Local and Stochastic Volatility

Johannes Ruf

Oxford-Man Institute of Quantitative Finance
johannes.ruf@oxford-man.ox.ac.uk

Slides available at
`http://www.oxford-man.ox.ac.uk/~jruf`

June, 2012

Some first remarks

- This mini course only touches on a few themes in the world of local and stochastic volatility models.
- The course puts more emphasis on models used for pricing and hedging than on models used for estimation.
- This presentation is partially based on notes by Michael Monoyios and Sergey Nadtochiy.

Definition: spot volatility

- In the Black-Scholes (BS) model,

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma S_t \mathrm{d}B_t,$$

  the constant σ is the *(spot) volatility* of S.

- Observe that

$$\lim_{\Delta t \downarrow 0} \frac{1}{\Delta t} \mathrm{Var}\left(\log \left(\frac{S_{t+\Delta t}}{S_t} \right) \right) = \frac{\mathrm{d}}{\mathrm{d}t} [\log(S)]_t = \sigma^2. \quad (1)$$

- We can use (1) as a definition for spot volatility in general models of the form

$$\mathrm{d}S_t = \mu_t S_t \mathrm{d}t + \sigma_t S_t \mathrm{d}B_t$$

  (in particular, without jumps) and have

$$\frac{\mathrm{d}}{\mathrm{d}t} [\log(S)]_t = \sigma_t^2.$$

Estimation: spot volatility

- Usually, the volatility is easier to estimate than the drift.
- Large amount of current research on high-frequency / *tick* data.
- In a simple setup: assume discrete price observations $s_i = S_{i\Delta t}(\omega)$ for $i = 0, \dots, n$ with $n\Delta t = T$.
- Define logarithmic returns $r_i := \log(s_i/s_{i-1})$ for $i = 1, \dots, n$.
- Itô's formula yields

$$\mathrm{d}\log(S_t) = \left(\mu_t - \frac{1}{2}\sigma_t^2\right)\mathrm{d}t + \sigma_t \mathrm{d}B_t.$$

- Thus,

$$R_i = \log\left(\frac{S_{i\Delta t}}{S_{(i-1)\Delta t}}\right) = \int_{(i-1)\Delta t}^{i\Delta t}\left(\mu_t - \frac{1}{2}\sigma_t^2\right)\mathrm{d}t + \int_{(i-1)\Delta t}^{i\Delta t}\sigma_t \mathrm{d}B_t.$$

Estimation: spot volatility II

- Assume, for a moment, that $\mu_t \equiv \mu$ and $\sigma_t \equiv \sigma$.
- Then,

$$R_i \sim \mathcal{N}\left(\left(\mu - \frac{1}{2}\sigma^2\right)\Delta t, \sigma^2 \Delta t\right).$$

- Thus, maximum likelihood estimator $\widehat{\sigma}^2$ of σ^2 is, with $\bar{r} = \sum_{i=1}^n r_i/n$,

$$\widehat{\sigma}^2 = \frac{1}{\Delta t} \cdot \frac{1}{n} \sum_{i=1}^n (r_i - \bar{r})^2 = \frac{1}{T} \sum_{i=1}^n (r_i - \bar{r})^2.$$

- How can we find an estimator for σ_t^2 if σ_t is not a constant?

Estimation: spot volatility III

- Remember:

$$\int_0^T \sigma_t^2 \mathrm{d}t = [\log(S)]_T.$$

- Thus,

$$\int_0^T \sigma_t^2 \mathrm{d}t \approx \sum_{i=1}^n r_i^2,$$

where the right-hand side is the approximate quadratic variation.

- If $\sigma_t \equiv \sigma$, this then yields an estimator $\widetilde{\sigma}^2$ of σ^2:

$$\widetilde{\sigma}^2 = \frac{1}{T}\sum_{i=1}^n r_i^2 \approx \frac{1}{T}\sum_{i=1}^n (r_i - \bar{r})^2 = \widehat{\sigma}^2.$$

Estimation: spot volatility IV

- General theory yields (remember, we did not allow for jumps) that

$$\sum_{i=1}^{n} r_i^2 \xrightarrow{\mathbb{P}} \int_0^T \sigma_t^2 \mathrm{d}t$$

as $\Delta t \downarrow 0$.

- Thus, we can estimate $\int_0^T \sigma_t^2 \mathrm{d}t$ consistently with high-frequency data.
- However, keep in mind that model assumptions (diffusion) do not describe well super-high frequency data (ticker size, ...).
- If we know $\int_0^T \sigma_t^2 \mathrm{d}t$ for all $T > 0$, then we can determine σ_t^2 Lebesgue-almost everywhere. (More cannot be expected.) However the integrated version seems to be the more natural quantity in any case.

Estimation: spot volatility V

- Consider again the case $\sigma_t \equiv \sigma$.
- One then needs to decide on a choice of T, often 30 or 180 days when using daily observations.
- Development of more sophisticated methods, e.g. (exponentially weighted) moving averages (RiskMetrics)
- See notes for an example with Dow Jones data.

Realized versus implied volatility

- Realized volatility estimate: based on historical data (past observations).
- Implied volatility: based on current market prices.
- Observe that BS-pricing formula (for calls and puts) implies, as a function of σ, an inverse function. For each price (in its range), there exists a unique σ, which, when put into the BS formula, yields that price.
- Given a market price $C_t^{\mathrm{MKT}}(T,K)$, the implied volatility $\Sigma_t(T,K)$ is the unique volatility, that solves

$$C_t^{\mathrm{BS}}(\Sigma_t(T,K)^2, S_t, T, K) = C_t^{\mathrm{MKT}}(T,K).$$

- The function Σ_t is called *volatility surface*.
- If BS model was correct, $\Sigma_t(\cdot,\cdot)$ would be constant.

Implied volatility of SP500 index options

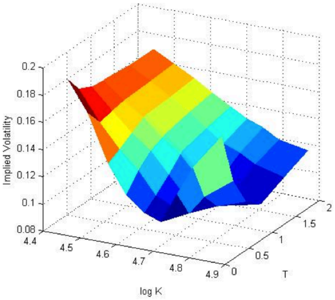

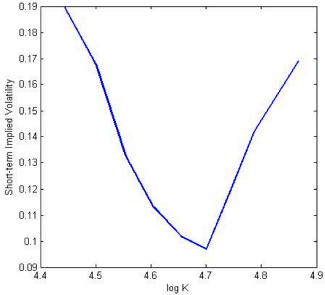

Black-Scholes and implied volatility

- BS assumption of constant implied volatility clearly does not hold in markets where calls and puts are liquidly traded (otherwise, implied volatilities cannot be observed).
- Graphs illustrate that implied volatilities today change as maturity and strike changes.
- Moreover, implied volatilities for fixed maturities and strikes also change over time; that is, $\Sigma_t(T, K)$ as a function of t is not constant.

VIX

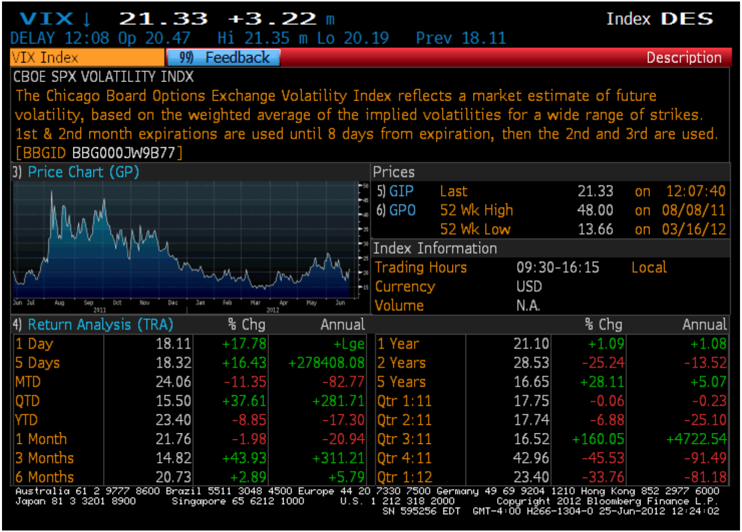
VIX ↓ 21.33 +3.22 m
Index DES
DELAY 12:08 Op 20.47 Hi 21.35 m Lo 20.19 Prev 18.11
VIX Index
99) Feedback
Description
CBOE SPX VOLATILITY INDX
The Chicago Board Options Exchange Volatility Index reflects a market estimate of future volatility, based on the weighted average of the implied volatilities for a wide range of strikes. 1st & 2nd month expirations are used until 8 days from expiration, then the 2nd and 3rd are used.
[BBGID BBG000JW9B77]
3) Price Chart (GP)
Prices
5) GIP Last 21.33 on 12:07:40
6) GPO 52 Wk High 48.00 on 08/08/11
52 Wk Low 13.66 on 03/16/12
Index Information
Trading Hours 09:30-16:15 Local
Currency USD
Volume N.A.
4) Return Analysis (TRA) % Chg Annual % Chg Annual
1 Day 18.11 +17.78 +Lge
5 Days 18.32 +16.43 +278408.08
MTD 24.06 -11.35 -82.77
QTD 15.50 +37.61 +281.71
YTD 23.40 -8.85 -17.30
1 Month 21.76 -1.98 -20.94
3 Months 14.82 +43.93 +311.21
6 Months 20.73 +2.89 +5.79
1 Year 21.10 +1.09 +1.08
2 Years 28.53 -25.24 -13.52
5 Years 16.65 +28.11 +5.07
Qtr 1:11 17.75 -0.06 -0.23
Qtr 2:11 17.74 -6.88 -25.10
Qtr 3:11 16.52 +160.05 +4722.54
Qtr 4:11 42.96 -45.53 -91.49
Qtr 1:12 23.40 -33.76 -81.18
Australia 61 2 9777 8600 Brazil 5511 3048 4500 Europe 44 20 7330 7500 Germany 49 69 9204 1210 Hong Kong 852 2977 6000
Japan 81 3 3201 8900 Singapore 65 6212 1000 U.S. 1 212 318 2000 Copyright 2012 Bloomberg Finance L.P.
SN 595256 EDT GMT-4:00 H266-1304-0 25-Jun-2012 12:24:02

VIX II


VIX ↑ 21.36 +3.25 m
Index GP
DELAY 12:15 Op 20.47 Hi 21.36 m Lo 20.19 Prev 18.11
VIX Index
95) Save As
96) Actions
97) Edit
98) Table
Line Chart
06/27/2007 - 06/25/2012 Last Price Line
11) Compare Mov. Avgs Volume USD
1D 3D 1M 6M YTD 1Y 5Y Max Daily
Security/Study Event
Australia 61 2 9777 8600 Brazil 5511 3048 4500 Europe 44 20 7330 7500 Germany 49 69 9204 1210 Hong Kong 852 2977 6000
Japan 81 3 3201 8900 Singapore 65 6212 1000 U.S. 1 212 318 2000 Copyright 2012 Bloomberg Finance L.P.
SN 595256 H266-1304-0 25-Jun-12 12:30:20 EDT GMT-4:00

Modeling of implied volatilities

- A word of warning: In this course, we shall not model implied volatilities;
- instead, we shall model the process σ_t.
- Implied volatilities correspond, in some sense, to averages of future realizations of paths of σ_t.
- The direct modeling of implied volatilities is highly complex; in particular to check whether these models satisfy standard no-arbitrage conditions, e.g.
  - call prices are convex functions in strike K,
  - call prices are increasing functions in maturity T.

Stylized facts

- Volatility clustering and persistence: small price moves follow small moves, large moves follow large moves (high autocorrelation of volatility measures).
- Thick tails: distribution of asset returns have heavier tails than normal distribution (*leptokurtic distribution*).
- Negative correlation between prices and volatility: when prices go down, volatility tends to rise (*leverage effect*).
- Mean reversion: volatility tends to revert to some long-run level.

BS model does not capture these stylized facts.

Autocorrelation of squared returns

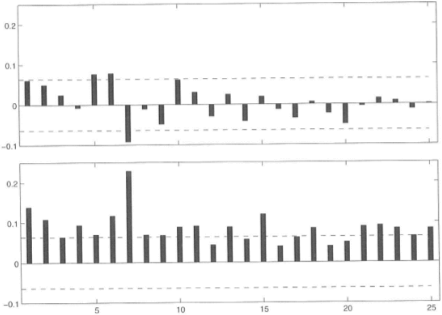

Fig. 6.3. Top: Autocorrelation function of NASDAQ log return series. Bottom: Autocorrelation function of the corresponding squared series.

From Lai and Xing, Statistical Models and Methods for Financial Markets (Springer).

Various volatility models

- Pricing and hedging models. In increasing generality:
  - deterministic models: $\sigma_t = \sigma(t)$ only function of t,
  - *local volatility models*: $\sigma_t = \sigma(t, S_t)$ function of t and S_t,
  - stochastic volatility models: additional stochastic factors, e.g. SABR, Heston.
- Econometric models (mainly for estimation / forecasting): ARCH, GARCH, EGARCH, IGARCH, ARMA-EGARCH, ...

Econometric models

- Usually formulated in discrete time, suitable for statistical estimation ("time series" analysis).
- Motivated by an attempt to model volatility clustering.
- Heteroskedastic = variance / volatility can change.
- *Autoregressive Conditional Heteroskedastic (ARCH)* models: With P_i denoting the log-price at times $i\Delta t$,

$$P_i = P_{i-1} + \alpha + \eta_i \varepsilon_i,$$

where α represents the trend, $\{\varepsilon_i\}_{i\in\mathbb{N}}$ is a family of i.i.d. standard normally distributed random variables and

$$\eta_i^2 = \beta_0 + \sum_{j=1}^{k} \beta_j (P_{i-j} - P_{i-j-1} - \alpha)^2.$$

Thus, large observed recent price changes increase volatility of next price change.

Econometric models II

- *Generalized ARCH (GARCH)* models depend also on past values of η_i.
- Usually a GARCH model needs less parameters. (ARCH models often need large k to get good fit.)
- Analytic expressions for maximum-likelihood estimators or forecasted volatilities are available.
- In EGARCH models, one distinguishes between positive and negative returns.
- For details, see for example books by Lai and Xing, Statistical Models and Methods for Financial Markets (Springer) or Tsay, Analysis of Financial Time Series (Wiley).

Outline from now on

Pricing and hedging models:

1. deterministic models: $\sigma_t = \sigma(t)$ only function of t
2. *local volatility models*: $\sigma_t = \sigma(t, S_t)$ function of t and S_t
3. stochastic volatility models: additional stochastic factors

Deterministic volatility models

- Simplest generalization of BS:

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma(t) S_t \mathrm{d}B_t,$$

where $\sigma(\cdot)$ is a deterministic function of t.

- Pricing and hedging of contingent claims is basically the same as in BS.
- Market is *complete* (all contingent claims can be replicated by trading in the underlying).
- Contingent claim price $v(t, S_t)$ corresponding to terminal payoff $h(S_T)$ usually satisfies BS PDE (assume no dividends)

$$v_t(t,s) + rsv_s(t,s) + \frac{1}{2}\sigma(t)^2 v_{s,s}(t,s) - rv(t,s) = 0,$$
$$v(T,s) = h(s).$$

Deterministic volatility models II

- By the Feynman-Kac theorem,

$$v(t,s) = E^{\mathbb{Q}}\left[e^{-r(T-t)}h(S_T)|S_t = s\right],$$

where S has $\mathbb{Q}$-dynamics

$$\mathrm{d}S_t = rS_t\mathrm{d}t + \sigma(t)S_t\mathrm{d}B_t^{\mathbb{Q}}.$$

- Now, solve SDE to obtain

$$\log(S_T) = \log(S_t) + r(T-t) - \frac{1}{2}\int_t^T \sigma(u)^2\mathrm{d}u + \int_t^T \sigma(u)\mathrm{d}B_u^{\mathbb{Q}}$$

to observe that distribution of S_T, given $S_t = s$, is normal:

$$\log(S_T) \sim \mathcal{N}\left(\log(s) + \left(r - \frac{1}{2}\bar{\sigma}_t^2\right)(T-t), \bar{\sigma}_t^2(T-t)\right)$$

where

$$\bar{\sigma}_t^2 = \frac{1}{T-t}\int_t^T \sigma(u)^2\mathrm{d}u.$$

Deterministic volatility models III

- Thus, in all BS pricing formulas for European, path-independent contingent claims, just replace σ by $\bar{\sigma}_t$.
- E.g., price of a call option at time t if $S_t = s$ is given by

$$C_t^{\mathrm{BS}}(\bar{\sigma}_t^2, s, T, K) = s\mathcal{N}(d_1) - e^{-r(T-t)}K\mathcal{N}(d_2),$$

where

$$d_1 = \frac{\log(s/K) + (r + \bar{\sigma}_t^2)(T-t)}{\bar{\sigma}_t\sqrt{T-t}},$$

$$d_2 = d_1 - \bar{\sigma}_t\sqrt{T-t},$$

$$\bar{\sigma}_t^2 = \frac{1}{T-t}\int_t^T \sigma(u)^2 \mathrm{d}u.$$

Definition: local volatility model

- Further generalization of BS:

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma(t, S_t) S_t \mathrm{d}B_t.$$

- The deterministic function

$$(t, s) \to \sigma(t, s)$$

  is called *local volatility*.
- Model is complete.
- Option price $v(t, S_t)$ for terminal payoff $h(S_T)$ usually satisfies BS PDE

$$v_t(t, s) + rsv_s(t, s) + \frac{1}{2}\sigma(t, s)^2 v_{s,s}(t, s) - rv(t, s) = 0,$$
$$v(T, s) = h(s).$$

Example: Constant Elasticity of Variance (CEV) model

- Important example:

$$\sigma(t, s) = \delta s^{\beta}$$

  for $\delta > 0$ and usually $\beta \leq 0$.
- Case $\beta > 0$ needs care (strict local martingality).
- $\beta = -1/2$: Cox-Ingersoll-Ross / square-root process (well-known from modeling interest rates).
- $\beta < 0$ yields *leverage effect*: spot volatility increases as asset price declines.
- Generally, be cautious concerning possibly positive probability of hitting zero.
- Process is analytically tractable (including analytic formulas for barrier and lookback options).
- Can be extend (by making δ stochastic) to SABR model.

Another example: Quadratic Normal Volatility model

- Risk-neutral dynamics described by

$$\mathrm{d}S_t = (aS_t^2 + bS_t + c)\mathrm{d}B_t.$$

- Here,

$$\sigma(t, s) = as + b + \frac{c}{s}.$$

- Used to price Foreign Exchange options.
- Strict local martingality is again an issue,
- but process is analytically tractable.

Calibration

- We shall assume that there is a liquid market for “vanilla” calls / puts.
- Then, these contingent claims are marked to market,
- and could be used as hedging instruments to price exotic contingent claims.
- E.g., vanillas written on major indices (SP500, SP100, DJ, DAX, FTSE, ...), large stocks, currencies.
- Throughout this section, we shall try to find a local volatility function σ, such that the model prices (expectations under the risk-neutral measure) agree with the observed market prices (*calibration*).
- In other words, we are trying to choose a certain distribution among a class of distributions.

Digression: Breeden-Litzenberger formula

- Assume that S is a Markov process with a density $p(t, s, T, \cdot)$ for S_T, conditioned on $S_t = s$.
- Then,

$$C_t(s, T, K) = e^{-r(T-t)} \int_0^\infty p(t, s, T, y)(y - K)^+ \mathrm{d}y.$$

- Thus,

$$\begin{aligned} \frac{\partial}{\partial K} C_t(s, T, K) &= -e^{-r(T-t)} \int_K^\infty p(t, s, T, y) \mathrm{d}y, \\ \frac{\partial^2}{\partial K^2} C_t(s, T, K) &= e^{-r(T-t)} p(t, s, T, K), \end{aligned}$$

- implying that observing all call prices yields the (marginal) density of S_T.

Digression: Breeden-Litzenberger formula II

Several practical issues:

- Not all call prices can be observed (as the corresponding calls are not traded).
- Call prices are only observed with "some error" (bid-ask-spread).
- Differentiation is numerically highly unstable.
- Interpolating observed call prices with a smooth function is very difficult (interpolation scheme might not be arbitrage-free [call prices monotone and convex in K]).

Kolmogorov equations

- Assume from now on that S is given by

$$\mathrm{d}S_t = rS_t\mathrm{d}t + \sigma(t, S_t)S_t\mathrm{d}B_t,$$

in particular, S is Markovian, and (assume that it) has a density $p(t, s, T, \cdot)$ for S_T, conditioned on $S_t = s$.

- Then $p(t, s, T, y)$ satisfies "BS PDE"

$$-p_t(t, s, T, y) = rsp_s(t, s, T, y) + \frac{1}{2}\sigma(t, s)^2 s^2 p_{s,s}(t, s, T, y),$$
$$p(T, s, T, y) = \delta(s - y)$$

where $\delta(\cdot)$ is Dirac delta function, satisfying

$$\int_{-\infty}^{\infty} h(y)\delta(x - y)\mathrm{d}y = h(x).$$

- This PDE is also called *Kolmogorov backward equation*.

Kolmogorov equations II

- Observed that T, y are constant in Kolmogorov backward equation.
- By a simple argument based on integrating by parts (see lecture notes), one can derive the following PDE for $p(t, s, T, y)$:

$$
\begin{aligned}
p_T(t, s, T, y) &= -r\frac{\partial}{\partial y}(yp(t, s, T, y)) + \frac{1}{2}\frac{\partial^2}{\partial y^2}\left(\sigma(T, y)^2 y^2 p(t, s, T, y)\right), \\
p(t, s, t, y) &= \delta(s - y)
\end{aligned}
$$

- This PDE is called *Kolmogorov forward equation* or *Fokker-Planck equation*.
- Now, t, s are constant.
- This PDE is quite useful as its solution corresponds to the whole surface $p(t, s, \cdot, \cdot)$.

Dupire's formula

- Multiply forward Fokker-Planck equation by $(y-K)^+$ and integrate over y to obtain

$$\frac{\partial}{\partial T}\int_K^\infty p(t,s,T,y)(y-K)\mathrm{d}y =$$
$$-\int_K^\infty r\frac{\partial}{\partial y}(yp(t,s,T,y))(y-K)\mathrm{d}y + \frac{1}{2}\int_K^\infty \frac{\partial^2}{\partial y^2}\left(\sigma(T,y)^2y^2p(t,s,T,y)\right)(y-K)\mathrm{d}y.$$

- Integrating by parts (under sufficient regularity assumptions), multiplying by $e^{-r(T-t)}$, and applying the Breeden-Litzenberger formula yield

$$e^{-r(T-t)}\frac{\partial}{\partial T}\int_K^\infty p(t,s,T,y)(y-K)\mathrm{d}y = e^{-r(T-t)}\left(\int_K^\infty ryp(t,s,T,y)\mathrm{d}y + \frac{1}{2}\sigma(T,K)^2K^2p(t,s,T,K)\right)$$
$$= rC_t(s,T,K) - rK\frac{\partial}{\partial K}C_t(s,T,K) + \frac{1}{2}\sigma(T,K)^2K^2\frac{\partial^2}{\partial K^2}C_t(s,T,K).$$

- We obtain Dupire's formula

$$\sigma(T,K)^2 = 2\frac{\frac{\partial}{\partial T}C_t(s,T,K) + rK\frac{\partial}{\partial K}C_t(s,T,K)}{K^2\frac{\partial^2}{\partial K^2}C_t(s,T,K)}.$$

Dupire's formula II

- Do not confuse Dupire's equation

$$\frac{\partial}{\partial T}C + ry\frac{\partial}{\partial y}C - \frac{1}{2}\sigma(T,y)^2 y^2 \frac{\partial^2}{\partial y^2}C = 0$$

with the BS equation

$$\frac{\partial}{\partial t}C + rs\frac{\partial}{\partial s}C + \frac{1}{2}\sigma(t,s)^2 s^2 \frac{\partial^2}{\partial s^2}C - rC = 0.$$

- Main advantage of Dupire's equation is that it treats call price as a function of strike and maturity.
- Dupire's formula can be used to calibrate a local volatility model to call prices.

Dupire's formula III

Uniqueness: Given a continuum of arbitrage-free call prices there exists at most one local vol surface which calibrates them:

$$\sigma(T,K)=\sqrt{2\frac{\frac{\partial}{\partial T}C_t(s,T,K)+rK\frac{\partial}{\partial K}C_t(s,T,K)}{K^2\frac{\partial^2}{\partial K^2}C_t(s,T,K)}}.$$

Existence: Not obvious as the SDE

$$\mathrm{d}S_t=rS_t\mathrm{d}t+\sigma(t,S_t)S_t\mathrm{d}B_t$$

needs to have a solution.

Dupire's formula IV

- Neither Dupire nor Derman-Kani (who developed a discrete-time version) thought of local volatility as a realistic model for the evolution of actual volatility.
- Local volatility can be interpreted as a "code-book", a translation of a call price surface, due to the one-to-one mapping of a (arbitrage-free) call price surface and a local volatility surface.
- Why do we need more complicated models? Local volatility models basically capture all marginal distributions.

Dupire's formula and implied volatility

- Remember: Implied volatility $\Sigma_t(T,K)$ defined via

$$C_t^{\mathrm{BS}}(\Sigma_t(T,K)^2, s, T, K) = C_t(s, T, K).$$

- Reparameterize (dimensionless variables):

$$w(T,x) = \Sigma_t(T, se^{r(T-t)}e^x)^2(T-t).$$

- Then,

$$\sigma(T, se^{r(T-t)}e^x)^2 = \frac{w_T}{1 - \frac{x}{w}w_x + \frac{1}{4}\left(-\frac{1}{4} - \frac{1}{w} + \frac{x^2}{w^2}\right)w_x^2 + \frac{1}{2}w_{x,x}}.$$

- For details, see for example Gatheral, The Volatility Surface (Wiley).

Dupire's formula: drawbacks

- Requires continuum of observed call prices.
- Prices are not exactly observed due to bid-ask spreads.
- Differentiation is numerically unstable.
- Interpolations are difficult as no-arbitrage conditions have to be guaranteed.
- More generally, local volatility models do not capture the correct dynamics. (E.g., calibrate local volatility model at times t_0 and t_1 and observe that parameters usually change completely.)

Outline for the rest of this section

Calibration of local volatility models

1. Inverse problems
2. Examples
3. Regularization
4. Application to calibration of local volatility models

Inverse problems

- A problem is called an *inverse problem* if it is defined as a inverse of some other, “more explicitly” stated problem.
- E.g., instead of going from a model (here described through the local volatility $\sigma(\cdot,\cdot)$) to option prices, going the other way around, from option prices to model parameters.
- A problem is called *well-posed* if
  1. a solution exists,
  2. the solution is unique, and
  3. the solution depends “continuously” on the data (input).
- Otherwise, the problem is called *ill-posed*.

Examples: well- and ill-posed problems

- Integration: given f, find $F(x) = \int_0^x f(z)\mathrm{d}z$. Consider an erroneous observation $\tilde{f} := f + g$ with $\sup_x |g(x)| \leq \epsilon$. Then,

$$\left| \int_0^x \tilde{f}(z)\mathrm{d}z - \int_0^x f(z)\mathrm{d}z \right| \leq \int_0^x |g(z)|\mathrm{d}z \leq \epsilon x.$$

  Thus, integration represents a well-posed problem.

- Differentiation: given F, find $f(x) = \frac{\partial}{\partial x}F(x)$. Consider an erroneous observation $\tilde{F} := F + \epsilon \sin(x/\epsilon^2)$. Then $|\tilde{F}(x) - F(x)| \leq \epsilon$, but

$$\left| \frac{\partial}{\partial x}\tilde{F}(x) - \frac{\partial}{\partial x}F(x) \right| = \frac{\left|\cos\left(\frac{x}{\epsilon^2}\right)\right|}{\epsilon} \to \infty$$

  as $\epsilon \downarrow 0$. Thus, differentiation represents an ill-posed problem.

Implied volatility of SP500 index options

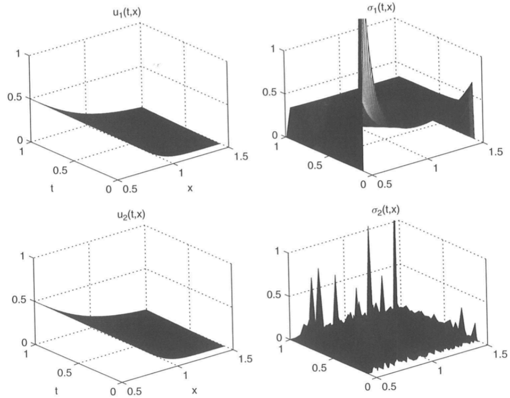

Figure 1 Extreme sensitivity of Dupire formula to noise in the data. Two examples of call price function (left) and their corresponding local volatilities (right). The prices differ through IID noise $\sim UNIF(0, 0.001)$, representing a bid–ask spread

From Cont, Encyclopedia of Quantitative Finance (Wiley).

Tikhonov regularization

- Consider an abstract ill-posed inverse problem

$$F(x) = y, x \in X, y \in Y.$$

- Solution may not exist, may not be unique, or may be unstable.
- Remedy: Solve a *regularized* optimization problem instead:

$$\min_{x \in X} \|F(x) - y\|_Y^2 + \alpha G(x),$$

  where G is some (convex) penalty function, e.g., $G(x) = \|x - x_0\|_X^2$ with x_0 a "prior" guess.
- *Regularization factor* α needs to be determined by cross-validation.

Tikhonov regularization applied to local vol. calibration

- Market prices $V^{\text{mkt}}(T_i, K_i)$ and model prices $V^{\sigma}(T_i, K_i)$ (produced in a model with local volatility $\sigma(\cdot,\cdot)$)
- V: usually out-of-the money calls / puts.
- Consider, for some weights w_i,

$$\min_{\sigma\in\mathfrak{S}} \sum_{i=1}^{n} w_i \left(V^{\text{mkt}}(T_i, K_i) - V^{\sigma}(T_i, K_i)\right)^2 + \alpha\|\sigma\|^2.$$

- For example,

$$\|\sigma\|^2 = \int_0^T \int_0^\infty \left(\frac{\partial^2\sigma}{\partial K^2}\right)^2 + \left(\frac{\partial\sigma}{\partial T}\right)^2 \mathrm{d}K\mathrm{d}t$$

  ("Flatter" volatility surfaces are preferred)
- Choice of weights w_i should depend on liquidity of corresponding options.
- Choice of $\mathfrak{S}$: often, spline-based representation (local volatility is parameterized by finitely many values, thus optimization problem is finite-dimensional problem).

Tikhonov reg. applied to local vol. calibration II

Advantages:

- No need to have continuum of observed strikes and maturities.
- No need to interpolate market prices.
- Convex penalization leads to numerical stability.
- Calibrated surface is smooth due to choice of penalization norm.

Disadvantages:

- Computationally demanding.
- Penalization criterion does not include weights to take into account distribution of S_t; thus, criterion overweights values with small probability of occurrence.

Problems of local volatility models

- Local volatility models can perfectly fit marginals (European-style path-independent options),
- but have problems with pricing path-dependent options, and
- their dynamics are not realistic.

Formal description: stochastic volatility models

- Allow now σ_t to be a stochastic process:

$$
\begin{aligned}
\mathrm{d}S_t &= \mu(t, S_t, Y_t)S_t\mathrm{d}t + \sigma(t, S_t, Y_t)S_t\mathrm{d}B_t, && (2)\\
\mathrm{d}Y_t &= a(t, S_t, Y_t)\mathrm{d}t + b(t, S_t, Y_t)\mathrm{d}W_t, && (3)\\
\mathrm{d}\langle B, W\rangle_t &= \rho\mathrm{d}t.
\end{aligned}
$$

- We can write $W_t = \rho B_t + \sqrt{1-\rho^2}Z_t$, where Z is a BM independent of B.
- If the underlying asset (with price process S_t) is the only hedging instrument, then the market is *incomplete* since, in general, contingent claims written on Y cannot be replicated by trading in S only.
- $\rho < 0$ corresponds to the *leverage effect*.

Risk-neutral measures

- Denote by $\mathbb{P}$ the *historical* (*physical*) probability measure under which S and Y have the dynamics of (2) and (3).
- If we assume *no-arbitrage* (*NFLVR* to be precise), then there exists a measure $\mathbb{Q}$, equivalent to $\mathbb{P}$, such that $e^{-rt}S_t$ is a local martingale under $\mathbb{Q}$. (*1st Fundamental Theorem of Asset Pricing*)
- However, $\mathbb{Q}$ might not need to be unique. Indeed, it is unique if and only if the model is complete (each contingent claim can be replicated perfectly by trading in S alone). (*2nd Fundamental Theorem of Asset Pricing*)

Risk-neutral measures II

- Denote the class of equivalent local martingale measures by $\mathcal{M}$.
- Any $\mathbb{Q} \in \mathcal{M}$ is characterized by its *stochastic discount factor / Radon-Nikodym derivative* with respect to $\mathbb{P}$:

$$Z_t^{\mathbb{Q}} := \left.\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right|_{\mathcal{F}_t} = \mathcal{E}(-\lambda \cdot B - \psi \cdot Z)_t,$$

where $\mathcal{E}$ is the *Doléans-Dade (stochastic) exponential*

$$\mathcal{E}(X)_t := e^{X_t - X_0 - \frac{1}{2}[X]_t},$$

and

$$\lambda_t := \frac{\mu(t, S_t, Y_t) - r}{\sigma(t, S_t, Y_t)},$$

and ψ is progressively measurable.

Risk-neutral measures III

- As we have seen, λ is fixed by the model (representing the *market price of risk*), but ψ is an (almost) arbitrary process.
- For any $\mathbb{Q} \in \mathcal{M}$, we have $\mathbb{E}[Z_T^{\mathbb{Q}}] = 1$.
- Thus, this yields a necessary condition on ψ. For this condition to hold, Novikov's condition is sufficient:

$$\mathbb{E}\left[e^{\frac{1}{2}\int_0^T (\lambda_t^2 + \psi_t^2)\mathrm{d}t}\right] < \infty.$$

Risk-neutral measures IV

- Under $\mathbb{Q}$, we obtain that $B^{\mathbb{Q}}, X^{\mathbb{Q}}$ are independent BM with

$$B_t^{\mathbb{Q}} = B_t + \int_0^t \lambda(u, S_u, Y_u)\mathrm{d}u,$$
$$Z_t^{\mathbb{Q}} = Z_t + \int_0^t \psi_u \mathrm{d}u.$$

- Furthermore, we obtain the dynamics

$$\begin{aligned}
\mathrm{d}S_t &= rS_t\mathrm{d}t + \sigma(t, S_t, Y_t)S_t\mathrm{d}B_t^{\mathbb{Q}},\\
\mathrm{d}Y_t &= \left(a(t, S_t, Y_t) - b(t, S_t, Y_t)\left(\rho\lambda(t, S_t, Y_t) + \sqrt{1-\rho^2}\psi_t\right)\right)\mathrm{d}t\\
&\quad + b(t, S_t, Y_t)\mathrm{d}W_t^{\mathbb{Q}},\\
\mathrm{d}\langle B^{\mathbb{Q}}, W^{\mathbb{Q}}\rangle_t &= \rho\mathrm{d}t.
\end{aligned}$$

Arbitrage-free prices

- Consider a claim that pays $h(S_T, Y_T)$ at time T.
- Denote

$$C_t^{\mathbb{Q}} = \mathbb{E}^{\mathbb{Q}}\left[e^{-r(T-t)}h(S_T, Y_T)|\mathcal{F}_t\right]$$

- If (S, Y) is Markovian under $\mathbb{Q}$, e.g. if $\psi_t = \psi(t, S_t, Y_t)$, we have $C_t^{\mathbb{Q}} = v^{\mathbb{Q}}(t, S_t, Y_t)$ for

$$v^{\mathbb{Q}}(t, s, y) = \mathbb{E}^{\mathbb{Q}}\left[e^{-r(T-t)}h(S_T, Y_T)|S_t = s, Y_t = y\right].$$

- Clearly, $C^{\mathbb{Q}}$ and $v^{\mathbb{Q}}$ depend on the choice of risk-neutral measure $\mathbb{Q}$.
- All possible, arbitrage-free prices of the claim with payoff $h(S_T, Y_T)$, are given by the interval

$$\left[\inf_{\mathbb{Q}\in\mathcal{M}} C_t^{\mathbb{Q}}, \sup_{\mathbb{Q}\in\mathcal{M}} C_t^{\mathbb{Q}}\right].$$

Choice of risk-free measure

- If $\inf_{\mathbb{Q}\in\mathcal{M}} C_t^{\mathbb{Q}} = \sup_{\mathbb{Q}\in\mathcal{M}} C_t^{\mathbb{Q}}$, then there is a unique price and the claim can be perfectly replicated, by the general theory. (e.g., usually [but not always] $h(S_T, Y_T) = S_T$).
- Otherwise, distinguish two cases:
  1. Only S is a liquidly traded asset, and there are no other hedging instruments available.
  2. There is another hedging instrument available (e.g., a call).
- In case 1, we have a problem. Possible approaches:
  - Reconsider, whether the model should be changed.
  - Find a hedging strategy that minimizes risk (e.g., VaR), quadratic hedging, ...
  - Take an ad-hoc measure: *Minimal Martingale (Entropy) Measure*, ...
  - Choose the superreplicating price: $\sup_{\mathbb{Q}\in\mathcal{M}} C_t^{\mathbb{Q}}$. This often is a very conservative approach.

Portfolio dynamics and lack of replication

- If an investor holds Δ_t units of S at each time t and keeps / borrows the remainder in / from the bank account (self-financing strategy), then her wealth process $X := \{X_t\}_{t\in[0,T]}$ satisfies

$$\begin{aligned}\mathrm{d}X_t &= rX_t\mathrm{d}t + \Delta_t S_t \sigma(t, S_t, Y_t)\left(\lambda(t, S_t, Y_t)\mathrm{d}t + \mathrm{d}B_t\right)\\ &= rX_t\mathrm{d}t + \Delta_t S_t \sigma(t, S_t, Y_t)\mathrm{d}B_t^{\mathbb{Q}}.\end{aligned}$$

- Price of a claim $C_t^{\mathbb{Q}} = v^{\mathbb{Q}}(t, S_t, Y_t)$ satisfies (under $\mathbb{P}$):

$$\begin{aligned}\mathrm{d}C_t^{\mathbb{Q}} = &(v_t^{\mathbb{Q}} + \mathcal{A}v^{\mathbb{Q}})\mathrm{d}t + (\sigma S_t v_s^{\mathbb{Q}} + \rho b v_y^{\mathbb{Q}})\mathrm{d}B_t\\ &+ \sqrt{1-\rho^2}\, b v_y^{\mathbb{Q}}\mathrm{d}Z_t,\end{aligned}$$

where

$$\mathcal{A}v = \frac{1}{2}\sigma^2 s^2 v_{s,s} + \frac{1}{2}b^2 v_{y,y} + \rho\sigma s b v_{s,y} + \mu s v_s + a v_y.$$

- It is simple to see that usually perfect hedging with S only is impossible.

Market completion

- Introduce another traded asset, for example, a contingent claim (that cannot be replicated by trading in S only) with payoff $g(S_{\widetilde{T}}, Y_{\widetilde{T}})$ at time $\widetilde{T} \geq T$. Denote its price at time t by $O_t = u(t, S_t, Y_t)$.
- Now,

$$\begin{aligned} \mathrm{d}O_t = &(u_t + \mathcal{A}u)\mathrm{d}t + (\sigma S_t u_s + \rho b u_y)\mathrm{d}B_t \\ &+ \sqrt{1-\rho^2} b u_y \mathrm{d}Z_t, \\ &= (u_t + \mathcal{A}u - \lambda(\sigma S_t u_s + \rho b u_y) - \psi\sqrt{1-\rho^2} b u_y)\mathrm{d}t \\ &+ (\sigma S_t u_s + \rho b u_y)\mathrm{d}B_t^{\mathbb{Q}} + \sqrt{1-\rho^2} b u_y \mathrm{d}Z_t^{\mathbb{Q}}. \end{aligned}$$

- Thus, there exists exactly one ψ (assuming $u_y \neq 0$) such that the drift term equals ru under the corresponding $\mathbb{Q}$.
- Thus, with S and O as traded assets, there exists only one risk-neutral martingale measure and thus, the market is complete.

Market completion II

- An investor can hold Δ_t units of S, N_t units of O, and the bank account. Her wealth process X satisfies

$$\begin{aligned}
\mathrm{d}X_t &= \Delta_t \mathrm{d}S_t + N_t \mathrm{d}O_t + r(X_t - \Delta_t S_t - N_t O_t)\mathrm{d}t \\
&= \left(r(X_t - N_t u) + \Delta_t \sigma\lambda S_t + N_t(u_t + \mathcal{A}u)\right)\mathrm{d}t \\
&\quad + (\Delta_t \sigma S_t + N_t(\sigma S_t u_s + \rho b u_y))\mathrm{d}B_t + N_t\sqrt{1-\rho^2}\, b u_y \mathrm{d}Z_t.
\end{aligned}$$

- In order to hedge a contingent claim perfectly, we need a wealth process with $X_T = h(S_T, Y_T)$.
- We continue by setting the $\mathrm{d}Z$ and $\mathrm{d}B$-terms equal in the dynamics of X and $C^{\mathbb{Q}}$:

$$\begin{aligned}
\mathrm{d}C_t^{\mathbb{Q}} &= (v_t^{\mathbb{Q}} + \mathcal{A}v^{\mathbb{Q}})\mathrm{d}t + (\sigma S_t v_s^{\mathbb{Q}} + \rho b v_y^{\mathbb{Q}})\mathrm{d}B_t \\
&\quad + \sqrt{1-\rho^2}\, b v_y^{\mathbb{Q}} \mathrm{d}Z_t,
\end{aligned}$$

Market completion III

- Equating the $\mathrm{d}Z$-terms yields

$$N_t u_y = v_y^{\mathbb{Q}}.$$

This component of the hedging strategy is called "Vega hedging".

- Now, equating the $\mathrm{d}B$-terms yields

$$\Delta_t = v_s^{\mathbb{Q}} - \frac{v_y^{\mathbb{Q}}}{u_y} u_s.$$

This component of the hedging strategy is called "Delta hedging".

- Thus, the "volatility risk" of the claim is offset by the vega hedge and the delta hedge gets adjusted by the delta provided through the vega hedge.

Market completion IV

- Equating now the $\mathrm{d}t$-terms yields

$$\frac{u_t + \mathcal{A}u - \sigma\lambda S_t u_s - ru}{u_y} = \frac{v_t^{\mathbb{Q}} + \mathcal{A}v^{\mathbb{Q}} - \sigma\lambda S_t v_s^{\mathbb{Q}} - rv^{\mathbb{Q}}}{v_y^{\mathbb{Q}}}.$$

- Observe that the left-hand side equals

$$(\rho\lambda + \sqrt{1-\rho^2}\psi)b.$$

- This comes from setting the drift equal to ru in

$$\begin{aligned}\mathrm{d}O_t = {} & (u_t + \mathcal{A}u - \lambda(\sigma S_t u_s + \rho b u_y) - \psi\sqrt{1-\rho^2}bu_y)\mathrm{d}t \\ & + (\sigma S_t u_s + \rho b u_y)\mathrm{d}B_t^{\mathbb{Q}} + \sqrt{1-\rho^2}bu_y\mathrm{d}Z_t^{\mathbb{Q}}.\end{aligned}$$

Market completion V

- Thus, both u and $v^{\mathbb{Q}}$ satisfy the PDE

$$f_t + \mathcal{A}^{\mathbb{Q}} f - rf = 0,$$

where

$$\mathcal{A}^{\mathbb{Q}} f = \frac{1}{2}\sigma^2 s^2 f_{s,s} + \frac{1}{2} b^2 f_{y,y} + \rho\sigma s b f_{s,y} + r s f_s + \left(a - b\left(\rho\lambda + \sqrt{1-\rho^2}\psi\right)\right) f_y .$$

- Observe that $\mathcal{A}^{\mathbb{Q}}$ is the generator of

$$\begin{aligned} \mathrm{d}S_t &= rS_t \mathrm{d}t + \sigma(t, S_t, Y_t) S_t \mathrm{d}B_t^{\mathbb{Q}}, \\ \mathrm{d}Y_t &= \left(a(t, S_t, Y_t) - b(t, S_t, Y_t)\left(\rho\lambda(t, S_t, Y_t) + \sqrt{1-\rho^2}\psi_t\right)\right)\mathrm{d}t \\ &\quad + b(t, S_t, Y_t)\mathrm{d}W_t^{\mathbb{Q}}, \\ \mathrm{d}\langle B^{\mathbb{Q}}, W^{\mathbb{Q}}\rangle_t &= \rho \mathrm{d}t. \end{aligned}$$

- Thus, the Feynman-Kac formula yields

$$v(t, s, y) = \mathbb{E}^{\mathbb{Q}}\left[e^{-r(T-t)} h(S_T, Y_T) | S_t = s, Y_t = y\right].$$

Special case: zero correlation

- Assume for the moment that $\rho = 0$.
- Then

$$\begin{aligned} \mathrm{d}S_t &= rS_t\mathrm{d}t + \sigma(t, S_t, Y_t)S_t\mathrm{d}B_t^{\mathbb{Q}}, \\ \mathrm{d}Y_t &= (a(t, S_t, Y_t) - b(t, S_t, Y_t)\psi_t)\,\mathrm{d}t + b(t, S_t, Y_t)\mathrm{d}W_t^{\mathbb{Q}}, \\ \mathrm{d}\langle B^{\mathbb{Q}}, W^{\mathbb{Q}}\rangle_t &= 0\mathrm{d}t. \end{aligned}$$

- A call with payoff $h(S_T) = (S_T - K)^+$ has a price

$$\begin{aligned} v(t, s, y) &= \mathbb{E}^{\mathbb{Q}}\left[e^{-r(T-t)}h(S_T)|S_t = s, Y_t = y\right] \\ &= \mathbb{E}^{\mathbb{Q}}\left[\mathbb{E}^{\mathbb{Q}}[e^{-r(T-t)}h(S_T)|\mathcal{F}_t, \{\sigma_t\}_{t\geq 0}]|S_t = s, Y_t = y\right] \\ &= \mathbb{E}^{\mathbb{Q}}\left[C_t^{\mathrm{BS}}(\bar{\sigma}_t^2, s, T, K)|Y_t = y\right], \end{aligned}$$

where

$$\bar{\sigma}_t^2 := \frac{1}{T-t}\int_t^T \sigma^2(s, Y_s)\mathrm{d}s.$$

Examples of stochastic volatility models

- Hull-White model: $\rho = 0$, $\mu(\cdot,\cdot,\cdot) = \mu$, $\sigma(\cdot,\cdot,y) = \sqrt{y}$ and Y GBM:

$$\begin{aligned} \mathrm{d}S_t &= \mu S_t \mathrm{d}t + \sqrt{Y_t} S_t \mathrm{d}B_t, \\ \mathrm{d}Y_t &= a Y_t \mathrm{d}t + b Y_t \mathrm{d}W_t, \\ \mathrm{d}\langle B, W\rangle_t &= 0 \mathrm{d}t. \end{aligned}$$

- Heston model: $\mu(\cdot,\cdot,\cdot) = \mu$, $\sigma(\cdot,\cdot,y) = \sqrt{y}$ and Y a square-root / Cox-Ingersoll-Ross process (mean-reverting!):

$$\begin{aligned} \mathrm{d}S_t &= \mu S_t \mathrm{d}t + \sqrt{Y_t} S_t \mathrm{d}B_t, \\ \mathrm{d}Y_t &= a(b - Y_t) \mathrm{d}t + b\sqrt{Y_t} \mathrm{d}W_t, \\ \mathrm{d}\langle B, W\rangle_t &= \rho \mathrm{d}t. \end{aligned}$$

- Standard models are mainly chosen due to their analytic tractability.

Another class of stochastic volatility models

- Sometimes, stochastic volatility models are specified through a time-change $V_t = \int_0^t v_s \mathrm{d}s$.
- The time change V_t has an interpretation as *business time* (in contrast to versus *calendar time*).
- E.g.,

$$S_t = e^{rt+B_{V_t}}.$$

- If B and V are independent, this is just another representation of the stochastic volatility model above, due to the self-similarity of BM ($B_{ct} \overset{d}{=} \sqrt{c}B_t$).
- Advantage: Interpretation as business time.
- Disadvantage: It is not clear, a priori, how to model the leverage effect, for example.

Hedging: stochastic volatility model vs. BS

- In theory, stochastic volatility models can help in hedging the "volatility risk", which is an improvement in comparison with the BS model.
- However, hedging ratios depend strongly upon the parameters, and are sensitive with respect to changes in parameters (recalibration).
- This is another example of an ill-posed problem.
- Often, a simple model does better since its parameters can be calibrated more efficiently, are more robust, and hedging errors can get "averaged out".