Jumps and Estimation Risk in Finance and Decision Making

George Dotsis
Department of Accounting, Finance and Management and Essex Finance Centre, University of Essex
gdotsis@essex.ac.uk

Abstract
This paper provides an overview of my thesis on the statistical and economic role of jumps and estimation risk in contingent claims analysis. Regarding estimation risk, by employing non parametric statistical methods the thesis derives sampling distributions of European and American options that are consistent with arbitrage conditions and reduces considerably the finite sample bias in GARCH option pricing models. Furthermore, a simple statistical framework for gauging the impact of estimation risk in a variety of optimal decision-making rules is applied and it is demonstrated that incorrectly omitting estimation risk has a considerable impact on optimal decision making. With respect to jump risk, the thesis examines its impact on the dynamics of volatility using data from the spot and derivatives markets. The analysis points towards the central role of jumps in the dynamics of volatility under both statistical and financial criteria. Finally, the thesis derives optimal policy rules when economic agents face uncertainty and irreversibility and the underlying variables are subject to abrupt changes. It is shown that the possibility of abrupt changes is an important feature that accelerates policy implementation.

1. Introduction

Traditional financial decision-making models builds upon the assumption that the underlying stochastic variables display local smooth changes that can be approximated by diffusion processes the parameters of which are known with certainty. However, one of the endemic hazards of financial decision-making is the

risk of a sudden large shock in the underlying variables. The crash of 1987 and the meltdown during the Russian crisis in 1998 are only two of the many recent examples of such events. Such phenomena cannot be captured by diffusion processes alone, but require mixture models of high activity and low activity events, termed as jump diffusion models (eg., Duffie et al. 2000;). Furthermore, in any empirical application the true values of the parameters of the underlying diffusion are never known with certainty and the investor must estimate them from a set of available data. Consequently, when statistical estimates are used in place of the true parameters investors face an additional uncertainty related to the so-called estimation risk (see Gibson et al., 1999; Barberis, 2002), because the parameters are estimated with errors.

In the thesis I examine the statistical and economic role of jumps and estimation risk in contingent claims analysis in finance. Contingent claims analysis concern a wide spectrum of problems embracing both financial derivatives and optimal decision making rules implied by real options analysis (see Merton, 1990; Dixit and Pindyck, 1994) and is undoubtedly considered as the state of the art methodology in financial modelling. In the thesis, jump risk is examined with respect to the dynamics of volatility and optimal timing rules for policy implementation problems. Estimation risk is examined in the context of pricing options and optimal timing of financial investment decisions.

Jumps and estimation risk in contingent claims analysis are particularly challenging in terms of econometric estimation and modelling. The econometric estimation of the dynamics of the underlying state variables is demanding because in continuous-time jump diffusion process the conditional density functions of discretely sampled returns are rarely known in closed form and hence conventional estimation methods cannot be applied. Moreover, parametric methods cannot adequately tackle estimation risk because contingent claims are highly non-linear functions of the underlying variables and hence the statistical errors on the input parameters are transmitted in a complex manner.

The paper is structured as follows. Section 2, describes the main results of the thesis on the statistical and economic role of estimation risk in option pricing and optimal timing of financial investment decisions. Section 3, presents the main findings of the thesis with respect to statistical and economic role of jump risk. Finally, the concluding section underlines the key conclusions of the thesis.

2. Estimation Risk

2.1 Inferring the distribution of Option Prices

Inference on the distribution of option prices has been a long-standing problem in the financial literature. In particular, the use of estimates in place of the true but unknown volatility in the Black-Scholes formula (hereafter, BSF), gives rise to an interesting statistical problem since the estimate of the variance rate affects the estimate of the corresponding option price. This problem involves estimation risk whereby the model is valid but the input parameters are uncertain.

A significant complication arises from the fact that even an unbiased estimate of the variance does not produce an unbiased estimate of the option price since the BSF is nonlinear with respect to volatility (see Boyle and Anathanarayanan, 1977). Moreover, it is well known that reducing estimation risk by increasing sample size, either by sampling frequency or time horizon, may not be always appropriate in view of microstructures and nonstationarities (e.g., see Boyle and Anathanarayanan, 1977, Campbell et al., 1997, Ch. 3).

Lo (1986) developed the most general framework based on asymptotic statistical theory for estimation and testing contingent-claim asset pricing models, such as the BSF. A computational advantage of this approach is that the limiting distribution of the contingent claim can be derived regardless of whether the claim is priced in closed form or numerically. The problem of statistical bias due to the use of historical variance has not received much attention in the empirical option pricing literature. Knight and Satchell (1997) argued that this could be due to the relatively small magnitude of biases reported in simulation studies, the development of more general formulae than the BSF and the widespread adoption of implied volatilities instead of historical variances. However, as discussed by the authors, the reported biases may appear so small due to the conservative volatility levels and the short expirations adopted in the simulation studies. Finally, the authors stress that implied volatilities are not always readily available, as for example, in real option pricing.

In the thesis I propose the use of a bootstrap methodology in inferring the properties of the finite sample distribution of option prices and hedge parameters when estimates of variance are used. Despite its computational intensity, it is argued that this methodology may have considerable advantages over the alternatives discussed previously, in terms of: small sample statistical properties, data requirements, computational simplicity and consistency against no-arbitrage bounds. As reported by the numerous applications of the bootstrap methodology in statistics and econometrics, it is particularly well suited for the problem underhand, which involves a highly nonlinear estimation with a limited amount of data. In an empirical application, I explore the merit of the proposed methodology in comparison to Lo's asymptotic approach using spot data from the S&P 500 index.

While the empirical literature in this area has been concerned with the distribution of European option prices, I also look into the distribution of a hedge parameter, delta, which is of great concern to investors, particularly to those selling options. Moreover, I examine the distribution of American options, something that also has not received attention in the empirical literature. Although American options present significant computational problems, since no analytical pricing model exists, they are by far the most widely traded derivatives.

The bootstrap is a resampling method of simulation for inferring the distribution of a statistic derived from a sample by treating the sample as the population. It is nonparametric in the sense that unlike alternatives such as Monte Carlo simulation methods, it does not draw repeated samples from assumed distributions. The bootstrap carries out conventional statistical calculations in an unconventional way: by purely computational means, rather than through the use of mathematical formulas. Bootstrap can be thought as a nonparametric maximum likelihood theory applied via the computer to a more complicated class of estimation problems. The development of this approach is relatively new (Efron, 1979) and has become increasingly popular with the widespread availability of powerful computers. Although bootstrapping is computationally intensive, it is intuitively appealing and particularly easy to implement. A variety of resampling schemes, standard error estimation algorithms and hypothesis testing procedures have been developed within the bootstrap methodology.

It has been widely demonstrated that under mild conditions, boostrapping provides more accurate approximations in small samples to the distribution of many statistics than classical large sample approximations. As demonstrated by Horowitz (2001), under mild regularity conditions, the bootstrap provides approximations to distributions of statistics, coverage probabilities of confidence intervals, and rejection probabilities of hypothesis tests that are at least accurate as the approximations of 1st order asymptotic distribution theory, without entailing the algebraic complexity of higher order expansions. Even if asymptotic and bootstrap standard errors are the same, the confidence intervals can be different if the bootstrap distribution is skewed. Although bootstrap estimates are asymptotically efficient, they are not necessarily unbiased, but tend to have small biases compared to the magnitude of their standard errors.

The bootstrap is especially well suited for cases where it is difficult to calculate the asymptotic distribution of an estimator or statistic. In the case of finite variance the bootstrap distribution converges weakly to normality. Singh (1981) has showed that under the existence of third moments, the bootstrap is asymptotically a better approximation to the true distribution than the normal on the basis of Edgeworth expansion. In the case where third moments do not exist then the bootstrap approximation is asymptotically equivalent to the normal (Hall, 1988).

I study the sampling distributions of European and American option prices and hedge parameters, respectively. The data employed consist of daily closing prices for the Standard and Poors 500 Index (S&P500) from 6/10/03 to 31/12/03, a total of 60 observations. The relatively short sample length was determined according to standard practice (e.g. see Hull, 2004), in order to avoid nonstationarities. Daily returns were calculated as usual via logarithmic differencing. A maximum likelihood estimate of variance is used to calculate volatility levels, and, subsequently, point estimates of option prices using the BSF and the Barrone-Adesi and Whaley (1987) approximation to American option prices, respectively. Lo's (1986) asymptotic standard errors (ASE) of prices are then calculated For the European options I also study the asymptotic distribution of the partial derivative of option price with respect to the underlying asset, the so-called hedge parameter delta. I then compare the results for European and American options with those obtained via bootstrapping, respectively.

I find that biases in the BSF option prices due to the use of estimates in historical variance are likely to be small. However, I show that asymptotic approaches produce confidence intervals for out of the money options that violate no-arbitrage bounds. In contrast, under the bootstrap approach prices are always consistent with no arbitrage conditions. I also extend the results to the hedge parameter delta and to American option prices, where I obtain similar results. However, I find that an increase in volatility and time to maturity, contrary to what has been conjectured in the literature, does not necessarily increase biases in option prices. Shifting volatility, time to maturity, risk-free rates and dividend yields appears to have significant, yet mixed, effects on the shape of the distribution of option prices and deltas, respectively.

2.1 Bias Reduction in Option Prices

There has been a growing interest recently within academia and industry for option pricing models assuming time-varying volatilities of financial asset returns. In these so-called GARCH option pricing models, the conditional variance of returns is modelled under various specifications of a Generalized Autoregressive Conditionally Heteroskedastic (GARCH) process. A great advantage they provide is that options can be priced solely on the basis of observables, without necessarily resorting to option market data. Moreover, GARCH option pricing models have been found able to fit well European option prices.

In the thesis I explore the finite sample properties of a popular GARCH option-pricing model proposed by Heston and Nandi (2000). Option prices in this model depend on the unknown GARCH parameters and the risk premium of the underlying asset process. Since these parameters are typically estimated via maximum likelihood, they may be biased in finite samples. Moreover, it is known that the assumptions underlying a GARCH process are not necessarily robust to the specification of the sampling frequency, ie., if the true frequency for the process is weekly, then the GARCH parameters estimated from monthly data are not unbiased estimates of the true parameters.

Using standard ML estimation, the parameter vector of the Heston Nandi model are estimated from the discrete log returns of the asset price. Limit theorems show that for a vector of parameters, the asymptotic distribution of a non-linear function of the parameters, as in an option pricing model, is given by:

$$\sqrt{n}\left(\hat{F}_{ML}(\hat{q}) - F(q)\right) \xrightarrow{A} N\left(0, V_F\right)$$

where, $V_F \equiv \dfrac{\partial F}{\partial q} V_q \dfrac{\partial F}{\partial q'}$

The parameter $V_?$ is the asymptotic variance-covariance matrix of ? and is vector of derivatives with respect to the input parameters. Therefore, asymptotic theory implies that as the sample grows larger, the option prices are normally distributed around the true price with variance. As noted by Lo (1986), the practical usefulness of asymptotic theory depends on the sample size required to approximately attain asymptotic normality. Since no general guidelines exist, this issue should be resolved in view of the particular problem in hand. Therefore, I rely on simulation results to examine the number of observations required for the estimated option price to converge to its "true" price.

The simulations are performed with the parameters estimated by Heston and Nandi (2000, Table 1, pp. 597) using 3 years of daily SP500 returns. These parameters imply an unconditional annual volatility of about 9.5%. I also run the simulations for hypothetical unconditional volatility levels of 15% and 20%. Unconditional volatility is "shifted" by arbitrarily changing the GARCH parameters. Although these parameters are not estimated from historical data, it is instructive to assess the effect of higher volatility on the finite sample properties of option prices. For each level of volatility I examine sample sizes of 500, 750, 1,500 and 3,000 observations, respectively. Following standard practise, in each simulated path I disregard the first 100 observations. I run 1,000 simulations and for each simulated path I estimate the parameter vector ? via ML. I then compute option prices for various moneyness and maturity levels, respectively. For pricing the options I employ the risk neutral unconditional variance implied by the parameters. The conditional variance cannot be used since it is path dependent and cannot be defined ex-ante in order to obtain the "true" option price. For both the GARCH parameters and the option prices I estimate the percentage bias defined as the percentage difference between the average estimate from the 1,000 simulations and the true value, respectively.

The simulations show that the parameters of the GARCH process vary considerably in terms of percentage bias when unconditional volatility is 9.5%. For all sample sizes, the risk premium exhibits the largest percentage bias. For example, for a sample size of 750 observations the risk premium estimate is biased upward by over 230%. It is important to note that biases in the parameters persist even when 3,000 returns, 12 years of daily data, are employed. The standard errors decrease roughly at the rate implied by asymptotic theory, i.e., by a factor of $\sqrt{2}$, when the sample observations increase from 750 to 1,500 and from 1,500 to 3,000, respectively. Interestingly enough, even though the bias in the GARCH parameters is substantial, the bias of the unconditional risk-neutral variance is roughly zero in all cases. The average risk neutral unconditional variance derived from the simulations implies an annual risk neutral volatility 9.53%, which matches very closely the one implied by the "true" parameters.

In terms of option valuation, I price call options with 30 and 60 days to maturity, respectively. In general, the results suggest that the finite sample bias in the option price estimates are small across the configurations examined. Only short term, out-of-the-money options are significantly affected by the biases in the GARCH parameters. In particular, the 10% out-of-the-money call option with maturity 30 days is overpriced by 38% for a sample size of 500 observations. Although mispricing tends to reduce as the sample size increases, even with 3,000 observations the option price is biased upwards by over 6%. Short-term at-the-money, 5%-out-of-the-money and 5%-in-the-money options display negligible bias. As we move to 60 days to maturity, all options are accurately priced for all sample sizes. As noted by Heston and Nandi (2000), the option prices depend only marginally on the value of risk premiumr and this explains the fact that its large bias is not fully transmitted into the option prices. It is reasonable to argue that the large biases observed for short-term options could be also due to the fact that the tails of the conditional distribution over short intervals rely heavily on the GARCH parameters. In contrast, options with longer maturities depend primarily on the risk neutral variance input. The simulation study showed that the risk neutral unconditional variance is symmetric around the true value. Hence, it appears that the symmetry of the volatility input is passed into the option prices. The decrease in the standard errors of the option prices found to be roughly consistent with asymptotic theory. At-the-money options display the largest standard errors since they have the largest sensitivity with respect to the variance (vega) and consequently are more influenced by its sampling variation.

In order to reduce the bias observed in short term out of the money options I employ the jackknife resampling method. The jackknife was originally proposed by Quenouille (1956) as a general method for bias reduction. In relation to other bias reduction methods, the jackknife is superior because of its generality and its small computational cost. Although it is possible to jackknife the parameter estimates and then plug the revised estimates into the option pricing formula, I apply the jackknife directly to the option prices themselves. This approach was also favoured by Phillips and Yu (2005), since it deals directly with the ultimate quantity of interest and avoids the problem discussed previously arising from the fact that an unbiased estimate of variance does not necessarily produce an unbiased option price. Moreover, it is not certain that jackknifed GARCH parameters will still satisfy the inequality constraint that ensures finite moments in the underlying process. I concentrate on those option pricing configurations that were found to exhibit the largest biases. In general, the results indicate that the bias reduction ability of the jackknife improves along with the size of the sample used. More specifically, for a sample size of 750 observations, the bias is decreased by 82% for the option with 30 days to maturity, from 25.1% to only 4.4% with a 25% increase in the standard error of the estimate. For the 15% unconditional volatility level and 1,500 observations, bias is reduced from 19.9% to -4.6%. The bias in the jackknifed option prices is negligible for the two lower volatility levels studied when 3,000 observations are used. For the 20% unconditional volatility level, the jackknife manages to reduce bias only when 3,000 observations are employed.

These results can be used to re-interpret previous findings in the literature with respect to the inability of GARCH option pricing models to reproduce empirically realistic price levels for certain option configurations (eg., Duan et al, 2003). The mispricing could also be due to the finite sample properties of the estimators employed rather than to a structural flaw of the option pricing model. In any case, I have shown empirical studies may be highly questionable if estimation risk remains unexplored and the necessary steps for bias reduction are not attempted.

2.1 Optimal Decision Making under Estimation Risk

Real option theory is a relatively new approach to investment evaluation, which stresses the irreversibility of investment decisions and the uncertainty of the surrounding economic environment. However, empirical applications have been rare since relevant market data for the required parameters are usually not available. As with financial options, the volatility level (s) is a key parameter which influences mostly the outcome of the investment evaluation process. A variety of approaches have been used in the literature to approximate project volatility (e.g., Davis, 1998). Similar to financial options, the use of estimates in place of the true but unknown variance gives rise to an interesting statistical problem since the estimate of the variance rate affects the estimates of the corresponding real option value and threshold. Such problems have been identified in the real options literature and ad-hoc sensitivity analysis is typically used as a way of assessing the impact of changes in volatility.

However, real option investment analysis has not been formally examined from the perspective of estimation risk, whereby the model is valid but the input parameters are uncertain. Estimation risk is particularly relevant for real options since the availability of data is, at best, very limited. In the thesis The methodology applied in the thesis draws from previous research on the impact of volatility estimation risk on financial options in finite samples (Boyle and Anathanarayanan, 1977) and asymptotically (Lo, 1986). Following the approach of Mc Donald and Siegel (1986), I study the option to invest under the assumption that the value of the project evolves according to a geometric Brownian motion.

The statistical results indicate that optimal decision making, even when reasonably large sample sizes are available, will be difficult since the critical project value confidence intervals are considerably wide and asymmetric. Uncertainty in optimal investment thresholds is considerably higher for the right wing of the confidence interval. Incorrectly assuming symmetry in the confidence interval, as typically implied by sensitivity analysis, means that the decision rule will tend to be too optimistic and overvalue projects.

In order to highlight further the implications of estimation risk in optimal decision making I examine its impact on the model of Paddock et al. (1988). This is a well established model that treats the valuation of an offshore oil tract as a three-stage investment problem comprising of the exploration, development and extraction of oil. Given that the development stage requires the largest capital expenditures, I focus on the valuation of the undeveloped reserve and the decision as to when to develop. Applying the statistical methodology, I show that that the option to develop the reserve entails significant estimation risk. The results reveal that incorrectly omitting estimation risk can lead to suboptimal decision making and significant economic losses.

3. Jump Risk

3.1 Volatility

The presence of jumps in asset returns has been the subject of many studies in finance. Existing results emphasize the economic and statistical role of jumps and highlight its fundamental importance in a variety of financial applications such as derivatives valuation (e.g., Bates, 1991; Naik and Lee, 1990, Broadie Chernov and Johannes, 2004), asset allocation (e.g., Liu, Longstaff and Pan, 2003), corporate debt (e.g., Chen and Kou, 2005), term structure modeling (e.g., Johannes, 2004) and econometric estimation (e.g., Eraker, Johannes and Polson, 2003;). Recently, the impact of jumps has been examined with respect to a new class of financial products termed as volatility derivatives. Volatility derivatives are written on some measure of volatility such as implied volatility indices and provide pure exposure to volatility changes. Given recent evidence that during times of market stress diversification fails (eg., Skintzi and Refenes, 2005) and volatility changes rapidly (e.g., Eraker, Johannes and Polson, 2003) these products are especially suited so as to act as hedging vehicles. The correct specification of the stochastic process that an implied volatility index follows is necessary in order to price and manage the risk of derivatives written on implied volatility, accurately. Moraux et al. (1999), estimated mean-reverting processes using data on implied volatility indices. Wagner and Szimayer (2004) and Psychogios et al (2005) investigate the presence of jumps in the implied volatility index by estimating autonomous mean reverting jump diffusion process. They found evidence of significant positive jumps in implied volatilities. However, these studies are restricted to statistical analysis and they do not employ data from the derivatives market. However, it is often the case that statistical results are not always compatible with financial results. In the thesis I explore the ability of alternative univariate diffusion and jump diffusion processes to capture the dynamics of implied volatility indices over time and to match the observed volatility derivative prices. In contrast to the previous literature, I also estimate non-stationary processes. The validity of the alternative stochastic processes is assessed under both econometric and financial metrics by using a rich data set. Data on a plethora of European and American implied volatility indices (VIX, VXO, VXD, VDAX, VX1 and VX6), and the rapidly evolving CBOE volatility futures market are employed. The use of data drawn from different countries/sectors within the same country also sheds light on whether the properties of the models depend on these factors. In addition, this facilitates the investigation of the performance of the alternative specifications under different construction settings, since some of these indices are based on a different construction method (e.g., VIX versus VXO). Within the econometric framework, conditional Maximum Likelihood Estimation (MLE) is used to estimate the parameters of the various volatility processes. In the case where the conditional density function does not have a closed-form, the characteristic function is derived. Then, Fourier inversion of the characteristic function is employed. Standard statistical tests are used to compare the alternative processes. Within the financial metric, the CBOE volatility futures market is used to rank the alternative processes within a futures pricing context. Under the risk-adjusted probability measure, the volatility futures price equals the conditional expected value of the future volatility. Hence, the valuation of volatility futures is not model-free; for any given process, the pricing performance of the corresponding volatility futures pricing model is examined. To the best of my knowledge, Zhang and Zhu (2006) is the only study that has investigated the pricing of volatility futures empirically. However, they do this for a specific model without conducting a horse race among different models.

The parameters of the various processes are estimated by the conditional Maximum Likelihood (ML) method (see Hamilton, 1994, for a description of the method). I rely on the conditional MLE because it takes into account the empirically documented dependence of the observations. MLE has been commonly used to estimate continuous time models in finance (see Sundaresan, 2000, for a review, and the references therein). This is because it has desirable statistical properties. The set $\hat{\Theta}$ of the ML estimators are consistent, asymptotically efficient achieving the Cramer-Rao lower bound for consistent estimators, and they are normally distributed, $\hat{\Theta} \sim N[\Theta, \{I(\Theta)^{-1}\}]$, where $[I(\Theta)]^{-1} = (-E[\frac{\partial^2 \log L(\Theta)}{\partial\Theta\partial\Theta'}])^{-1}$. Moreover, Ait-Sahalia (2004) has shown that alternative estimation methods such as the generalized method of moments cannot attain the efficiency of the ML estimators. Moreover, he has found that MLE can disentangle the diffusion from the jump component.

The conditional MLE requires the conditional (transition) density function $f[V(t+t)|V(t),\Theta]$ (t>0) of the process V_t, where t denotes the sampling frequency of observations and T is the set of parameters to be estimated. For the processes where the conditional density function does not exist in a closed form, the corresponding conditional characteristic function is derived. The required conditional density function is obtained by Fourier inversion of the characteristic function. Maximizing the likelihood function via Fourier inversion, though computationally intensive, provides asymptotically efficient estimates of the unknown parameters (see Singleton, 2001, for a discussion and applications). Moreover, even though the likelihood function of jump diffusion processes may be unbounded, its Fourier transform is always bounded.

After estimating the various processes I derive the corresponding futures prices model. Let $G_t(V,T)$ denote the futures price at time t for a futures contract on V with maturity T. Under the risk-adjusted equivalent martingale measure Q, $G_t(V,T)$ equals the conditional expected value of V_T at time T (see Cox et al., 1981); the expected value is conditional on the information up to time t. The pricing performance of the various models depends on the remaining time-to-maturity. As the time-to-maturity increases the futures price that corresponds to the mean-reverting processes (non-mean-reverting processes) tends to a constant (grows exponentially). Therefore, the "mean-reverting" volatility futures models are expected to perform worse as the time-to-maturity increases since they cannot capture the stochastic evolution of implied volatility. The assessment of the pricing performance of the volatility futures models is done as follows. The parameters of the pricing models are the risk-adjusted parameters. These can be obtained either by calibration or by making an assumption about the market price of volatility/jump risk and then apply Girsanovs' theorem. In the case of calibration, a natural approach would be to calibrate each model to the market futures prices of the shortest futures series. Then, the pricing performance of each model would be investigated for the remaining futures series. However, calibration (e.g., by means of a generalized non-linear least squares regression method) is an ill-posed problem in the case of the jump-diffusion processes under consideration due to an under-identification problem. Therefore, I use the parameters estimated from the time series data and resort to an assumption about the market price of the volatility/jump risk.

The econometric analysis finds discontinuities in implied volatility while mean reversion is of second order importance. Most importantly, the results are also consistent under the financial metric, where jump diffusion processes display the lowest pricing errors, regardless of the market price of risk employed. These results provide further support to the fact that jump risk is an important aspect in financial markets.

3.2 Optimal Investment

Optimal investment rules derived from real options theory have a broad applicability that goes beyond pure financial investment decisions. Recently, optimal decision making rules implied by real options analysis have also been applied to optimal timing problems of environmental policies.

The real options approach is though to be superior since it takes explicitly into account three important characteristics of environmental problems, namely, uncertainty, irreversibility and flexibility. Firstly, there is uncertainty over both the magnitude environmental damage (environmental uncertainty) and the future costs and benefits of a certain policy to reduce this damage (economic uncertainty). Secondly, there are at least two kinds of irreversibility inherent in environmental problems: environmental damage itself is at least partially irreversible, and, the costs of adopting policies to reduce this damage are largely sunk costs. Finally, policy adoption is clearly not a “now-or-never” proposition since it can be delayed in anticipation of new information.

Motivated by the recent empirical evidence on the possibility of jumps in emission levels or social costs from pollution related phenomena, the thesis examines the resulting effects on the optimal timing of environmental policy decisions. More specifically, by extending the approach of Pindyck (2000), I allow for unexpected abrupt changes in the evolution of the stock of pollutant. I then study ceteris paribus the effects of jumps in the future costs and benefits of environmental damage and its reduction. In contrast to Pindyck (2000) where the analysis is carried out under the assumption of diffusion processes, I let the state variables follow jump diffusion processes. Smooth changes are governed by a Brownian motion while the arrival of extreme sudden events is controlled by a Poisson process with a fixed jump size. In order to preserve simplicity and analytical tractability, I do not model economic and environmental uncertainty simultaneously (as in Pindyck, 2002). The issue of jumps in the state variables underlying policy decisions is an issue that has not been thoroughly examined in the literature. Baranzini et al. (2003) have modelled the net discounted benefit, defined as the ratio of expected discounted benefits to expected discounted costs, as a Brownian motion with jumps.

However, their model is different, because it assumes that uncertainty regarding costs and benefits is strictly exogenous. In order to make uncertainty endogenous, I allow the benefits from emissions reduction to depend on the current level of the stock of pollutant, which is assumed to follow a jump diffusion process. The Baranzini et al. (2003) approach presents problems in its implementation since the ratio of discounted benefits to expected discounted costs is difficult to measure empirically. In the present chapter I allow for jumps in the pollutant level itself for which a wealth of empirical data and structural models exist.

I find that the optimal threshold were policy should be adopted?is a strictly decreasing function of the occurrence of extreme events. If the stock of pollutant is subject to large unexpected changes then policy should be adopted earlier. Moreover, an increase in the probability of a large change in the future flow of social costs will cause environmental policies to be adopted earlier. An increase in jump size will have a similar effect.

The results of this study show that neglecting the possibility of jumps in the stock of a pollutant may lead to suboptimal delays in the adoption of environmental policy decisions. Similar results are obtained when the social cost per unit of stock of pollutant is allowed to rise sharply. Overall, the presence of discontinuities is a factor that should be treated with great caution since it may affect significantly the optimal timing of environmental policies.

Conclusions

The thesis examines the statistical and economic role of jumps and estimation risk in contingent claims analysis. For tackling estimation risk in derivatives pricing, the thesis proposes the use of non parametric resampling methods. I studied the sampling distribution of European and American option prices and deltas when historical estimates of variance are used using statistical bootstrapping. I showed that in contrast to conventional methods such as large sample theory, bootstrapping always produces sampling distributions for prices and hedge ratios that do not violate no-arbitrage bounds and hence can be reliably adapted in empirical applications. Moreover, I examined the finite sample properties and estimation risk involved in the implementation of the Heston and Nandi (2000) GARCH option pricing model. I found that the model overprices options with respect to their "true" value only in the case of short-term, out-of-the-money configurations. I provided evidence that for adequate sample sizes, the jackknife resampling method can be used effectively and with small computational cost to reduce finite sample bias in the GARCH option prices. These results can be used to re-interpret previous findings in the literature with respect to the inability of GARCH option pricing models to reproduce empirically realistic price levels for certain option configurations (eg., Duan et al, 2003). The mispricing could also be due to the finite sample properties of the estimators employed rather than to a structural flaw of the option pricing model.

I examined the impact of volatility estimation risk on real options and optimal decision rules. Using a simple statistical framework, I inferred the distribution of the optimal investment threshold both in finite samples and asymptotically. In both case I showed that the biases are small even in very small samples. However, I demonstrated that volatility estimation risk induced wide confidence intervals for the true value. In an empirical application I applied the statistical framework to the Paddock et al. (1988) study. I derived the confidence intervals and I showed that uncertainty of the volatility estimate induces wide confidence intervals for the true value of the investment. The results of indicated that incorrectly omitting volatility estimation risk might have a considerable impact on optimal decision making and result in significant economic losses.

Jump risk was shown to be an important characteristic in financial market and in particular with respect to the dynamics of volatility. I explored the ability of alternative univariate diffusion and jump diffusion processes to capture the dynamics of implied volatility indices over time and to match the prices of observed volatility option prices. The validity of the alternative implied volatility processes was assessed under both econometric and financial metrics. To this end, data were employed from major European and American implied volatility indices and the rapidly growing CBOE volatility futures market. Within the econometric metric, maximum likelihood estimation has been carried out. The characteristic functions have been derived in the cases where they were required. Under the financial metric, the processes were ranked according to the pricing performance in the CBOE market of the corresponding volatility futures pricing models. I showed that under both metrics the simplest jump diffusion model à la Merton (1976) performed best followed very closely by the simplest diffusion specification (geometric Brownian motion). I showed that the econometric results hold regardless of the implied volatility index under scrutiny

Finally, I derived optimal timing rules for environmental policies when the underlying state variables follow jump diffusion processes. I considered two cases: environmental uncertainty and economic uncertainty. I showed that the presence of jumps in emission levels or in social costs from pollution related phenomena is an important characteristic that must be taken into consideration. I showed that in both cases the presence of jumps accelerates the optimal timing of environmental policy decisions. Moreover, I demonstrated that when the probability or size of jumps in either state variable grows arbitrarily large then policies should be implemented immediately.

The thesis leads to a variety of research questions to be addressed in the future. In particular, the impact of estimation risk on asset allocation strategies still remains an unresolved issue in the finance literature. Recent studies compare the out-of-sample performance of simple allocation rules (eg., 1/N) to models of optimal asset-allocation (static and dynamic models). Surprisingly, they find that simple allocation strategies perform better. The reason is that is that the loss from sub-optimal diversification is smaller relative the loss from the estimation errors in the parameters required to implement optimal asset allocation strategies. Moreover, the impact of jumps on asset returns is also an unresolved issue. An interesting direction of future research would be to examine the relationship between extreme events in financial markets and the equity premium puzzle.

References

Ait-Sahalia, Y. (2004). "Disentangling Diffusion from Jumps", *Journal of Financial Economics*, 74, 487-528.

Baranzini, A., Chesney, M., and Morisset, J. (2003). "The Impact of Possible Climate Catastrophes on Global Warming Policy", *Energy Policy*, 31, 691-701.

Barberis, N. (2000). “Investing for the Long Run when Returns Are Predictable”, *Journal of Finance*, 55, 225-264.

Barone-Adesi, G., and Whaley, R.E. (1987). “Efficient Analytic Approximation of American Option Values”, *Journal of Finance*, 42, 301-20.

Bates, D. (1991). “The Crash of '87 -- Was It Expected? The Evidence from Options Markets”, *Journal of Finance,* 46, 1009-1044.

Boyle, P.P., and Ananthanarayanan, A.L. (1977). “The Impact of Variance Estimation in Option Valuation Models”, *Journal of Financial Economics*, 5, 375-387.

Broadie, M., Chernov, M., and Johannes, M. (2004). “Model Specification and Risk Premiums: The Evidence from the Futures Options”, Working Paper, Columbia Business School.

Campbell, J.Y., Lo, A., and MacKinlay, C.A. (1997). The Econometrics of Financial Markets, Princeton NJ: Princeton University Press.

Chen, N., and Kou, S. (2005). “Credit Spreads, Optimal Capital Structure, and Implied Volatility with Endogenous Default and Jump Risk”, Working Paper, Columbia University.

Davis, G. (1998). “Estimating Volatility and Dividend Yield When Valuing Real Options to Invest or Abandon”, *Quarterly Review of Economics and Finance*, 38, 725-754.

Dixit, A. and Pindyck, R. (1994). Investment under Uncertainty. Princeton, New Jersey: Princeton University Press.

Duffie, D., Pan, J., and Singleton, K. (2000). “Transform Analysis and Asset Pricing for Affine Jump-Diffusions”, *Econometrica*, 68, 1343-1376.

Efron, B. (1979) “Bootstrap Methods: Another look at the Jacknife”, *Annals of Statistics*, 7, 1-26.

Eraker, B., Johannes, M., and Polson, N. (2003). “The Impact of Jumps in Volatility and Returns”, *Journal of Finance,* 53, 1269-1300.

Gibson, R., Lhabitant, F.S., Pistre, N., and Talay, D. (1999). “Interest Rate Model Risk: An Overview”, *The Journal of Risk*, 1, 37-62.

Hall, P. (1988). “Rate of Convergence in Bootstrap Approximations”, *Annals of Probability*, 16, 1665-1684.

Heston, S.L., and Nandi, S. (2000). “A Closed-Form GARCH Option Valuation Model”, Review of Financial Studies, 13, 585-625.

Horowitz, J.L. (2001). “The Bootstrap”, in: Heckman, J.J. and Leamer, E.E (Eds.) Handbook of Econometrics ,Amsterdam: Elsevier Science.

Johannes, M. (2004). “The Economic and Statistical Role of Jumps to Interest Rates”, Journal of Finance, 59, 227-260

Knight, J., and Satchell, S.E. (1997). “Existence of Unbiased Estimators of the Black/Scholes Option Price, Other Derivatives and Hedge Ratios”, Econometric Theory, 13, 791-807.

Liu, J., Longstaff, F., and Pan, J. (2003). “Dynamic Asset Allocation with Event Risk”, *Journal of Finance,* 58, 231-259.

Lo, A.W. (1986). “Statistical Tests of Contingent Claims Asset Pricing Models”, *Journal of Financial Economics*, 17, 143-173.

McDonald R., and Siegel, D. (1986), “The Value of Waiting to Invest”, *The Quarterly Journal of Economics*, 101, 707-728.

Merton, R.C. (1990). Continuous-Time Finance Oxford, U.K.: Blackwell.

Moraux, F., Navatte, P., and Villa, C. (1999). “The Predictive Power of the French Market Volatility Index: A Multi Horizons Study”, *European Finance Review*, 2, 303-320.

Naik, V., and Lee, M. (1990). “General Equilibrium Pricing of Options on the Market Portfolio with Discontinuous Returns”, *Review of Financial Studies,* 3, 493–521.

Paddock, J., Siegel, D., and Smith, J. (1988). “Option Valuation of Claims on Real Assets: The Case of Offshore Petroleum Leases”, *Quarterly Journal of Economics*, 103, 479-508.

Phillips, P.C.B., and Yu, J. (2005). “Jackknifing Bond Option Prices”, *Review of Financial Studies*, 18, 707-742.

Pindyck, R. (2000). “Irreversibilities and the Timing of Environmental Policy”, *Resource and Energy Economics*, 22, 233-259.

Pindyck, R. (2002). “Optimal Timing Problems in Environmental Economics”, *Journal of Economic Dynamics and Control*, 26, 1677-1697.

Psychoyios, D., Dotsis, G., and Markellos R. (2005). “Does VIX Jump? Implication for Pricing and Hedging Volatility Risk, MSL Working Paper.

Quenouille, M.H. (1956). “Notes on Bias in Estimation”, *Biometrika*, 43, 353–360.

Singh, K. (1981). “On the Asymptotic Accuracy of Efron’s Bootstrap”, *Annals of Statistics*, 9, 1187-1195.

Singleton, K. (2001). "Estimation of Affine Asset Pricing Models Using the Empirical Characteristic Function”, *Journal of Econometrics*, 102, 111-141.

Skintzi, V., and Refenes, A. (2005). “Implied Correlation Index: A New Measure of Diversification”, *Journal of Futures Markets*, 25, 171-197.

Sundaresan, S.M. (2000). “Continuous Time Methods in Finance: A Review and an Assessment”, *Journal of Finance,* 55, 1569-1622.

Wagner, N., and Szimayer, A. (2004). “Local and Spillover Shocks in Implied Market Volatility: Evidence for the U.S. and Germany”, *Research in International Business and Finance*, 18, 237-251.

Zhang, J.E., and Zhu, Y. (2006). "VIX Futures", *Journal of Futures Markets*, forthcoming.