Model uncertainty and its impact on the pricing of derivative instruments.

Rama CONT

Centre de Mathématiques Appliquées, Ecole Polytechnique

June 2004.*

Abstract

Model uncertainty, in the context of derivative pricing, results in mispricing of contingent claims due to uncertainty on the choice of the pricing models. We introduce here a quantitative framework for measuring model uncertainty in option pricing models. After discussing some properties which a quantitative measure of model uncertainty should verify in order to be useful and relevant in the context of risk management of derivative instruments, we propose two methods for measuring model uncertainty which verify these properties, yield numbers which are comparable to other risk measures and compatible with observations of market prices of a set of benchmark derivatives. The first method is based on a coherent risk measure compatible with market prices of derivatives, while the second method is based on a convex risk measure. Both methods lead to a decomposition of risk measures into a market value and a premium for model uncertainty. Finally, we discuss some implications for management of “model risk”.

*This project was supported by a research grant from the Europlace Institute of Finance. Part of this work was done in the framework of a research project on model uncertainty at HSBC-CCF, Division of Market and Model Risk. Previous versions of this work were presented at the Hermes Center (Nicosia), 3rd METU Economics Research Conference (2003) and the European conference on Arbitrage theory and applications (Paris, 2003). We thank Sana BenHamida, Joël Bessis, Jean-François Boulier, Rémi Bourrette, Larry Epstein, Walter Schachermayer and Franck Viollet for helpful discussions.

Contents

1	Introduction	3
2	Risk, uncertainty and ambiguity	5
2.1	Bayesian model averaging	6
2.2	Worst case approaches	7
2.3	Risk measures	8
3	Model uncertainty in the context of derivative valuation	10
3.1	Statistical uncertainty vs uncertainty on pricing rules	10
3.2	Benchmark instruments vs illiquid products	11
3.3	Some requirements for a measure of model uncertainty	11
4	A quantitative framework for measuring model uncertainty	13
4.1	An axiomatic setting for model uncertainty	13
4.2	A “coherent” measure of model uncertainty	15
4.3	Examples	16
5	A convex measure of model uncertainty	21
6	Discussion	25
6.1	Summary	26
6.2	Specifying the class of models	26
6.3	Updating with new information	27
6.4	Conclusion	28
A	Proof of Proposition 1	31
B	Proof of Proposition 2	32

1 Introduction

In March 1997, Bank of Tokyo/Mitsubishi announced that its New York-based derivatives unit had suffered a \$ 83 million loss because their internal pricing model overvalued a portfolio of swaps and options on U.S. interest rates. A few weeks later, NatWest Capital Markets announced a £50 million loss because of a mispriced portfolio of German and U.K. interest rate options and swaptions run by a single derivatives trader in London. According to observers having followed these events, many “of the situations [...] that led to (recent) derivatives losses were attributable to model risk” [Elliott (1997)].

With the dissemination of quantitative methods in risk management and advent of complex derivative products, mathematical models have come to play an increasingly important role in financial decision making, especially in the context of pricing and hedging of derivative instruments. While the use of models has undeniably led to a better understanding of market risks, it has in turn given rise to a new type of risk, known as “model risk” or “model uncertainty”, linked to the uncertainty on the choice of the model itself. According to a recent report [Williams (1999)], \$ 5 billion in derivatives losses in 1999 were attributable to “model risk”.

Uncertainty on the choice of the pricing model can lead to the mispricing of a derivative products. While model uncertainty is acknowledged by most operators who make use of quantitative models, most of the discussion on this subject has stayed at a qualitative level and a quantitative framework for measuring model uncertainty is still lacking. As noted by [Williams (1999)], “there are no packaged, off-the-shelf systems for model risk management”. Some questions for which one would like quantitative answers are:

- How sensitive is the value of a given derivative to the choice of the pricing model?
- Are some instruments more model-sensitive than others?
- How large is the model uncertainty of a portfolio compared with its market risk?
- Can one provision for “model risk” in the same way as one provisions for market risk and credit risk?

One could wonder whether model uncertainty deserves a separate treatment from other sources of uncertainty in financial markets. Indeed, the classical approach to decision under uncertainty [Savage] does not distinguish between different sources of risk: “model uncertainty” should be indistinguishable from market risk, credit risk,...which would imply that “model uncertainty” simply amounts to weighting various models with probabilities and representing all sources of uncertainty using a probability distribution on the enlarged space comprising “models” + scenarios. Indeed, such “model averaging” approaches have been proposed in the Bayesian literature [Hoeting et al., 1999] (see Section 2). However, this approach is in strong contrast with the current practices in

risk management: as noted by [Routledge & Zin (2001)], market participants typically use different criteria to measure “market risk” and “model risk”, the former being valued by using a probabilistic model while the latter is approached through a worst case approach, for instance by stress testing of portfolios.

This had led to the distinction between *risk* -uncertainty on outcomes for which the probabilities are known- and *ambiguity* or model uncertainty-when several specifications are possible for such probabilities [Knight]. [Ellsberg] has shown that aversion to ambiguity clearly plays a role in decision making. A growing body of literature has explored decision under ambiguity, its axiomatic foundations [Gilboa & Schmeidler, Epstein] and implications for the behavior of security prices [Epstein & Wang, Routledge & Zin (2001)]. Some of these ideas have resurfaced in the recent literature on coherent risk measures [Artzner et al(1999)] and their extensions [Föllmer & Schied (2002)].

Although general in nature, these approaches do not take into account some specific features of the use of probabilistic models in the pricing of derivatives. The notion of coherent risk measure does not distinguish hedgeable and non-hedgeable risks, nor does it differentiate between market risk and model uncertainty. And, although coherent measures of risk are expressed in monetary units, when applied to traded options they may lead to numbers which are not necessarily comparable to the mark-to-market value of these options. Also, in the context of derivative pricing, models are often specified not in terms of objective probabilities but “risk-neutral” probabilities so, in incomplete markets, ambiguity can prevail on pricing criteria even when there is no ambiguity on the underlying price process itself. These remarks show that model uncertainty in option pricing cannot be reduced to the classical setting of decision under ambiguity and merits a specific treatment.

We attempt to address these issues by defining a framework for quantifying model uncertainty in option pricing models. We first discuss, at an intuitive level, some properties which a quantitative measure of model uncertainty should possess in order to qualify as a measure of model uncertainty in the context of risk measurement and management. We then propose two methods for measuring model uncertainty which verify these properties and are compatible with observations of market prices of a set of benchmark derivatives. Both methods lead to a decomposition of risk measures into a market value and a premium for model uncertainty.

The paper is structured as follows. We start by discussing some existing approaches to decision-making in presence of multiple probability measures in Section 2. Some specific features of the use of models in the valuation of derivative instruments are not taken into account in these general frameworks; these issues are discussed in section 3, where we give an intuitive definition of model uncertainty in the context of derivative markets and enumerate some properties a measure of model uncertainty must have in order to be meaningful for risk management of derivative instruments. In section 4 we formulate these requirements in mathematical terms and present a methodology for measuring model uncertainty which verifies these requirements. This method requires to specify a set of pricing models and calibrate them to a set of market option prices;

this requirement is relaxed in Section 5, where a more general approach based on convex risk measures is proposed. Section 6 concludes by summarizing our main contributions, discussing some open questions and pointing out possible implications of our work for the measurement and management of “model risk”. We have attempted to motivate the mathematical notions introduced through examples which illustrate their relevance.

2 Risk, uncertainty and ambiguity

In the standard framework in mathematical finance and, more generally, in decision theory, the starting point is the specification of a stochastic model: a set of future scenarios $(\Omega, \mathcal{F})$ and a probability measure $\mathbb{P}$ on these outcomes. However there are many circumstances, especially in financial decision making, where the decision maker or risk manager is not able to attribute a precise probability to future outcomes. This situation has been called “uncertainty” by [Knight], by contrast with “risk”, when we are able to specify a unique probability measure on future outcomes.¹ More precisely, we speak of *ambiguity* when we are facing several possible specifications $\mathbb{P}_1, \mathbb{P}_2, \dots$ for probabilities on future outcomes [Epstein].

In his 1961 thesis, [Ellsberg] established a distinction between aversion to risk—related to lack of knowledge of future outcomes— and aversion to ambiguity², related to the lack of knowledge of their probabilities, and showed that aversion to ambiguity can strongly affect decision makers behavior and resolve some paradoxes of classical decision theory. More recently, ambiguity aversion has shown to have important consequences in macroeconomics [Hansen et al., 1999, Hansen et al., 2002] and for price behavior in capital markets [Chen & Epstein, Epstein & Wang, Routledge & Zin (2001)].

Two different paradigms have been proposed for evaluating uncertain outcomes in presence of ambiguity. The first one, which consists of averaging over possible models, has been used in the statistical literature [Raftery (1993), Hoeting et al., 1999]. The other one is based on worst-case or “maxmin” approach and has been axiomatized by [Gilboa & Schmeidler] and studied in the context of asset pricing by [Epstein & Wang, Routledge & Zin (2001)] and others. Related to this worst-case approach is the recent literature on coherent measures of risk [Artzner et al(1999)]. We review in this section these approaches and their possible implications and shortcomings for quantifying model uncertainty for portfolios of derivatives.

¹This distinction appeared in [Knight], hence the term “Knightian uncertainty” sometimes used to designate the situation where probabilities are unknown. We remark here that the term “model risk” sometimes used in the financial literature is somewhat confusing in this respect and the term “model uncertainty” should be preferred.

²We use here the terminology of [Epstein].

2.1 Bayesian model averaging

A lot of attention has been devoted to model uncertainty in the context of statistical estimation, using a Bayesian approach. [Hoeting et al., 1999] note that “data analysts typically select a model from some class of models and then proceed as if the selected model had generated the data. This approach ignores the uncertainty in model selection, leading to over-confident inferences and decisions that are more risky than one thinks they are”. Bayesian model averaging is one way to incorporate model uncertainty into estimation procedures.

Let $\mathcal{M} = \{M_1, \dots, M_J\}$ be a family a of candidate models whose parameters (not necessarily in the same sets) are denoted by $\theta_1 \in E_1, \dots, \theta_J \in E_J$. Consider a Bayesian observer with two levels of prior beliefs:

- Priors on model parameters: $p(\theta_j|M_j)$ is a prior density (on E_j) that summarizes our views about the unknown parameters of model j , given that M_j holds.
- Prior “model weights”: $\mathbb{P}(M_j), j = 1..J$, the prior probability that M_j is the “true” model.

Given a set of observations y , the posterior probability for the model M_j is given by

$$\mathbb{P}(M_j|y) = \frac{p(y|M_j)\mathbb{P}(M_j)}{\sum_{k=1}^J p(y|M_k)\mathbb{P}(M_k)} \quad (1)$$

where $p(y|M_j)$ is the integrated likelihood of the data under model M_j :

$$p(y|M_j) = \int_{E_k} \mathbb{P}(y|\theta_j, M_j)p(\theta_j|M_j) d\theta_j \quad (2)$$

Suppose now we want to compute a model dependent quantity, given by the expectation of a random variable X : we only have the observations y but we are uncertain about the model to use. The Bayesian model averaging approach suggests to compute this quantity in each model and average over the models, weighting each model by its posterior probability given the observations:

$$\hat{E}[X|y] = \sum_{j=1}^m E[X|y, M_j] \mathbb{P}(M_j|y) \quad (3)$$

If M_j are alternative option pricing models, this would amount to computing option prices in each model and taking a weighted average across models. Similarly one can use the following quantity to measure dispersion across models:

$$\hat{D}[X|y] = \sum_{j=1}^m \{E[X|y, M_j] - \hat{E}[X|y]\}^2 \mathbb{P}(M_j|y) \quad (4)$$

Averaging across models, whether or not it is done in a Bayesian way, provides a higher stability of the estimates obtained. However, several obstacles appear when trying to apply this approach in the framework of option pricing.

First, this method not only requires specifying –as in any Bayesian method– a prior $p(\theta_j|M_j)$ on parameters of each model, but also a prior probability $\mathbb{P}(M_j)$ on possible models, which is more delicate. How does one weigh a stochastic volatility model with respect to a jump-diffusion model? How should prior weights vary with the number of factors in interest rate models? While such questions might be ultimately reasonable to ask, not much experience is available in assigning such prior weights.³ In other words, this approach requires too much probabilistic sophistication on the part of the end user.

The second obstacle is computational: the posterior distributions involved in the formulas above are not explicit and sampling from them requires the use of Markov Chain Monte Carlo algorithms, which are computationally intensive. Such an approach has been attempted in the case of Black-Scholes model [Jacquier & Jarrow (2000)] but seems less feasible as soon as we move to more complex models. It should also be noted that, because of these computational difficulties, the Bayesian model averaging literature deals with relatively simple model structures (linear and regression-type models).

Overall, the main justification for averaging over models is that it improves predictive ability [Hoeting et al., 1999] of some target quantity (say, an option price). However, the main concern of risk management is not to predict prices but to quantify the risk associated with them so model averaging seems less relevant in this context.

2.2 Worst case approaches

The model averaging procedure described above, whether or not it is done in a Bayesian way, is in fact consistent with the classical approach to decision under uncertainty [Savage], which does not distinguish between different sources of risk: in this approach, “model uncertainty” should be indistinguishable from market risk, credit risk,... “Model uncertainty” then simply amounts to weighting various models with probabilities and representing all sources of uncertainty using a probability distribution on the enlarged space comprising “models” + scenarios. However, this approach is in strong contrast with the current practices in risk management: market participants do not specify probabilistic beliefs over models and, as noted by [Routledge & Zin (2001)], typically use different criteria to measure “market risk” and “model risk”, the former being valued by risk neutral pricing (averaging across scenarios) while the latter is approached through a worst case approach, for instance by stress testing of portfolios.

Aside from being observed in the practice of risk managers, the worst-case approach also has a firm axiomatic foundation: [Gilboa & Schmeidler] propose a system of axioms under which an agent facing ambiguity chooses among a set A of feasible alternatives by maximizing a “robust” version of expected utility (also called “maxmin” expected utility), obtained by taking the worst case over all models:

$$\max_{X \in A} \min_{\mathbb{P} \in \mathcal{P}} E^{\mathbb{P}}[U(X)] \quad (5)$$

³In the statistical literature, uniform priors on models are often used.

Here the *risk* aversion of the decision maker is captured by the utility function U , while the aversion to ambiguity (model uncertainty) is captured by taking the infimum over all models in $\mathcal{P}$. It can be shown that the worst case approach actually corresponds to an “infinite aversion to model uncertainty” [Adam]. With respect to model averaging procedures described in section 2.1, worst case approaches are more conservative, more robust and require less sophisticated inputs on the part of the user. Thus, they are more amenable to the design of a robust, systematic approach for measuring model uncertainty.

2.3 Risk measures

Related to the worst case approach described above is the notion of *coherent risk measure*. A risk measurement methodology is a way of associating a number (“risk measure”) $\rho(X)$ with a random variable X , representing the payoff of an option, a structured product or a portfolio. More precisely, if we define a payoff as a (bounded measurable) function $X : \Omega \rightarrow \mathbb{R}$ defined on the set Ω of market scenarios and denote the set of payoffs as E , then a risk measure is a map $\rho : E \rightarrow \mathbb{R}$. [Artzner et al(1999)] enumerate a set of properties that ρ needs to possess in order to be useful as a measure of risk in a risk management context:

1. Monotonicity: if a portfolio X dominates another portfolio Y in terms of payoffs then it should be less risky: $X \geq Y \Rightarrow \rho(X) \leq \rho(Y)$.
2. Risk is measured in monetary units: adding to a portfolio X a sum a in numeraire reduces the risk by a : $\rho(X + a) = \rho(X) - a$.
3. Sub-additivity: this is the mathematical counterpart of the idea that diversification reduces risk.

$$\rho(X + Y) \leq \rho(X) + \rho(Y) \quad (6)$$

4. Positive homogeneity : the risk of a position is proportional to its size.

$$\forall \lambda > 0, \rho(\lambda X) = \lambda \rho(X) \quad (7)$$

A risk measure $\rho : E \rightarrow \mathbb{R}$ verifying these properties is called a coherent risk measure. [Artzner et al(1999)] show that any coherent measure of risk can be represented as the highest expected payoff in a family $\mathcal{P}$ of models:

$$\rho(X) = \sup_{\mathbb{P} \in \mathcal{P}} E^{\mathbb{P}}[-X] \quad (8)$$

Interestingly, this representation is a *result* of the “axioms” of risk measures: it shows that any coherent risk measure is representable as a worst case expected utility with a zero “risk aversion” (i.e. a linear “utility”) and an infinite aversion to ambiguity. It remains to specify the family $\mathcal{P}$ and different choices will yield different measures of risk. Many familiar examples of risk measures can be represented in this form [Artzner et al(1999), Föllmer & Schied (2002b)].

Coherent risk measures were generalized in [Föllmer & Schied (2002)] by relaxing the positive homogeneity hypothesis: if conditions (3) and (4) above are replaced by

$$\forall \lambda \in [0, 1], \quad \rho(\lambda X + (1 - \lambda)Y) \leq \rho(\lambda X) + \rho((1 - \lambda)Y) \quad (9)$$

then ρ is called a *convex risk measure*. Under an additional continuity condition, a convex risk measure can be represented as

$$\rho(X) = \sup_{\mathbb{P} \in \mathcal{P}} \{E^{\mathbb{P}}[-X] - \alpha(\mathbb{P})\} \quad (10)$$

where $\alpha(\mathbb{P})$ is a positive “penalty” function. Allowing α to take infinite values, one can always extend $\mathcal{P}$ to the set of all probability measures on $(\Omega, \mathcal{F})$. A coherent risk measure as defined above then corresponds to the special case where α only takes the values 0 or ∞ .

Several remarks can be made at this stage on the possible use of this approach for derivatives. First, since $\rho(X)$ is specified in monetary units, one can attempt to compare it to the market price of X if it is traded in the market. For example, $-\rho(X)$ and $\rho(-X)$ (risk of a short position in X) could be used to derive a price interval and be compared to the market bid-ask spread for the derivative. In fact some authors have used the term “risk-adjusted” value for $-\rho(X)$. However there is no ingredient in the axioms above guaranteeing that such a comparison will be meaningful. Indeed, the elements $\mathbb{P} \in \mathcal{P}$ represent alternative choices for the “objective” evolution of the market: they are not risk-neutral measures and the quantities $E^{\mathbb{P}}[X]$ should not be interpreted as “prices”. For example the “risk-adjusted value” of a forward contract on USD/EUR, which has a unique model-free valuation compatible with arbitrage constraints, is not equal in general to this arbitrage value. Therefore, while $\rho(X)$ is measured in monetary units it cannot be directly related to prices. In fact, it is obvious from the properties above that, if $\rho(\cdot)$ is a coherent (resp. convex) risk measure, then for *any* scaling factor $\lambda > 0$, $\lambda\rho$ is *also* a coherent (resp. convex) risk measure: there is no natural *normalization* involved in the definition.⁴ [Föllmer & Schied (2002)] propose an additive normalization for a convex risk measure ρ by setting $\rho(0) = 0$, in which case the representation (10) must be modified but this still leaves us the freedom of a multiplicative scaling.

Second, a coherent risk measure $\rho(\cdot)$ does not distinguish in general between hedgeable and unhedgeable risks. For example, $\rho(X)$ may be the same for a position in futures or for a path-dependent option whereas the risks involved in the case of the call option are of different nature: in one case they can be replicated in a model-free way by taking positions in the underlying whereas in the other case hedging requires assumptions on the future stochastic behavior of the underlying and is model-dependent.

In order to better situate these issues, we will now discuss some requirements one would like to impose on a measure of model uncertainty in the context of

⁴Which is reminiscent of the ad-hoc scaling factors often used in computing regulatory capital.

derivative pricing (Section 3) then proceed to formalize them in mathematical terms (Section 4). The relation with coherent and convex of measures of risk will then become clear.

3 Model uncertainty in the context of derivative valuation

Stochastic models of financial markets usually represent the evolution of the price of a financial asset as a stochastic process $(S_t)_{t \in [0, T]}$ defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$. An option on S with maturity T then corresponds to a random variable H_T , whose value is revealed at T and depends on the behavior of the underlying asset S between 0 and T . For example, a call option with maturity T_i and strike price K_i is then a financial contract which pays out $\max(0, S_{T_i} - K_i) = (S_{T_i} - K_i)^+$ to the holder at the maturity date T_i . The main focus of option pricing theory has been to define a notion of value for such options and compute this value.

In an arbitrage-free market, the assumption of linearity of prices leads to the existence of a probability measure $\mathbb{Q}$ equivalent to $\mathbb{P}$ such that the value $V_t(H)$ of an option with payoff H is given by:

$$V_t(H) = B(t, T)E^{\mathbb{Q}}[H|\mathcal{F}_t] \quad (11)$$

where $B(t, T)$ is a discount factor. In particular the discounted asset price is a martingale under $\mathbb{Q}$. Here the probability measure $\mathbb{Q}$ does not describe “objective probabilities”: for an event $A \in \mathcal{F}$, while $\mathbb{P}(A)$ represents its probability of occurrence, $\mathbb{Q}(A)$ represents the value of an option with terminal payoff equal to $1/B(t, T)$ if A occurs and zero otherwise. A pricing model, specified by such a “risk-neutral” probability measure $\mathbb{Q}$, therefore encodes market consensus on values of derivative instruments rather than any “objective” description of market evolution: it should be seen as a market-implied model.

3.1 Statistical uncertainty vs uncertainty on pricing rules

When speaking of stochastic models and model uncertainty, one should therefore distinguish *econometric* models, where one specifies a probability measure $\mathbb{P}$ in an attempt to model the historical evolution of market prices, from *pricing* models where a risk-neutral probability measure $\mathbb{Q}$ is used to specify a pricing rule whose role is to relate prices of various instruments in an arbitrage free manner.

If $\mathbb{P}$ corresponds to a complete market model (for example, a one dimensional diffusion model for a single asset) then the pricing rule $\mathbb{Q}$ is uniquely defined by $\mathbb{P}$. Uncertainty on $\mathbb{Q}$ can then only result from uncertainty on $\mathbb{P}$ - which results from the lack of identification of $\mathbb{P}$ from historical data- so we are in the classical case of ambiguity or Knightian uncertainty described in [Knight, Ellsberg, Epstein, Routledge & Zin (2001)]. However if $\mathbb{P}$ corresponds

to the more realistic case of an *incomplete* market model (for example, a jump-diffusion or stochastic volatility models for a single asset or a multifactor diffusion model with more factors than tradable assets) then the knowledge of $\mathbb{P}$ does not determine the pricing rule $\mathbb{Q}$ in a unique way. Therefore, even if $\mathbb{P}$ is known with certainty we still face uncertainty in the choice of the pricing model $\mathbb{Q}$. Thus, the notion of model uncertainty in the context of option pricing extends beyond the traditional framework of statistical uncertainty on the evolution of the underlying. While the literature mentioned in section 2 has focused on “statistical uncertainty”, we will focus here on uncertainty on pricing rules.

We also note that in existing works on model uncertainty [Chen & Epstein, Epstein & Wang, Gundel, Karatzas & Zamfirescu] all probability measures $\mathbb{P} \in \mathcal{P}$ are assumed to be equivalent to a reference probability $\mathbb{P}_0$.⁵ While this is often considered to be a “technical” hypothesis, it can actually quite restrictive: it means that all models agree on the universe of possible scenarios and only differ on their probabilities. For example, if $\mathbb{P}_0$ defines a complete market model, this hypothesis entails that there is no uncertainty on option prices! A fundamental example such as a diffusion model with uncertain volatility [Avellaneda et al, Lyons] does not verify this hypothesis. We will not assume this hypothesis in the sequel.

3.2 Benchmark instruments vs illiquid products

When discussing the role of mathematical models in derivative markets, one should also distinguish between liquidly traded options, for which a market price is available, and exotic or illiquid options, which are issued over-the-counter and for which a market price is often unavailable. For the former, which includes call and put options on major indices, exchange rates and major stocks, the price is determined by supply and demand on the market. Pricing models are therefore not used to price such options; their market prices are rather used as inputs in order to “calibrate” (mark-to-market) option pricing models. For exotic, over the counter or illiquid options, there is no market price available: the value of the option is computed using a pricing model. In order to guarantee coherence (in the sense of absence of arbitrage) between these two categories of instruments, the pricing rule chosen should be consistent with the observed market prices of the traded options. Thus a pricing model acts as an arbitrage-free “interpolation/ extrapolation” rule, extending the price system from market-quoted instruments to non-quoted ones.

3.3 Some requirements for a measure of model uncertainty

We now translate the above remarks into a set of requirements that any measure of model uncertainty in derivative valuation should take into account. Hereafter by a “model” we mean an arbitrage-free option pricing rule, represented by a

⁵With the notable exception of [Avellaneda et al] and [Lyons].

(risk-neutral) probability measure $\mathbb{Q}$ on $(\Omega, \mathcal{F})$ such that $(S_t)_{t \in [0, T]}$ is a martingale under $\mathbb{Q}$.

Consider now a (model-dependent) value $V(\mathbb{Q})$, a typical example of which is the value at $t = 0$ of a random terminal payoff X : $V(\mathbb{Q}) = B(0, T)E^{\mathbb{Q}}[X]$. Other examples are provided by values of options with early exercise features, such as an American put $V(\mathbb{Q}) = \sup_{\tau} E^{\mathbb{Q}}[B(\tau, T)(K - S_{\tau})^+]$ where the supremum is taken over all non-anticipating (random) exercise times $\tau \in [0, T]$. Since these quantities depend on the choice of the pricing rule $\mathbb{Q}$, it is natural to ask what the impact of this choice on their value is. The “model uncertainty” of $V(\cdot)$ is defined as the uncertainty on the value of $V(\mathbb{Q})$ resulting from the uncertainty in the specification of $\mathbb{Q}$. Based on the above discussion, here are some natural requirements that a measure of model uncertainty should verify:

1. For liquidly traded options, the price is determined by the market: there is no model uncertainty on the value of a liquid option.
2. Any measure of model uncertainty must take into account the possibility of setting up (total or partial) hedging strategies in a model-free way. If an instrument can be replicated in a model-free way, then its value involves no model uncertainty. If it can be partially hedged in a model-free way, this should also reduce the model uncertainty on its value.
3. When some options (typically, call or put options for a short maturities and strikes near the money) are available as liquid instruments on the market, they can be used as hedging instruments for more complex derivatives. A typical example of a model-free hedge using options is of course a *static* hedge using liquid options, a common approach for hedging exotic options.
4. If one intends to compare model uncertainty with other, more common, measures of (market) risk of a portfolio, the model uncertainty on the value of a portfolio should be expressed in monetary units and normalized to make it comparable to the market value of the portfolio.
5. As the set of liquid instruments becomes larger, the possibility of setting up static hedges increases which, in turn, should lead to a decrease in model uncertainty on the value of a typical portfolio.

In order to take the above points into account, we therefore need to specify not only the class of (alternative) models considered but also the set of hedging instruments which are available. While the usual approach is to allow dynamic trading in the underlying, it is common market practice to use static or semi-static positions in call and put (“vanilla”) options to hedge exotic options [Allen & Padovani] so we will also include this possibility in the framework presented below.

4 A quantitative framework for measuring model uncertainty

Let us now define a quantitative setting taking into account the above remarks. Consider a set of market scenarios $(\Omega, \mathcal{F})$. We stress that there is no reference probability measure defined on Ω . The underlying asset is represented by a measurable mapping: $S : \Omega \mapsto D([0, T])$ where $D([0, T])$ denotes the space of right continuous functions with left limit (this allows for jumps in prices) and $S(\omega)$ denotes the trajectory of the price in the market scenario $\omega \in \Omega$. A contingent claim will be identified with the terminal value at T of its payoff, represented by a random variable H revealed at T . In order to simplify notations, we will omit discount factors: all payoffs and asset values are assumed to be discounted values.

4.1 An axiomatic setting for model uncertainty

In order to define a meaningful methodology for measuring model uncertainty we need the following ingredients:

1. Benchmark instruments: these are options written on S whose prices are observed on the market. Denote their payoffs as $(H_i)_{i \in I}$ and their observed market prices by $(C_i^*)_{i \in I}$. In most cases a unique price is not available; instead, we have a range of prices $C_i^* \in [C_i^{\text{bid}}, C_i^{\text{ask}}]$.
2. A set of arbitrage-free pricing models $\mathcal{Q}$ consistent with the market prices of the benchmark instruments: the (discounted) asset price $(S_t)_{t \in [0, T]}$ is a martingale under each $\mathbb{Q} \in \mathcal{Q}$ with respect to its own history⁶ $\mathcal{F}_t$ and

$$\forall \mathbb{Q} \in \mathcal{Q}, \forall i \in I, \quad E^{\mathbb{Q}}[|H_i|] < \infty \quad E^{\mathbb{Q}}[H_i] = C_i^* \quad (12)$$

In a realistic setting the market price C_i^* is only defined up to the bid-ask spread so one may relax the consistency constraint (12) to:

$$\forall \mathbb{Q} \in \mathcal{Q}, \forall i \in I, \quad E^{\mathbb{Q}}[H_i] \in [C_i^{\text{bid}}, C_i^{\text{ask}}] \quad (13)$$

A typical example that the reader can keep in mind is when a finite number of pricing models are available as alternatives: $\mathcal{Q} = \{\mathbb{Q}_k, k = 1..n\}$ and we will use this case in most examples. However the assertions below will also hold for infinite convex sets $\mathcal{Q}$ with a finite number of extremal elements.

Remark 4.1 (Parameter uncertainty vs uncertainty on model type) Some authors have distinguished between “parameter” uncertainty and “model uncertainty” [Kerkhof et al.]. We find this distinction to be irrelevant: if $(\mathbb{Q}_\theta)_{\theta \in E}$ is a parametric family of (pricing) models, different values $\theta_1, ..\theta_n$ of the parameter will define probability measures $\mathbb{Q}_i, i = 1..n$ and this is the only ingredient we need here. The fact that they can be embedded in a “single” parametric

⁶Defined as the natural filtration of S completed by the null sets of $\frac{1}{|\mathcal{Q}|} \sum_{\mathbb{Q} \in \mathcal{Q}} \mathbb{Q}$.

family is purely conventional and depends on the arbitrary definition of a “parametric family”. In fact by embedding all model in $\mathcal{Q}$ in a single super-model one can always represent model uncertainty as “parameter uncertainty”.

The set of contingent claims can be defined as those with a well defined price in all models:

$$\mathcal{C} = \{H \in \mathcal{F}_T, \sup_{\mathbb{Q} \in \mathcal{Q}} E^{\mathbb{Q}}[|H|] < \infty\} \quad (14)$$

When $\mathcal{Q}$ is finite this is simply the set of terminal payoffs which have a well-defined value under any of the alternative pricing models: $\mathcal{C} = \bigcap_{k=1}^n L^1(\Omega, \mathcal{F}_T, \mathbb{Q}_k)$.

For a simple (i.e. piecewise constant and bounded) predictable process $(\phi_t)_{t \in [0, T]}$ representing a self-financing trading strategy, the stochastic integral $\int_0^t \phi_u \cdot dS_u$ corresponds to the gain from trading between 0 and t and its discounted value is a $\mathbb{Q}$ -martingale. Note that the usual construction of this stochastic integral depends on the underlying measure $\mathbb{Q}$. Following [Doléans-Dade (1971)], one can construct a stochastic integral with respect to the whole family $\mathcal{Q}$: there exists a process $G_t(\phi)$ such that for every $\mathbb{Q} \in \mathcal{Q}$ the equality

$$G_t(\phi) = \int_0^t \phi_u \cdot dS_u$$

holds $\mathbb{Q}$ -almost surely. $G_t(\phi)$ is then a $\mathbb{Q}$ -martingale and defines a model-free version of the gain of the strategy ϕ . In the case where $\mathcal{Q}$ is finite, this construction coincides with stochastic integral constructed with respect to $\overline{\mathbb{Q}} = \frac{1}{|\mathcal{Q}|} \sum_{\mathbb{Q} \in \mathcal{Q}} \mathbb{Q}$. However it is more natural to refer to the models $\mathbb{Q} \in \mathcal{Q}$ instead of $\overline{\mathbb{Q}}$.

The set of simple predictable processes can be enlarged in various ways⁷ to include more complex strategies; we will denote by $\mathcal{S}$ be the set of admissible trading strategies and require that the stochastic integral $G_t(\phi) = \int_0^t \phi \cdot dS$ is well-defined and is a $\mathbb{Q}$ -martingale bounded from below $\mathbb{Q}$ -a.s. for each $\mathbb{Q} \in \mathcal{Q}$. Note that we have made no assumption about market completeness or incompleteness, nor do we require that the probability measures $\mathbb{Q} \in \mathcal{Q}$ be equivalent with each other.

Consider now a mapping $\mu : \mathcal{C} \mapsto [0, \infty[$ represent the model uncertainty on the contingent claim X . The properties enumerated in section 3 can be stated as follows:

1. For liquid (benchmark) instruments, model uncertainty reduces to the uncertainty on market value:

$$\forall i \in I, \quad \mu(H_i) \leq |C_i^{\text{ask}} - C_i^{\text{bid}}| \quad (15)$$

2. Effect of hedging with the underlying:

$$\forall \phi \in \mathcal{S}, \quad \mu(X + \int_0^T \phi_t \cdot dS_t) = \mu(X). \quad (16)$$

⁷For various definitions of admissible strategies see [Kabanov] or [Delbaen & Schachermayer].

In particular the value of a contingent claim which can be replicated in a model free way by trading in the underlying has no model uncertainty:

$$[\exists x_0 \in \mathbb{R}, \exists \phi \in \mathcal{S}, \forall \mathbb{Q} \in \mathcal{Q}, \quad \mathbb{Q}(X = x_0 + \int_0^T \phi_t dS_t) = 1] \Rightarrow \mu(X) = 0. \quad (17)$$

3. Convexity: model uncertainty can be decreased through diversification.

$$\forall X_1, X_2 \in \mathcal{C}, \forall \lambda \in [0, 1] \quad \mu(\lambda X_1 + (1 - \lambda)X_2) \leq \lambda \mu(X_1) + (1 - \lambda)\mu(X_2). \quad (18)$$

4. Static hedging with traded options:

$$\forall X \in \mathcal{C}, \forall u \in \mathbb{R}^K, \quad \mu(X + \sum_{i=1}^k u_i H_i) \leq \mu(X) + \sum_{i=1}^k |u_i| (C_i^{\text{ask}} - C_i^{\text{bid}}) \quad (19)$$

In particular the value of any instrument which can be statically replicated with traded options, model uncertainty reduces to the uncertainty on the cost of replication:

$$[\exists u \in \mathbb{R}^K, \quad X = \sum_{i=1}^K u_i H_i] \Rightarrow \mu(X) \leq \sum_{i=1}^K |u_i| |C_i^{\text{ask}} - C_i^{\text{bid}}|. \quad (20)$$

Remark 4.2 Contrarily to the conditions defining coherent risk measures, condition (15) defines a scale for μ : if μ verifies the above properties then $\lambda\mu$ still verifies them for $0 < \lambda \leq 1$ but not necessarily for $\lambda > 1$. This allows to construct a maximal element among all mappings proportional to μ , defined as the one which saturates the range constraint (15):

$$\mu_{\max} = \lambda_{\max} \mu \quad \lambda_{\max} = \sup\{\lambda > 0, \lambda\mu \text{ verifies (15)}\} \quad (21)$$

As long as the set of benchmark instruments is non-empty, then $0 < \lambda_{\max} < \infty$. Then for any $0 < \lambda \leq 1$, $\lambda\mu_{\max}$ still verifies (15)-(16)-(17)-(18)-(19)-(20) and now λ can now be interpreted as the proportion of the bid-ask spread which is attributed to model uncertainty.

4.2 A “coherent” measure of model uncertainty

Given the ingredients above, we now construct a measure of model uncertainty verifying the above properties. A payoff $X \in \mathcal{C}$ has a well-defined value in all the pricing models $\mathbb{Q} \in \mathcal{Q}$. Define the upper and lower price bounds by:

$$\bar{\pi}(X) = \sup_{i=1..n} E^{\mathbb{Q}_i}[X] \quad \underline{\pi}(X) = \inf_{i=1..n} E^{\mathbb{Q}_i}[X] = -\bar{\pi}(-X) \quad (22)$$

$X \mapsto \bar{\pi}(-X)$ then defines a coherent risk measure. Any of the pricing models $\mathbb{Q} \in \mathcal{Q}$ will give a value for X falling in the interval $[\underline{\pi}(X), \bar{\pi}(X)]$. For a payoff whose value is not influenced by model uncertainty, $\bar{\pi}(X) = \underline{\pi}(X)$. We propose to measure the impact of model uncertainty on the value of a contingent claim X by

$$\mu_{\mathcal{Q}}(X) = \bar{\pi}(X) - \underline{\pi}(X) \quad (23)$$

Proposition 1 (A coherent measure of model uncertainty)

1. $\bar{\pi}, \underline{\pi}$ assign values to the benchmark derivatives compatible with their market bid-ask prices:

$$\forall i \in I, \quad C_i^{\text{bid}} \leq \underline{\pi}(H_i) \leq \bar{\pi}(H_i) \leq C_i^{\text{ask}} \quad (24)$$

2. $\mu_{\mathcal{Q}} : \mathcal{C} \mapsto \mathbb{R}^+$ defined by (23) is a measure of model uncertainty verifying the properties (15)-(16)-(17)-(18)-(19)-(20).

Proof: see Appendix.

Taking the difference between $\bar{\pi}(X)$ and $\underline{\pi}(X)$ decomposes the coherent risk measure $\bar{\pi}(-X)$ into two components: the model uncertainty, $\mu_{\mathcal{Q}}(X)$ and the market value of the derivative, which may be computed using one of the pricing models (say, $E^{\mathbb{Q}_1}[X]$). $\mu_{\mathcal{Q}}(X)$ can then be used either to compute a margin (for an OTC instrument) or to provision for model uncertainty on this trade.

The computation of the worst case bounds $\bar{\pi}, \underline{\pi}$ is similar to the superhedging approach [El Karoui & Quenez]. If all models in $\mathcal{Q}$ correspond to complete market models, then $\bar{\pi}(X)$ can be interpreted as the cost of the cheapest strategy dominating in the worst case model. However in the usual superhedging approach $\mathcal{Q}$ is taken to be the set of *all* martingale measure equivalent to $\mathbb{P}$ so price intervals produced by superhedging tend to be quite large and sometimes coincide with maximal arbitrage bounds [Eberlein & Jacod], rendering them useless when compared with market prices. Using the approach presented above, if X is the terminal payoff of a traded option our construction interval $[\underline{\pi}(X), \bar{\pi}(X)]$ is compatible with bid-ask intervals for this option. This remark that the calibration condition (13) is essential to guarantee that our measure of model uncertainty is both nontrivial and meaningful.

Remark 4.3 For ease of presentation we have considered a finite number of alternative models. However it is easily observed that the family of $\{\mathbb{Q}_i, i = 1..n\}$ can be replaced by its convex hull $\{\sum \alpha_i \mathbb{Q}_i, \alpha_i \geq 0, \sum \alpha_i = 1\}$ without changing the value of $\bar{\pi}, \underline{\pi}$ and μ . $\sum \alpha_i \mathbb{Q}_i$ is interpreted as a statistical mixture of the models $(\mathbb{Q}_i)$, obtained by picking at random one of the models $i = 1..n$ with probability $(\alpha_i)_{i=1..n}$ and pricing the option under $\mathbb{Q}_i$. More generally using the Choquet representation theorem [Choquet] one can show that the above results hold for any convex set $\mathcal{Q}$ with a finite number of extremal points.

4.3 Examples

The following example shows that a given payoff can be highly exposed to model uncertainty while its "market risk" is estimated as being low.

Example 4.1 (Uncertain volatility) Consider a market where there is a riskless asset with interest r , a risky asset S_t and a call option on S with maturity T , trading at price C^* at $t = 0$. Consider the alternative diffusion models:

$$\mathbb{Q}_i : \quad dS_t = S_t[r dt + \sigma_i(t) dW_t] \quad (25)$$

where $\sigma_i : [0, T] \rightarrow]0, \infty[$ is a bounded deterministic volatility function with and W a standard Brownian motion under $\mathbb{Q}_i$. Then the calibration condition (12) reduces to

$$\frac{1}{T} \int_0^T \sigma_i(t)^2 dt = \Sigma^2 \quad (26)$$

where Σ is the Black-Scholes implied volatility associated to the call price C^* . Obviously (26) has many solutions, each of which corresponds to a different scenario for the evolution of market volatility. Examples of such solutions which are widely used by practitioners are piecewise constant or piecewise linear functions of t ; consider for example

$$\sigma_1(t) = \Sigma \quad (27)$$

$$\sigma_i(t) = a_i 1_{[0, T_1]} + \sqrt{\frac{T^2 \Sigma^2 - T_1 a_i^2}{T - T_1}} 1_{]T_1, T]} \quad i = 2..n \quad (28)$$

with $\Sigma < a_i < \Sigma \sqrt{T/T_1}$ for $i = 2..n$. Set $a_1 = \Sigma$ and let

$$\bar{a} = \max\{a_i, i = 1..n\} \quad \underline{a} = \min\{a_i, i = 1..n\}$$

Now consider the issue of a call option X with maturity $T_1 < T$ (with a possibly different strike). For each $i = 1..n$, $(\Omega, \mathcal{F}, \mathcal{F}_t^S, \mathbb{Q}_i)$ defines a complete market model (i.e. the martingale representation property holds) so under each $\mathbb{Q}_i$ the call can be perfectly hedged. However the corresponding Δ -hedging strategy depends on the volatility structure: it is not model free. Therefore, while the P&L of the delta-hedged position is almost surely zero according the model $\mathbb{Q}_i$ used to compute the Δ , it is a random variable with non-zero variance under any $\mathbb{Q}_j, j \neq i$. In fact using the monotonicity of the Black Scholes formula with respect to volatility it is easy to show that

$$\bar{\pi}(X) = C^{BS}(K, T_1; \bar{a}) \quad \underline{\pi}(X) = C^{BS}(K, T_1; \underline{a}) \quad (29)$$

This example also shows that, when all the alternative pricing models considered are (one dimensional) diffusion models, model uncertainty reduces to "Vega risk", that is, uncertainty on volatility.

The next example shows that, conversely, a position in derivatives can have a considerable exposure to market risk but no exposure to model uncertainty:

Example 4.2 (Butterfly position) Consider a market where options are liquidly traded at strike levels $K_1 < K_2 < K_3$, at market prices $C_t^*(T, K_i), i = 1, 2, 3$ where K_2 is at the money. A butterfly position consists in taking a short position in two units of the at-the-money call option $C_t(T, K_2)$ and a long position in each of the calls $C_t(T, K_1), C_t(T, K_3)$. This position has an exposure to "gamma" risk but since it can be synthesized using market-traded options in a model-free way the model uncertainty on its value reduces to the uncertainty on the cost of the static hedge:

$$\mu_Q(V_t) \leq |C^{\text{bid}}(K_1) - C^{\text{ask}}(K_1)| + |C^{\text{bid}}(K_2) - C^{\text{ask}}(K_2)| + 2|C^{\text{bid}}(K_3) - C^{\text{ask}}(K_3)|.$$

A typical portfolio of derivatives will be exposed both to market risk and model uncertainty, but the above examples illustrate the difference between the two concepts.

The above examples are theoretical. In the case of index options, one disposes of more than a hundred prices and a simple model such as (25) is insufficient to reproduce their smile and skew features: more sophisticated models such as local volatility models [Dupire (1994)], stochastic volatility models or models with jumps have to be used. Given an empirical data set of option prices, how can a family of pricing models compatible with market prices of options be specified in an effective way? Can one implement an algorithm capable of generating such a class of models verifying (13) and subsequently computing $\mu(X)$ for any given payoff X ? [BenHamida & Cont (2004)] have given an example of such a procedure in the case of diffusion ("local volatility") models:

Example 4.3 Using an evolutionary algorithm, [BenHamida & Cont (2004)] construct a family $\{\mathbb{Q}_i, i = 1..n\}$ of local volatility models

$$\mathbb{Q}_i : \quad dS_t = S_t[r dt + \sigma_i(t, S_t) dW_t]$$

compatible with a given set of call option prices $(C_i^{\text{bid}}, C_i^{\text{ask}})$ in the following manner: we start with a population of candidate solutions $(\sigma_i(\cdot, \cdot), i = 1..N)$ and evolve them iteratively through random search / selection cycles until the prices generated for the benchmark options by the local volatility functions $(\sigma_i(\cdot, \cdot), i = 1..N)$ become compatible with their bid-ask spreads. Denoting by E the set of admissible local volatility functions, this algorithm defines a Markov chain in E^N , which is designed to converge to a set of model parameters which minimize the difference between model and market prices of benchmark options.

Here is an empirical example, obtained by applying this procedure to DAX index options on June 13, 2001. The benchmark instruments are European call and put options traded on the market, numbering at around 150 quoted strikes and maturities. The implied volatility surface is depicted in figure 4.3. Figure 4.3 illustrates two local volatility functions obtained by applying the algorithm described above to the data: while these two volatility functions look different, they are both compatible with the market prices of quoted European call options and this cannot be distinguished on the sole basis of market information. However, they will not give rise to the same values for American or exotic options for which we face model uncertainty. Other examples are shown in figure 4.3 which illustrates several such local volatility functions: note the high level of uncertainty on short term volatility, due to the fact the value of short term options is not affected very much by the volatility and thus the information implied by these options on volatility is imprecise. The diffusion models defined by these local volatility functions can then be used to price a given exotic option, leading to a range of prices as in Section 4.

We present now another example where market risk and model risk are both present, which allows to compare the two; it also illustrates that our approach to measuring model uncertainty is not tied to the class of diffusion models and can incorporate more general specifications:

Figure 1: DAX options implied volatilities: 13 June 2001.

Figure 2: Two examples of local volatility functions compatible with market prices of DAX options, June 13 2001.

Figure 3: A family of local volatility functions $\sigma(t, S)$ compatible with the market prices of DAX options on June 13, 2001 computed using an evolutionary algorithm [BenHamida & Cont (2004)].

Example 4.4 (Uncertainty on model type: local volatility vs jumps)

Consider the following jump-diffusion model, used in many cases to reproduce implied volatility skews and smiles in short term options:

$$\mathbb{Q}_1 : \quad S_t = S_0 \exp[\mu t + \sigma W_t + \sum_{j=1}^{N_t} Y_j] \quad (30)$$

In this example we choose $\sigma = 10\%$, $\lambda = 1$ and Y_j are IID random variables with the probability density shown in Figure 4.4. Figure 4.4 shows the implied volatilities for call options, computed using (30) as a model for risk neutral dynamics.

The user, uncertain whether $\mathbb{Q}_1$ is the right model to use, decides to price it using a more familiar diffusion model

$$\mathbb{Q}_2 : \quad \frac{dS_t}{S_t} = r dt + \sigma(t, S_t) dW_t \quad (31)$$

where $\sigma(t, S)$ is calibrated to the implied volatilities in figure (4.4). The resulting volatility function $\sigma(t, S)$ is shown in figure 4.4 (right).

These two models give exactly the same prices for all call options with maturities between 0.1 and 1 year and all strikes between $\pm 10\%$ of the money. Using these options as benchmark instruments, $\mathcal{Q} = \{\mathbb{Q}_1, \mathbb{Q}_2\}$ verifies (12). However,

Figure 4: Left: Density of jump sizes in the jump-diffusion model (30). Right: The local volatility, $\sigma(t, S)$, as function of underlying asset and time in (31).

Table 1: Model uncertainty on a barrier option.

	Local volatility	Black-Scholes+jumps	μ_Q
At the money call T=0.2	3.5408	3.5408	0
Knock out call			
K=105, T=0.2, Barrier B=110	2.73	1.63	1.1

as figure 4.4 shows, the typical scenarios they generate are completely different: $\mathbb{Q}_1$ generates discontinuous price trajectories with stationary returns while $\mathbb{Q}_2$ generates continuous trajectories with highly non-stationary, time dependent returns!

Consider now the pricing of a barrier option, say a knock out call with strike at the money, maturity $T = 0.2$ and a knock-out barrier $B = 110$. Due to the high short term volatilities needed in the diffusion model to calibrate the observed call prices, the price is higher than in the model with jumps. As shown in table 1, model uncertainty on the value of this exotic, yet quite common, derivative represents 40% of its selling price! This example clearly illustrates that, even for common derivatives, model uncertainty does not represent a small correction to the price but a major factor of risk, as important as market risk.

5 A convex measure of model uncertainty

The approach discussed above is quite intuitive to understand, but its implementation requires to “calibrate” various models to a set of benchmark instruments, a task which can be more or less difficult depending on the complexity of the models and the payoff structures of the benchmark instruments. We will now see that this difficulty can be overcome by using the notion of convex risk measure, due to [Föllmer & Schied (2002)]. As noted in Section 2.3, a convex risk measure can be represented in the form (32), where the penalty function α is a rather abstract object whose value does not have a clear financial interpre-

Figure 5: Implied volatilities for European call and put options, produced by model (30) or (31).

Figure 6: Left: typical sample path of the jump-diffusion model (30). Right: sample path of the local volatility model (31). Both models give rise to the same implied volatility skew and term structure, shown in Figure 4.4.

tation. Also, the representation (32) lacks a normalization which could allow to compare it to marked-to-market values of portfolios. We show here that a special choice of the penalty function can resolve these issues.

Consider as in Section 4, a family of pricing rules $\mathcal{Q}$ and a set of benchmark options with payoffs $(H_i)_{i \in I}$ with market prices $(C_i^*)_{i \in I}$. However, we do not require anymore that these pricing models match the market prices (12). The idea is to relax the calibration constraint (12) on the pricing models $\mathbb{Q} \in \mathcal{Q}$: instead of requiring the pricing models to reproduce the market prices of benchmark instruments exactly as in (12) or within bid-ask spreads as in (13), we can consider a much larger class of pricing models but *penalize* each model according to the precision with which it prices the benchmark instruments:

$$\rho(X) = \sup_{\mathbb{Q} \in \mathcal{Q}} \{ E^{\mathbb{Q}}[-X] - \sum_{i \in I} |C_i^* - E^{\mathbb{Q}}[H_i]| \} \quad (32)$$

In the language of [Föllmer & Schied (2002)], ρ is a convex risk measure associated with the penalty function

$$\begin{aligned} \alpha(\mathbb{Q}) &= \sum_{i \in I} |C_i^* - E^{\mathbb{Q}}[H_i]| \quad \text{if } \mathbb{Q} \in \mathcal{Q} \\ &= +\infty \quad \text{if } \mathbb{Q} \notin \mathcal{Q} \end{aligned} \quad (33)$$

Define now the following valuation measures:

$$\pi^*(X) = \rho(-X) = \sup_{\mathbb{Q} \in \mathcal{Q}} \{ E^{\mathbb{Q}}[X] - \sum_{i \in I} |C_i^* - E^{\mathbb{Q}}[H_i]| \} \quad (34)$$

$$\pi_*(X) = -\rho(X) = \inf_{\mathbb{Q} \in \mathcal{Q}} \{ E^{\mathbb{Q}}[X] + \sum_{i \in I} |C_i^* - E^{\mathbb{Q}}[H_i]| \} \quad (35)$$

This means we price the payoff X with all the pricing models $\mathbb{Q} \in \mathcal{Q}$ but we penalize each model price by the pricing error it commits on the benchmark instruments. In other words, we take more or less seriously the prices produced by any of the pricing models according to the precision with which they are capable of reproducing the market prices of benchmark instruments. We define now, by analogy with (23):

$$\forall X \in \mathcal{C}, \quad \mu_*(X) = \pi^*(X) - \pi_*(X) \quad (36)$$

The following proposition shows that π^* and π_* are concave versions of $\bar{\pi}, \underline{\pi}$ with similar properties and allow to define a measure of model uncertainty with the required properties:

Proposition 2 *Let π^*, π_* and μ_* be defined by (34)-(35)-(36).*

1. π^* assigns to any benchmark instrument a value lower than its market price:

$$\forall i \in I, \quad \pi^*(H_i) \leq C_i^* \quad (37)$$

π_* assigns to any benchmark instrument a value higher than its market price:

$$\forall i \in I, \quad \pi_*(H_i) \geq C_i^* \quad (38)$$

2. Assume the class of pricing models contains at least one model compatible with the market prices of the benchmark options:

$$\exists \mathbb{Q} \in \mathcal{Q}, \quad \forall i \in I, \quad E^{\mathbb{Q}}[H_i] = C_i^*. \quad (39)$$

Then for any payoff $X \in \mathcal{C}$, $\pi^*(X) \geq \pi_*(X)$.

3. Under assumption (39), μ_* defined by (36) is a measure of model uncertainty verifying the properties (15)-(16)-(17)-(18).
4. Static hedging reduces model uncertainty: under assumption (39), diversifying a position using long positions in benchmark derivatives reduces model uncertainty:

$$[1 \geq \lambda_k \geq 0, \sum_{k=0}^K \lambda_k = 1] \Rightarrow \mu_*(\lambda_0 X + \sum_{k=1}^K \lambda_k H_k) \leq \mu_*(X) \quad (40)$$

In particular, any position which can be replicated by a convex combination of available derivatives has no model uncertainty:

$$[\exists (\lambda_i)_{i \in I}, 1 \geq \lambda_i \geq 0, \sum_{i \in I} \lambda_i = 1, X = \sum_{i \in I} \lambda_i H_i] \Rightarrow \mu_*(X) \leq 0. \quad (41)$$

Proof: see Appendix.

Remark 5.1 More generally we could have considered as penalty function a weighted average of pricing errors relative to benchmark instruments:

$$\alpha(\mathbb{Q}) = \sum_{i=1}^n w_i |C_i^* - E^{\mathbb{Q}}[H_i]|.$$

However, it is interesting to note that requiring (37) for any specification of $\mathcal{Q}$ implies that $w_i \leq 1$. Similarly, requiring (38) for any specification of $\mathcal{Q}$ implies $w_i \geq 1$ so equal weighting seems to be the only possibility which satisfies both constraints. In fact, H_i should be interpreted as the payoff of i -th benchmark option, the nominal being determined by the (maximal) quantity of the i -th option *available* to the investor.

Remark 5.2 Properties (40)-(41) replace properties (19)-(20) verified by the coherent measure of model uncertainty $\mu_{\mathcal{Q}}$: this is due to the fact that a convex measure of risk cannot extrapolate the risk of a portfolio to larger, proportional portfolio. A closer look at (41) shows that it is the only reasonable definition of a static hedge: in line with remark 5.1, H_i represents, in nominal terms, the

maximal position in the i -th derivative so feasible positions in this derivative are of the type $\lambda_i H_i$ with $0 \leq \lambda_i \leq 1$. By contrast in (19)-(20) we implicitly allowed unlimited short and long positions in all derivatives and H_i was defined up to a constant.

Remark 5.3 The constraint (39) of including at least one model which calibrates the market prices is quite weak. In practice μ_* can be computed in the following manner. Given a set of prices for the benchmark options, we first choose a pricing model $\mathbb{Q}_1$ which can reproduce these prices and is easily calibrated to option prices. Typical examples are one dimensional diffusion models (local volatility models) used for equity and index derivatives and the SABR model [Hagan et al., 2002] for European options on interest rates. Such models are typically used for their ability to calibrate market prices, but may not generate realistic dynamics for future market scenarios. We have then the freedom to add to such a calibrated model $\mathbb{Q}_1$ other pricing models $\mathbb{Q}_2, \mathbb{Q}_3, \dots$ with more realistic features but which may more complex to calibrate. The procedure above does not require to calibrate these models precisely but simply to penalize their pricing errors, thus their inclusion in the set of pricing models is not a computational burden. In this way, the easy-to-calibrate model $\mathbb{Q}_1$ anchors our measure of model uncertainty in the market prices while more realistic models $\mathbb{Q}_2, \mathbb{Q}_3, \dots$ can be incorporated without having to set up heavy numerical procedures for their calibration. Finally, we remark that (39) is equivalent to requiring that $\rho(0) = 0$ (see the proof of Proposition 2), which is the normalization condition proposed in [Föllmer & Schied (2002)].

It is interesting to examine what happens if more liquid instruments become available. This has the effect of increasing π_* and of decreasing π^* : thus, the measure of model uncertainty becomes smaller. This can be interpreted in the following way: the addition of more liquid options allows a wider range of model-free (static) hedging strategies which allow to reduce exposure to model uncertainty on a given portfolio.

Also, from the expression of the penalty function (33) it is clear that models with lower pricing errors will be more and more favored as the number of benchmark instruments $|I|$ increases. As $|I| \rightarrow \infty$, the sum defining $\alpha(\mathbb{Q})$ will stay finite only if the calibration error is bounded independently of $|I|$; this happens for instance if the pricing model $\mathbb{Q}$ misprices only a finite number of benchmark options, all others being calibrated. Conversely, when there are no options available ($I = \emptyset$), ρ is a coherent risk measure defined by the set $\mathcal{Q}$.

6 Discussion

We now summarize the main contributions of this work, discuss some open questions and possible implications for the risk management of derivative instruments.

6.1 Summary

We have proposed a quantitative framework for measuring the impact of model uncertainty on derivative pricing. Starting from a set of traded benchmark options and a family $\mathcal{Q}$ of option pricing models, we associate a measure of model uncertainty $\mu(X)$, expressed in monetary units, with any derivative with payoff X in two ways. The first method (Section 4) requires the models in $\mathcal{Q}$ to be pricing models calibrated to the benchmark options and computes the range $\mu_{\mathcal{Q}}(X)$ of prices for X over all of these calibrated models. The second method (Section 5) does not require any calibration but penalizes the lack of calibration in the spirit of convex risk measures. In both cases, the specification of a set of benchmark instruments constitutes a key ingredient in our procedure, which was missing in preceding approaches to model risk: without it, the measures of model risk may range between zero and infinity and be meaningless when compared to market values of portfolios.

Both of these approaches verify the intuitive requirements, outlined in Section 3, that a measure of a model uncertainty should have in order to be meaningful in the context of risk management. They are both compatible with market values of traded options and take into account the possibility of model-free hedging with options. They lead to a decomposition of the risk of a position into the sum of a first term, which is of the same order of magnitude as its nominal value and a second term, which can be interpreted as a component of the bid-ask spread due to model uncertainty. Measures of model uncertainty computed in this manner are realistic enough to be considered as bid-ask values and directly compared with market prices and other, more common measures of risk.

Our approach does not require the set of pricing models considered to define equivalent measures on scenarios. When all the models considered are one dimensional diffusions, model uncertainty reduces to uncertainty on future volatility and the approach adopted here is similar to the Lagrangian Uncertain Volatility model of [Avellaneda & Paras]. The notion of model uncertainty proposed here reaches beyond the concept of uncertain volatility and can encompass other types of models (jumps, stochastic volatility) or sources of model uncertainty (number of factors in multifactor models, jump sizes, ..).

Finally, it is important to note that our measures of model risk are not defined “up to a normalization constant” as in [Artzner et al(1999)] or [Föllmer & Schied (2002)]: they directly produce numbers consistent with mark-to-market valuation of portfolios, when available, and do not require an ad-hoc scaling factor in order to be meaningfully used to provision for model risk.

6.2 Specifying the class of models

The relevance of our measure of model risk partly hinges on the specification of the class $\mathcal{Q}$ of models. As noted by [Hansen & Sargent, 2001], “the development of computationally tractable tools for exploring model misspecification [...] should focus on what are the interesting classes of candidate models for

applications”. This issue seems less difficult in option pricing than, say, in macroeconomics, since a market consensus has emerged on a set of standard pricing models (though not a unique model!) for each type of underlying asset in the last decade.

In the approach described in section 4, a further requirement is the ability to calibrate the models to market observations. Standard model calibration algorithms yield a single solution/ pricing model. A first way out is to specify different model classes and perform calibration separately in each model class, yielding a calibrated set of parameters from each class. This approach takes into account uncertainty on model type. Another approach, an example of which was given in Example 4.3, is to consider a single model class but recognize that the calibration problem may have multiple solutions and use a stochastic search algorithm [BenHamida & Cont (2004)]. The two approaches are not exclusive and may be combined. The availability of efficient numerical procedures will ultimately orient market practice in one direction or the other.

The approach of section 5, based on convex risk measures, has the advantage of relaxing this calibration requirement and thus is potentially more flexible from a computational point of view. The advantages and drawbacks of the two approaches remain to be studied in specific settings.

6.3 Updating with new information

In the above discussion, we have considered a market viewed at time $t = 0$. How does the procedure described above apply as time evolves? By analogy with (23), one could define a dynamic bid-ask interval by replacing expectations by conditional expectations:

$$\bar{\pi}_t(X) = \sup_{k=1..n} E^{\mathbb{Q}_k}[X|\mathcal{F}_t^S] \quad \underline{\pi}_t(X) = \inf_{k=1..n} E^{\mathbb{Q}_k}[X|\mathcal{F}_t^S] \quad (42)$$

Then, for each t , $\bar{\pi}_t$ and $\underline{\pi}_t$ define coherent risk measures that make use of the information on the evolution of the underlying up to time t and a natural candidate for building a measure of model uncertainty would then be $\mu_t(X) = \bar{\pi}_t(X) - \underline{\pi}_t(X)$. However there are at least two objections to the formulation (42). The first objection is that it does not guarantee dynamic consistency (i.e. a dynamic programming principle). This problem has been studied in the framework of a Brownian filtration in [Chen & Epstein, Peng (2002)] where it is shown that a special structure has to be imposed on $\mathcal{Q}$ in order for dynamic consistency to hold. These authors examine the case where the family $\mathcal{Q}_0$ is kept fixed while the market evolves through the evolution of the asset price. However the pricing models $\mathbb{Q} \in \mathcal{Q}_0$ have only been calibrated to the value of the benchmark options at $t = 0$ and simply conditioning them on the evolution of the underlying asset clearly does not exploit the information given by the evolution of the market prices of the benchmark options. This is due to the fact that, in a realistic framework, one cannot assume that the “true” model describing the joint evolution of the benchmark derivatives is included in the set $\mathcal{Q}_0$. In line with this remark, the market practice is to re-calibrate pricing

models as prices of options evolve through time. This re-calibration procedure implies that the set of pricing rules $\mathcal{Q}_0$ cannot be used at a later date but has to be replaced by a set $\mathcal{Q}(t)$ of pricing rules verifying:

$$\forall \mathbb{Q} \in \mathcal{Q}(t), \forall i \in I, \quad E^{\mathbb{Q}}[H_i] = C_i^*(t). \quad (43)$$

or, in the bid-ask version:

$$\forall \mathbb{Q} \in \mathcal{Q}(t), \forall i \in I, \quad E^{\mathbb{Q}}[H_i] \in [C_i^{\text{bid}}(t), C_i^{\text{ask}}(t)] \quad (44)$$

This leads to a time-dependent set $\mathcal{Q}(t)$ of forward-looking models $\{\mathbb{Q}_i(t), i = 1..n\}$ where $\mathbb{Q}_i(t)$ is an updated version of $\mathbb{Q}_i$ and defines a risk-neutral measure on the future paths $D([t, T])$ verifying (44). In other words, since the result of model calibration procedure at time t depends on the prices of benchmark instruments $(C_i^*(t), i \in I)$, the measures $(\mathbb{Q}_i(t))_{i=1..n}$ are random measures, whose evolution depends on the market prices of benchmark options. Hence the updating procedure implied by re-calibration procedures is more subtle than conditioning on the past evolution of the underlying asset: the updating procedure must also reflect the evolution of prices in the options market.

A related practical question is that of sensitivity of measures of model uncertainty to market conditions.⁸ This question is already present in the case of market risk measures such as VaR, which can fluctuate in a non-negligible manner as market conditions (e.g. prices of underlying assets) vary. A strong sensitivity would blur the distinction between market risk and model uncertainty. Case studies remain to be done in order to clarify the impact of this sensitivity in practical examples of derivatives portfolios.

6.4 Conclusion

Quantitative risk management took off in the 1990s with the availability of simple tools such as Value-at-Risk for measuring market risk: notwithstanding its technical imperfections, Value-at-Risk convinced practitioners that it is possible in practice to quantify market risk, had a great impact on risk management practices and motivated many researchers to improve this methodology in various ways.

In the recent years, various case studies have indicated the importance of “model risk” in the derivative industry and some spectacular failures in risk management of derivatives have emphasized the consequences of neglecting model uncertainty. Many large financial institutions are conscious of this issue and have been developing methods to tackle it systematically.

We have provided in this paper a simple methodology which can be used to quantify model uncertainty and provides meaningful figures compatible with mark-to-market values of portfolios, when they are available. Our approach can serve as a basis for provisioning for model uncertainty or simply as a decision aid for risk managers and regulators.

⁸We thank Joël Bessis for pointing out this issue.

We hope this work will stimulate further case studies using the methodology presented here, in order to better understand the impact of model uncertainty in various contexts.

References

- [Adam] Adam, K. (2004) On the relation between robust and Bayesian decision making, *Journal of economic dynamics and control*, 28, 2105–2117.
- [Allen & Padovani] Allen, S. and Padovani, O. (2002) Risk management using quasi-static hedging, *Economic Notes*, 31(2), 277–236.
- [Artzner et al(1999)] Artzner P., Delbaen F., Eber J.M., Heath D. (1999) Coherent measures of risk, *Mathematical Finance*, **9**, No. 3, pp 203–228.
- [Avellaneda et al] Avellaneda M., Levy A. & Paras, A. (1995): Pricing and hedging derivative securities in markets with uncertain volatilities, *Applied Mathematical Finance*, 2, 73–88.
- [Avellaneda & Paras] Avellaneda, M. & Paras, A. (1996): Managing the Volatility Risk of Portfolios of Derivative Securities: the Lagrangian Uncertain Volatility Model, *Applied Mathematical Finance*, 3, 31–52.
- [BenHamida & Cont (2004)] Ben Hamida, S. & Cont, R. (2004) Recovering volatility from option prices by evolutionary optimization, Rapport Interne CMAP No 534, Ecole Polytechnique.
- [Buff] Buff, R. (2001) Uncertain volatility models, Springer Finance.
- [Carr et al. 2001] Carr, P.; Geman, H.; Madan, D.B. (2001) Pricing and hedging in incomplete markets, *Journal of Financial Economics* pp. 131–167.
- [Chen & Epstein] Chen, Z. & Epstein, L.G. (2002) Ambiguity, risk and asset returns in continuous time, *Econometrica*, .
- [Choquet] Choquet, G. (1955) Theory of capacities, *Annales de l’Institut Fourier*, Vol 5, pp 131–295.
- [Delbaen & Schachermayer] Delbaen, F. & Schachermayer, W. (1998) The fundamental theorem of asset pricing for unbounded stochastic processes. *Math. Ann.*, 312(2):215–250, 1998.
- [Doléans-Dade (1971)] Doléans-Dade, C. (1971) Intégrales stochastiques par rapport à une famille de probabilités, in: *Séminaire de Probabilités V*, Lecture Notes in mathematics Vol. 191, pp. 141–146, Berlin: Springer.
- [Dupire (1994)] Dupire, B. (1994) Pricing with a smile, *RISK*, 7 (1), 18–20.
- [Eberlein & Jacod] Eberlein, E. and Jacod, J. (1997) On the range of option prices, *Finance and Stochastics*, 1, 131–140.

- [Ellsberg] Ellsberg, D. (1961) Risk, Ambiguity, and the Savage Axioms, *Quarterly Journal of Economics*, 75, 643-669.
- [Elliott (1997)] Elliott, M. (1997) Controlling Model Risk, *Derivatives Strategy*, June.
- [El Karoui & Quenez] El Karoui, N. & Quenez, M.C. (1995) Dynamic programming and the pricing of contingent claims in incomplete markets, *SIAM Journal of Control and Optimization*, Vol 33, No 1, 29-66.
- [El Karoui & Rouge] El Karoui, N. & Rouge, R. (2000) Pricing via utility maximization and entropy, *Mathematical Finance*, 10, pp. 259-276.
- [Epstein] Epstein, L.G. (1999) A definition of uncertainty aversion, *Review of Economic Studies*, 65, 579-608.
- [Epstein & Wang] Epstein, L.G. & Wang, L. (1995) Uncertainty, risk neutral measures and asset price booms and crashes, *Journal of Economic Theory*, 67, 40-80.
- [Föllmer & Schied (2002)] Föllmer, H. & Schied, A. (2002) Convex measures of risk and trading constraints, *Finance and Stochastics*, 6, No. 4, 429-447.
- [Föllmer & Schied (2002b)] Föllmer, H. & Schied, A. (2002) *Stochastic Finance*, Walter de Gruyter.
- [Frittelli & Gianin (2002)] Frittelli, M. & Gianin, E.R. (2002) Putting order in risk measures, *Journal of Banking and Finance*, Volume 26, Issue 7, Pages 1247-1561.
- [Gilboa & Schmeidler] Gilboa & Schmeidler (1989) Maxmin expected utility with non-unique prior, *Journal of mathematical economics*, 18, 141-153.
- [Gundel] Gundel, A. (2004) Robust utility maximization for complete and incomplete market models, Working paper, Humboldt Universität, Berlin.
- [Hagan et al., 2002] Hagan P., Kumar D., Lesniewski A., Woodward D. (2002) Managing smile risk, *Wilmott Magazine*.
- [Hansen et al., 1999] Hansen L., Sargent Th., Tallarini T. (1999) Robust permanent income and pricing, *Review of Economic Studies*, 66, 872-907.
- [Hansen & Sargent, 2001] Hansen, L. & Sargent, Th. (2001) Time inconsistency of robust control, Working paper.
- [Hansen et al., 2002] Hansen L., Sargent Th., Turluhambetova G.A., Williams N. (2002) Robustness and uncertainty aversion, Working Paper.
- [Hoeting et al., 1999] Hoeting J.A., Madigan D., Raftery A.E., Volinsky C.T. (1999) Bayesian model averaging: a tutorial, *Statistical Science*, 14, 4, 382-417.

- [Jacquier & Jarrow (2000)] Jacquier, E. & Jarrow, R. (2000): Bayesian analysis of contingent claim model error, *Journal of Econometrics*, Vol. 94, 145–180.
- [Kabanov] Kabanov, Yuri (2001) Arbitrage theory, in: Jouini et al (eds.): *Option pricing, interest rates and risk management*, Cambridge University Press.
- [Karatzas & Zamfirescu] Karatzas, I. and Zamfirescu, I.M. (2004) Optimal stopping with model uncertainty, Working Paper.
- [Kerkhof et al.] Kerkhof J., Melenberg B. and Schumacher, H. (2002) Model risk and regulatory capital, Working Paper.
- [Knight] Knight, F. (1921) *Risk, uncertainty and profit*, Boston: Houghton Mifflin.
- [Lyons] Lyons, T.J. (1995) Uncertain volatility and the risk free synthesis of derivatives, *Applied mathematical finance*, 2, 117–133.
- [Lyons & Smith] Lyons, T.J. & Smith, A.T. (1999) Uncertain volatility, *RISK*, September, 106–109.
- [Peng (2002)] Peng, S. (2002) Nonlinear expectations and dynamic risk measures, CIME Lectures.
- [Raftery (1993)] Raftery, A. E. (1993) Bayesian model selection in structural equation models. In: K. Bollen and J. Long (eds.): *Testing Structural Equation Models*, 163180, Sage, Newbury Park, CA.
- [Routledge & Zin (2001)] Routledge, B.R. & Zin, S. (2001) Model uncertainty and liquidity, NBER Working Paper 8683.
- [Savage] Savage, L.J. (1954) *The foundations of statistics*, Wiley.
- [Williams (1999)] Williams, D. (1999) Models vs. The Market: Survival of the Fittest, Report FIN514, Meridien Research.

A Proof of Proposition 1

1. Each $\mathbb{Q} \in \mathcal{Q}$ verifies the calibration condition (13). By taking supremum (resp. infimum) over $\mathbb{Q} \in \mathcal{Q}$ we obtain: $C_i^{\text{ask}} \geq \bar{\pi}(H_i) \geq \underline{\pi}(H_i) \geq C_i^{\text{bid}}$.
2. (15) follows from the above inequality. To show (16), note that for any $\phi \in \mathcal{S}$ and any $\mathbb{Q} \in \mathcal{Q}$ the gains process $G_t(\phi)$ is a $\mathbb{Q}$ -martingale so $E^{\mathbb{Q}}[X + \int_0^T \phi_t \cdot dS_t] = E^{\mathbb{Q}}[X]$ so $\bar{\pi}(X + \int_0^T \phi_t \cdot dS_t) = \bar{\pi}(X)$ and $\underline{\pi}(X + \int_0^T \phi_t \cdot dS_t) = \underline{\pi}(X)$ so $\mu_{\mathcal{Q}}(X + \int_0^T \phi_t \cdot dS_t) = \mu_{\mathcal{Q}}(X)$. Choosing in particular $X = x_0 \in \mathbb{R}$ we obtain (17).

To show the convexity property (18), consider $X, Y \in \mathcal{C}$ and $\lambda \in [0, 1]$. For each $\mathbb{Q} \in \mathcal{Q}$ we have

$$\begin{aligned} \lambda \inf_{\mathbb{Q} \in \mathcal{Q}} E^{\mathbb{Q}}[X] + (1 - \lambda) \inf_{\mathbb{Q} \in \mathcal{Q}} E^{\mathbb{Q}}[Y] &\leq E^{\mathbb{Q}}[\lambda X + (1 - \lambda)Y] \\ &\leq \lambda \sup_{\mathbb{Q} \in \mathcal{Q}} E^{\mathbb{Q}}[X] + (1 - \lambda) \sup_{\mathbb{Q} \in \mathcal{Q}} E^{\mathbb{Q}}[Y] \end{aligned}$$

By taking supremum (resp. infimum) over $\mathbb{Q}$ we obtain:

$$\begin{aligned} \lambda \underline{\pi}(X) + (1 - \lambda) \underline{\pi}(Y) &\leq \underline{\pi}(\lambda X + (1 - \lambda)Y) \\ &\leq \bar{\pi}(\lambda X + (1 - \lambda)Y) \leq \lambda \bar{\pi}(X) + (1 - \lambda) \bar{\pi}(Y) \end{aligned}$$

from which (18) is easily derived.

Let us now show (19). Consider a portfolio composed of a long position in X , and positions $u_i, i = 1..k$ in k benchmark options $H_i, i = 1..k$. Assume without loss of generality that the first $u_1, \dots, u_{k_1}$ are long positions, the others being short positions. Since any $\mathbb{Q} \in \mathcal{Q}$ verifies (13), we have:

$$\begin{aligned} E^{\mathbb{Q}}[X] + \sum_{i=1}^{k_1} u_i C_i^{\text{bid}} + \sum_{i=k_1+1}^k u_i C_i^{\text{ask}} &\leq E^{\mathbb{Q}}[X + \sum_{i=1}^k u_i H_i] \\ &\leq E^{\mathbb{Q}}[X] + \sum_{i=1}^{k_1} u_i C_i^{\text{ask}} + \sum_{i=k_1+1}^k u_i C_i^{\text{bid}} \end{aligned}$$

By taking the supremum (resp. the infimum) over $\mathbb{Q} \in \mathcal{Q}$ we obtain:

$$\begin{aligned} \underline{\pi}(X) + \sum_{i=1}^{k_1} u_i C_i^{\text{bid}} + \sum_{i=k_1+1}^k u_i C_i^{\text{ask}} &\leq \underline{\pi}(X + \sum_{i=1}^k u_i H_i) \\ \bar{\pi}(X + \sum_{i=1}^k u_i H_i) &\leq \bar{\pi}(X) + \sum_{i=1}^{k_1} u_i C_i^{\text{ask}} + \sum_{i=k_1+1}^k u_i C_i^{\text{bid}} \end{aligned}$$

Adding the last two inequalities and taking into account the signs of u_i we obtain:

$$\bar{\pi}(X + \sum_{i=1}^k u_i H_i) - \underline{\pi}(X + \sum_{i=1}^k u_i H_i) \leq \bar{\pi}(X) - \underline{\pi}(X) + \sum_{i=1}^k |u_i (C_i^{\text{bid}} - C_i^{\text{ask}})|$$

which yields (19). Substituting $X = 0$ yields (20).

B Proof of Proposition 2

Let π^*, π_* and μ_* be defined by (34)-(35)-(36).

1. Noting that

$$\forall \mathbb{Q} \in \mathcal{Q}, \forall i \in I, -|C_i^* - E^{\mathbb{Q}} H_i| \leq C_i^* - E^{\mathbb{Q}}[H_i] \text{ we obtain}$$

$$\begin{aligned} E^{\mathbb{Q}} H_i - \sum_{j \in I} |C_j^* - E^{\mathbb{Q}} H_j| &\leq E^{\mathbb{Q}} H_i - |C_i^* - E^{\mathbb{Q}} H_i| \\ &\leq E^{\mathbb{Q}} H_i + C_i^* - E^{\mathbb{Q}}[H_i] \leq C_i^* \end{aligned}$$

Taking the supremum over $\mathbb{Q} \in \mathcal{Q}$ we obtain $\pi^*(H_i) \leq C_i^*$.

Similarly, starting from

$$\forall \mathbb{Q} \in \mathcal{Q}, \forall i \in I, |C_i^* - E^{\mathbb{Q}} H_i| \geq C_i^* - E^{\mathbb{Q}}[H_i] \text{ we obtain}$$

$$\begin{aligned} E^{\mathbb{Q}}[H_i] + \sum_{j \in I} |C_j^* - E^{\mathbb{Q}}(H_j)| &\geq E^{\mathbb{Q}}[H_i] + |C_i^* - E^{\mathbb{Q}} H_i| \\ &\geq E^{\mathbb{Q}}[H_i] + C_i^* - E^{\mathbb{Q}}[H_i] \geq C_i^* \end{aligned}$$

Taking the infimum over $\mathbb{Q} \in \mathcal{Q}$ we obtain $\pi_*(H_i) \geq C_i^*$.

2. Since ρ defined by (32) is a convex risk measure, applying (9) to $Y = -X$ and $\lambda = 1/2$ yields:

$$\rho(0) = \rho\left(-\frac{X}{2} + \frac{X}{2}\right) \leq \frac{1}{2}(\rho(X) + \rho(-X))$$

Since $\pi^*(X) = \rho(-X)$ and $\pi_*(X) = -\rho(X)$ we obtain

$$\forall X \in \mathcal{C}, \quad \pi^*(X) \geq \pi_*(X) + 2\rho(0)$$

Now remark that $-\rho(0) = \inf_{\mathbb{Q} \in \mathcal{Q}} \sum_{i \in I} |C_i^* - E^{\mathbb{Q}} H_i|$ is simply the smallest calibration error achievable using any of the pricing models in $\mathcal{Q}$. If we assume the class of pricing models contains at least *one* model compatible with the market prices of the benchmark options:

$$\exists \mathbb{Q} \in \mathcal{Q}, \forall i \in I, \quad E^{\mathbb{Q}}[H_i] = C_i^*.$$

then $\rho(0) = 0$ and we obtain $\forall X \in \mathcal{C}, \pi^*(X) \geq \pi_*(X)$.

3. From the above inequality $\pi^*(H_i) \geq \pi_*(H_i)$ with (37) and (38) we obtain:

$$\forall i \in I, C_i^* \geq \pi^*(H_i) \geq \pi_*(H_i) \geq C_i^*$$

hence $\mu_*(H_i) = 0$ for all $i \in I$. To show (17), note that for any $\phi \in \mathcal{S}$ and any $\mathbb{Q} \in \mathcal{Q}$ the gains process $G_t(\phi)$ is a $\mathbb{Q}$ -martingale so $E^{\mathbb{Q}}[X] =$

$x_0 + E^{\mathbb{Q}}[\int_0^T \phi_t dS] = x_0$. Therefore $\rho(X) = -x_0$ and $\pi^*(X) = \pi_*(X) = 0$ hence $\mu_*(X) = 0$. More generally,

$$\begin{aligned} \rho(X + \int_0^T \phi_t dS_t) &= \sup_{\mathbb{Q}} \{ E^{\mathbb{Q}}[X + \int_0^T \phi_t dS_t] + \sum_{i \in I} |C_i^* - E^{\mathbb{Q}}(H_i)| \} \\ &= \sup_{\mathbb{Q}} \{ E^{\mathbb{Q}}[X] + \sum_{i \in I} |C_i^* - E^{\mathbb{Q}}(H_i)| \} \end{aligned}$$

since $\int_0^t \phi_t dS$ is a martingale under each $\mathbb{Q} \in \mathcal{Q}$, which implies (16).

Using the convexity property (9) of ρ we have, for any $X, Y \in \mathcal{C}$ and $\lambda \in [0, 1]$:

$$\begin{aligned} \mu_*(\lambda X + (1 - \lambda)Y) &= \rho(\lambda X + (1 - \lambda)Y) + \rho(-\lambda X - (1 - \lambda)Y) \\ &\leq \lambda \rho(X) + (1 - \lambda) \rho(Y) + \lambda \rho(-X) + (1 - \lambda) \rho(-Y) \\ &= \lambda [\rho(X) + \rho(-X)] + (1 - \lambda) [\rho(Y) + \rho(-Y)] \\ &= \lambda \mu_*(X) + (1 - \lambda) \mu_*(Y) \end{aligned}$$

which shows the convexity property (18) for μ .

4. To show (40), consider $\lambda_k \geq 0, k = 0..K$ with $\sum_{k=0}^K \lambda_k = 1$. Using the convexity of μ_*

$$\mu_*(\lambda_0 X + \sum_{k=1}^K \lambda_k H_k) \leq \lambda_0 \mu_*(X) + \sum_{k=1}^K \lambda_k \mu_*(H_k)$$

As shown above, $\mu_*(H_i) = 0$ for all $i \in I$ and $0 \leq \lambda_0 \leq 1$ we obtain (40). Choosing $X = 0$ and noting that $\mu_*(0) = 0$, we obtain (41).