tag{2.6}$$

In this formulation, *a priori* information for the model parameters in the form of simple bounds can be included. Problem (2.6) is ill-posed because the prices of liquid vanilla options are relatively insensitive to the tails of the price distribution. We illustrate this with an example. Suppose that the current stock price $S_0 = 100$, the interest rate $r = 0.05$, and the dividend yield $q = 0.02$. Assume that the risk-adjusted price dynamics are governed by the constant volatility jump model (2.1), with lognormally distributed jump sizes. The synthetic market parameters are assumed to be

$$\lambda^{\mathbb{Q}} = 0.1,\ \mu^{\mathbb{Q}} = -0.92,\ \gamma^{\mathbb{Q}} = 0.425,\ \text{and } \sigma = 0.2. \tag{2.7}$$

Andersen and Andreasen (2000) calibrate Merton's model to a set of implied volatilities from S&P 500 index options, and determine the best-fit pricing parameters to be $[\lambda_{AA}^{\mathbb{Q}} = 0.089, \mu_{AA}^{\mathbb{Q}} = -0.8898, \gamma_{AA}^{\mathbb{Q}} = 0.4505, \sigma_{AA} = 0.1765]$. Clearly, our choice of synthetic market parameters in (2.7) are close to these values. A set of American option prices $\{\bar{V}_j\}_{j=1}^{30}$ is computed from the assumed jump model using techniques described in d'Halluin, Forsyth, and Labahn (2004) and d'Halluin, Forsyth, and Vetzal (2005). 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 (2.6) is a nonlinear minimization problem with simple bounds, and we compute the parameters using a trust region method (Coleman and Li, 1996).

Table 2.1 displays the estimated model parameters and their corresponding calibration errors using three alternative starting points. The table shows that there are significantly different parameter estimates 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 in some cases 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 is ill-posed.

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 Cont and Tankov (2004b). 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

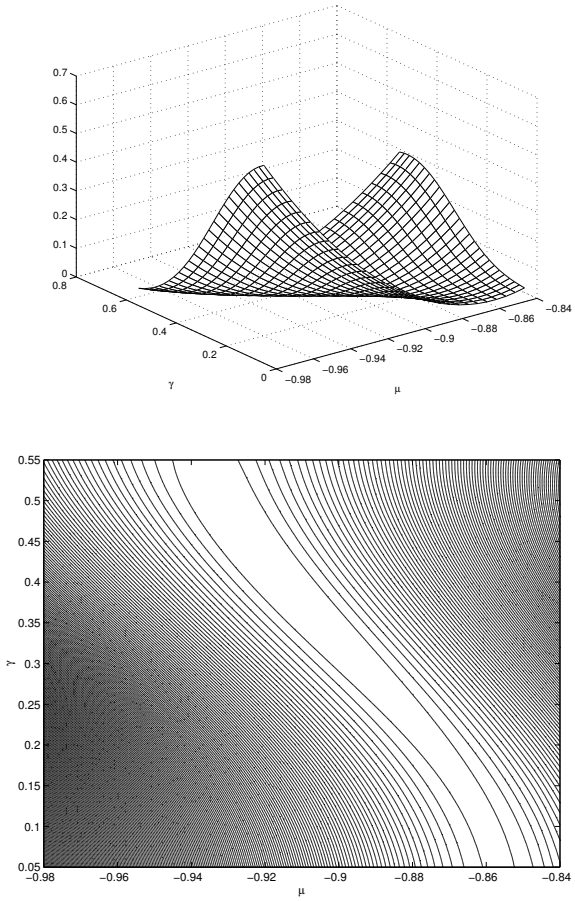

FIGURE 2.1: *Calibration error and associated contours, with* $\lambda^{\mathbb{Q}} = 0.1$ *and* $\sigma = 0.20$ *for Merton's constant volatility jump diffusion model. In this case, for fixed* $\lambda^{\mathbb{Q}}$ *and* σ*, the objective function has a very flat region, suggesting that the calibration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

intensity $\lambda^{\mathbb{Q}'}$ and volatility σ' are uniquely defined as the objective function has a clearly defined minimum. This indicates that the calibration problem (2.6) is well-posed if the parameters specifying the distribution of the jump amplitude are given. Note that it may be possible in some cases to estimate the standard deviation parameter $\gamma^{\mathbb{Q}}$ based on historical price data. This is because in an economic equilibrium model of the type described in, e.g. Bates (1988, 1991), or Naik and Lee (1990), $\gamma^{\mathbb{Q}} = \gamma^{\mathbb{P}}$ (where $\mathbb{P}$ is the objective probability measure). Of course, this would likely require a fairly lengthy time series of historical observations: we would need to have enough jumps in the price data to identify the features of the jump size distribution. However, even if such data were available, it would not be of much use in estimating the mean of $\log(J)$ under the risk-adjusted measure 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. Baker et al., 2004), path-dependent option values such as barrier and forward starting options also depend on risk-adjusted conditional transition densities from stock prices at one future time to those at another future time. 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 (Schoutens et al., 2004). Independent of a model, vanilla option

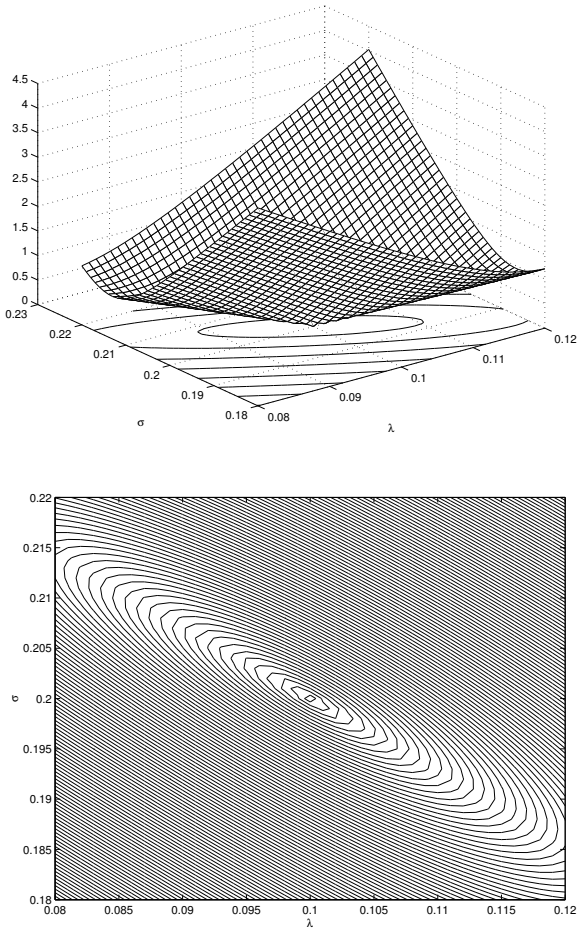

FIGURE 2.2: *Calibration error and associated contours, with* $\mu^{\mathbb{Q}} = -0.92$ *and* $\gamma^{\mathbb{Q}} = 0.425$*, for Merton's constant volatility jump diffusion model. In this case, for fixed* $\mu^{\mathbb{Q}}$ *and* $\gamma^{\mathbb{Q}}$*, the objective function has a well defined minimum.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

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 model, all of the $p(S_{t_1}; S_{t_2})$ conditional transition densities are of the same form as $p(S_0; S_T)$, and so analysis of $p(S_0; S_T)$ offers insight into the ill-posedness of the calibration. We compare the transition 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 model can be computed analytically (see, e.g. Labahn, 2003). 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 estimated 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 mainly affect the tail of the density function. This has minimal effect in pricing vanilla options, whose values are relatively insensitive to the tail of the price distribution. Moreover, we note the following:

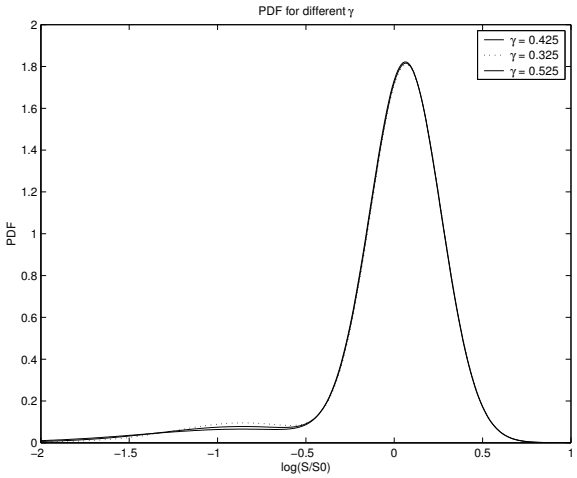

FIGURE 2.3: *Transition densities for a 1 year time horizon, with assorted values of γ and other model parameters fixed ($\mu^{\mathbb{Q}} = -0.92$, $\lambda^{\mathbb{Q}} = 0.1$, $\sigma = 0.20$).*

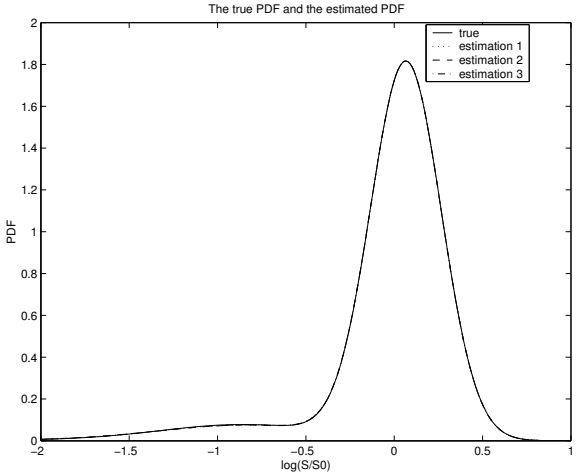

FIGURE 2.4: *Transition densities for a 1 year time horizon, with the synthetic market parameters (2.7) and the estimated parameters from 30 American options. The different estimates are obtained using various starting points in the calibration algorithm (see Table 2.1).*

Starting Points				Final Estimates				Calibration Error
$\lambda_0^{\mathbb{Q}'}$	$\mu_0^{\mathbb{Q}'}$	$\gamma_0^{\mathbb{Q}'}$	σ_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	7.6520e-06

TABLE 2.2: *Parameter estimation for Merton's jump diffusion model using 30 vanilla American puts and calls and three forward start options in a synthetic market. Bounds setting:* $\lambda^{\mathbb{Q}'} \in [0,1]$, $\mu^{\mathbb{Q}'} \in [-2,2]$, $\gamma^{\mathbb{Q}'} \in [0,1]$, $\sigma' \in [0,1]$. *True values:* $\lambda^{\mathbb{Q}} = 0.1$, $\mu^{\mathbb{Q}} = -0.92$, $\gamma^{\mathbb{Q}} = 0.425$, $\sigma = 0.2$.

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, implying 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 a highly accurate estimation of the minimizer.

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 which have values that 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 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 possibility is to add more market price information, e.g. prices of exotic options. To investigate this approach, three one-year maturity forward start put options were added to the given set of American option prices. These forward start options begin at 1 month, 6 months, and 9 months from the current time. 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 a few 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 option values and three forward start prices, are similar to Figure 2.4.

The inclusion of additional forward start options somewhat mitigates the ill-posedness of the jump model calibration problem because their values are more sensitive to the tail of the underlying price distribution. If accurate, deep out-of-the-money option prices are available, e.g. American or digital puts with strikes in the range $[30, 60]$ for the problem considered, preliminary tests indicate that the calibration problem becomes even better posed and accurate jump parameter estimates can be more easily obtained. This suggests that to ensure accurate jump parameter estimation, option instruments which have values that are quite sensitive to the corresponding tail of the underlying price distribution are necessary. An equity default swap (see, e.g. Wolcott, 2004) is an example of a type of option which contains path-dependent information far from at-the-money. This is essentially a one touch digital put with strike typically set at 30% of the underlying asset price at contract inception. If a liquid market develops in these instruments, then we can expect the calibration problem to be better posed.

An alternative approach to overcome the ill-posedness of the calibration problem 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 are sufficiently close to market prices. Cont and Tankov (2004b) gauge closeness to the prior using the relative entropy criterion, assuming that the return of the underlying price follows a Lévy process. Challenges with this approach include determination of an appropriate prior and an appropriate 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, in our experience it seems that the most difficult parameter to estimate in Merton's model with lognormal jumps is the standard deviation $\gamma^{\mathbb{Q}}$ of $\log(J)$. However, it is feasible to impose some reasonable range restriction on this parameter based on historical information because, as noted previously, in certain economic equilibrium models $\gamma^{\mathbb{P}} = \gamma^{\mathbb{Q}}$. This can potentially ease some of the difficulties in the jump model calibration problem.

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 calibration 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 (2.6) to account for the fact that the magnitude of the noise typically depends on the moneyness of the option. In particular, we consider instead the calibration problem

$$\min_{y\in\mathbb{R}^3,\sigma'\in\mathbb{R}} \quad \left(\sum_{j=1}^{m} w_j\big(V(S_0;0;K_j,T_j,y,\sigma') - \bar{V}_j\big)^2\right) \quad \text{subject to} \quad l_\sigma \le \sigma' \le u_\sigma,\ l_y \le y \le u_y, \tag{2.8}$$

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 with 30 vanilla American option prices calculated in a synthetic market. Let V^* denote the exact synthetic market option prices under the assumed model with parameters given in (2.7). We assume that the market prices $\bar{V}$ are noisy; specifically $\bar{V}$ is set as

$$\bar{V}_j = V_j^* + \left(0.01 + \left(\frac{K}{S_0} - 1\right)^2\right) \cdot \phi \cdot V_j^*, \tag{2.9}$$

where ϕ is drawn from a standard normal distribution. Hence, at-the-money options have relatively smaller noise. The maximum percentage price errors in the test samples are given in Table 2.3. In the calibration problem (2.8), we set the weights as

$$w_j = \exp\left(-20\left|\frac{K_j}{S_0} - 1\right|\right).$$

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 σ consistently has relative high accuracy, even when calibrating from noisy prices. However, different sets of noisy prices lead to significantly different estimates for the jump parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$. This is particularly true 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 transition 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 Section 2.2, we assumed that volatility σ was constant and focused on analyzing the calibration problem with respect to the jump parameters $(\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \gamma^{\mathbb{Q}})$ and σ. Unfortunately, such a simple model is generally incapable of accurately calibrating to market option prices of different maturities and strikes. In this section, we consider the calibration problem of a more complex jump model (2.1) in which the local volatility is a function of the underlying price S and time t, as has been suggested in Andersen and Andreasen (1999,

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 the four tests.*

Test	$\lambda^{\mathbb{Q}'}$	$\mu^{\mathbb{Q}'}$	$\gamma^{\mathbb{Q}'}$	σ'	$\frac{1}{2}\|V - V^*\|_2^2$	$\frac{1}{2}\|V - \bar{V}\|_2^2$	$\|V - V^*\|_\infty$	$\|V - \bar{V}\|_\infty$
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

TABLE 2.4: *Parameter estimation for Merton's jump diffusion model from noisy prices. Bounds are set as in Table 2.1, except that here* $\gamma^{\mathbb{Q}'} \in [0.3, 0.55]$. *True values:* $\lambda^{\mathbb{Q}} = 0.1$, $\mu^{\mathbb{Q}} = -0.92$, $\gamma^{\mathbb{Q}} = 0.425$, $\sigma = 0.2$.

2000). Similar to the calibration of a local volatility function in a generalized Black-Scholes model (Coleman et al., 1999), we represent the 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.

It has been illustrated in Coleman et al. (1999, 2001) 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 (2.5) is ill-posed even when the jump parameters $(\lambda^{\mathbb{Q}'}, \mu^{\mathbb{Q}'}, \gamma^{\mathbb{Q}'})$ are fixed, and the effects of this can be significant since the option prices are more sensitive to volatilities. Thus it is important to regularize the estimation problem with respect to the local volatility function estimation and to compute a parsimonious model for accurate pricing and hedging.

Smoothness has long been used as a regularization condition for a function approximation problem with limited observation data (Tikhonov and Arsenin, 1977; Vapnik, 1982; Wahba, 1990). In addition, smoothness of the local volatility function can be important in computational option valuation schemes. Lagnado and Osher (1997) propose the use of smoothness as a regularization condition to approximate the local volatility function in a generalized Black-Scholes model when market vanilla European option prices are given. For the regularized optimization problem proposed by Lagnado and Osher, 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 Coleman et al. (1999) for a more detailed discussion.

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 $\varsigma(\bar{S}_j, \bar{t}_j) = \bar{\sigma}_j, j = 1, \ldots, \ell$, 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, $\bar{\sigma} = (\bar{\sigma}_1, \ldots, \bar{\sigma}_\ell)$, 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

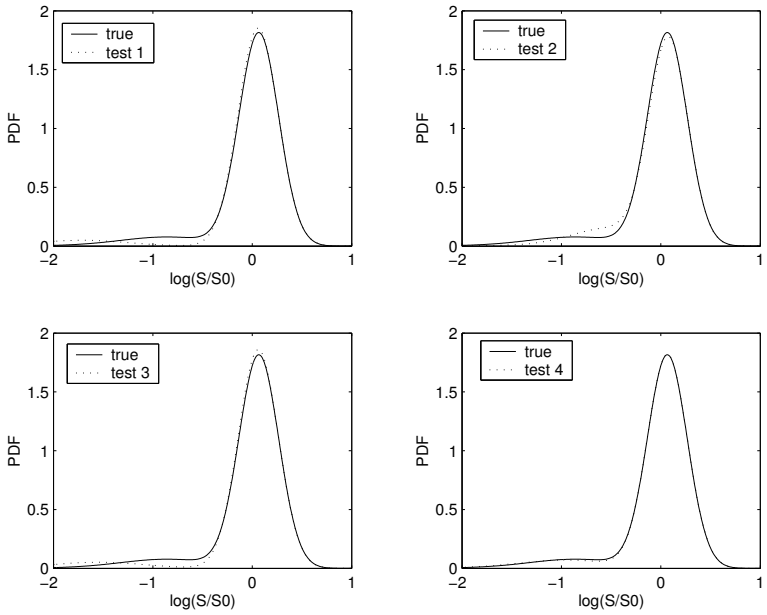

FIGURE 2.5: *Comparison of the synthetic (true) probability density function and the estimated probability density functions using the calibration results in Table 2.4. Compared to the true density, the most pronounced effect of the noisy data can be seen in the left tail of the estimated density.*

(y, ς), where $\varsigma = \varsigma(S, t; \bar{\sigma})$, i.e.,

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

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 $(\bar{S}_1, \bar{t}_1), \ldots, (\bar{S}_\ell, \bar{t}_\ell)$, solve for the ℓ-vector $\bar{\sigma}$

$$\min_{y \in \mathbb{R}^3, \bar{\sigma} \in \mathbb{R}^\ell} \quad \left(f(y, \bar{\sigma}) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{j=1}^{m} w_j \left(V_j^0(y, \bar{\sigma}) - \bar{V}_j \right)^2 \right) \quad \text{subject to} \quad l_\sigma \le \bar{\sigma} \le u_\sigma, \; l_y \le y \le u_y, \tag{2.10}$$

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, but the problem becomes even more computationally difficult.

Problem (2.10) 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 locations, $\{(\bar{S}_j, \bar{t}_j)\}_{j=1}^\ell$. The choice of the number of knots and their placement in spline approximation is generally a complicated issue (Wahba, 1990; Dierckx, 1993); we simply choose the minimum number of knots placed around $(S_0, 0)$ which leads to a sufficiently accurate calibration of the market prices.

Compared to the calibration problem for a generalized Black-Scholes model with a deterministic local volatility function in Coleman et al. (1999), the calibration problem under the jump diffusion model (2.1) is substantially more expensive in terms of computation. 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 (2.4), while in the European case, $V(S, t; K, T)$ (notational dependence on y, σ

is omitted here) satisfies the forward PIDE (Andersen and Andreasen, 2000)

$$-V_T + \big(q - r + \kappa^{\mathbb{Q}}\lambda^{\mathbb{Q}}\big)KV_K + \frac{1}{2}\sigma^2K^2V_{KK} + \big(1+\kappa^{\mathbb{Q}}\big)\lambda^{\mathbb{Q}}\left(\int_0^\infty V\big(S,t;K/z,T\big)g'\big(z\big)\,dz - V\big(S,t\big)\right) = qV, \quad (2.11)$$

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

The model estimation problem (2.10) is typically solved by an iterative method which requires at least function and gradient (Jacobian of $[V_1,\ldots,V_j]$) evaluations at each iteration. For European options, the forward equation (2.11) 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 (Coleman et al., 2002; d'Halluin et al., 2004).[3] 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 $\ell \times m$ option valuations, where ℓ is the total number of spline knots and m is the total number of option prices given.

The calibration problem (2.6) for Merton's jump model is far less costly since there are only four variables. 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 (2.6) and using it as a starting point for the more general model (2.10), 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 (2.8) 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 Andersen and Andreasen (2000). Recall that this first approximation also makes sense from the viewpoint of computational efficiency.

Stage 2: Estimating a local volatility function. In the second stage, the spline calibration problem (2.10) is solved with the estimation from the first stage calibration problem used as the starting point.

To illustrate, we assume that the diffusive volatility of the underlying price has the form of a constant elasticity of variance model as described in Cox and Ross (1976). 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 $\sigma(S) = 25/S$ for $S > 0$. 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 Section 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. Similar to Merton's jump model calibration in Section 2.2, different starting points produce somewhat different $(\lambda^{\mathbb{Q}},\mu^{\mathbb{Q}},\gamma^{\mathbb{Q}})$ estimates at the end of the second stage.

Figure 2.6 compares the estimated local volatility function with the true $\sigma(S)$. In spite of the difficulty in determining an accurate estimation of the jump parameters, a fairly accurate estimation of the local volatility function is obtained from a limited set of American option prices with appropriate smooth regularization using

[3] In certain special cases efficiency gains are possible. For example, if the volatility function is independent of S and t, then, as noted above, it is possible to solve a single forward equation for American options (Carr and Hirsa, 2003).

	Starting Points				Final Estimates				Calibration Error
Stage	$\lambda_0^{\mathbb{Q}'}$	$\mu_0^{\mathbb{Q}'}$	$\gamma_0^{\mathbb{Q}'}$	σ_0'	$\lambda_p^{\mathbb{Q}'}$	$\mu_p^{\mathbb{Q}'}$	$\gamma_p^{\mathbb{Q}'}$	σ_p'	$\frac{1}{2}\|V-\bar{V}\|^2$
1	0.4	0.4	0.4	0.4	0.1621	-0.7713	0.3056	0.2201	
2					0.0955	-0.9761	0.4535		4.6104e-05
1	0.1	0.8	0.1	0.1	0.1620	-0.7717	0.3071	0.2201	
2					0.0983	-0.9528	0.5043		2.8060e-05
1	0.2	-0.8	0.2	0.2	0.1618	-0.7724	0.3050	0.2201	
2					0.0931	-1.0023	0.4488		4.2284e-05

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

splines. In addition, we note that different starting points yield very similar volatility function estimates. While the ill-posedness of the constant volatility jump diffusion model calibration problem with respect to the jump parameters can be improved by including options far from at-the-money, adding option prices with more strikes and maturities will also further improve the accuracy the estimated local volatility function.

The test case above is a relatively simple one in that the given local volatility function $\sigma(S,t)$ is independent of t. As a further test, we calibrate a jump diffusion model with a volatility function to American option market price data for Brocade Communications Systems Inc. on April 21, 2004. The spot price is $S_0 = 6.19$, the dividend yield is zero, and the risk free rate r is time-dependent. This rate is derived from the yields to maturity of various zero-coupon government bonds (the values used 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 \$2.50 to \$20 and maturities range from one month to two years. Table A.2 in the Appendix provides the market option prices. 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 $[0.1S_0; 1.55S_0; 3S_0] \times [0; 1; 2]$. 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 corresponding objective function value is 0.0029 and the maximum calibration error is 0.0357. This largest calibration error occurs when 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. In this case 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 a fairly accurate calibration to vanilla option prices.

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

1. Even calibrating a simple Merton jump model from vanilla 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 consistently is relatively accurate. In addition, accurate model parameters can be eventually estimated after extensive optimization iterations, assuming that there are no errors in the price data and that the modelling assumptions are correct.

2. Relatively accurate estimation of the transition density function (in the case of constant volatility) can be obtained even when the fitted jump parameters 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.

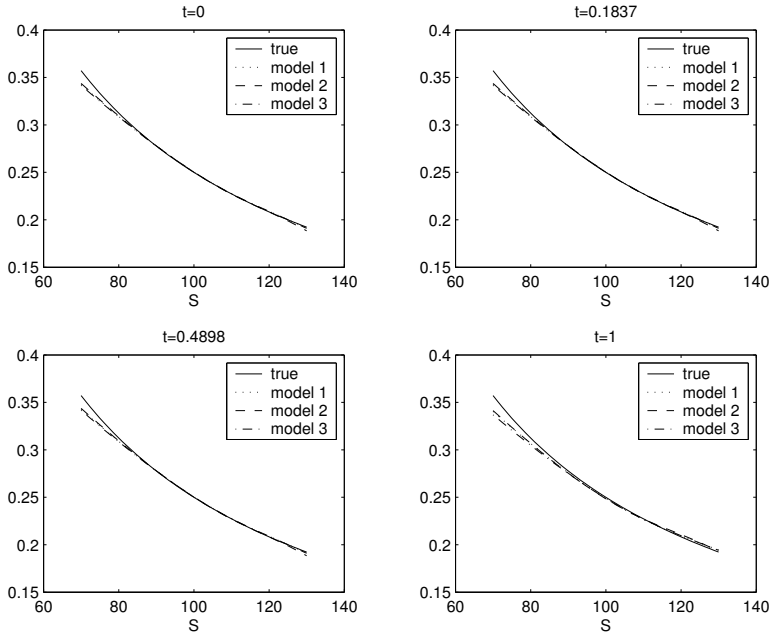

FIGURE 2.6: *Estimation of the local volatility function $\sigma(S,t)$ for four different values of t from 30 American options. Models 1-3 refer to the local volatility function obtained using different starting values in the calibration algorithm, as shown in Table 2.5.*

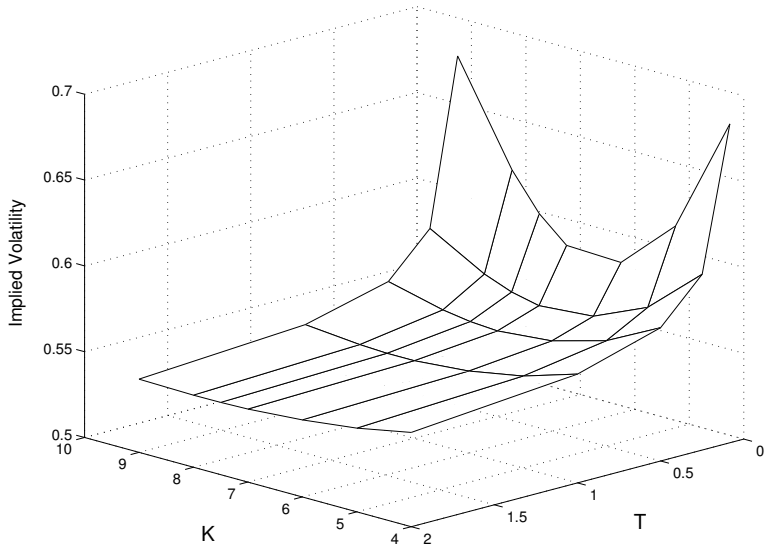

FIGURE 2.7: *Implied volatilities for options on Brocade Communications Systems Inc. on April 21, 2004.*

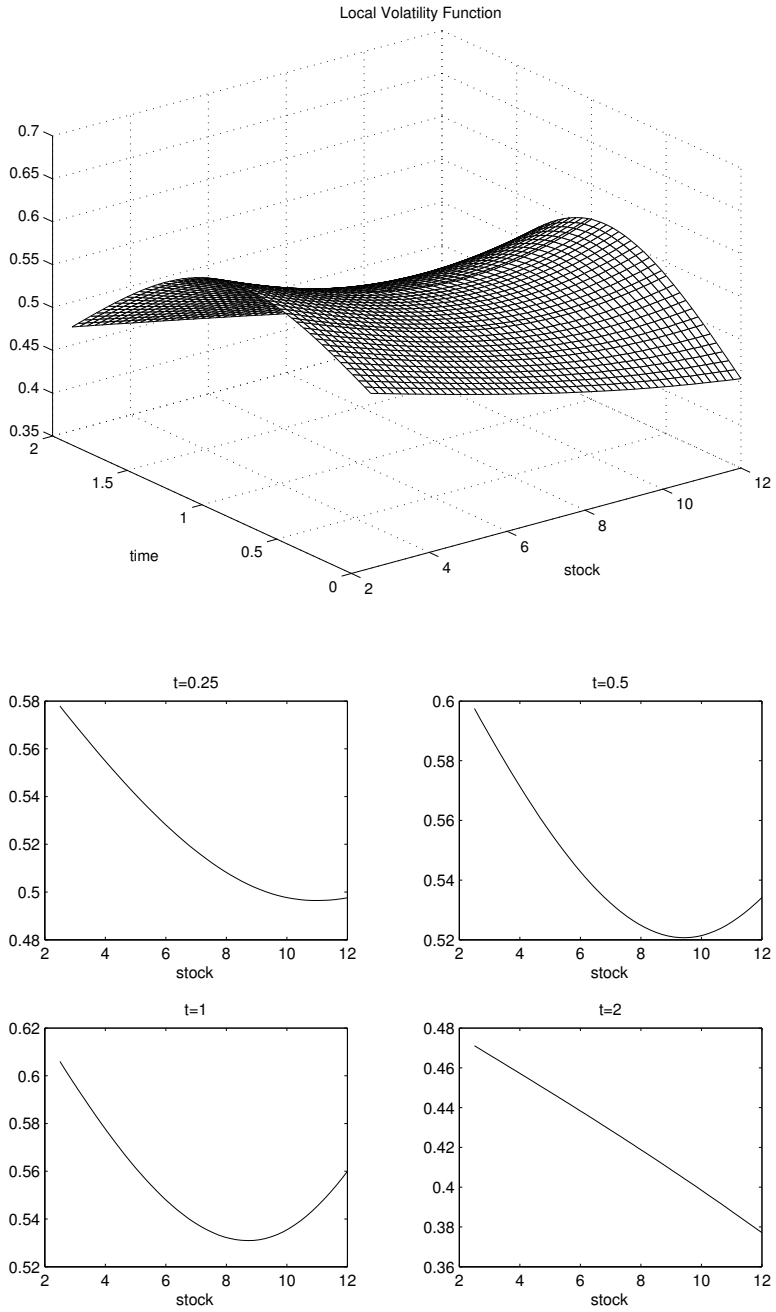

FIGURE 2.8: *The calibrated local volatility function $\sigma(S,t)$ for Brocade Communications Systems Inc. on April 21, 2004.*

3. When the available market prices are noisy, it is even more difficult to calibrate the jump parameters accurately. However, the estimation of volatility achieves relatively high accuracy, and the calibrated model yields relatively accurate pricing for American options with values that are relatively insensitive to the tail of the price distribution.

4. For a more general jump diffusion model (2.1), a relatively accurate estimate of 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

We now investigate the effect of jump model calibration error on hedging and evaluate the performance of two approaches to hedging 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 trading 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 an estimate 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, the dynamic hedging of a contingent claim under a jump diffusion process is a far greater challenge than it otherwise would be in the (complete) Black-Scholes universe (see, e.g. Naik and Lee, 1990). 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 errors.[4] However, the risk inherent in the jump diffusion model (with a continuum of possible jump sizes) can only be removed completely by an infinite number of hedging instruments. With a finite number of instruments, the diffusion risk can still be eliminated by imposing delta neutrality, but the compound Poisson process governing the arrival and magnitude of 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. Moreover, jump risk may be substantial. Consider a bank which is short an option with a convex payoff (e.g. a call). Suppose it constructs a continuously rebalanced delta neutral hedge using the underlying, so that it only has protection against the diffusion risk. If a jump occurs, the convexity of the option essentially guarantees the bank's position will lose money regardless of the size or direction of the movement (Merton, 1976). Even if a jump is deemed highly unlikely, the potentially disastrous consequences provide strong incentives for option writers to at least partially 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.

[4] Of course, even for pure diffusion cases dynamic hedging strategies applied to contracts such as barrier options become problematic if the spot approaches the barrier, since the option gamma becomes infinite. This has led some researchers to suggest static hedging for barrier options (Derman et al., 1995; Carr et al., 1998; Andersen et al., 2002).

3.1.1 Jump Risk

This section derives a mathematical representation of jump risk. The hedge portfolio contains an amount B in cash, is long e units of the underlying asset S, and long N additional hedging instruments $\vec{I} = [I_1, I_2, \ldots, I_N]^T$ (written on the underlying) with weights $\vec{\phi} = [\phi_1, \phi_2, \ldots, \phi_N]^T$. For notational simplicity, we assume that $\vec{I}$ denotes all of 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 a short position in the primary option $-V$, the overall hedged position has value

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

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

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

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

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

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

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

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

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

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

$$d\Pi = -\left[V_t + \frac{\sigma^2 S^2}{2} V_{SS} \right] dt + \vec{\phi} \cdot \left[\vec{I}_t + \frac{\sigma^2 S^2}{2} \vec{I}_{SS} \right] dt + rB\, dt + \left[-\Delta V + \vec{\phi} \cdot \Delta \vec{I} + e \Delta S \right] d\pi^{\mathbb{P}}\ , \tag{3.3}$$

indicating that $d\Pi$ is now a pure jump process with drift. For expository reasons, we assume here that the option V and none of the hedging instruments in $\vec{I}$ are exercised early.[5] Using an elementary rearrangement,

[5] In our implementation, if the short position in V is optimally exercised, then the hedge portfolio is liquidated to cover the short position. Also, if any of the hedging instruments are optimally exercised, then they are replaced by another instrument.

the pricing PIDEs for V and $\vec{I}$ in the continuation region of the form (2.2) may be written as

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

where $\mathbb{E}^{\mathbb{Q}}(\Delta S) = \mathbb{E}^{\mathbb{Q}}\left(S[J-1]\right) = S\mathbb{E}^{\mathbb{Q}}(J-1) = S\kappa^{\mathbb{Q}}$. Substituting (3.4) into (3.3) yields

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

Using the delta neutral constraint (3.2) gives

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

Therefore equation (3.5) indicates that the value of the overall hedged position grows at the risk free 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 component of the jump risk is deterministic, while the second part is stochastic as it depends on whether a jump occurs over the instant dt and what its size is. 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

If a jump happens ($d\pi^{\mathbb{P}} = 1$), the size of the stochastic part of the jump risk in (3.6) 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 overall hedged position 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 via a suitable linear constraint, i.e. $\mathbb{E}^{\mathbb{Q}}\big[\Delta H\big] = 0$.

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 the 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 can be eliminated by introducing M hedging instruments into $\vec{I}$. The linear system

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

ensures that for the current time t and asset price S_t, the diffusion risk is removed and the overall hedged position is invariant to jumps of size $J^* \in \{J_1, J_2, \ldots, J_M\}$. 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 reasonable to expect that 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 that an infinite number of instruments are needed to eliminate the jump risk. Therefore the goal is to reduce the 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 $W(J)$:

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

Note that $W(J)$ has the properties of a probability density, i.e. $W(J) \geq 0$, $\int_0^\infty W(J)\, dJ = 1$. 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 remaining input for the optimization problem (3.8) is the weighting function. One logical choice is setting $W(J) = g^{\mathbb{P}}(J)$, the jump size distribution under the objective measure (Bates, 1988; Andersen and Andreasen, 2000). In this case the optimization problem is similar to a local variance minimization. To see this, note that the instantaneous change in the delta neutral hedged portfolio (3.5) may be written as

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

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

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

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 size, the resulting hedging strategy is that suggested in Andersen and Andreasen (2000).

Unfortunately, the density $g^{\mathbb{P}}(J)$ is usually not explicitly known. Recall that it is the risk-adjusted distribution $g^{\mathbb{Q}}(J)$ that is required for pricing, and is implicitly determined by calibration to market prices. One may hazard a guess for the real world distribution. For example, as noted below (see equation (4.1)), in Naik and Lee (1990)'s general equilibrium version of 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 its unknown real world counterpart $g^{\mathbb{P}}(J)$. On the other hand, a scarcity of information may be encapsulated by 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 here 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

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

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 numerical techniques. The expectations are first transformed into correlation integrals and then the fast Fourier transform is used to compute these integrals in an efficient manner (as in d'Halluin et al., 2005). 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 adopt the standard approach of setting small singular values to zero after the SVD is completed, and then use this modified decomposition to compute the portfolio weights.

This strategy is essentially that proposed in Bates (1988), though approached from a different point of view. Bates 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 match 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 reiterate: 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 wants to form a hedge portfolio at time zero that imposes delta neutrality and reduces the jump risk exposure. 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—imposing continuity results in better numerical behavior.

The left panel of Figure 3.2 presents the jump risk profile $\Delta H(J)$ for the optimized hedge portfolios resulting from various hedging strategies. These curves represent the change in the value Π of the overall hedged position as a function of jump size. There is a large potential loss to a portfolio that ignores jumps, which is evident in the *Delta Hedge Only* curve. 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 intercepts where the curves cross the zero line are fortuitous points as there is no change in the overall hedged position for jumps corresponding to these amplitudes. This is exactly what is desired. As more hedging instruments are included, the curves more closely hug the zero line and the number of intercepts increases. In fact, for the case of 10 hedging instruments, the curve remains so close to zero that its structure is not clearly seen; an enlargement is provided in the right panel 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

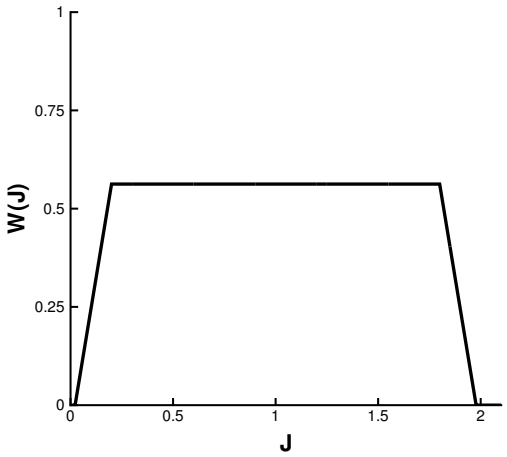

FIGURE 3.1: *Uniform-like weighting function $W(J)$ (as in equation (3.8)) used in some of the simulations. Experiments showed that use of the ramp tails resulted in better numerical behavior.*

given that the straddle costs \$21.74. Both Bates (1988) and Andersen and Andreasen (2000) contain similar graphs. Overall, it is clear from this example that the outlined procedure can be used to reduce the jump risk exposure for an option seller.

3.1.3 Implementing the Jump Risk Minimization Dynamic Hedging Strategy

Merton (1976) considered the characteristics of a dynamic delta neutral hedge that uses the underlying. Since this 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, so as noted above, 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, as noted in Cont and Tankov (2004a), the hedge accounts for the average effect of jumps though it is obviously not perfect due to the random arrival of jumps. However, in general jump risk will not be completely diversifiable, so the jump processes 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 model) is studied by Naik and Lee (1990). It is initially assumed the volatility used for pricing is σ, 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 for pricing.

Delta hedging alone is clearly insufficient, so it is desirable 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 $t = 0$ 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

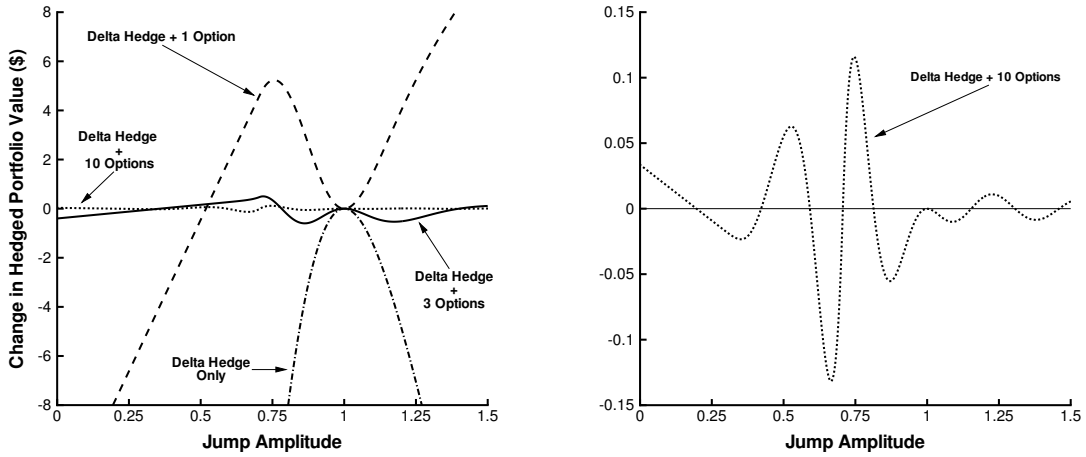

FIGURE 3.2: *Dollar change in portfolio value resulting from a jump. The left panel shows the performance of several hedging strategies, while the right panel is an enlargement of the curve when 10 options are used to hedge. The option being hedged is a one year American straddle.*

makes the initial value of the overall hedged position zero, i.e. $B(0) = V(S_0, 0) - e(0)S_0 - \vec{\phi}(0) \cdot \vec{I}(S_0, 0)$. 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 rebalancing contains

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

An instant after rebalancing the overall hedged position 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 sub-optimal (this will be discussed further below).

We are interested in the value of the overall hedged position when the option expires. Ideally, the portfolio has a value of zero at expiry (or when the primary option is exercised), 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):

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

3.2 Semi-Static Hedging

As an alternative to the dynamic hedging approach described above, a semi-static hedging strategy can be implemented by minimizing directly the difference between the value of the primary option and the value of 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 ideas described in Carr and Wu (2004), which use 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 is not instantaneously delta neutral 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.

Assume that we have an existing primary option maturing at T, which has value at time t denoted by $V(S_t, t)$. Also suppose that the hedge portfolio is rebalanced at a set of discrete times $0 = t_0 < t_1 < \cdots < t_M = T$, where $t_k = k\delta t$. Following Carr and Wu (2004), it can be shown that the value of the primary option may be computed based on the option replication principle

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

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 $S_{t_{M-1}}$. This equation shows 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. By no-arbitrage, 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 through times $t_{M-1}, \ldots, t_0$, 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 vanilla options are restricted to a few strike prices. Consequently, instead of using the integration equation, we compute an optimal replication in a weighted least squares sense. Ideally, the weights are given by the transition probability density function for the underlying price under the objective probability measure $\mathbb{P}$. When this is not available, an approximation (e.g. the transition 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 calculated 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

$$\underbrace{\vec{\phi}(t_k) \cdot \vec{I}(t_k) + e(t_k)S(t_k) + B(t_k)}_{\text{instant after rebalancing}} = \underbrace{\vec{\phi}(t_{k-1}) \cdot \vec{I}(t_k) + e(t_{k-1})S(t_k) + B(t_{k-1}) \exp\{r\delta t\}}_{\text{instant before rebalancing}}. \tag{3.14}$$

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

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

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

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

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

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

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

and $\mathbb{E}_k^{\mathbb{P}}$ denotes the conditional expectation under the probability measure $\mathbb{P}$, or a suitable approximation. Note that $\mathbb{E}_k^{\mathbb{P}}$ effectively 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 a range of S values, then the precise estimation of this weighting function (the $\mathbb{P}$ measure density) should not be 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 transition density function. This is used below 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 Section 3.1.3. Note, however, that the semi-static and dynamic hedging strategies have 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 one year European straddle with strike \$100. As a 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 and put options with strikes in intervals of \$5 are used as hedging instruments. These options are assumed to 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 may be linked to the risk-adjusted measure in (2.7) via utility-based equilibrium relations

$$\begin{aligned}\sigma^{\mathbb{P}} &= \sigma^{\mathbb{Q}} \\ \gamma^{\mathbb{P}} &= \gamma^{\mathbb{Q}} \\ \mu^{\mathbb{P}} &= \mu^{\mathbb{Q}} + (1-\beta)(\gamma^{\mathbb{Q}})^2 \\ \lambda^{\mathbb{P}} &= \lambda^{\mathbb{Q}} \exp\Big\{\big(1-\beta\big)\big(\mu^{\mathbb{Q}} + \frac{1}{2}(1-\beta)(\gamma^{\mathbb{Q}})^2\big)\Big\}\end{aligned} \tag{4.1}$$

where $1-\beta$ is the coefficient of relative risk aversion.[6] With $1-\beta = 2$, the parameter set of Table 4.1 is obtained. For simplicity in the following, the $\mathbb{Q}$ volatility is taken to be a constant, since we are focusing on hedging strategies. 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 $\alpha^{\mathbb{P}} = 0.1779$.

[6] The relations given in (4.1) are for the case of power utility, as in Naik and Lee (1990) and Bates (1991). One possible generalization is to the case of uncertainty aversion, as in Liu et al. (2005), which includes power utility as a special case. As our intention here is to simply obtain an approximate transformation to the $\mathbb{P}$ measure parameters, we use only the simple power utility case.

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 real world $\mathbb{P}$ measure that characterize the synthetic market (Merton's constant volatility jump diffusion model). The coefficient of relative risk aversion is $1-\beta=2$.*

In addition to assessing the efficacy of the hedging strategies, the simulations also explore the consequences of using an estimated pricing measure. Recall from Section 2 that many different pricing measure parameter sets produce close agreement to the values of vanilla options generated by the $\mathbb{Q}$ measure of Table 4.1. We will refer to the $\mathbb{Q}$ measure in Table 4.1 as the true pricing measure. It may be thought of as the pricing rule responsible for all option values observed in the market.

In the experiments to follow, 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 $(\lambda^{\mathbb{P}'}, \mu^{\mathbb{P}'}, \gamma^{\mathbb{P}'}, \sigma') = (0.0310, -0.3825, 0.4906, 0.1991)$. Five simulation sets are 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 $W(J)$ is set to the lognormal density suggested by the objective measure. For semi-static hedging, the conditional expectation in (3.18) is taken with respect to the transition 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 estimate the objective measure $\mathbb{P}'$. In turn, this estimated objective measure is used to establish the weighting function (for dynamic hedging) or to calculate the expectation in (3.18) (for semi-static hedging).

The intention of the final two simulation sets (#4 and #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 both to calculate the option values and to 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

For the moment, we will only use the underlying and call options as instruments in the dynamic hedge, since put-call parity allows any desired position in a put to be replicated by positions in a call (and the underlying). This is reasonable here since all options are European—Section 4.2 below considers American options, and in that context 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

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 ($\mu = 0.5, \gamma = 0.1$)

TABLE 4.2: *Summary of the simulation experiments. For all simulations, the true objective $\mathbb{P}$ measure is used to generate asset price paths. Moreover, all observable prices in the synthetic market are derived using the true pricing $\mathbb{Q}$ measure.*

			Percentiles			
Simulation Set	Mean	Std. Dev.	0.02%	0.2%	99.8%	99.98%
Delta hedging only	0.1172	0.2377	-4.6026	-3.0158	0.1907	0.2201
1	0.0003	0.0311	-0.3403	-0.1166	0.1494	0.3947
2	0.0003	0.0299	-0.3640	-0.1173	0.1447	0.3517
3	0.0004	0.0249	-0.2860	-0.0964	0.1180	0.2915
4	-0.0001	0.0464	-0.5027	-0.1897	0.2321	0.4987
5	0.0195	0.2191	-2.7946	-1.6948	0.4301	3.9394
Uniform weighting	-0.0004	0.0229	-0.2413	-0.0879	0.1061	0.2821

TABLE 4.3: *Dynamic hedging: statistics of the relative P&L distributions. Primary option: European straddle. For the α% percentile, approximately α% of the 500,000 simulations resulted in a relative P&L less than the reported amount.*

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 used.

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 of the 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, 500,000 simulations are carried out (using the same random numbers across the different sets). With a mean arrival rate of $\lambda^{\mathbb{P}} = 0.0228$, there are expected to be about 5,700 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 indicated 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 they near maturity. However, the protection against jump risk should not be adversely affected.

The distributions of relative P&L are presented in Figure 4.1, with statistical measures provided in Table 4.3. The outliers of these distributions are very important, as it is these potentially catastrophic events one tries to prevent by hedging. Various outliers (in the form of the 0.02%, 0.2%, 99.8%, and 99.98% percentiles) are given in the table. Results are presented for two simulation experiments particular to dynamic hedging: (i) a simple delta neutral hedge that only uses the underlying; and (ii) use of the uniform-like weighting function $W(J)$ plotted in Figure 3.1.

The first row of Table 4.3 gives the summary statistics for the case where only delta hedging is used, with the corresponding distribution plot of the relative P&L shown in the bottom left panel of Figure 4.1. It appears that the seller of the option is almost always making money. In fact, the 99% value-at-risk is

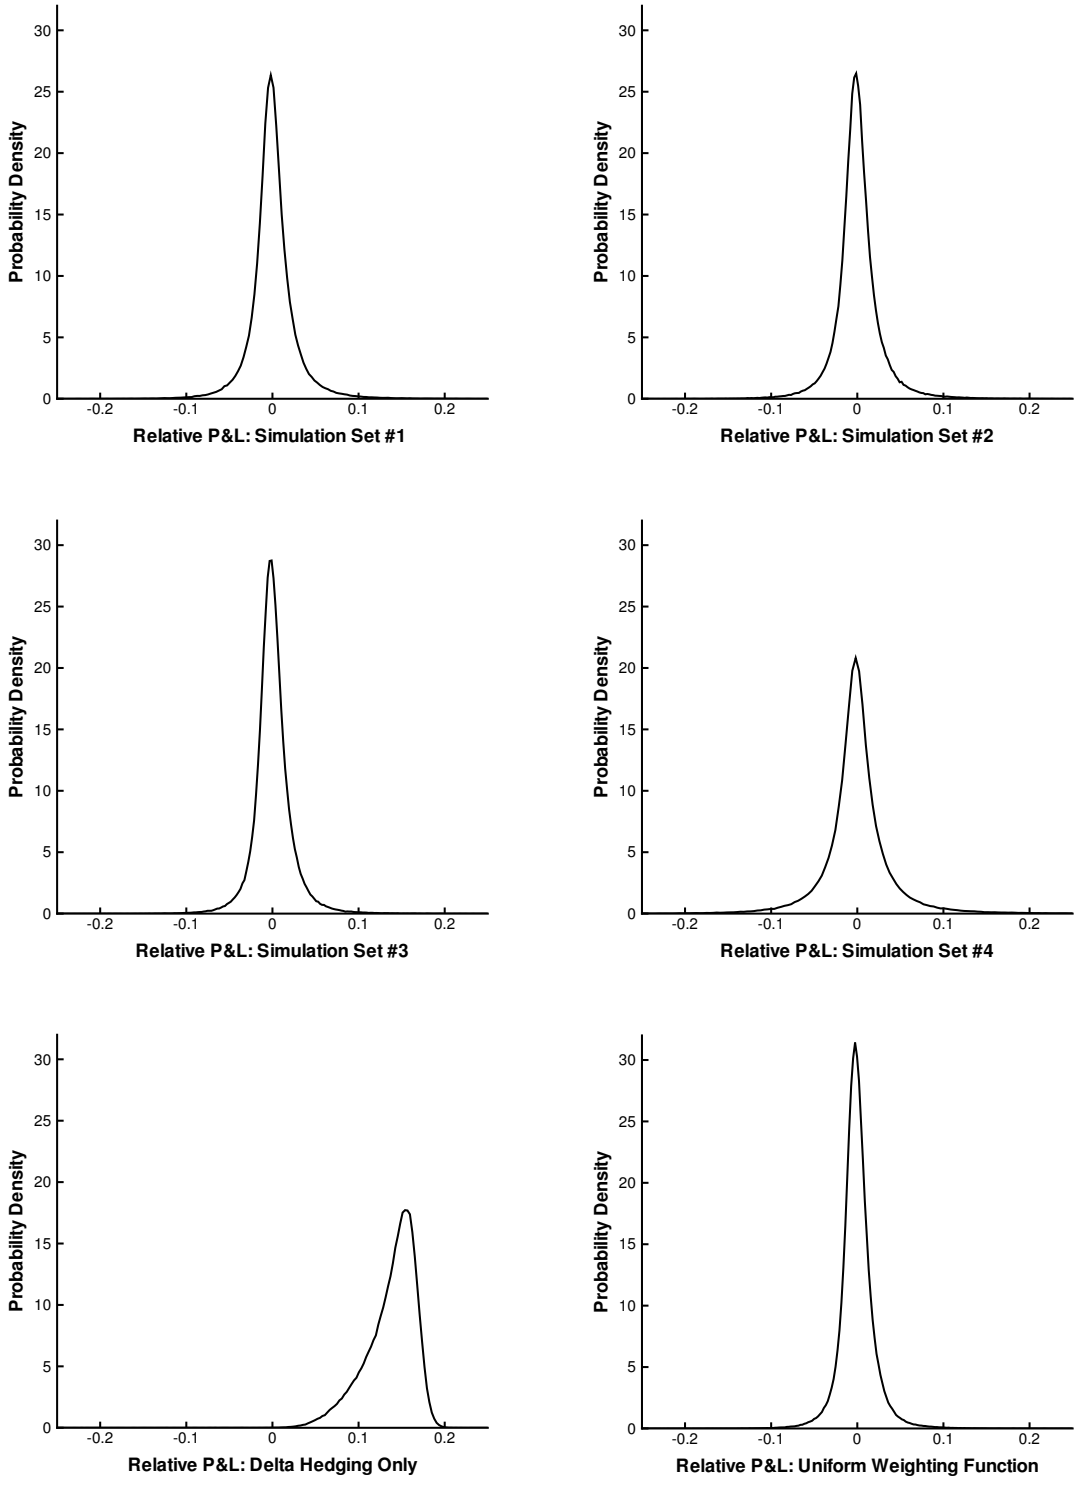

FIGURE 4.1: *Dynamic hedging: relative P&L distributions. Primary option: European straddle.*

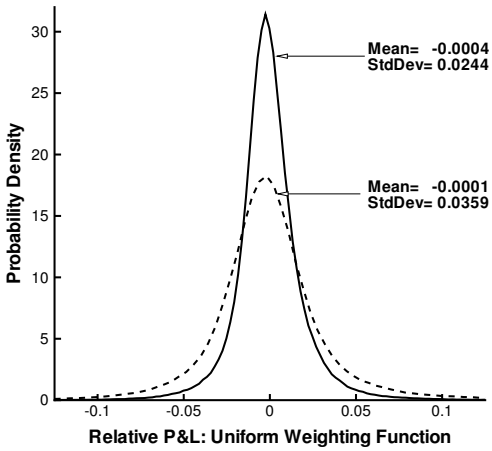

FIGURE 4.2: *Dynamic hedging: relative P&L distributions when uniform-like weighting is used. Primary option: European straddle. The solid curve corresponds to the case of 40 complete rebalances, with 3 adjustments of the underlying between each complete rebalance. The dashed curve represents the case when only the 40 complete rebalances are carried out.*

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 the $\mathbb{Q}$ measure is more pessimistic than the $\mathbb{P}$ measure. 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 overall hedged position, often quite a substantial one.

Simulation sets #1 through #4 demonstrate that the dynamic hedging strategy of jump risk minimization works 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 expected shape of a symmetric, mean-zero density having a limiting form of 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 $\mathbb{Q}'$ produce 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 $\mathbb{P}'$ gives similar results, though there is a slight degradation for the weighting function suggested by $\mathbb{Q}$. The salient conclusion from simulation set #5, which has poor summary statistics, 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 three 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 the relative P&L distribution are presented in Figure 4.2. The solid curve represents the trading regime previously mentioned, i.e. 40 complete rebalances with three adjustments of the underlying between each. The dashed curve corresponds to the case where only the 40 complete rebalances are carried out. Incorporating frequent updating of the underlying to impose delta neutrality should lead to 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.

			Percentiles			
Simulation Set	Mean	Std. Dev.	0.02%	0.2%	99.8%	99.98%
1	0.0008	0.0116	-0.1235	-0.0463	0.0529	0.1889
2	0.0007	0.0116	-0.1238	-0.0463	0.0528	0.1881
3	0.0009	0.0116	-0.1204	-0.0463	0.0526	0.1841
4	-0.0001	0.0121	-0.1169	-0.0543	0.0579	0.1996
5	0.0019	0.0179	-0.4525	-0.0583	0.0898	0.3404

TABLE 4.4: *Semi-static hedging: statistics of the relative P&L distributions. Primary option: European straddle.*

4.1.2 Semi-Static Hedging Experiments

We now provide a study of semi-static hedging, as discussed in Section 3.2. We use the same scenarios as for dynamic hedging above in Section 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 three months until maturity and strikes closest to $[0.8, 0.9, 1.0, 1.1, 1.2]S_t$ are employed. 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. We hold this portfolio until t_1, at which point it is rebalanced 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 three months until expiry and strikes in intervals of \$5 that are 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 500,000 underlying paths used for each simulation set.

Table 4.4 presents the statistics of the relative P&L distributions for each simulation set. The results for sets #1 and #2 are very close, due to the fact the estimated measure $\mathbb{Q}'$ leads to fairly accurate pricing of vanilla options. In addition, the results of simulation set #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 vanilla options. Simulation set #4 assumes that the objective measure $\mathbb{P}$ is not available, and the risk-adjusted measure is instead used in the hedging computation. The results are close to those found using the $\mathbb{P}$ measure. In simulation set #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 a very poor choice for the transition probability density function. This is an encouraging result, suggesting that this hedging strategy is quite robust to misspecification of the objective transition density. This observation is consistent with the results obtained in Carr and Wu (2004).

4.2 Dynamic Hedging of an American Option

This section applies the jump risk minimization procedure to the problem of dynamically hedging an American option. The primary contract to be hedged is an American put with a half-year maturity. The hedging horizon is three months, and out-of-the-money American puts and calls are used as the hedging instruments (these all have the same maturity as the hedging horizon). 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 put in the hedging portfolio may be optimally exercised before the end of the hedging horizon. 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 parameters as in the European case are used, except that the mean arrival rate of jumps in the market is set about four times higher, to $\lambda^{\mathbb{P}} = 0.1$. This provides a higher stress on the hedging system, as over

Number of Rebalances	Hedging Options	Mean	Std. Dev.	Percentiles 0.02%	0.2%	99.8%	99.98%
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	-0.0822	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 4.5: *Dynamic hedging: summary statistics of the relative P&L distributions. Primary option: American straddle. Hedging horizon: three months.*

the course of 500,000 simulations about 12,500 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 4.5 provides statistics for the distributions of relative P&L. Figure 4.3 plots the distributions for the cases 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. The use of other instruments within our dynamic hedge mitigates this jump risk, as is evident with the .02% and 0.2% percentiles. 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 (i.e. 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 4.5 that using the hedging strategy based on minimization of the infinitesmal variance (the dynamic strategy) does a remarkable job of reducing the standard deviation even if the rebalancing interval is monthly. In this case, the distinction between dynamic and semi-static becomes blurred.

5 Conclusion

Andersen and Andreasen (1999, 2000) have suggested that a jump diffusion model coupled with a local volatility function may be able to capture observed market cross-sectional implied volatilities. In this paper, we investigate two crucial challenges associated with such a jump model, model calibration and hedging of jump risk. Previous authors (e.g. Cont and Tankov, 2004b) have shown that the calibration problem for a 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 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 vanilla options lack information regarding tails of the price distribution, there are insufficient option prices to completely determine a local volatility function. We generate synthetic market prices where the true market follows a jump diffusion process with known parameters. We then illustrate that, while it is difficult to obtain accurate parameter estimates for

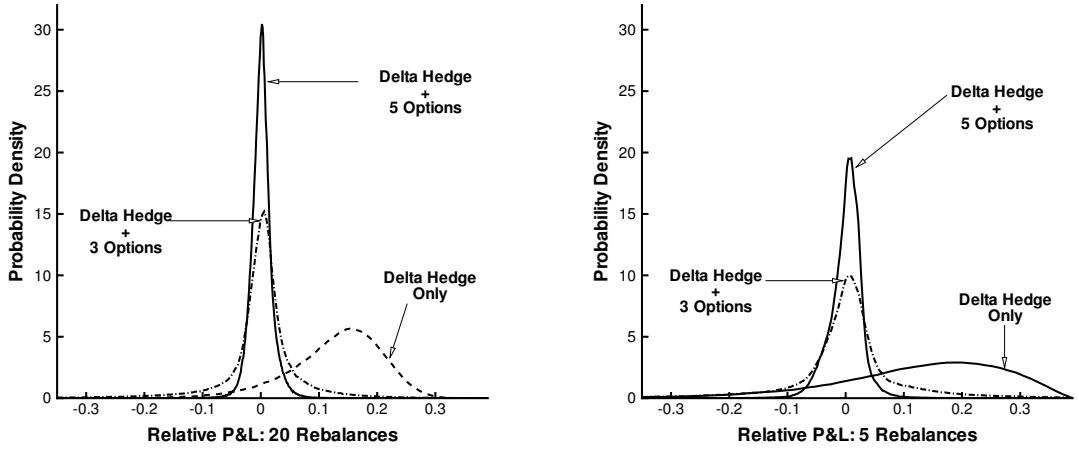

FIGURE 4.3: *Dynamic hedging: relative P&L distributions. Primary option: American straddle. Hedging horizon: three months.*

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, in general 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 with a portfolio containing 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 relative profit and loss distributions of these strategies. The experiments included an investigation of estimation errors in the parameters, assuming that a hedger would clearly 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 Carr and Wu (2004), we assume that the market is incomplete with a limited set of options available for hedging. 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 probability density function of the objective price process seems 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 transition density 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 transition density was used. It appears that the transition density is dominated by the

diffusion volatility, which is comparatively easy to estimate. The jump parameters mainly affect the tails of the distribution. Since the transition density 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.

The main 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, at least in the straddle examples we considered.
- Both dynamic and semi-static hedging strategies perform very well with hedging portfolios containing five options. The semi-static approach is robust to poor estimation of the $\mathbb{P}$ measure transition density function.

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. We defer a detailed investigation of these topics to future research, merely noting here that it is straightforward to include a penalty term in our objective function so that trading of options at each rebalance date is minimized. It is also worth noting that most of our results regarding the ill-posedness of the calibration problem and the hedging of jump risk should be expected to carry over to more complicated models which incorporate stochastic volatility as well as jumps (though obviously volatility risk would have to be managed in addition to jump risk). Again, however, a detailed investigation of this is left for subsequent research. 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

Andersen, L. and J. Andreasen (1999). Jumping smiles. *Risk* 12(11), 65–68.

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

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

Baker, G., R. Beneder, and A. Zilber (2004). FX barriers with smile dynamics. Working paper, Erasmus University.

Bakshi, G. and C. Cao (2002). Risk-neutral kurtosis, jumps, and option pricing: Evidence from 100 most actively traded firms on the CBOE. Working paper, Smith School of Business, University of Maryland.

Bakshi, G., C. Cao, and Z. Chen (1997). Empirical performance of alternative option pricing models. *Journal of Finance* 52, 2003–2049.

Bates, D. S. (1988). Pricing options under jump-diffusion processes. Working paper 37-88, Rodney L. White Center for Financial Research, The Wharton School, University of Pennsylvania.

Bates, D. S. (1991). The crash of '87: Was it expected? The evidence from options markets. *Journal of Finance* 46, 1009–1044.

Bates, D. S. (2000). Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics* 94, 181–238.

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

Carr, P. and A. Hirsa (2003). Why be backward? *Risk* 16(1), 103–107.

Carr, P. and L. Wu (2004). Static hedging of standard options. Working paper, Courant Institute, New York University.

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

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

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

Coleman, T. F., Y. Li, and A. Verma (2002). A Newton method for American option pricing. *The Journal of Computational Finance* 5(3), 51–78.

Cont, R. and P. Tankov (2004a). *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, Boca Raton.

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

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

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

d'Halluin, Y., P. Forsyth, and K. Vetzal (2005). Robust numerical methods for contingent claims under jump diffusion processes. *IMA Journal of Numerical Analysis* 25, 65–92.

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

Dierckx, P. (1993). *Curve and Surface Fitting with Splines*. Oxford University Press, Oxford.

Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.

Dumas, B., J. Fleming, and R. E. Whaley (1998). Implied volatility functions: Empirical tests. *Journal of Finance* 53, 2059–2106.

Dupire, B. (1994). Pricing with a smile. *Risk* 7(1), 18–20.

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

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 327–343.

Hull, J. and A. White (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance* 42, 281–300.

Labahn, G. (2003). Closed form PDF for Merton's jump diffusion model. Technical report, School of Computer Science, University of Waterloo.

Lagnado, R. and S. Osher (1997). Reconciling differences. *Risk* 10(4), 79–83.

Liu, J., J. Pan, and T. Wang (2005). An equilibrium model of rare-event premia and its implication for option smirks. *Review of Financial Studies* 18, 131–164.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.

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

Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance* 49, 771–818.

Schoutens, W., E. Simons, and J. Tistaert (2004). A perfect calibration! Now what? *Wilmott Magazine* March, 66-78.

Tikhonov, A. N. and V. Y. Arsenin (1977). *Solutions of Ill-Posed Problems*. W. H. Winston, Washington.

Vapnik, V. N. (1982). *Estimation of Dependences Based on Empirical Data*. Springer-Verlag, Berlin.

Wahba, G. (1990). *Spline Models for Observational Data*. SIAM, Philadelphia.

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

A Appendix: Market Data

This appendix provides the market data used in the calibration test in Section 2.3 for Brocade Communications Systems Inc. on April 21, 2004.

Zero Coupon Government Bonds

Maturity	1 month	3 months	6 months	1 year	2 years
Yield	1.09%	1.17%	1.28%	1.64%	2.43%

TABLE A.1: *Yields to maturity for zero coupon government bonds on April 21, 2004.*

American Option Prices ($)

K	$T = .0833$	$T = 0.25$	$T = 0.5$	$T = 1$	$T = 2$
2.5				0.05	0.16
3			0.03	0.11	0.28
4		0.05	0.14	0.33	0.61
5	0.06	0.21	0.40	0.68	1.05
6	0.32	0.58	0.83	1.17	1.59
7	0.15	0.40	0.67	1.08	1.69
7.5	0.08	0.28	0.54	0.94	1.55
8	0.04	0.20	0.43	0.82	1.42
9	0.02	0.11	0.28	0.62	1.20
10		0.06	0.19	0.48	1.02
11		0.04	0.13	0.37	0.88
12		0.03	0.09	0.30	0.76
12.5		0.02	0.08	0.27	0.71
15		0.01	0.04	0.17	0.53
17.5		0.01	0.03	0.12	0.40
20			0.02	0.09	0.32
25			0.02	0.06	0.22
30			0.02	0.04	0.17
40			0.01	0.03	0.12
45				0.03	0.10

TABLE A.2: *American option prices for Brocade Communications Systems Inc. on April 21, 2004. The spot price* $S_0 = \$6.19$. *Call prices are reported when the strike* $K > S_0$. *Put prices are given when* $K < S_0$.