The integrated correlated variance as a statistical inverse problem

Ruth Kaila
Aalto University, Helsinki

6th World Congress of the Bachelier Finance Society
June 2010, Toronto

Ranunculus glacialis

Outline of the presentation

- Statistical inverse problems
- Correlated Hull-White formula for option pricing
- Inverse problem of integrated correlated variance

Inverse problems and ill-posedness

Let us have an unknown hidden variable x and an observed variable y, related to x by the model function f:

$$y = f(x, e),$$

where e accounts for the measurement noise.

The inverse problem is to get information of x by measuring y - for example to estimate the implied volatility, local volatility, or risk-neutral price density from option prices.

Often inverse problems are ill-posed: they have no unique solution and small errors in the data propagate to large errors in x.

- traditional regularization: the ill-posed problem is replaced by a nearby problem that is well-posed

- instead of applying traditional regularization, we recast the inverse problem in the form of statistical inference on the distribution of the unknown

  – allows us to integrate additional prior knowledge to our estimation process

  – even if the deterministic ill-posed problem does not have a unique solution, there always exists a probability density of the unknown, the variance of which may be large or small

Stochastic inverse problem; Bayesian approach

X is an unknown hidden random variable; an observed random variable Y is related to X by the model function f, so that

$$Y = f(X, \epsilon),$$

where also the noise is modelled as a random variable ϵ.

We apply a Bayesian approach. The solution of the inverse problem, the posterior density, is then given by

$$P_{\mathsf{posterior}}(x \mid y) \propto P_{\mathsf{likelyh}}(y \mid x) P_{\mathsf{prior}}(x),$$

where x and y are realizations of X and Y.

Bayesian approach with hyperprior

We apply a Bayesian approach with a hyperprior, so that

$$P_{\text{posterior}}(x \mid y) \propto P_{\text{likelyh}}(y \mid x) P_{\text{prior}}(x \mid \theta) P_{\text{hyper}}(\theta),$$

where θ is part of the inverse problem and will be defined by the data.

Maximum-A-Posteriori and Conditional Mean estimates

Two common Bayesian point estimates:

Maximum-A-Posteriori (MAP)

$$x_{\text{MAP}} = \arg\max P(x \mid y),$$

- tells which value of x maximizes the posterior distribution of this unknown. Leads to an optimization problem.

Conditional mean (CM)

$$x_{\text{CM}} = \mathbf{E}\{x \mid y\},$$

- provides information of the point of mass of the posterior distribution. Leads to Markov Chain Monte Carlo (MCMC) sampling and an integration problem.

Correlated stock price and volatility processes

Model the stock price process as

$$dX_t = \sqrt{1-\rho^2}\sigma_t X_t dW_{1,t} + \rho \sigma_t X_t dW_{2,t},$$

where the σ dynamics are independent on the Brownian motion W_1 but dependent on the Brownian motion W_2, and $\rho \in [-1, 1]$ is the correlation between price and volatility shocks.

The integrated variance $\bar{\sigma}_t^2$ is given by

$$\bar{\sigma}_t^2 = \frac{1}{T-t} \int_t^T \sigma_s^2 ds,$$

and the integrated correlated variance by $\bar{\sigma}_{\rho,t}^2 = (1-\rho^2)\bar{\sigma}_t^2$.

Now $\log(X_T/X_t)$ is normally distributed conditional on $\bar{\sigma}^2_{\rho,t}$ and ξ_t, with

$$\mathbb{E}\Big\{\log(\frac{X_T}{X_t}) \mid \bar{\sigma}^2_{\rho,t}, \xi_t\Big\} = \log(\xi_t) - \frac{1}{2}\bar{\sigma}^2_{\rho,t}(T-t),$$
$$\mathrm{Var}\Big\{\log(\frac{X_T}{X_t}) \mid \bar{\sigma}^2_{\rho,t}, \xi_t\Big\} = \bar{\sigma}^2_{\rho,t}(T-t).$$

where the integrated correlated variance is given by

$$\bar{\sigma}^2_{\rho,t} = \frac{1-\rho^2}{T-t}\int_t^T \sigma_s^2 ds$$

and

$$\xi_t = \exp\left(-\frac{1}{2}\rho^2\int_t^T \sigma_s^2 ds + \rho\int_t^T \sigma_s dW_{2,s}\right).$$

The correlated Hull-White price for a European call is then

$$U_t^{\mathsf{HW},\rho}(x; K, T; \sigma_t^2) = \mathbb{E}\Big\{U_t^{\mathsf{BS}}(x\xi_t; K, T; \bar{\sigma}^2_{\rho,t}) \mid \bar{\sigma}^2_{\rho,t}, \xi_t, x = X_t\Big\}.$$

(Hull-White 1987, correlated Hull-White by Willard 1996)

We assume that

$$u_t^{\mathsf{obs}} = U_t^{\mathsf{HW},\rho} + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathsf{Var_t})$$

and u_t^{obs} is a realization of $U_{t,\mathsf{true}} \sim \mathcal{N}(\widehat{U}_t, \mathsf{Var_t})$, with

$$\begin{aligned} \widehat{U}_t &= (u_{t,\mathsf{bid}} + u_{t,\mathsf{offer}})/2, \\ \mathsf{Var_t} &= V(u_{t,\mathsf{bid}} - u_{t,\mathsf{offer}}), \end{aligned}$$

where V is positive finite constant.

Joint density and marginalized densities related to correlated Hull-White formula

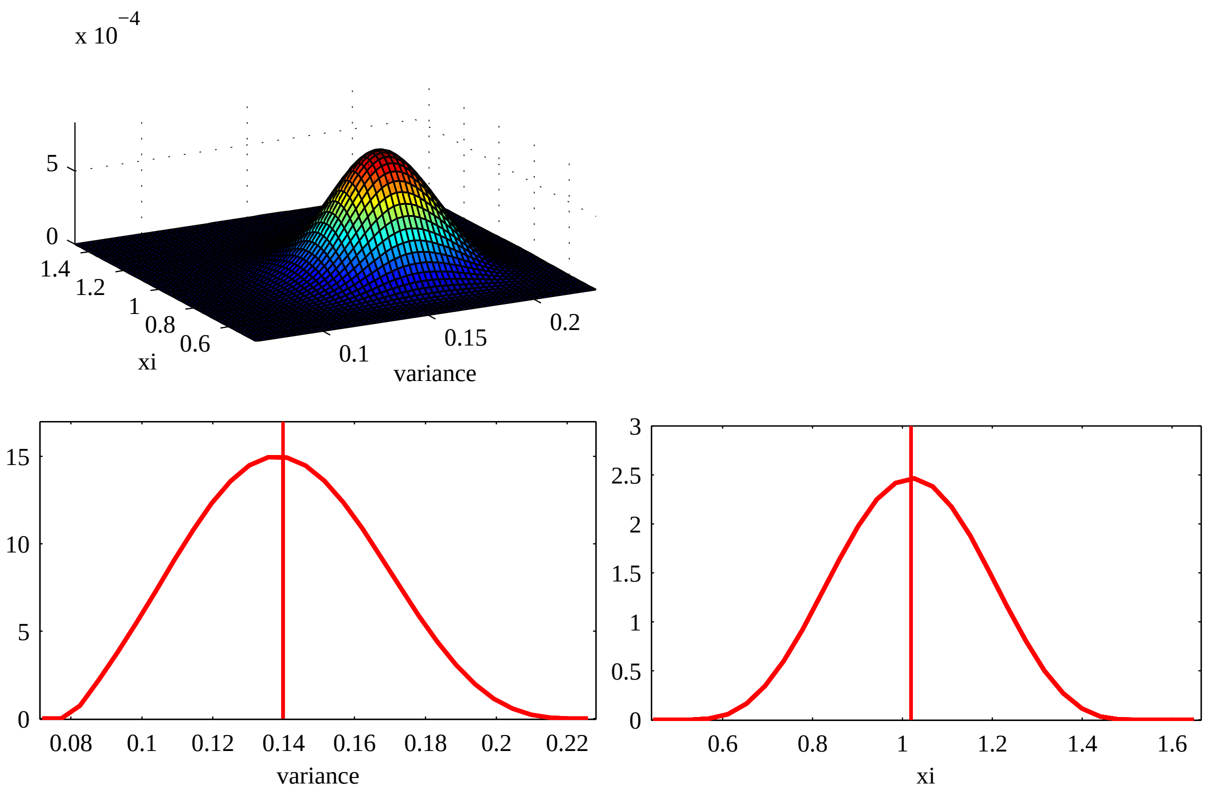

Joint probability density and marginalization

The joint probability density of two real valued random variables X and Y is defined as

$$P\{X \in A, Y \in B\} = \iint_{A \times B} \pi(x, y) dx dy,$$

and, assuming that such a density exists, the marginal density of X is recovered as

$$\pi(x) = \int_{\mathbb{R}} \pi(x, y) dy.$$

We estimate the joint probability distribution $(\bar{\sigma}^2_{\rho,t}, \xi_t)$. When estimating the integrated correlated variance, we marginalize the correlated Hull-White formula with respect to an estimate of ξ_t, and when estimating ξ_t, we marginalize the formula with respect to an estimate of $\bar{\sigma}^2_{\rho,t}$.

Integrated correlated variance and marginalization

We alternate two problems: estimate one of the unknown densities at time, marginalizing with respect to the other one.

$$
\begin{aligned}
U_t^{\mathsf{HW},\rho}(x;K,T;\sigma_t^2) &= \mathbb{E}\Big\{U_t^{\mathsf{BS}}(x\xi_t;K,T;\bar\sigma_{\rho,t}^2) \mid \bar\sigma_{\rho,t}^2, \xi_t = \xi_t^{\mathsf{est}}\Big\}.\\
U_t^{\mathsf{HW},\rho}(x;K,T;\sigma_t^2) &= \mathbb{E}\Big\{U_t^{\mathsf{BS}}(x\xi_t;K,T;\bar\sigma_{\rho,t}^2) \mid \bar\sigma_{\rho,t}^2 = \bar\sigma_\rho^{2,\mathsf{est}}, \xi_t\Big\}.
\end{aligned}
$$

Information content of the densities

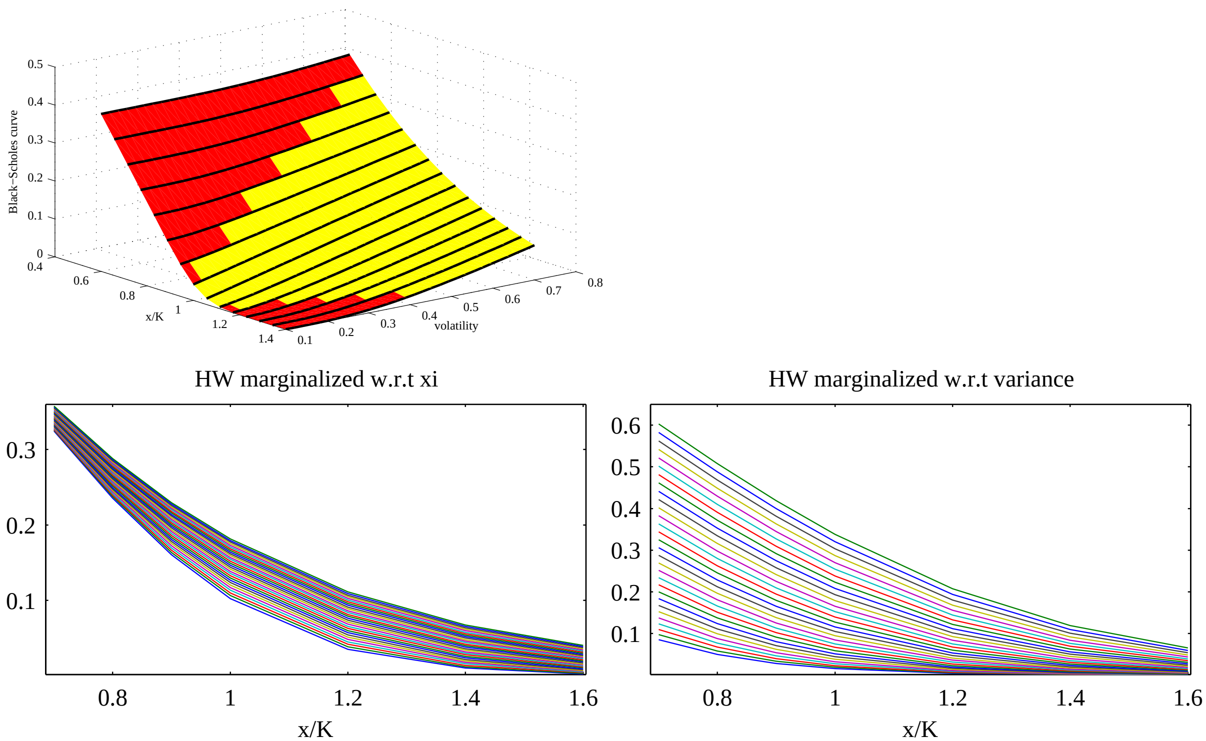

Prior information and assumptions for $\bar{\sigma}^2_{\rho,t}$ and ξ_t

We approximate the true probability density $\bar{\sigma}^2_{\rho,t}$ with a discrete density z, $z \in \mathbb{R}^n$.

- we know that the density of interest is non-negative:

$$\pi(z) \geq 0.$$

- we know that the cumulative density equals one:

$$\sum z = 1.$$

  We denote by α the cumulative density corresponding to z.

- we assume that the density is smooth and takes on positive values on the interval $[a, a + M]$.

- the likelihood function is given by

$$P_{\text{li}}(u^{\text{obs}} \mid \alpha, \alpha_n = 1) \propto P_<(\alpha) \exp\left(-\frac{1}{2\text{Var}}\|u^{\text{obs}} - AV\alpha\|^2\right)$$

with

$$\begin{aligned} a_{ij} &= \frac{M}{n} U^{\text{BS}}(x; K_i, T; \bar{\sigma}_j^2), \quad A = a_{ij}, \\ z &= V\alpha, \end{aligned}$$

so that V is a first order difference matrix. Also,

$$P_<(\alpha) = \begin{cases} 1, & \text{if } \alpha_{j+1} \geq \alpha_j, \quad 1 \leq j \leq n-1 \\ 0 & \text{elsewhere.} \end{cases}$$

- the likelihood function is given by

$$P_{\text{li}}(u^{\text{obs}} \mid \alpha, \alpha_n = 1) \propto P_{<}(\alpha) \exp\Big(-\frac{1}{2\text{Var}}\|u^{\text{obs}} - AV\alpha\|^2\Big)$$

- the prior density is given by

$$P_{\text{prior}}(\alpha) \propto \exp\Big(-\frac{1}{2}\|\theta^{-1/2}L\alpha\|^2\Big),$$

where L is a first order finite difference matrix and θ its variance

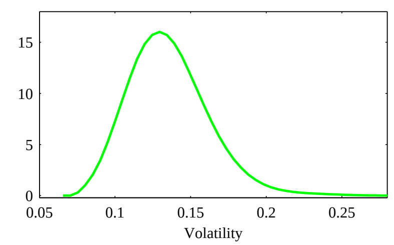

- the prior with a hyperprior is given by

$$P_{\text{prior}}(\alpha \mid \theta) \propto \exp\Big(-\frac{1}{2}\|D^{-1/2}L\alpha\|^2 - \sum_{j=1}^{n}\frac{\theta_j}{\theta_0} + (\beta - \frac{3}{2})\log\theta\Big),$$

where $D^{1/2} = \text{diag}(\theta_1^{1/2}, \theta_2^{1/2} \ldots$ and $\theta_n^{1/2}) \in \mathbb{R}^{n\times n}$ is the random variance with prior based on the gamma distribution

$$P_{\text{hyper}}(\theta) \propto \prod_{j=1}^{n} \exp\left(-\frac{\theta_j}{\theta_0} + (\beta - 1)\log\theta_j\right)$$

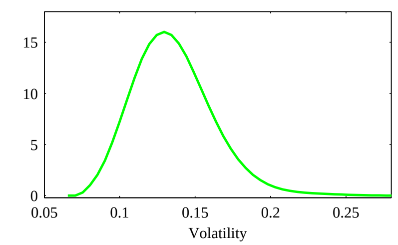

- the posterior density is given by

$$
\begin{aligned}
&P_{\mathrm{post}}(\alpha, \theta \mid u^{\mathrm{obs}}, \alpha_n = 1) \\
&\qquad \propto P_<(\alpha) \exp\Bigg(-\frac{1}{2\mathrm{Var}} \| u^{\mathrm{obs}} - AV\alpha \|^2 \\
&\qquad - \frac{1}{2} \| D^{-1/2} L \alpha \|^2 - \sum_{j=1}^{n} \frac{\theta_j}{\theta_0} + (\beta - \frac{3}{2}) \log \theta \Bigg)
\end{aligned}
$$

ideas from Somersalo, Calvetti 2007 and imaging inverse problems

MAP and MCMC/CM estimates for $\bar{\sigma}^2_{\rho,t}$ and ξ_t

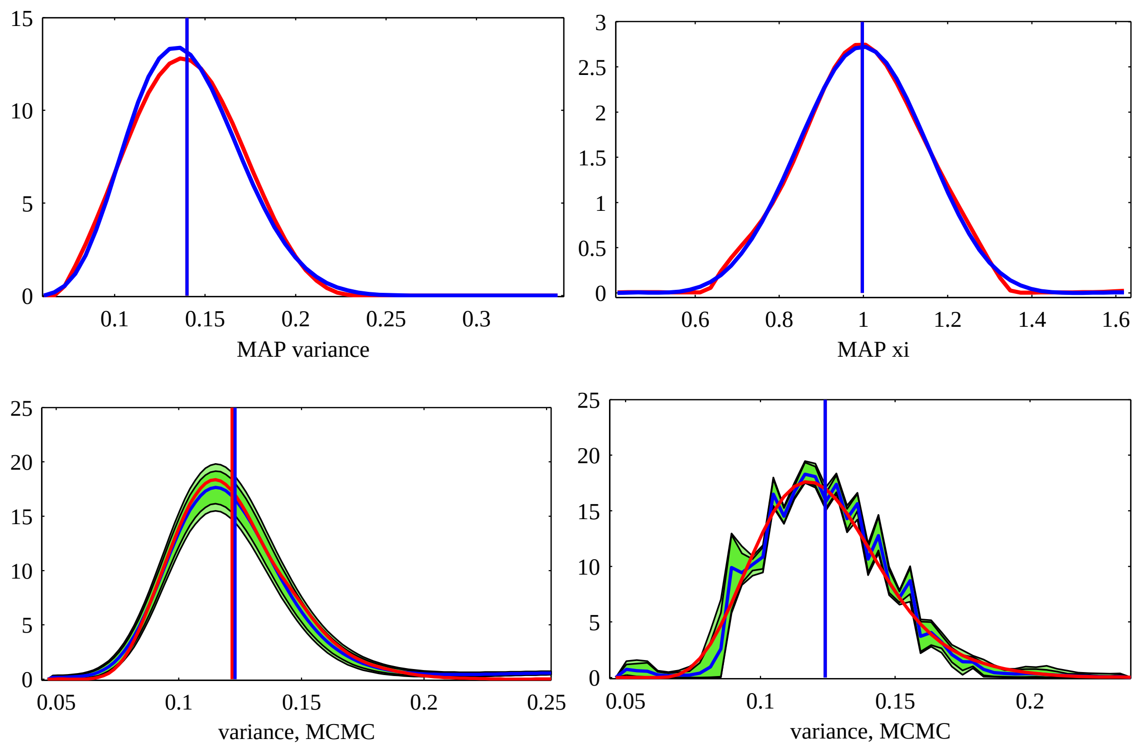

Uncorrelated integrated variance

We assume that the distribution of $\bar{\sigma}_t^2$ is log-normal:

$$\bar{\sigma}^2_{\rho,t} \sim \mathcal{N}(e^{\mu_t}, e^{\varsigma_t^2}),$$

and $\mu_t \in \mathcal{I}[\mu_{\mathsf{min}}, \mu_{\mathsf{max}}]$, $\varsigma_t^2 \in \mathcal{I}[\varsigma^2_{\mathsf{min}}, \varsigma^2_{\mathsf{max}}]$ are unknown. We compute which lognormal discrete distribution z gives the maximum likelihood:

$$P_{\mathsf{li}}(u_t^{\mathsf{obs}}) \propto \exp\Big(-\frac{1}{2\mathsf{Var_t}}\|u_t^{\mathsf{obs}} - U_t^{\mathsf{HW}}(x, z_t(\mu_t, \varsigma_t))\|^2\Big).$$

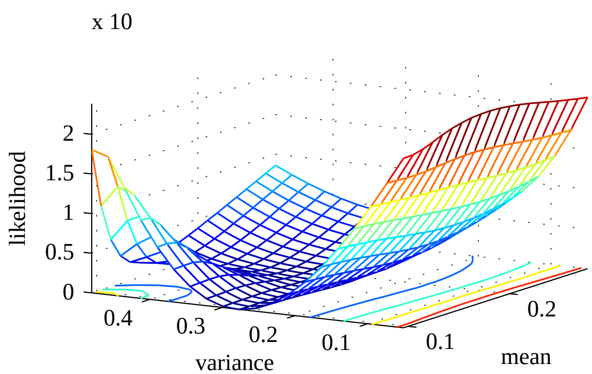

Applications of the integrated correlated variance

- to price variance and volatility derivatives,
- to use H-W and $\pi(\bar{\sigma}_t^2)$ when computing hedging ratios
- to price digital options

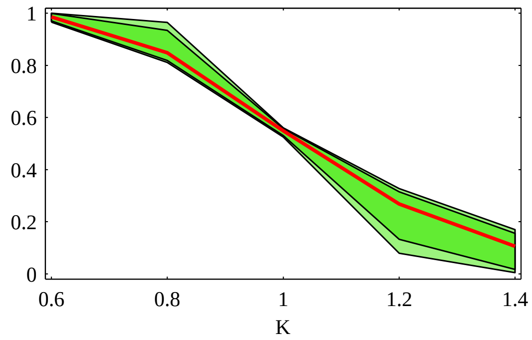

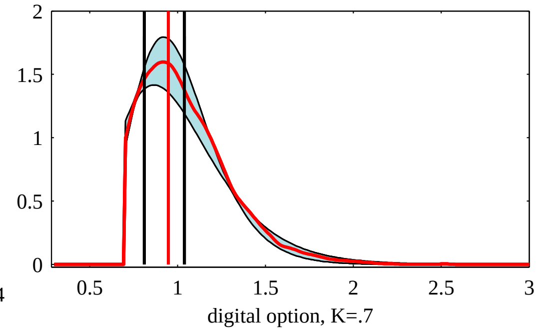

Kiitos !