

## *Model Uncertainty and Its Impact on Derivative Pricing*

**Alok Gupta, Christoph Reisinger,  
Alan Whitley**

University of Oxford

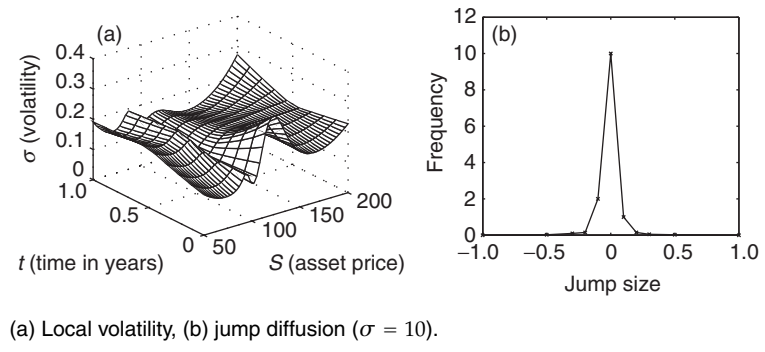
Financial derivatives written on an underlying can normally be priced and hedged accurately only after a suitable mathematical model for the underlying has been determined. This chapter explains the difficulties in finding a (unique) realistic model: model uncertainty. If the wrong model is chosen for pricing and hedging, unexpected and unwelcome financial consequences may occur. By “wrong model” we mean either the wrong model type (specification uncertainty) or the wrong model parameter (parameter uncertainty). In both cases, the impact of model uncertainty on pricing and hedging is significant. A variety of measures are introduced to value the model uncertainty of derivatives and a numerical example again confirms that these values are a significant proportion of the derivative price.

In this introductory section, we will look at various ways in which the model selection problem may manifest itself and the consequences of this for derivative pricing.

### **Motivating examples**

To demonstrate the prevalence of model uncertainty in derivative pricing, we look at calibrating different derivative pricing models to a set of observed prices. We consider a set of 60 European call prices for six maturities varying between one month and one year, and 10 strikes varying between 90% and 110% of the spot value. We look at two examples of trying to calibrate models to these prices. In the first we do not assume we know the model type (thus investigating specification uncertainty), and in the second we fix the model

**Figure 21.1** Local volatility model and jump diffusion model fitted to 60 European call prices




**Table 21.1** Heston stochastic volatility model fitted to 60 European call prices

| Parameter                | Value   |
|--------------------------|---------|
| Rate of reversion        | 0.0745  |
| Long-run variance        | 0.1415  |
| Volatility of volatility | 0.1038  |
| Correlation              | -0.2127 |
| Initial variance         | 0.0167  |

type and try to identify the model parameter (therefore considering parameter uncertainty).

**Example 21.1 (specification uncertainty).** To the same set of 60 observed European call prices, we fit a local volatility model (Dupire 1994), a jump-diffusion model (Cont and Tankov 2004) and a Heston stochastic volatility model (Heston 1993). The models are all very different: a one-factor continuous process, a two-factor continuous process and a discontinuous process respectively. But each model is nevertheless fitted to within an average of three basis points (bp) of the same quoted call prices. Two of the calibrated models are displayed in Figure 21.1, where part (a) shows a local volatility surface that reproduces the 60 prices to within an average of 3bp, part (b) shows the jump density (of an exponential Lévy process) that does the same. Table 21.1 gives the parameters for a Heston model that also fits the observed prices.

**Figure 21.2** Four very different local volatility surfaces fitted to the same set of 60 calibration prices



Surfaces reproduce prices on average to within 3bp of the observable prices. European call options have maturities varying between one month and one year and strikes varying between 90% and 110% of the spot value.

**Example 21.2 (parameter uncertainty).** For the same set of 60 European call prices, we now assume the model is known to be local volatility and try to find different local volatility surfaces that fit the prices. This is now a problem related to the uncertainty involved when fitting a pricing model. Observe in Figure 21.2 the variety of differently shaped local volatility surfaces that arise. Note that each local volatility surface reproduces all 60 prices to within an average of 3bp. We have only plotted four surfaces to demonstrate the disparity between the shapes.

Now imagine trying to price another option on the same underlying. Below are the prices given by the different models and the different model calibrations in Examples 21.1 and 21.2, respectively, for a three-month up-and-out barrier call option with strike 90% of the spot and barrier 110% of the spot.

As Table 21.2 indicates, the derivative price variation is noticeable for different local volatility surfaces – up to 26 basis points. The disparity in pricing is even larger for different model types – by up to 177 basis points.

**Table 21.2** Barrier prices found by different fitted models

| Model                 | Price | Local volatility | Price |
|-----------------------|-------|------------------|-------|
| Local volatility      | 7.41  | A                | 7.51  |
| Jump-diffusion        | 9.18  | B                | 7.32  |
| Stochastic volatility | 7.77  | C                | 7.48  |
|                       |       | D                | 7.58  |

Barrier prices found by different fitted models according to Example 21.1 (in the two left-hand columns) and Figure 21.2 (the two right-hand columns). We price an up-and-out barrier call option with strike 90% of the spot and barrier 110% of the spot value.

### Risk and Knightian uncertainty

The cost of selecting the wrong model is most commonly referred to as model risk, but the more accurate terminology is model uncertainty. The reason for this is as follows. Suppose the underlying  $S$  takes different values  $S(\omega)$  depending on the future scenario  $\omega \in \Omega$  that occurs. And let  $\mathbb{P}$  be the probability measure corresponding to the set of future scenarios  $\Omega$ . Then risk corresponds to not knowing which future scenario  $\omega \in \Omega$  will occur, whereas uncertainty corresponds to lack of knowledge of the probability measure  $\mathbb{P}$ .

The distinction between risk and uncertainty was highlighted by Knight (1921).<sup>1</sup> Although subtle, the difference between not knowing the future state and not knowing the probability of the possible future states is important, and investors and risk-managers are likely to have differing aversions to both.

It is important to also note the difference between model uncertainty and market incompleteness (Branger and Schlag 2004). In an incomplete market, the true data-generating process may be known, but not all contracts are attainable, so the pricing measure is not unique. However, under model uncertainty, we do not even know the true data-generating process.

### Sources and types of model uncertainty

The investigation of model uncertainty applied to financial models is in its infancy and only began to receive attention when, for example, Derman (1996) published his research notes on model uncertainty. Derman identifies several sources of model uncertainty,

which we group into three distinct classes, with examples, and give some extensions.

**Incorrect model.** On a fundamental level, known mathematical models might not be capable of projecting stock movements, such that brinkmanship or psychology might play a more important role. Furthermore, some factors might have been forgotten, or factors incorrectly modelled (eg, as deterministic when they are stochastic or vice versa). The model for a price process in one market may be inappropriate for another in a different market with differing levels of interest rates or volatility. A model suitable in a stable market might become inappropriate in a time of financial crisis. Once market frictions such as transaction costs and illiquidity are factored in, a model might no longer be applicable.

**Incorrect solution.** The model might be correct but the final analytical solution found could be wrong. For example, Li's famous copula formula (Li 2000), on which billions of dollars were invested, had an important right bracket missing. Numerical approximations may not be accurate enough. For example, Monte Carlo methods often need many simulations to converge. Software and hardware can be faulty so, because a lot of trading and pricing platforms use many thousands of lines of program code, difficult-to-detect errors can cause incorrect solutions. For example, the R2009b release of MATLAB incorrectly solved a linear system for a transposed  $2 \times 2$  matrix.<sup>2</sup>

**Incorrect calibration.** Instationarity of the underlying process may cause previous calibrations to no longer be applicable. Instability of the solution might imply that the wrong model (parameter) is chosen. Furthermore, there may be a lack of robustness of the solution, ie, pricing and hedging is non-robust with respect to the modelling assumptions.

Whatever the source of model uncertainty, the consequences can be dramatic and costly. We highlight some of these effects in the following section, with a focus on the first and last points above.

### **Effect of uncertainty on derivative pricing**

Uncertainty in the modelling of the underlying will manifest itself as risk in derivative pricing and hedging. There is a broad spectrum of model uncertainty ranging, at one extreme, from situations where

very little, if anything, is known about the detailed structure of the model of the underlying (specification uncertainty), to the opposite extreme where the structure of the model is specified in detail but there is uncertainty over the parameters of the model (parameter uncertainty).

At the one extreme, even though we do not have a full specification of the model, we may still be able to draw conclusions about derivative pricing. There is a substantial body of literature that discusses the range of arbitrage-free prices for certain kinds of derivatives where only the most general properties of the model of the underlying are assumed. For example, elementary bounds on the prices of vanilla derivatives can be given that are completely model-free (Chance 2008). As a fully parameterised description of the model is not assumed in such cases, we will refer to this kind of uncertainty as specification uncertainty.

Later in this chapter, as examples of specification uncertainty we will describe two model uncertainty frameworks where the underlying follows a stochastic volatility process. Model uncertainty is present in these frameworks because we have only limited information on the volatility process driving the model; no particular form of volatility process is assumed and it remains unparameterised.

At the other extreme, model uncertainty may reduce to the complete description of a family of models in which the 'true' model is believed to lie and where each family member is fully specified as a parameterised model. For example, we may have a family of classic Black-Scholes models which differ only in the value of the (constant) volatility used to define each member; the volatility is then the parameter describing the uncertainty. We will refer to this later as parameter uncertainty. This kind of uncertainty lends itself to analysis by Bayesian methods, as we can begin by assigning a prior probability distribution to the set of parameters describing the family of models and then use observed data (such as the behaviour of the underlying or prices of derivatives on the underlying) to derive a posterior distribution for the parameters. This approach is described in detail later. Other methods of describing the uncertainty in this framework exist; for example, the uncertainty might be described using the language of fuzzy set theory, where "possibility" replaces the concept of probability, eg, as illustrated by Swishchuk *et al* (2008).

Typically, we will have some prior beliefs about the model uncertainty and, in the absence of any other information, will take these beliefs into account in pricing and hedging. For example, in the situation where we do not specify the form of the volatility, we might only begin with the prior belief that the spot volatility always lies between 0.2 and 0.3. But, in the case of parameter uncertainty, we might begin with the prior belief that the appropriate family of models consists of Black–Scholes models with (constant) volatility between 0.2 and 0.3.

Analysis of derivative pricing and hedging using only this limited prior information is likely to result in an unacceptably wide range of possible prices and hedging strategies. It is standard practice to then make use of additional data to calibrate (or “constrain”) the family of possible models in order to reduce the uncertainty range. As an excellent example of this, Cox and Obłój (2008, 2009) describe price bounds for certain kinds of barrier options that make only mild assumptions about the model of the underlying, but assume knowledge of European option prices for the same maturity and all strikes.

Calibration may use data for only the underlying to reduce uncertainty about the real world measure or it may use observable prices of traded derivatives to reduce uncertainty about the risk-neutral measure. A key feature of the Bayesian framework applied to parameter uncertainty (detailed later) is that calibration yields a posterior probability distribution for the family of allowable models which can then be used to inform the pricing of other derivatives. A posterior distribution results because many models may give a “sufficient” fit to the calibration data, although some will be more probable than others.

In this chapter, the main focus will be on the problem of incorrect model choice. As described above, specification uncertainty and parameter uncertainty are two ways of describing the situation where we believe we have some (perhaps limited) prior information about the “true” model, eg, that it is a member of a particular family of models that we can describe in some way. We must acknowledge this and understand the implications for the pricing and hedging of derivatives.

### MODEL-FREE PRICING

In this section we will present examples of model uncertainty where the family of models cannot or shall not be parameterised. The first example to be presented is the uncertain volatility model introduced by Avellaneda *et al* (1995) and we will follow this example with a variant of this problem as described by Mykland (2009). These examples will illustrate that, although we have a situation where the model for the underlying is uncertain and we cannot fully specify the members of the family of allowable models, we can nevertheless obtain useful information about derivative pricing and hedging. Furthermore, just as with the Bayesian techniques to be described later, it is possible to use the prices of traded financial instruments in a form of calibration that reduces the effect of the model uncertainty on the pricing and hedging of derivatives.

#### Bounded volatility assumptions

Avellaneda *et al* (1995) described a pricing problem where the model for the underlying,  $S_t$ , is taken to be the stochastic differential equation

$$dS_t = S_t(\mu_t dt + \sigma_t dZ_t)$$

where  $Z_t$  is a standard Brownian motion,  $\mu_t$  is the drift and the volatility,  $\sigma_t$ , is a stochastic process satisfying the condition

$$\sigma_t \in [\sigma_{\min}, \sigma_{\max}] \quad \text{for all } t \in [0, T] \quad (21.1)$$

for some non-negative constants,  $\sigma_{\min}$  and  $\sigma_{\max}$ , over a time interval  $[0, T]$ . Avellaneda *et al* (1995) mention that the bounds could, in fact, be taken to be functions of the price of the underlying and of time. This uncertainty framework has also been studied by Lyons (1995).

Avellaneda *et al* consider the pricing and hedging of a European derivative with payoff  $h(S_T)$  at maturity,  $T$ , in this framework (in fact, they consider the pricing and hedging of a portfolio of European claims with various maturities) and give expressions for upper and lower bounds for the derivative price,  $W^+(S_t, t)$  and  $W^-(S_t, t)$ .

The initial upper bound,  $W^+(S_0, 0) = W_0^+$ , is the lowest price that can be charged for the derivative such that, by following an appropriate hedging strategy, the seller can be sure to avoid making a loss on hedging. The initial lower bound,  $W^-(S_0, 0)$ , is the highest price that can be paid for the derivative such that, by following an



appropriate hedging strategy, the buyer can be sure to avoid making a loss on hedging.

The authors show that these price bounds satisfy two versions of a non-linear partial differential equation, the Black–Scholes–Barenblatt (BSB) equation. The BSB equation is

$$\frac{\partial W}{\partial t} + r\left(S\frac{\partial W}{\partial S} - W\right) + \frac{1}{2}\hat{\sigma}^2\left[\frac{\partial^2 W}{\partial S^2}\right]S^2\frac{\partial^2 W}{\partial S^2} = 0$$

In the equation for  $W^+$ ,  $\hat{\sigma}$  is defined by

$$\hat{\sigma}[\Gamma] = \begin{cases} \sigma_{\max} & \text{if } \Gamma \geq 0 \\ \sigma_{\min} & \text{if } \Gamma < 0 \end{cases}$$

and, in the equation for  $W^-$ ,  $\sigma_{\min}$  and  $\sigma_{\max}$  are interchanged. This equation must be solved numerically using the portfolio payoffs as time boundary conditions.

In the case of a derivative with a convex payoff, the BSB ask price is the same as the standard Black–Scholes price with the volatility set identically equal to  $\sigma_{\max}$  and the bid price is the standard Black–Scholes price with the volatility set identically equal to  $\sigma_{\min}$ . For a concave payoff, these results apply with  $\sigma_{\max}$  and  $\sigma_{\min}$  interchanged. For payoffs of mixed convexity, the BSB ask and bid prices may lie outside the range of the standard Black–Scholes prices computed with the volatility set identically equal to  $\sigma_{\max}$  or  $\sigma_{\min}$ .

Under certain conditions on the stochastic volatility process, as  $\sigma_{\max}$  tends to  $\infty$  and  $\sigma_{\min}$  tends to 0, the ask price is given by the smallest concave super-majorant of the payoff function, ie, the smallest concave function that is always at least as large as the payoff (Cvitanic *et al* 1999). So, for example, in the case of a European call option when the volatility is unbounded, the ask price is given by the spot price of the asset, which is an elementary model-free bound.

The span of prices between  $W^+$  and  $W^-$  is a measure of the effect of model uncertainty on the price of the derivative. If the derivative were sold for a price outside this range, an arbitrage opportunity would arise for either the buyer or the seller. When the volatility range shrinks to zero, the upper and lower price bounds both equal the Black–Scholes price for the then fixed volatility, as would be expected.

The solution of the BSB equation also provides the appropriate hedging strategy to use if the derivative is sold for, say,  $W^+(S_0, 0)$ .

The hedge ratio is then simply given by the delta of the value process, ie,  $\partial W^+ / \partial S$ .

Frey and Sin (2001) point out that, in practice, it may not be possible to determine a finite upper bound on the volatility process to use as  $\sigma_{\max}$  (nor may it be possible to precisely specify the lower bound  $\sigma_{\min}$ ). They suggest a practical approach in which volatility bounds are estimated such that the probability of the volatility process straying outside the bounds during the time interval  $[0, T]$  is  $1 - \alpha$ , ie, they define a prediction set for the volatility process where  $\alpha$  is the probability that the conditions assumed in the calculation of the ask price do not hold. If the ask price is calculated using these bounds and the derivative is then continuously hedged using the corresponding hedge ratios, the probability of incurring a hedging error will not exceed  $\alpha$ .

In practice, it would be convenient if the volatility bounds could be chosen so that the probability of incurring a hedging error were equal to some agreed tolerance level,  $\beta$ . Unfortunately, the link between the  $\alpha$  as described above and the resulting  $\beta$  is rather weak and  $\beta$  can be much smaller than  $\alpha$ . As a result, too high a value of  $\alpha$  may be chosen, leading to an excessively high ask price, for example.

This behaviour is illustrated by the results of a pricing and hedging simulation carried out by the authors. In this simulation, the underlying volatility process follows a Heston model with known parameters, and  $\sigma_{\max}$  and  $\sigma_{\min}$  were chosen to be 0.40 and 0.10, respectively. Simulation over a period of six months showed that around 50% of the volatility paths strayed outside these limits. The BSB equation with these volatility bounds was used to calculate the ask price for a particular butterfly spread and a hedging simulation was then run using the hedge ratios from the BSB solution with the asset price trajectories being again derived from the Heston model. Analysis of the hedging errors from the simulation showed that only 5% of the paths incurred hedging errors.

Furthermore, when the volatility bounds were changed to 0.50 and 0.03, it was found that about 5% of the paths violated the constraints but none of the 50,000 paths sampled resulted in hedging errors.

In this case, if the error tolerance had been set at 5%, it would have been acceptable to set the volatility bounds at 0.40 and 0.10. However, if the bounds had been set, in a rather natural way, to

ensure that the model conditions were only violated 5% of the time, the bounds could have been set at 0.50 and 0.03. This choice would have resulted in a higher, and thus less competitive, ask price, albeit with a much lower probability of incurring a hedging error.

### **Super-replication**

In an incomplete market, the price for a claim  $X$  cannot in general be uniquely identified by no-arbitrage arguments. In this case the “super-replication” price of  $X$  is an indicator of the (maximum selling) value of the claim. Under certain conditions, the super-replication price is equal to  $\sup_{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}}[X]$ , where  $\mathbb{Q}$  belongs to the set of pricing measures.

The pricing problem in the uncertain volatility model framework is an example of super-replication. It addresses the problem of identifying an initial wealth and a self-financing trading strategy that will almost surely achieve the derivative payoff. In this case, the super-replication must be achieved whenever the volatility process satisfies the stated condition. The terminology “super-hedging” is sometimes used instead of super-replication.

The uncertain volatility model has been explicitly studied as a super-replication problem by Frey (2000), who identifies a process that super-replicates a given European claim given the conditions satisfied by the volatility process (Frey’s method also works for a restricted set of path-dependent claims, eg, some special types of barrier derivatives). The super-replicating process he identifies has a value process given by the price of an American derivative where the underlying is a normalised geometric Brownian motion (ie, it has zero drift and unit volatility), the payoff is a modified version of the original derivative payoff and the volatility bounds determine the exercise window for the American derivative.

Although Frey finds a super-replicating process and hence the initial worst-case ask price, the super-replication is not necessarily minimal, ie, it may be possible to charge a lower price and hedge appropriately to guarantee hedging losses do not occur. He goes on to discuss when it is possible to formulate the problem as the solution of a partial differential equation (recovering the BSB equation) which does give a minimal super-replication and demonstrates the circumstances under which the two approaches give the same result.

Although Frey deals only with the super-replicating price, a similar method could be used to compute a lower bound on the derivative price, the sub-replication price, thus producing an uncertainty range for the derivative price.

Frey presents his results in detail for a single derivative of maturity  $T$  and his method clearly extends to a portfolio of European derivatives with the same maturity. He also describes how to calculate the super-replication price for a portfolio of European derivatives with different maturities.

**Bounded total-variance assumptions**

Mykland (2009) describes an uncertain volatility model which differs from the framework by Avellaneda *et al*, Lyons and Frey by placing an alternative constraint on the stochastic volatility process. Mykland assumes that the volatility process,  $\sigma_t$ , satisfies the following total-variance condition

$$\mathcal{E}^- \leq \int_0^T \sigma_t^2 dt \leq \mathcal{E}^+ \tag{21.2}$$

for two constants  $\mathcal{E}^-$  and  $\mathcal{E}^+$ . The quantity being constrained here is the quadratic variation of the log process for the underlying.

Mykland shows that the worst-case (“conservative”) ask price for a European derivative with maturity  $T$  can be described as the price of a suitable American derivative in this framework. In other words, there is a starting price,  $A_0$  and a super-replication process,  $V$ , with  $V_0 = A_0$ , whose value at time  $t$ ,  $V_t$ , is given by an American derivative price and there is an associated hedging strategy that ensures no loss will be made as long as the volatility process satisfies the condition given above.

Given a volatility process that satisfies the volatility bounds in Equation 21.1, we will have

$$\sigma_{\min}^2 T \leq \int_0^T \sigma_t^2 dt \leq \sigma_{\max}^2 T$$

so the range of initial prices that results from the Mykland framework cannot be wider than the range derived from the framework by Avellaneda *et al*, Lyons and Frey, given the same volatility process.

An important feature of the Mykland model is that the hedging strategy relies on the continuous estimation of the realised variance of the underlying, ie

$$\int_0^t \sigma_u^2 du$$

as this quantity is used to adjust the exercise window for the American derivative being calculated. This quantity can be estimated from the history of the log returns of the underlying but the accuracy of the super-replication will depend on the accuracy with which the realised variance is estimated.

Mykland provides a further interpretation of his model where the condition satisfied by the quadratic variation is viewed as the definition of a prediction set, ie, a probability  $(1 - \alpha)$  is attributed to the outcome that a realisation of the volatility process will satisfy the constraints. With this interpretation, the conservative ask price becomes the price at which the derivative can be sold, so that, by following the appropriate hedging strategy, the seller can be sure that the probability of incurring a hedging error is at most  $\alpha$ . The idea is that the seller can choose a tolerable level for the risk of incurring hedging errors and price and hedge the derivative accordingly.

This approach can be compared to the use of a prediction set with the bounded volatility uncertainty framework as described earlier. It should again be noted that a hedging error does not necessarily result from a particular volatility path violating the total-variance bounds,  $\mathcal{E}^-$  and  $\mathcal{E}^+$ , so  $\alpha$  may overstate the probability of a hedging error occurring.

### Using calibration to reduce model uncertainty

An important property of the price found by super-replication is that the pricing mechanism is non-linear, ie, if we have two European derivatives maturing at times  $T_1$  and  $T_2$ , with payoffs  $h_1(S_{T_1})$  and  $h_2(S_{T_2})$ , then the super-replication price for a linear combination of these two derivatives will not in general be the same linear combination of the super-replication prices. This property of the pricing mechanism opens up the possibility of reducing the worst-case ask price for a European claim by setting up a static hedge using traded derivatives in the same underlying.

If we let the holdings in the  $n$  hedging instruments be given by  $\lambda_1, \dots, \lambda_n$ , and instrument  $i$  have payoff  $h_i(S_{T_i})$  and traded price  $p_i$ , then the worst-case ask price for another derivative with payoff  $h(S_T)$  is given by

$$W_0^+(h + \lambda_1 h_1 + \lambda_2 h_2 + \dots + \lambda_n h_n) - (\lambda_1 p_1 + \lambda_2 p_2 + \dots + \lambda_n p_n)$$

ie, we calculate the worst-case “ask” price for the enlarged (statically hedged) portfolio and then subtract the initial cost of the hedge. We

then find the static hedge coefficients  $\lambda_1, \dots, \lambda_n$  which minimise this price. Due to the non-linearity of the pricing mechanism, the result may be lower than the ask price in the absence of the hedge, so a reduced ask price can be charged, safe in the knowledge that hedging the enlarged portfolio will not incur a hedging loss if the volatility satisfies the assumed condition.

This method of reducing the uncertainty range of the derivative price is a form of calibration and, in the absence of a bid–ask spread for the traded derivatives, the calibration would be exact, as the value of a traded derivative calculated by this method will simply deliver the traded derivative price. In practice there will be bid–ask spreads and a simple approach might be to use the average of the bid and ask prices to represent the price of a traded derivative in the equation. However, a more sophisticated approach would be to include the bid or ask prices in the formula depending on the signs of the hedging coefficient,  $\lambda_i$ , which determine whether the hedges in the traded derivatives are held long or short. This will lead to a more complicated optimisation problem to determine the hedging coefficients.

For the Avellaneda *et al* uncertain volatility model, this method of static hedging can easily be applied, as the pricing mechanism described in Avellaneda *et al* (1995) can handle a portfolio of mixed maturity European derivatives. Similarly, as Frey (2000) describes how to calculate the super-replication price for a portfolio of mixed maturity European derivatives, his method could also be used for calibration against a set of mixed maturity European options.

In the case of Mykland’s uncertainty volatility model, a method of determining the ask price for a mixed maturity portfolio would be required. This would probably require constraints on the quadratic variation to be specified for time intervals other than  $[0, T]$ .

Model uncertainty measures will be discussed in more detail later but, in passing, we can observe that one measure of model uncertainty can be defined by setting

$$\mu(X) = \text{conservative ask price}(X) - \text{conservative bid price}(X)$$

where  $X$  is a claim. In the case of both the Frey and Mykland super-replication approaches, this is because the conservative ask price can be written as

$$\bar{\pi}(X) = \sup_{P \in \mathcal{P}} \{\mathbb{E}^P[X]\}$$

where  $\mathcal{P}$  is a suitable family of probability measures. Cont (2006) describes how an upper bound on price of this type can be combined with a lower bound given by

$$\underline{\pi}(X) = -\bar{\pi}(-X)$$

to give a coherent measure of model uncertainty. The lower bound on price is the conservative bid price.

Although calibration (using a static hedge with an appropriate choice of calibrating instruments) can be expected to reduce the uncertainty in the price of a derivative, the lack of a detailed specification of the underlying model may mean that the uncertainty range cannot be reduced to an acceptable level. For this reason, it may be necessary to attempt a more detailed description of the underlying model and use the Bayesian techniques described in the next section to reduce the range of prices and better quantify the model uncertainty.

However, it should be borne in mind that moving to a more precise description of the underlying model carries the risk that the family of candidate models will be made too narrow, so that uncertainty is underestimated. Consequently, selection of the family of candidate models and the associated choice of the prior distribution should be made with some care.

### CALIBRATION AND PARAMETER UNCERTAINTY

If a model class has been chosen for the underlying, a decision still has to be made on how to calibrate the model for subsequent derivative pricing and hedging. The problem is non-trivial and a wrong decision can be costly for the decision maker. In this section, we focus on the problem of choosing a suitable parameter for a model, and explain why it is often difficult and unstable. We then present a robust Bayesian solution.

#### The inverse problem

Let the underlying asset price process be  $S = (S_t)_{t \geq 0}$  and suppose it depends on the time  $t$ , stochastic process(es)  $Z = (Z_t)_{t \geq 0}$  and model parameter  $\theta \in \Theta$ , ie

$$S_t = S(S_0, t, (Z_u)_{0 \leq u \leq t}; \theta) \quad (21.3)$$

where  $S_0$  is the value of the underlying at time 0.

Suppose we wish to price an option on  $S(\theta)$  that has maturity  $T$  and a payoff function  $h$ . Let  $f_t(\theta)$  be the price of this option at  $t$  when the model parameter is  $\theta$ . If we assume there is a risk-neutral pricing measure  $\mathbb{Q}$ , then we can explicitly write this price as

$$f_t(\theta) = \mathbb{E}^{\mathbb{Q}}[B(t, T)h(S(\theta)) \mid (S_u)_{0 \leq u \leq t}]$$

where  $B(t, T)$  is a discount factor between  $t$  and  $T$ .

Now we observe many such option prices  $\{f_t^{(i)}(\theta) : i \in I_t\}$  at time  $t \in [0, T]$ , where  $I_t$  is an index set of data. Usually, prices are only observed to within bid–ask spreads so there exists a noise component  $\{e_t^{(i)} : i \in I_t\}$ , ie

$$V_t^{(i)} = f_t^{(i)}(\theta^*) + e_t^{(i)} \quad (21.4)$$

for  $i \in I_t$ .  $\theta^*$  denotes the true parameter. The calibration problem is then to select  $\theta$  which best reproduces the observed prices  $\{V_t^{(i)} : i \in I_t, t \in Y_n([0, T])\}$ , where  $Y_n([0, T]) = \{t_1, \dots, t_n : 0 = t_1 < t_2 < \dots < t_n \leq T\}$  is an  $n$ -partition of the interval  $[0, T]$ . The calibration problem is an example of an “inverse problem”, since we know the forward function  $f_t$  which enables us to compute the price if we know  $\theta$ , but we do not know how to explicitly recover  $\theta$  if we know the price.

However, before attempting to find the solution  $\theta$ , it is first necessary to ascertain whether a stable solution exists at all.

### Well-posedness

We call a mathematical problem “well posed” if it satisfies Hadamard’s criteria:<sup>3</sup>

- (i) for all admissible data, a solution exists;
- (ii) for all admissible data, the solution is unique;
- (iii) the solution depends continuously on the data.

If, on the other hand, a mathematical problem violates one or more of the above criteria, then we call it “ill posed”. Parameter identification problems are often ill posed. In the context of calibration, we start by assuming we can find a solution fitting the data to within an acceptable error tolerance,  $\delta$  say, and hence satisfying (i). The classical example where this is not given is the Black–Scholes model, where a single volatility parameter cannot be chosen to simultaneously fit options with different strikes and maturities to an acceptable level.



Such models would not be adequate in practice, and we assume that the class of models is sufficiently rich to contain models which fit the prices to an acceptable degree, in which case we often cannot guarantee properties (ii) and (iii). The effects of violating either of these two properties will be seen for pricing and hedging.

A good example for this is the local volatility model, which extends the Black–Scholes model by allowing the volatility to depend on both time and the spot price. Dupire’s Formula (Dupire 1994) provides an explicit expression of this volatility function in terms of option values for a continuum of strikes and maturities. However, finitely many data – as are observable in practice – do not suffice to pin down the infinite-dimensional parameter.

If there is more than one possible solution, ie, more than one calibrated parameter, then we call the inverse problem “underdetermined”. This happens when we do not have enough market prices to restrict the value of the calibrated parameter. In this situation, choosing the wrong calibrated parameter will lead to incorrect pricing and hedging of other options, which can result in losses for a trading agent.

Furthermore, the admissible data is almost always noisy (the values are only observed with added error as in Equation 21.4), so we assume the true values to lie within some confidence interval around the observed value. In our context, prices are never observed exactly, but only to within a bid–ask spread. So the bid–ask spread can be thought of as this error confidence interval.

If a solution does not depend continuously on the data, ie, market prices, then a small mis-pricing in the market of one of the observed prices can lead to a disproportionately large error in the chosen calibrated parameter. This is again drastically exemplified by the local volatility model, which is extremely ill conditioned with respect to noisy observations even when the surface is discretised. And again, this results in incorrect pricing and hedging of other (exotic) contracts.

### **Regularisation**

We call the process of approximating an ill-posed problem by a well-posed problem “regularisation”. A vast literature<sup>4</sup> exists on handling ill-posed problems and especially ill-posed inverse problems.

Let us consider a general inverse problem in which we know the forward function  $f$  and want to solve

$$f(\theta) = V, \quad \theta \in \Theta, \quad V \in \mathcal{V} \quad (21.5)$$

for finite-dimensional  $\theta$ , but do not know the inverse function  $f^{-1}$ .  $\Theta$  is the parameter set, eg, the set of discretised local volatility functions in the context of the previous section, and  $\mathcal{V}$  is the image set, eg, the set of quoted prices for vanilla instruments. Suppose further that we can only observe an approximation  $V^\delta$  for  $V$ ,  $\|V^\delta - V\|_{\mathcal{V}} \leq \delta$  with some observation error norm, and are instead trying to solve  $f^{-1}(V^\delta) = \theta^\delta$ . Assume that  $f^{-1}$  does not satisfy Hadamard's condition (ii) and/or (iii) from the previous section.

The most widely used approach to regularisation is to replace  $f^{-1}$  by a regularisation operator  $f_\lambda^{-1}$  with regularisation parameter  $\lambda > 0$  which depends on  $\delta$  and/or  $V^\delta$ . The operator and parameter are chosen so that

$$\begin{aligned} \lambda &= \lambda(\delta, V^\delta) > 0 \\ f_\lambda^{-1}: \mathcal{V} &\rightarrow \Theta \text{ is bounded for all } \lambda \in (0, \lambda_0) \\ \limsup_{\lambda \rightarrow 0} \{ \|f_\lambda^{-1}(V^\delta) - f^{-1}(V^\delta)\|_\Theta \} &= 0 \end{aligned}$$

This ensures that  $(f_\lambda^{-1}(V^\delta) =: \theta_\lambda^\delta) \rightarrow \theta^\delta$  as  $\lambda \rightarrow 0$  (Engl *et al* 1996).

It still remains, however, to find a regularisation operator and parameter. There are several methods for doing so:<sup>5</sup> using the spectrum of operator  $f$ , using Fourier, Laplace and other integral transformations.<sup>6</sup>

A common way to address the potential non-existence of a solution is to replace Equation 21.5 by a minimisation (least-squares) problem for the calibration error

$$\|g(\theta^\delta)\|_{\mathcal{V}} = \|f(\theta^\delta) - V^\delta\|_{\mathcal{V}}$$

and to address the non-uniqueness and stability by adding a stabilising function  $h: \Theta \rightarrow \mathbb{R}$ . Hence, our original problem (Equation 21.5) becomes

$$\text{find the } \theta^\delta \text{ which minimises } \|g(\theta)\|_{\mathcal{V}}^2 + \lambda h(\theta) \quad (21.6)$$

An appropriate choice for  $h$  varies from problem to problem, but common practice is to take a Tikhonov functional (Tikhonov *et al* 1977). The Tikhonov functional favours solutions with smaller  $h(\theta)$ .

In the context of calibrating local volatility, different functionals, composed of (potentially higher-order) Sobolev semi-norms of the local volatility function, have been proposed by various authors, originating with Jackson *et al* (1999), and including Crepey (2003), Egger and Engl (2005) and Achdou (2005).

However, this formulation for  $h$  usually has no immediate financial meaning, but rather is taken for mathematical convenience in terms of proving regularisation properties. A noteworthy exception in several respects is the work of Berestycki *et al* (2002), where the regularisation functional is motivated by an asymptotic analysis of the implied volatility for short time-to-maturity and far in-the-money and out-of-the-money options. The Bayesian framework, in contrast, offers a way to attach financial meaning to the regularisation term.

### Bayesian framework

Bayesian theory can be used to estimate the value of an unknown parameter and to quantify its uncertainty.<sup>7</sup> It provides a rigorous framework for combining prior information with observations to calculate likely values. It provides a natural way of “smoothing” inverse and regression problems (Green and Silverman 1994). Suppose we wish to estimate the value of some (finite-dimensional) parameter  $\theta$ . Assume we have some prior information for  $\theta$  (for example, that it is positive or represents a smooth function), summarised by a prior density  $p(\theta)$  for  $\theta$ . And suppose we observe some noisy data  $V = \{V_t: t \in Y_n\}$  (which in our case usually represents the observed option prices)

$$V_t = f_t(\theta^*) + e_t$$

for all  $t \in Y_n$ , where  $\theta^*$  is the true parameter,  $e_t$  is some random noise,  $Y_n$  is an index set of size  $n$  and  $f_t(\theta)$  is the option price at time  $t$  given model parameter  $\theta$ . Note that this is a special case of Equation 21.4 with one observation per time  $t$ , ie,  $|I_t| = 1$  for all  $t$ . In what follows, by abuse of notation, the function  $p$  will depend upon its argument. Then  $p(V | \theta)$  is the probability of observing the data  $V$  given  $\theta$ , and is called the likelihood function.

Now, an application of Bayes’s Rule implies that the posterior density of  $\theta$  is given by

$$p(\theta | V) = \frac{p(V | \theta)p(\theta)}{p(V)} \tag{21.7}$$

where the normalising constant  $p(V)$  is given by

$$p(V) = \int p(V | \theta)p(\theta) d\theta$$

and it is assumed that  $\theta \in \mathbb{R}^M$  for finite  $M$ .

**Definition 21.3.** A function  $L: \mathbb{R}^{2M} \rightarrow \mathbb{R}$  is a loss function if and only if

$$\begin{aligned} L(\theta, \theta') &= 0 & \text{if } \theta' = \theta \\ L(\theta, \theta') &> 0 & \text{if } \theta' \neq \theta \end{aligned}$$

where  $\theta, \theta' \in \mathbb{R}^M$ .

**Definition 21.4.** Given data  $V$  and loss function  $L$ , a corresponding Bayes estimator  $\theta_L(V)$  is a value of  $\theta$  which minimises the expected loss with respect to the posterior, ie

$$\theta_L(V) = \operatorname{argmin}_{\theta'} \left\{ \int_{\Theta} L(\theta, \theta') p(\theta | V) d\theta \right\} \quad (21.8)$$

Note that the minimiser  $\theta_L(V)$  is not necessarily unique. However, Gupta and Reisinger (2007) show that, for a certain class of loss functions and suitable calibration options, the Bayesian estimator can be proved to be consistent; that is, as more data is observed and the estimator updated, the estimate converges to the true value.

It is worth remarking that, for particular combinations of prior and likelihood function, both Gaussian for example, and a 0-1 loss function (Lehmann and Casella 1998), the minimisation formulation of Equation 21.8 is equivalent to Equation 21.6. In this sense, the Bayesian approach can be seen as a reformatting of the regularisation framework presented in the previous section (Fitzpatrick 1991).

**Example 21.5.** Consider claims  $C_i$  with corresponding observations

$$V^{(i)} = \frac{1}{2}(V^{(i) \text{ bid}} + V^{(i) \text{ ask}})$$

pricing functions

$$f^{(i)}(\theta) = \mathbb{E}^{\mathbb{Q}_\theta}[C_i]$$

noises

$$e^{(i)} \sim N\left(0, \frac{S_0^2}{10^8} \delta_i^2\right)$$

and weights

$$w_i = \left[\frac{1}{4}|V^{(i) \text{ bid}} - V^{(i) \text{ ask}}|^2\right]^{-1}$$

The likelihood function is then fixed. For the prior we use a Gaussian density  $\exp\{-\frac{1}{2}\lambda\|\theta\|^2\}$ , where  $\|\cdot\|$  is some norm of the finite-dimensional parameter  $\theta$  which summarises model  $\theta$  (with abuse of notation to simplify it). The posterior equation (Equation 21.7) then becomes

$$\begin{aligned} p(\theta | V) &\propto \exp\{-\frac{1}{2}\lambda\|\theta\|^2\} \\ &\times k \exp\left\{-\frac{1}{2} \frac{10^8}{\delta^2 S_0^2} \sum_{i \in I} \frac{4|\mathbb{E}^{\mathbb{Q}_\theta}[C_i] - V^{(i)}|^2}{|V^{(i)} \text{ bid} - V^{(i)} \text{ ask}|^2}\right\} \\ &= k \exp\left\{-\frac{1}{2} \left[ \frac{10^8}{\delta^2 S_0^2} \sum_{i \in I} \frac{4|\mathbb{E}^{\mathbb{Q}_\theta}[C_i] - V^{(i)}|^2}{|V^{(i)} \text{ bid} - V^{(i)} \text{ ask}|^2} + \lambda\|\theta\|^2 \right]\right\} \end{aligned}$$

where  $k$  is a normalising constant,  $\delta^2 = \sum_i w_i \delta_i^2$  and  $\lambda$  is a pre-defined constant indicating how strongly we believe in our prior assumptions. Observe that the Bayesian prior takes the role of the regularisation term and gives an interpretation for the regularisation parameter  $\lambda$  in Equation 21.6. Under the Bayesian framework,  $\lambda$  is viewed as the confidence parameter, ie, the strength of our belief in the prior assumptions.

### Bayesian pricing and hedging

In the previous section, we described how to find the Bayesian posterior  $p(\theta | V)$  and Bayesian estimator  $\theta_L(V)$  for the unknown parameter  $\theta$ . We now consider how these two quantities can be used to price and hedge contracts. There are three obvious approaches that can be taken, as detailed below: a naive Bayesian method, a partial Bayesian method and a full Bayesian method. Although the first is most commonly used in practice, and the second is conceptually simpler, the third method better uses the full power of the Bayesian approach and gives robust results. To clarify the three methods we reference the Black–Scholes delta hedge, but a similar method can be used for any hedge parameter.

- (i) Naive Bayesian method: only the Bayesian estimator  $\theta_L(V)$  is used. Examples of this method are least squares or the maximum *a posteriori* (MAP) estimator (Coleman *et al* 2001). The price of a different contract  $X$  with payoff  $h^X$  on  $S$  is taken to be

$$f_t^X(\theta_L(V)) = \mathbb{E}^{\mathbb{Q}}[B^{-1}(t, T)h^X(S(\theta_L(V))) | (S_u)_{0 \leq u \leq t}] \quad (21.9)$$

Similarly, the Black–Scholes delta hedge for  $X$  at time  $t$  is taken to be

$$\Delta(\theta_L(V)) = \frac{\partial f_t^X(\theta_L(V))}{\partial S_t}$$

This is a naive approach because the full information of the problem, as captured by the posterior distribution  $p(\theta | V)$ , is ignored.

- (ii) Partial Bayesian method (Branger and Schlag 2004): this method uses the full Bayesian posterior  $p(\theta | V)$  to average over the prices and deltas, so instead of finding the price in the Bayesian average model, it finds the Bayesian average model price and delta. More specifically, the price of contract  $X$  is given by

$$\int_{\theta} f_t^X(\theta) p(\theta | V) d\theta \quad (21.10)$$

and the hedge is given by

$$\int_{\theta} \frac{\partial f_t^X(\theta)}{\partial S_t} p(\theta | V) d\theta$$

This is only a partial approach because the price should correspond to the strategy which most closely hedges the contract  $X$  and there is no guarantee or intuition for why the above hedge should do this.

- (iii) Full Bayesian method (Gupta and Reisinger 2009): this method uses the precise formulation of Equation 21.8 with the posterior  $p(\theta | V)$  and for a suitable choice of loss function  $L$ . Suppose that  $\tilde{L}(\theta, \theta')$  corresponds to a measure of the hedging error caused by hedging contract  $X$  using parameter  $\theta'$  when the correct hedge is found using parameter  $\theta$ . Then

$$\theta_L(V) = \operatorname{argmin}_{\theta'} \left\{ \int_{\theta} \tilde{L}(\theta, \theta') p(\theta | V) d\theta \right\} \quad (21.11)$$

gives the optimal parameter to use for hedging and pricing  $X$ . In particular, we would take

$$f_t^X(\theta_L(V)) = \mathbb{E}^{\mathbb{Q}}[B^{-1}(t, T)h^X(S(\theta_L(V))) | (S_u)_{0 \leq u \leq t}]$$

for the price of the contract  $X$  and

$$\Delta_t(\theta_L(V)) = \frac{\partial f_t^X(\theta_L(V))}{\partial S_t}$$


---

as the delta hedge ratio. In this way we use the full information of the Bayesian posterior and the power of the loss function. Moreover, the hedge and price we use actually corresponds to a calibrated model, unlike for the previous method.

### **Advantages and disadvantages of Bayesian approaches**

For solving inverse problems in derivative pricing, the Bayesian framework offers some advantages over the regularisation method introduced earlier. Point estimates  $\theta_L(V)$  are useful, but of limited use without some measure of their correctness. The Bayesian approach offers a formal and consistent way to attach confidence to estimates. Equally, the approach provides a rigorous way to incorporate all available information regarding the unknown parameter, clearly differentiating between the *a priori* and observed information.

With special choices for the prior and likelihood, we can actually recover the regularisation operator in Equation 21.6 and the MAP estimator is equivalent to the solution of Equation 21.6. However, the advantage of the Bayesian approach is that we also discover a natural value for the regularisation parameter  $\lambda$ . As remarked earlier (see page 643 onwards),  $\lambda$  can be thought of as the confidence in the prior beliefs. This is important because in the regularisation method  $\lambda$  is often found through trial and error. The choice of stabilising term is often *ad hoc* or non-rigorous and therefore unsatisfactory. In the Bayesian framework, however, each term is meaningful and non-arbitrary.

Opponents of the Bayesian approach to data analysis often argue that it is fundamentally wrong to treat an unknown model parameter as a random variable and attach a distribution to it. They argue that the model parameter is unknown but not random. However, in some cases it is as important to be able to measure the uncertainty of a model parameter as it is to find the model parameter. One method of measuring the potential error is precisely to put a distribution on the model parameter and regard it as a random variable.<sup>8</sup> A second argument against the use of Bayesian theory is that the prior is inappropriate and meaningless, that scientists should not analyse data with any preconceptions or bias. However, in the mathematics of this chapter, the prior is a neat method of formally incorporating underlying assumptions. For example, no-arbitrage assumptions can be incorporated into the prior by attaching zero prior probability to

parameters which introduce arbitrage opportunities for calibration instruments.

Other, more practically minded, opponents of the Bayesian methodology sometimes argue that the assignment of probabilities to different parameters is too arbitrary, subjective and difficult. For example, Cont (2006) argues that assigning weights to models “requires too much probabilistic sophistication on the part of the end user”. However, the view here is that, because the calibration problem is ill-posed, we must draw on additional information not reflected in the prices of calibration instruments, and a prior naturally and unavoidably arises. Whether we choose to call the regularisation adjustment a roughness penalty function or smoothing term or prior is, in the opinion of the authors, a preference more of terminology than philosophy. Moreover, it is not even important that the prior should be very accurate or very carefully deliberated over; in typical option-pricing problems, the choice of a particular prior has less impact on the result. Indeed, if the estimator is updated by new observations, Gupta and Reisinger (2007) show that this estimator is consistent.

Given a Bayesian posterior distribution has been found, a variety of useful analyses can be performed:

- **Credible sets** (also known as confidence intervals) can be generated by finding sets of the parameter space which capture a certain proportion of the distribution. For example, if  $\theta$  is scalar, then taking an interval holding 95% of the distribution, with 2.5% in each tail, gives a centred confidence interval for the unknown parameter.
- **Marginal distributions** of a component of  $\theta$  can be found by integrating the joint posterior with respect to the other components. Viewing the marginal distribution of each component is useful in understanding how sensitive the joint posterior is to each of the components of  $\theta$  and also how much each component can vary.
- **Inferences** can be made about another quantity of interest,  $W$  say, that is a function of  $\theta$ . The spread of  $W$  can be measured and hence the errors associated with using a single point estimate for  $\theta$  can be calculated.



With respect to the third item, we can make inferences regarding the model uncertainty of a claim; this is the subject of the following section.

### MODEL UNCERTAINTY MEASURES

Having identified how model uncertainty can arise in the parameter-estimation problem, we now study measures of valuing this uncertainty. Risk measures are used in practice to determine the amount of capital to be held in reserve to make a risky position acceptable. Market risk measures like value-at-risk (VaR) are constructed on the implicit premise that a model for the market has been identified, and a risk measure  $\rho(X)$  for a contract (or future net worth)  $X$  is calculated within this model. Further examples are coherent and convex measures as introduced in Artzner *et al* (2002), Frittelli and Gianin (2002) and Föllmer and Schied (2002), which form the motivation for the following.

In market risk, the random variable  $X$  is understood to be random through its dependence on a “state of nature”  $\omega \in \Omega$ , ie,  $X = X(\omega)$ . In the context of the previous section, this would indicate the dependence on the path realised by the standard Brownian motion  $Z$ . If the model and all its parameters are known, this determines the law of  $X$ . However, the philosophy of this chapter is to acknowledge that the model/parameter  $\theta \in \Theta$  is not known, and we make this explicit by writing  $X(\omega, \theta)$ , and  $\rho^{\theta}$  for the corresponding market risk measure under model  $\theta$ .

In this section, instead of referring to different model types and different model parameters, we simply refer to different models. This is to emphasise that the measures presented can be applied very generally to either competing model types and/or a fixed model type with competing parameters.

The approaches described in this section all propose ways of accounting for this model uncertainty, but differ in the way they aggregate market risk and model uncertainty into a combined risk measure, or conversely how they separate out a model uncertainty measure from an overall measure of risk. A further divide can be drawn between measures based on worst-case scenarios within the assumed set of models, and those incorporating distributional information on the probabilities imposed on models/parameters, inferred, eg, from Bayesian analysis.

When it comes to measuring the model uncertainty of a derivative contract, it appears both logically consistent and practically relevant to associate with  $X$  a hedging portfolio with a hedge parameter  $\Delta$ , and to then apply one of the measures of the next two subsections (see pages 650 and 651). Then  $X = X(\omega, \theta, \Delta)$ , where  $\theta$  is the true parameter and  $\Delta$  is, for example, determined by hedging according to a model with parameter  $\theta'$ , in which case we write  $X = X(\omega, \theta, \theta')$  for clarity. A simple example would be a European option which is priced and hedged under a Black–Scholes model with assumed volatility  $\sigma'$ , when the true volatility is actually  $\sigma$  (so in this problem the unknown parameter “ $\theta$ ” is the scalar Black–Scholes volatility parameter “ $\sigma$ ”).

The measures of the final three subsections of this section are aimed at measuring the model uncertainty reflected in the spread of derivative prices more directly. Taking a view on hedging, they implicitly make the assumption that the distribution of (model) option prices  $f_i(\theta)$ , derived from a (posterior) distribution for the unknown parameter  $\theta$ , is a good indicator for the model uncertainty present in the subsequent hedging strategy. This will be justified for vanilla options but can underestimate the uncertainty for more exotic derivatives.

### Risk-averaging measures

Branger and Schlag (2004) consider Bayesian market risk measures which are close in spirit to the philosophy of this book. For a set  $\Theta$  of candidate models  $\theta$ , denote as above the probability of model  $\theta$  by  $p(\theta)$ .  $\mathbb{P}_\theta$  is the probability measure for the set of future scenarios corresponding to model  $\theta$  and  $\rho^{\mathbb{P}_\theta}(X)$  is the market risk measure of contract  $X$  under  $\mathbb{P}_\theta$ . Then Branger and Schlag define two different Bayesian methods of integrating market and model risk: model integration and risk integration.

In the first method, model integration, Branger and Schlag (2004) define the weighted market measure  $\mathbb{P}$  and the consequent market risk measure as follows

$$\rho(X) = \rho^{\mathbb{P}}(X), \quad \text{where } \mathbb{P} = \sum_{\theta \in \Theta} p(\theta) \mathbb{P}_\theta \quad (21.12)$$

Observe that there is a degree of symmetry in the above expression:  $\mathbb{E}^{\mathbb{P}}[X]$  can be viewed as a double sum (or double integral in

the infinite model and scenario case) over the different models and scenarios.

For the second method, risk integration, Branger and Schlag (2004) define the weighted market risk measure by

$$\rho(X) = \sum_{\theta \in \Theta} p(\theta) \phi(\rho^{\mathbb{P}_\theta}(X)) \quad (21.13)$$

for some model risk aversion function  $\phi$ .  $\phi$  is increasing and taken as convex if the decision maker is model risk averse, linear if the decision maker is model risk neutral and concave if the decision maker is model risk preferring.  $\phi(x) = x^n$  for different  $n \geq 1$  are proposed as possible convex functions. Note that Branger and Schlag (2004) measure market risk and model risk together, whereas an agent might find it useful to have a value for each separately.

### Risk-differencing measures

Kerkhof *et al* (2002) look to quantify model uncertainty with a view to determining how much regulatory capital should be set aside. They specify model uncertainty  $\mu$  as the difference between the worst-case market risk measure  $\rho$  and some reference market risk measure corresponding to reference model  $\alpha \in \Theta$

$$\mu(X) = \sup_{\theta \in \Theta} \rho^{\mathbb{P}_\theta}(X) - \rho^{\mathbb{P}_\alpha}(X) \quad (21.14)$$

where each model  $\theta$  in  $\Theta$  corresponds to measure  $\mathbb{P}_\theta$  and so gives a different market risk  $\rho^{\mathbb{P}_\theta}(X)$  for claim  $X$ .

This model uncertainty measure is interpretable as a conservative premium to be allocated in addition to the market risk measure in the assumed model  $\alpha$ , to account for uncertainty of the true model. For the actual form of  $\rho$ , Kerkhof *et al* (2002) suggest a number of alternatives: VaR, the coherent market risk measures introduced by Artzner *et al* (2002), worst conditional expectation and tail conditional expectation.

### Worst-case measures

Suppose we observe claims  $C_i$ , with corresponding observable bid-ask spreads  $[V^{(i) \text{ bid}}, V^{(i) \text{ ask}}]$  for  $i \in I$ , that we use as a calibration set and a set of models  $\Theta$ . As before, let  $\mathbb{Q}_\theta$  represent the risk-neutral probability measure for asset price process  $S$  corresponding to the model  $\theta$  for  $S$ . Now assume that

$$\text{for all } \theta \in \Theta, \quad \mathbb{E}^{\mathbb{Q}_\theta}[C_i] \in [V^{(i) \text{ bid}}, V^{(i) \text{ ask}}] \quad \text{for all } i \in I \quad (21.15)$$

ie, all measures  $\theta \in \Theta$  reproduce benchmark options to within their bid–ask spreads.

Let  $\mathcal{X} = \{X: \forall \theta \in \Theta, \mathbb{E}^{\mathbb{Q}_\theta}[|X|] < \infty\}$  be the set of all contingent claims that have a well-defined price in every model. Define  $\Phi$  to be the set of admissible trading strategies  $\phi$  such that

$$\mathbb{E}^{\mathbb{Q}_\theta} \left[ \int_0^t \phi_u dS_u \right] = 0$$

for all  $\theta \in \Theta$ . For simplicity, we assume the risk-free rate of growth is zero, so there is no discounting.

Cont (2006) defines a function  $\mu: \mathcal{X} \rightarrow [0, \infty)$  to be a model uncertainty measure if it satisfies Equation 21.15 and the following four axioms.

- A1. For benchmark options, the model uncertainty is no greater than the uncertainty of the market price

$$\text{for all } i \in I, \quad \mu(C_i) \leq |V^{(i) \text{ bid}} - V^{(i) \text{ ask}}|$$

- A2. Model-dependent dynamic hedging with the underlying does not reduce model uncertainty, since the hedge is model dependent

$$\text{for all } \phi \in \Phi, \quad \mu\left(X + \int_0^T \phi_t dS_t\right) = \mu(X)$$

But if the value of a claim can be totally replicated in a model-free way using only the underlying, then the claim has zero model uncertainty

if  $\exists x \in \mathbb{R}, \phi \in \Phi$  such that, for all  $\theta \in \Theta$ ,

$$X = x + \int_0^T \phi_t dS_t, \quad \theta\text{-a.s.}, \quad \text{then } \mu(X) = 0$$

- A3. Diversification does not increase the model uncertainty of a portfolio

for all  $X_1, X_2 \in \mathcal{X}$  and for all  $\lambda \in [0, 1]$ ,

$$\mu(\lambda X_1 + (1 - \lambda)X_2) \leq \lambda \mu(X_1) + (1 - \lambda) \mu(X_2)$$

- A4. Static hedging of a claim with traded options is bounded by the sum of the model uncertainty of that claim and the uncertainty
-

in the cost of replication

for all  $X \in \mathcal{X}$  and for all  $a \in \mathbb{R}^d$ ,

$$\mu\left(X + \sum_{i=1}^d a_i C_i\right) \leq \mu(X) + \sum_{i=1}^d |a_i| |V^{(i) \text{ bid}} - V^{(i) \text{ ask}}|$$

Cont (2006) shows that the function

$$\mu_0(X) = \sup_{\theta \in \Theta} \{\mathbb{E}^{\mathbb{Q}_\theta}[X]\} - \inf_{\theta \in \Theta} \{\mathbb{E}^{\mathbb{Q}_\theta}[X]\} \quad (21.16)$$

is a measure of model uncertainty, ie, it satisfies Equation 21.15 and the axioms (A1)–(A4). The measure finds the difference between the highest and lowest prices in  $\Theta$ . It is called the “worst-case” measure because it finds the largest difference amongst the collection of prices  $\mathbb{E}^{\mathbb{Q}_\theta}[X]$  for contract  $X$ .

Cont (2006) generalises the above to the case when not all the models  $\theta$  satisfy Equation 21.15, and instead assumes only that there exists at least one model  $\theta$  that satisfies Equation 21.15. Under subtle modification of the axioms (A1)–(A4), Cont proposes the function

$$\mu_0^*(X) = \sup_{\theta \in \Theta} \{\mathbb{E}^{\mathbb{Q}_\theta}[X] - \alpha_0(\theta)\} - \inf_{\theta \in \Theta} \{\mathbb{E}^{\mathbb{Q}_\theta}[X] + \alpha_0(\theta)\} \quad (21.17)$$

with the convex penalty functional  $\alpha_0$  defined by

$$\alpha_0(\theta) = \|\mathbb{E}^{\mathbb{Q}_\theta}[C] \triangleright V\|$$

where  $\|\cdot\|$  is a vector norm on  $\mathbb{R}^{|I|}$  and

$$(\mathbb{E}^{\mathbb{Q}_\theta}[C] \triangleright V)_i = \max\{V^{(i) \text{ bid}} - \mathbb{E}^{\mathbb{Q}_\theta}[C_i], \mathbb{E}^{\mathbb{Q}_\theta}[C_i] - V^{(i) \text{ ask}}, 0\}$$

as a model uncertainty measure. It is a “penalised worst-case” measure because it finds the largest difference amongst the collection of penalised prices  $\mathbb{E}^{\mathbb{Q}_\theta}[X] - \alpha_0(\theta)$  for contract  $X$ . Note the penalisation  $\alpha_0(\theta)$  reflects the calibration error, ie, the difference between market prices  $V$  and corresponding model prices  $\mathbb{E}^{\mathbb{Q}_\theta}[C]$  of the benchmark options  $C$ .

### Coherent measures

Motivated by the coherent model uncertainty measures introduced by Cont (2006), Gupta and Reisinger (2008) look at defining coherent model uncertainty measures which are more in spirit with the coherent market risk measures introduced by Artzner *et al* (2002).

Gupta and Reisinger measure the distribution of  $-2|\mathbb{E}^{\mathbb{Q}_\theta}[X] - x|$  in  $\Theta$  (we have already taken the expectation over the scenarios  $\omega$ ) for some fixed point  $x$ . The authors cite the following properties that should be expected of model uncertainty measures.

- Function of spreads: the “spread” of a claim  $X$  is the set of the prices for  $X$  found by all the different models.
- Monotonicity: if the spread of prices for  $Y$  is greater than that for  $X$ , then the model uncertainty for  $Y$  should be greater than for  $X$ .
- Sub-additivity: claims  $X$  and  $Y$  have a combined spread less than or equal to the sum of the individual spreads so the uncertainty measure should reflect this.
- Homogeneity: the spread for  $X$  should scale linearly with the number of claims  $X$ .

An example of a coherent measure is the “average-value” coherent measure given by

$$\mu_1(X) = E^Q[2|\mathbb{E}^{\mathbb{Q}_\theta}[X] - M^Q[\mathbb{E}^{\mathbb{Q}_\theta}[X]]|]$$

where  $M^Q[\mathbb{E}^{\mathbb{Q}_\theta}[X]]$  is the median value of  $\mathbb{E}^{\mathbb{Q}_\theta}[X]$  with respect to some measure  $Q$  on the set of models  $\theta$  and the expectation  $E$  is also taken with respect to measure  $Q$  on  $\Theta$ .


### Convex measures

Motivated by the convex model uncertainty measures introduced by Cont (2006), Gupta and Reisinger (2008) construct a set of axioms for convex model uncertainty measures which more closely follow those introduced by Frittelli and Gianin (2002) for convex market risk measures. Gupta and Reisinger drop the assumption that all or any models  $\theta$  in the model set  $\Theta$  satisfy Equation 21.15. A variety of convex model uncertainty measures is likely to be far more applicable than coherent measures, since it is atypical to find a large set of perfectly calibrated models.

Suppose  $p(\theta | V)$  is as given in Example 21.5. Then one example of a convex measure is

$$\mu_\lambda^*(X) = \sup_{\theta \in \Theta} \{\mathbb{E}^{\mathbb{Q}_\theta}[X] - \alpha_\lambda(\delta_\theta)\} - \inf_{\theta \in \Theta} \{\mathbb{E}^{\mathbb{Q}_\theta}[X] + \alpha_\lambda(\delta_\theta)\}$$

**Figure 21.3** 600 local volatility surface samples from the Bayesian posterior



where the convex penalty functional  $\alpha$  is given by

$$\begin{aligned} \alpha_\lambda(\delta_\theta) &= \frac{S_0}{10^4} \left[ -2\delta^2 \log \frac{p(\theta | V)}{k} \right]^{1/2} \\ &= \left[ \sum_{i \in I} w_i |\mathbb{E}^\theta[C_i] - V^{(i)}|^2 + \lambda \|\theta\|^2 \right]^{1/2} \end{aligned}$$

Observe that the confidence parameter,  $\lambda$ , plays a crucial role in determining the size of the model uncertainty values for contract  $X$ .


**Worked example**

Each surface is a sample of the posterior so is equally opaque.  $t$ , time in years;  $\sigma$ , volatility;  $S$ , asset price.

We continue with the numerical example presented in Examples 21.1 and 21.2. Using the Markov chain Monte Carlo Metropolis sampling algorithm described in Gupta and Reisinger (2007), we can sample the Bayesian posterior  $p(\theta | V)$  constructed in Example 21.5, where the unknown parameter  $\theta$  represents a discretised local volatility surface. Figure 21.3 shows a sample of 600 local volatility surfaces calibrated to an average of 3bp. Note the variety of shapes of local volatility surfaces.

With these samples, we can construct a distribution for the prices of derivatives. Recall the three-month up-and-out barrier call option we priced in Table 21.2 with strike  $0.9S_0$  and barrier  $1.10S_0$ . The Bayesian posterior  $p(\theta | V)$  gives the distribution of prices plotted

**Figure 21.4** Distribution of prices for barrier option, found using Bayesian posterior



From the 600 surfaces plotted in Figure 21.3 we construct the probability density function of the price of the barrier option and display the associated MAP and Bayes price. The Bayes price is calculated by an approximation to Equation 21.10 using the 600 sample surfaces. Also shown are the minimum and maximum prices calculated from the 600 surfaces.

**Table 21.3** Different model uncertainty values for the barrier option

| Model uncertainty measure | Value | Percentage of Bayes price |
|---------------------------|-------|---------------------------|
| Penalised worst-case      | 0.80  | 10.6                      |
| Convex                    | 0.77  | 10.3                      |

The model uncertainty measures are calculated using the Bayesian posterior  $p(\theta | V)$ .

in Figure 21.4. In this graph we have shown the MAP price 7.62, computed using Equation 21.9 and the Bayesian price 7.47, computed by Equation 21.10.

Next, using this distribution of prices, we can compute different model uncertainty measures for derivatives as described above, and we give examples for the penalised worst-case and convex measures. We use the Bayesian posterior  $p(\theta | V)$  to construct these model uncertainty measures as follows.



1. Penalised worst-case measure: price the claim  $X$  in models  $\theta \in \Theta$  and take supremums and infimums. Note that  $\Theta$  is the set of parameters  $\theta$  with positive posterior density.
2. Convex measure: take the Bayesian posterior  $p(\theta | V)$  and select confidence parameter  $\lambda = 1$  in order to get an even trade-off between the penalisation from the prior and the likelihood function.

The results are shown in Table 21.3.

We observe from Figure 21.4 that both these values, 0.80 and 0.77, capture the size of the interval over which the majority of the probability density function is concentrated. The convex measure is simply a generalisation of the penalised worst-case and allows the decision maker to tune the measure using  $\lambda$  according to their confidence in prior beliefs.

## CONCLUSION

### Decision rules for regulators and risk-managers

In practice, a “yes–no” rule-of-thumb for judging whether a contract  $X$  has a model uncertainty value  $\mu(X)$  that is too high would be useful for regulators and risk-managers. The VaR measure for market risk is widely used to make the following kind of decision

$$\text{reject } X \text{ if } \text{VaR}_{0.01}(X) > c_\rho$$

ie, do not buy  $X$  if the smallest loss of  $X$  in the worst 1% of scenarios is greater than  $c_\rho$ . Thus we could identically construct a yes–no rule for model uncertainty measure  $\mu$ , such as

$$\text{reject } X \text{ if } \mu(X) > c_\mu$$

ie, do not buy  $X$  if the model uncertainty value  $\mu(X)$  is greater than  $c_\mu$ . The value of  $c_\mu$  should depend on the investor’s risk preferences or there could be industry-standard values set by regulators. For example,  $c_\mu$  could be taken as 10% of the time-0 price  $V_0^X$  of  $X$ . The obvious generalisation is to aggregate market risk and model uncertainty with a combined rule such as

$$\text{reject } X \text{ if } \text{VaR}_{0.01}(X) + \mu(X) > c$$

for some  $c$ . Then note that, if  $\text{VaR}_{0.01}(X) + \mu(X) < 0$ , we would always buy  $X$ , since positive returns are made in all combinations of market scenarios and models.

### Summary

At the start of this chapter the distinction between (market) risk and (model) uncertainty was clarified. Possible sources of model uncertainty were detailed and explained. The model selection problem for derivative pricing was formally presented, first in the case of model-free pricing and then in the case of parameter estimation. For the second case, a robust Bayesian solution was detailed. We next studied five classes of model uncertainty measures. An example using a discretised local volatility surface was given to demonstrate the use of model uncertainty measures and the consequence on pricing derivatives. Finally, decision rules were proposed for risk-managers and regulators.

### Recommended further reading

We briefly highlight a few publications of interest related to pricing and hedging under model uncertainty, measures of model uncertainty, Bayesian approaches and the topics presented in this chapter.

The application of Bayesian theory to calibration problems in mathematical finance, although not a novel idea, is something that has only gathered weight since the early 1990s; Jacquier *et al* (1994) showed that Bayes estimators for a particular class of stochastic volatility models outperform the widely used method of moments and quasi-maximum likelihood estimators. More recently, Bhar *et al* (2006) and Para and Reisinger (2007) have considered dynamic Bayesian approaches to calibrating instantaneous spot and forward interest rates respectively.

Recently, attention has turned to using the Bayesian framework to examine the implications of parameter uncertainty in financial models. Jobert *et al* (2006) consider a Bayesian approach to explain the consistently large observed excess return earned by risky securities over the return on Treasury bills. They argue that, by dropping the assumption that the parameters of the dividend process are known to an agent but instead the agent only has some prior beliefs of these parameters, the excess rates of return are a natural consequence. Similarly, Monoyios (2007) examines the effects of drift parameter uncertainty in an incomplete market in which claims on non-traded assets are optimally hedged by a correlated traded asset. Using Bayesian learning, Monoyios (2007) concludes that terminal hedging errors

are often very large. Jacquier and Jarrow (2000) look at the effect on parameter uncertainty and model error in the Black–Scholes framework. They use Bayes estimators to infer values for option prices and hedge ratios and assess non-normality of the posterior distributions.

Closer to the example of the local volatility model used in this chapter are the works by Darsinos and Satchell (2007a,b). Darsinos and Satchell (2007a) formulate a joint prior for the asset price  $S_t$  and the Black–Scholes implied constant volatility  $\sigma$  using historical log returns of the asset price. The prior is updated using newly observed returns to give the posterior. The posterior is then transformed to a function of the asset price  $S_t$  and Black–Scholes European call price  $c$  and marginalised to give the probability density function for the option price  $c$ . Darsinos and Satchell (2007b) use this density to forecast European call option prices one day ahead and numerical experiments show substantial improvement to benchmark mean implied volatility procedures, especially in terms of hedging profits.

Green and Figlewski (1999) conduct an empirical study into the market and model risk exposures faced by an agent trading European calls and puts. They consider different volatility forecasting methods based on historical data and apply the methods to four underlyings: S&P 500 Index, three-month US dollar Libor, 10-year Treasury yield and the Deutschmark Exchange Rate. Their first finding is that the strategy of writing and holding option positions without hedging produces very large risk exposures, even over long horizons, and diversification does not significantly reduce this risk exposure. After daily delta re-hedging was added to the portfolios, Green and Figlewski (1999) found that the standard deviation, mean and worst-case returns were all reduced. However, worst-case losses were still several times the initial premium, particularly for out-of-the-money contracts. They conclude that writing options with volatility mark-ups (of up to 50%) turns a very risky trading strategy into a profitable one. By writing an option with a volatility mark-up we mean that the value of the volatility used in the calculation of the price is greater (ie, “marked up”) than the volatility actually estimated from the data. This gives a price greater than would have been found with the original estimated volatility and is thus a safer price for an agent to sell the option for. This finding, Green

and Figlewski conclude, indicates that the model risk from mis-estimating volatility in trading and hedging derivatives positions is very large.

In contrast, Hull and Suo (2002) look at the model risk from mis-specification of the model rather than mis-estimation. They consider the pricing errors arising from a continually recalibrated local volatility model. They price a compound option, a European call option on a European call option and a barrier option. They find that the continually recalibrated local volatility model always correctly prices European style options, where the payoff is contingent on the asset price at just one time. However, for exotic options dependent on the distribution of the asset price at two or more times, the model can perform badly.

Hull and Suo (2002) argue that this failure of the local volatility surface is to be expected. They explain that the local volatility model is designed to match European options correctly but not options dependent on the value of the underlying asset at multiple times. Let  $\phi_n(t_1, \dots, t_n)$  be the joint probability distribution of the asset price at times  $t_1, \dots, t_n$  and  $\phi_1(t_1), \dots, \phi_n(t_n)$  be the marginal distributions of the asset price at times  $t_1, \dots, t_n$ , respectively. Then Hull and Suo (2002) point out that the local volatility model is designed so that all the marginals  $\phi_1(t_1), \dots, \phi_n(t_n)$  are correct but in no way correctly reproduce  $\phi_n(t_1, \dots, t_n)$  or any other joint probability distribution. And this is fundamental to why different local volatility surfaces can be fitted to the same calibration prices (marginals) as we saw in Chapter 20, but why these surfaces give very different prices for exotic and path dependent options (joint distributions). This point is further clarified by Britten-Jones and Neuberger (2000), who show how very different volatility processes can be adjusted to fit the same observed option prices exactly; hence the prevalence of high model uncertainty.

Contreras and Satchell (2003) use a Bayesian approach to construct confidence intervals for the VaR measure. They design priors for the mean  $\mu$  and standard deviation  $\sigma$  of  $\text{VaR}(X)$  for some claim  $X$ , and update these statistics using the observed data. However, because VaR is not sub-additive or convex, it can lead to anomalous values for a portfolio of options (Artzner *et al* 2002). For example, we can easily find two options such that the VaR of the portfolio of two options is greater than the sum of the individual VaRs.

The authors are grateful for financial support from the UK Engineering and Physical Sciences Research Council (EPSRC).

- 1 See also the discussion in Chapter 30.
- 2 See *NA Digest* v.09 n.48 (<http://www.netlib.org/na-digest-html/09/v09n48.html>) for details.
- 3 See, for example, Engl *et al* (1996).
- 4 See, for example Engl *et al* (1996) and Tikhonov *et al* (1977).
- 5 See Tikhonov *et al* (1977) for details.
- 6 Bouchouev and Isakov (1999) offer a good overview of applications in financial markets.
- 7 An introduction into Bayesian statistics in general and its computational methods is given in Chapters 1 and 2; other valuable references are Lehmann and Casella (1998) and Gelman *et al* (2004).
- 8 See Lindley (2006) for a very readable introduction to uncertainty.

#### REFERENCES

- Achdou, Y., 2005, "An Inverse Problem for Parabolic Variational Inequalities in the Calibration of American Options", *SIAM Journal of Control* 43(5), pp. 1583–1615.
- Artzner, P., F. Delbaen, J. M. Eber and D. Heath, 2002, "Coherent Measures of Risk", in *Risk Management: Value at Risk and Beyond*, p. 145 (Cambridge University Press).
- Avellaneda, M., A. Lévy and A. Paras, 1995, "Pricing and Hedging Derivative Securities in Markets with Uncertain Volatilities", *Applied Mathematical Finance* 2, pp. 73–88.
- Berestycki, H., J. Busca and I. Florent, 2002, "Asymptotics and Calibration of Local Volatility Models", *Quantitative Finance* 2(1), pp. 61–9.
- Bhar, R., C. Chiarella, H. Hung and W. J. Runggaldier, 2006, "The Volatility of the Instantaneous Spot Interest Rate Implied by Arbitrage Pricing: A Dynamic Bayesian Approach", *Automatica* 42(8), pp. 1381–93.
- Bouchouev, I., and V. Isakov, 1999 "Uniqueness, Stability and Numerical Methods for the Inverse Problem that Arises in Financial Markets", *Inverse Problems* 15, pp. 95–116.
- Branger, N., and C. Schlag, 2004, "Model Risk: A Conceptual Framework for Risk Measurement and Hedging", Technical Report, Goethe University.
- Britten-Jones, M., and A. Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility", *The Journal of Finance* 55(2), pp. 839–66.
- Chance, D. M., 2008, "Rational Rules and Boundary Conditions for Option Pricing", Teaching Note 99-05.
- Coleman, T. F., Y. Li and A. Verma, 2001, "Reconstructing the Unknown Local Volatility Function", in *Quantitative Analysis in Financial Markets: Collected Papers of the New York University Mathematical Finance Seminar*, p. 192 (World Scientific).
- Cont, R., 2006, "Model Uncertainty and Its Impact on the Pricing of Derivative Instruments", *Mathematical Finance* 16(3), pp. 519–47.
- Cont, R., and P. Tankov, 2004, *Financial Modelling with Jump Processes* (London: Chapman & Hall/CRC Press).
- Contreras, P., and S. E. Satchell, 2003, "A Bayesian Confidence Interval for Value-at-Risk", Cambridge Working Papers in Economics.

- Cox, A. M. G., and J. K. Obłój, 2008, "Robust Hedging of Double Touch Barrier Options", Arxiv Preprint arXiv:0808.4012.
- Cox, A. M. G., and J. K. Obłój, 2009, "Robust Pricing and Hedging of Double No-Touch Options", Arxiv Preprint arXiv:0901.0674.
- Crepey, S., 2003, "Calibration of the Local Volatility in a Trinomial Tree Using Tikhonov Regularization", *Inverse Problems* 19(1), pp. 91–127.
- Cvitanic, J., H. Pham and N. Touzi, 1999, "Super-Replication in Stochastic Volatility Models under Portfolio Constraints", *Journal of Applied Probability* 36(2), pp. 523–45.
- Darsinos, T., and S. Satchell, 2007a, "Bayesian Analysis of the Black–Scholes Option Price", in S. Satchell (ed), *Forecasting Expected Returns in the Financial Markets*, p. 117 (Elsevier).
- Darsinos, T., and S. Satchell, 2007b, "Bayesian Forecasting of Options Prices: Natural Framework for Pooling Historical and Implied Volatility Information", in S. Satchell (ed), *Forecasting Expected Returns in the Financial Markets*, p. 151 (Elsevier).
- Derman, E., 1996, "Model Risk", *Risk Magazine* 9(5), pp. 34–7.
- Dupire, B., 1994, "Pricing with a Smile", *Risk Magazine* 7(1), pp. 18–20.
- Egger, H., and H. W. Engl, 2005, "Tikhonov Regularization Applied to the Inverse Problem of Option Pricing: Convergence Analysis and Rates", *Inverse Problems* 21(3), pp. 1027–45.
- Engl, H. W., M. Hanke and A. Neubauer, 1996, *Regularization of Inverse Problems* (Kluwer Academic).
- Fitzpatrick, B. G., 1991, "Bayesian Analysis in Inverse Problems", *Inverse Problems* 7(5), pp. 675–702.
- Föllmer, H., and A. Schied, 2002, "Convex Measures of Risk and Trading Constraints", *Finance and Stochastics* 6(4), pp. 429–47.
- Frey, R., 2000, "Superreplication in Stochastic Volatility Models and Optimal Stopping", *Finance and Stochastics* 4(2), pp. 161–87.
- Frey, R., and C. A. Sin, 2001, "Bounds on European Option Prices under Stochastic Volatility", *Mathematical Finance* 9(2), pp. 97–116.
- Frittelli, M., and E. R. Gianin, 2002, "Putting Order in Risk Measures", *Journal of Banking and Finance* 26(7), pp. 1473–86.
- Gelman, A., J. B. Carlin, H. S. Stern and D. B. Rubin, 2004, *Bayesian Data Analysis*, Second Edition (London: Chapman & Hall/CRC).
- Green, P. J., and B. W. Silverman, 1994, *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach* (London: Chapman & Hall/CRC).
- Green, T. C., and S. Figlewski, 1999, "Market Risk and Model Risk for a Financial Institution Writing Options", *The Journal of Finance*, pp. 1465–99.
- Gupta, A., and C. Reisinger, 2007, "Calibrating Financial Models Using a Consistent Bayesian Estimator", Preprint.
- Gupta, A., and C. Reisinger, 2008, "Model Uncertainty Measures in Finance Using Bayesian Posteriors", Working Paper.
- Gupta, A., and C. Reisinger, 2009, "Optimal Bayesian Hedging", Working Paper.
- 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(2), pp. 327–43.

- Hull, J., and W. Suo**, 2002, "A Methodology for Assessing Model Risk and Its Application to the Implied Volatility Function Model", *The Journal of Financial and Quantitative Analysis* 37(2), pp. 297–318.
- Jackson, N., E. Suli and S. Howison**, 1999, "Computation of Deterministic Volatility Surfaces", *Journal of Computational Finance* 2(2), pp. 5–32.
- Jacquier, E., and R. Jarrow**, 2000, "Bayesian Analysis of Contingent Claim Model Error", *Journal of Econometrics* 94, pp. 145–80.
- Jacquier, E., N. G. Polson and P. E. Rossi**, 1994, "Bayesian Analysis of Stochastic Volatility Models", *Journal of Business and Economic Statistics* 12(4), pp. 371–89.
- Jobert, A., A. Platania and L. C. G. Rogers**, 2006 "A Bayesian Solution to the Equity Premium Puzzle", Report, Statistical Laboratory, University of Cambridge.
- Kerkhof, F. L. J., B. Melenberg and J. M. Schumacher**, 2002, "Model Risk and Regulatory Capital", CentER Discussion Paper 2002-27, pp. 1–57.
- Knight, F. H.**, 1921, *Risk, Uncertainty and Profit* (New York: Houghton Mifflin).
- Lehmann, E. L., and G. Casella**, 1998, *Theory of Point Estimation* (Springer).
- Li, D. X.**, 2000, "On Default Correlation: A Copula Approach", *Journal of Fixed Income* 9(4), pp. 43–54.
- Lindley, D. V.**, 2006, *Understanding Uncertainty* (Chichester: John Wiley & Sons).
- Lyons, T. J.**, 1995, "Uncertain Volatility and the Risk-Free Synthesis of Derivatives", *Applied Mathematical Finance* 2, pp. 117–33.
- Monoyios, M.**, 2007, "Optimal Hedging and Parameter Uncertainty", *IMA Journal of Management Mathematics* 18(4), p. 331.
- Mykland, P. A.**, 2009, "Option Pricing Bounds and Statistical Uncertainty: Using Econometrics to Find an Exit Strategy in Derivatives Trading", Technical Report, University of Chicago.
- Para, H., and C. Reisinger**, 2007, "Calibration of Instantaneous Forward Rate Volatility in a Bayesian Framework", Working Paper.
- Swishchuk, A., A. Ware and H. Li**, 2008, "Option Pricing with Stochastic Volatility Using Fuzzy Sets Theory", Northern Finance Association 2008 Conference Papers.
- Tikhonov, A. N., V. Y. Arsenin and F. John**, 1977, *Solutions of Ill-Posed Problems* (Washington, DC: V.H. Winston).