Model Uncertainty Measures Using Bayesian Posteriors

27-28th March 2009, Princeton University

Alok Gupta
D.Phil Mathematical Finance
alok.gupta@maths.ox.ac.uk

MCFG, Mathematical Institute, Oxford University
Nomura Bank
EPSRC

Supervised By Christoph Reisinger

Contents

- Financial Modelling
  - Calibration Problem
  - Bayesian Approach
  - Model Uncertainty (Motivation)
- Coherent Model Uncertainty Measures
  - Review: Coherent Market Risk Measures
  - Cont's Axioms
  - Worst-Case Measure
  - Generalised Bayesian Measures
- Convex Model Uncertainty Measures
  - Bayesian Posterior Penalty
  - Numerical Example
- Related Work
  - Optimal Bayesian Hedging
  - Consistency Of Bayesian Estimators
- Questions

Financial Modelling

Calibration Problem

Suppose we have an asset price process $S = (S_t)_{t \geq 0}$. Choose a model $\mathbb{M}_\theta \in \mathcal{M}$ for S so that, for any time $t \geq 0$,

$$S_t = \mathbb{M}_\theta(S_0, t, (P_u)_{0 \leq u \leq t})$$

where θ is the model parameter (or vector of parameters or function) and $P = (P_t)_{t \geq 0}$ is a random process.

To price claims on $S(\mathbb{M}, \theta)$ we first choose $\mathbb{M}$. Then the *calibration problem* is to find suitable θ which reproduces some benchmark market prices $V^* = \{V_i^* : i \in I\}$ of claims $C = \{C_i : i \in I\}$ written on S.

Bayesian Approach

Bayesian theory is used to estimate the value of an unknown parameter(s). Suppose we observe noisy prices $V^* = \{V_i^* : i \in I\}$ related to θ by

$$V_i^* = V_i(\theta) + e_i$$

where e_i for $i \in I$ are (not necessarily independent) noises and $V_i(\cdot)$ is the known model price dependent on θ. Then the *Bayesian posterior* density of θ is

$$p(\theta \mid V^*) \propto p(V^*|\theta)\, p(\theta)$$

The Bayesian posterior is a measure of the confidence in any particular value of the parameter θ based on prior assumptions and observed data.

Model Uncertainty (Motivation)

Motivation: What is the risk associated with choosing the wrong model (parameter) for the pricing of a financial derivative?

Literature: Referred to as *model risk* or *model uncertainty*.

- R.Cont, *Model Uncertainty And Its Impact On The Pricing Of Derivative Instruments*
- F.L.J.Kerkhof, B.Melenberg and H.Schumacher, *Model Risk and Regulatory Capital*
- S.Figlewski and T.C.Green, *Market Risk and Model Risk For a Financial Institution Writing Options*

Coherent Model Uncertainty Measures

Review: Market Risk Measures

In 1998 Artzner et al. introduced coherent market risk measure as a mapping from $\mathcal{X}$ into $\Re$ satisfying:

1. translational invariance
2. monotonicity
3. subadditivity
4. positive homogeneity

Examples of coherent market risk measures are:

i) worst-case $\rho_1(X) = \sup_{\omega \in \Omega}\{-X(\omega)\}$.

ii) average value $\rho_0(X) = \mathbb{E}^{\mathbb{P}}[-X(\omega)]$.

iii) expected shortfall $\rho_\beta(X) = \frac{1}{1-\beta}\int_\beta^1 -X(\omega)\,\mathbb{P}(\mathrm{d}\omega)$ for some $\beta \in (0,1)$.

Cont's Axioms

Assume that

$$\forall\, \mathbb{M}_\theta \in \mathcal{M}(I),\ \forall i \in I, \qquad \mathbb{E}^{\mathbb{M}_\theta}[C_i] \in [V_i^{bid}, V_i^{ask}].$$

Then $\mu : \mathcal{X} \to [0, \infty)$ is a *coherent model uncertainty measure* if it satisfies:

- For benchmark options, $\mu(C_i) \leq |V_i^{bid} - V_i^{ask}|$.
- Dynamic hedging with the underlying only reduces model uncertainty if the hedging strategy is model-free. Moreover, if the claim can be totally replicated in a model-free way then $\mu(X) = 0$.
- Diversification: $\mu(\lambda X_1 + (1-\lambda)X_2) \leq \lambda\mu(X_1) + (1-\lambda)\mu(X_2)$ for $\lambda \in [0, 1]$.
- If payoff can be statically replicated in a model-free way with traded options then its model uncertainty is bounded by the uncertainty on the cost of this replication:

$$\left[\exists a \in \Re^d, \forall \mathbb{M}_\theta \in \mathcal{M} \quad X = \sum_{i=1}^{d} a_i C_i \quad \mathbb{M}_\theta\text{-a.s.}\right] \Rightarrow \mu(X) \leq \sum_{i=1}^{d} |a_i||V_i^{bid} - V_i^{ask}|.$$

Worst-Case Measure

Cont's worst-case measure is

$$\mu_1(X) = \overline{\pi}(X) - \underline{\pi}(X)$$

with $\underline{\pi}(X) = \inf_{\mathbb{M}_\theta \in \mathcal{M}} \mathbb{E}^{\mathbb{M}_\theta}[X]$, $\overline{\pi}(X) = \sup_{\mathbb{M}_\theta \in \mathcal{M}} \mathbb{E}^{\mathbb{M}_\theta}[X]$.

$V_K(\sigma) =$ Black-Scholes European call value with strike K.
$\hat{V}_K(\sigma) = \frac{V_K(\sigma) - V_K(0.1)}{V_K(0.15) - V_K(0.1)}$ for $K = S_0, S_0/2$ and $\sigma \in [0.1, 0.15]$.

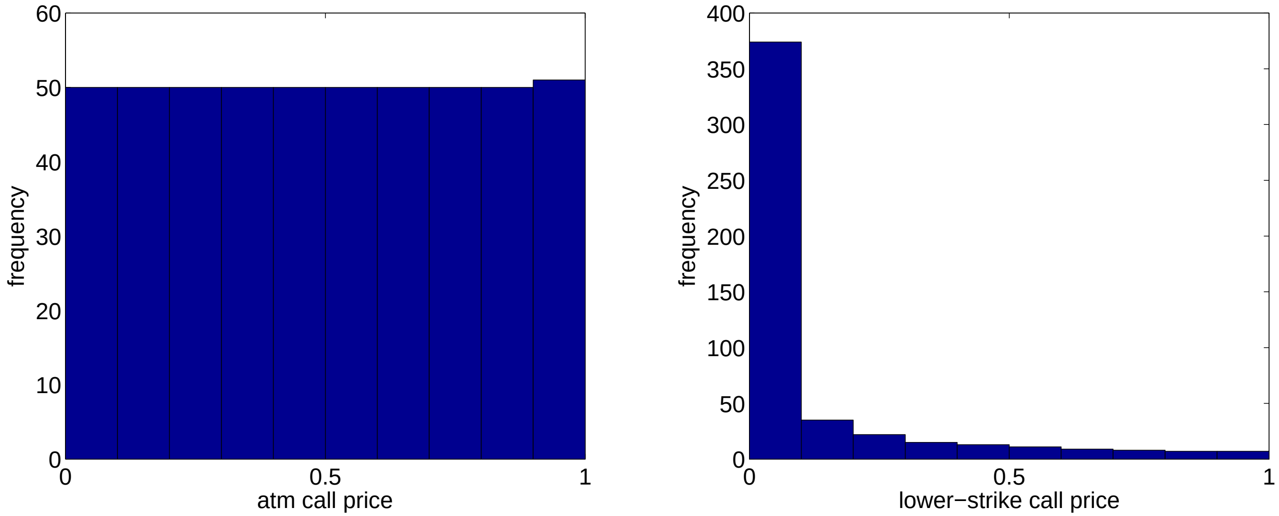

Generalised Bayesian Measures

Theorem 0.1 $\mu(X) = \rho(-\mathbb{E}^{\mathbb{M}_\theta}[X] + \underline{\pi}(X))$ *defines a coherent model uncertainty measure (where* $\mathbb{P}$ *corresponding to* ρ *is given by* $p(\theta|V^*)$*).*

Examples (c.f. market risk measures):

i) w-c $\mu_1(X) = \rho_1(-\mathbb{E}^{\mathbb{M}_\theta}[X] + \underline{\pi}(X)) = \overline{\pi}(X) - \underline{\pi}(X)$

ii) average value $\mu_0(X) = \mathbb{E}[\mathbb{E}^{\mathbb{M}_\theta}[X] - \underline{\pi}(X)|V^*]$

iii) expected shortfall $\mu_\beta(X) = \rho_\beta(-\mathbb{E}^{\mathbb{M}_\theta}[X] + \underline{\pi}(X))$

For the previous numerical example:

$$\mu_1(\hat{V}_{S_0}) = 1.00 \qquad \mu_{0.95}(\hat{V}_{S_0}) = 0.97 \qquad \mu_0(\hat{V}_{S_0}) = 0.50$$

$$\mu_1(\hat{V}_{S_0/2}) = 1.00 \qquad \mu_{0.95}(\hat{V}_{S_0/2}) = 0.83 \qquad \mu_0(\hat{V}_{S_0/2}) = 0.11$$

Convex Model Uncertainty Measures

Bayesian Posterior Penalty

Convex Measure: Idea is to penalise model according to how well it calibrates. Construction:

$$\begin{aligned}
\varepsilon &= \inf_{\mathbb{M}_\theta \in \mathcal{M}} \alpha(\mathbb{M}_\theta), \\
\pi^\alpha(X) &= \sup_{\mathbb{M}_\theta \in \mathcal{M}} \{\mathbb{E}^{\mathbb{M}_\theta}[X] - \alpha(\mathbb{M}_\theta)\} + \varepsilon, \\
\pi_\alpha(X) &= \inf_{\mathbb{M}_\theta \in \mathcal{M}} \{\mathbb{E}^{\mathbb{M}_\theta}[X] + \alpha(\mathbb{M}_\theta)\} - \varepsilon, \\
\mu_\alpha(X) &= \pi^\alpha(X) - \pi_\alpha(X)
\end{aligned}$$

for some function $\alpha > 0$ (e.g. a function of $p(\theta|V^*)$) satisfying:

$$\forall i \in I, \quad \pi^\alpha(C_i) \leq V_i^{ask} + \varepsilon \text{ and } \pi_\alpha(C_i) \geq V_i^{bid} - \varepsilon.$$

Theorem 0.2 *μ_α is a convex model uncertainty measure.*

Numerical Example

Local volatility surface was calibrated to 20 European call prices using 27 parameters. 5 options are priced with the Bayesian posterior. $S_0 = 5000$.

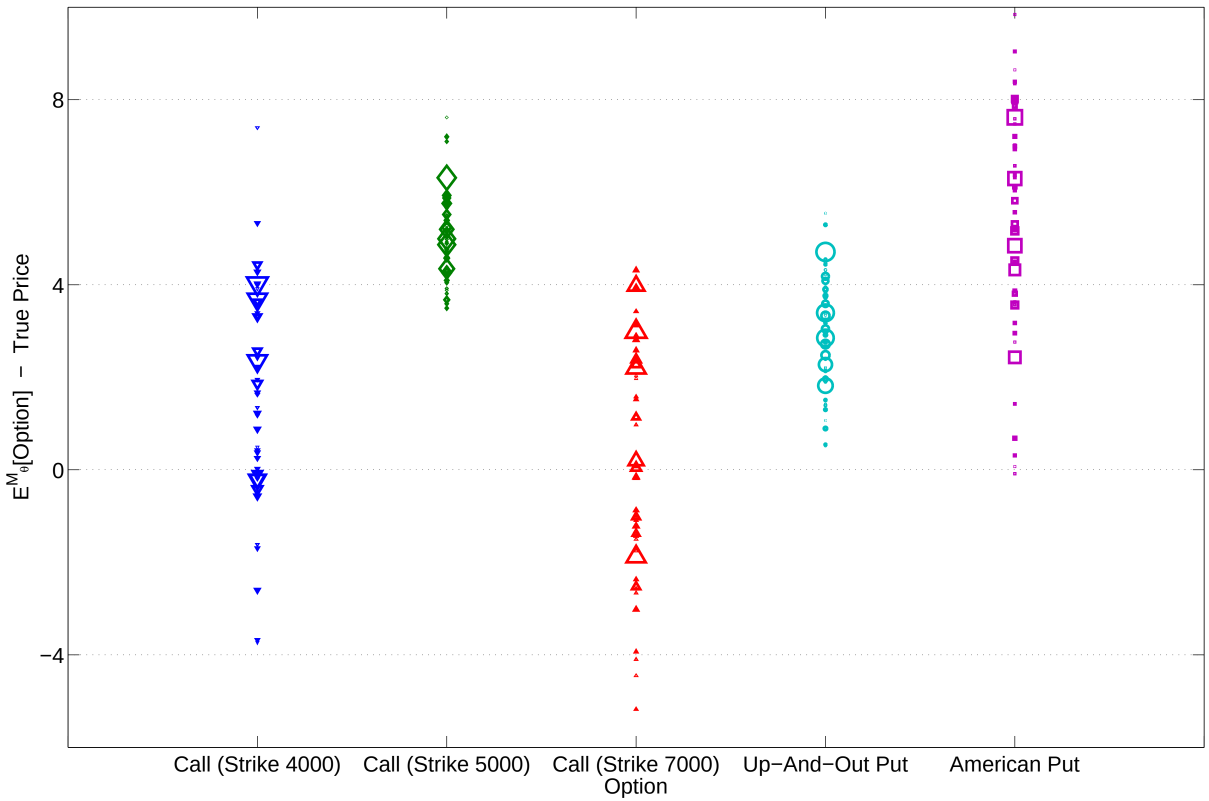

Worst-Case: ↑ 7.9 ↑ 1.8 ↑ 7.4 ↑ 2.9 ↑ 6.4

Related Work

Optimal Bayesian Hedging

Can use the Bayesian posterior to find model parameter θ which minimises the expected value of some hedging performance indicator

$$\theta_L(V^*) = \arg\min_{\theta'} \left\{ \int L(\theta, \theta')\, p(\theta|V^*)\, \mathrm{d}\theta \right\} .$$

where the loss function L is a function f of the hedging error,

$$L(\theta, \theta') = \mathbb{E}[f(\Pi_T(\theta') - V_T(\theta))]$$

and Π_t is the value of the hedging strategy for option with value V_t.

Optimal Bayesian Hedging

Example: Histogram comparing the hedging profits for usual MLE hedge and Bayesian (BEM) hedge for European call option with maturity 1 year and strike $1.1S_0$. Option value is 31.29.

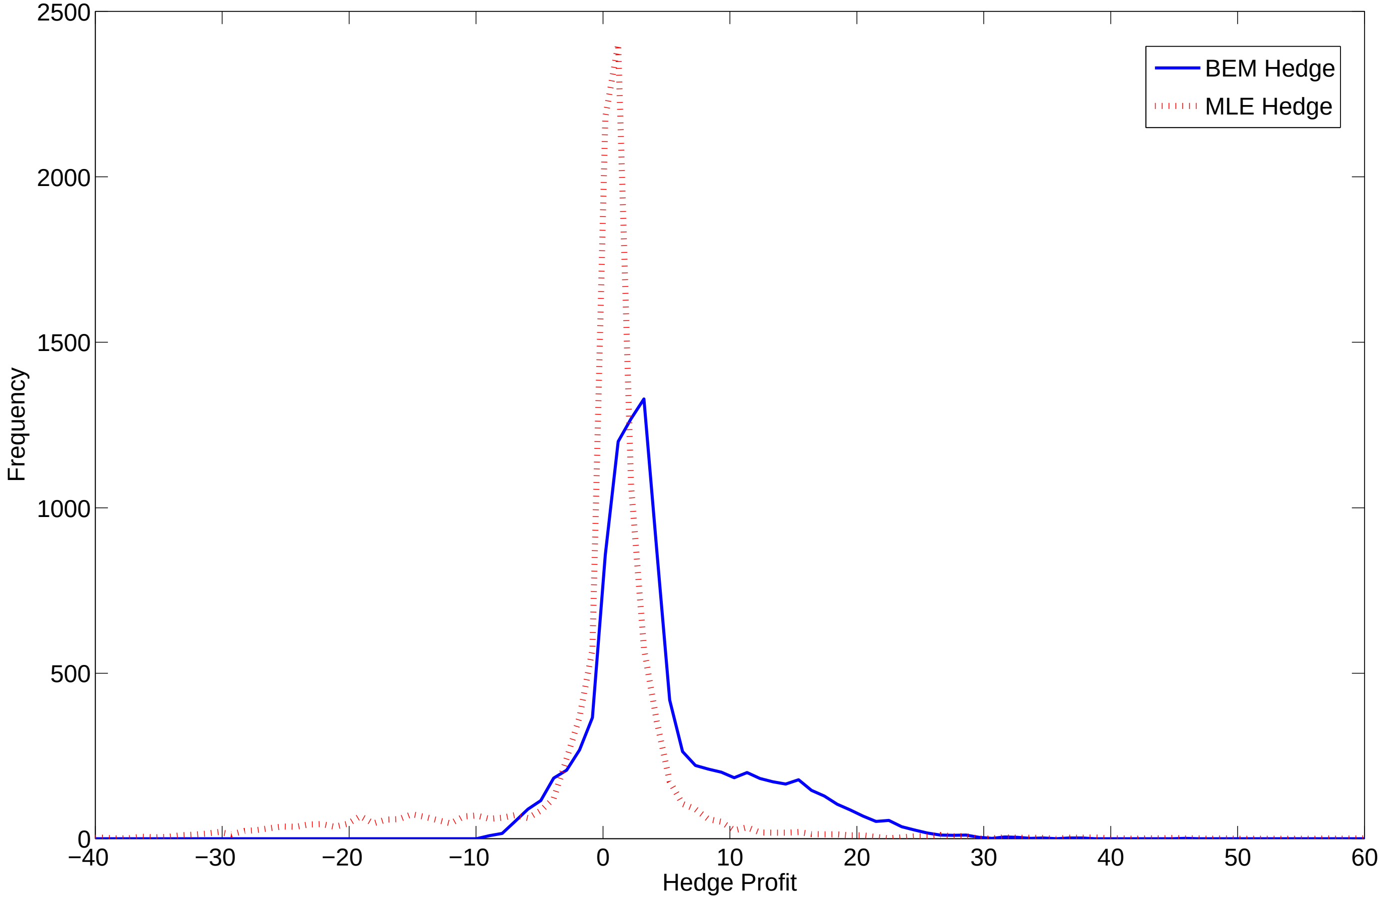

Consistency Of Bayesian Estimators

A sequence of estimators $\hat{\theta}_n(V^*)$ is *strongly consistent* if $\hat{\theta}_n \longrightarrow \theta^*$ almost surely where θ^* is the true value.

We have shown for Black-Scholes model with daily observation(s) and (non-)Gaussian noise that

$$\lim_{n \to \infty} p_n(\sigma | V^*) = \delta(\sigma - \sigma^*)$$

so the Bayesian estimators given by

$$\hat{\sigma}_n = \arg\min_{\sigma' \in \Sigma} \left\{ \int_{\Sigma} L(\sigma, \sigma') p_n(\sigma | V^*) \mathrm{d}\sigma \right\}$$

are consistent for all loss functions L.

Questions?