

Consistent Modeling of SPX and VIX options

Jim Gatheral



The Fifth World Congress of the Bachelier Finance Society
London, July 18, 2008

Disclaimer

The opinions expressed in this presentation are those of the author alone, and do not necessarily reflect the views of of Merrill Lynch, its subsidiaries or affiliates.

Motivation and context

- We would like to have a model that prices consistently
 - ① options on SPX
 - ② options on VIX
 - ③ options on realized variance
- We believe there may be such a model because we can identify relationships between options on SPX, VIX and variance. For example:
 - ① Puts on SPX and calls on VIX both protect against market dislocations.
 - ② Bruno Dupire constructs an upper bound on the price of options on variance from the prices of index options.
 - ③ The underlying of VIX options is the square-root of a forward-starting variance swap.
- VIX options give us the risk-neutral distribution of forward-starting variance swaps.
 - ① How do they constrain volatility dynamics in stochastic volatility models?

Outline

- 1 Historical development
 - Problems with one-factor stochastic volatility models.
 - Historical attempts to add factors.
- 2 Variance curve models
 - Bergomi's variance curve model.
 - Buehler's consistent variance curve functionals.
 - Double Lognormal vs Double Heston
- 3 The Double CEV model
 - Model calibration
 - Parameter stability
- 4 Time series analysis
 - Statistics of model factors
- 5 Options on realized variance
- 6 Conclusion

Problems with one-factor stochastic volatility models

- All volatilities depend only on the instantaneous variance v
 - Any option can be hedged perfectly with a combination of any other option plus stock
 - Skew, appropriately defined, is constant
- We know from PCA of volatility surface time series that there are at least three important modes of fluctuation:
 - level, term structure, and skew
- It makes sense to add at least one more factor.

Other motivations for adding another factor

- Adding another factor with a different time-scale has the following benefits:
 - One-factor stochastic volatility models generate an implied volatility skew that decays as $1/T$ for large T . Adding another factor generates a term structure of the volatility skew that looks more like the observed $1/\sqrt{T}$.
 - The decay of autocorrelations of squared returns is exponential in a one-factor stochastic volatility model. Adding another factor makes the decay look more like the power law that we observe in return data.
 - Variance curves are more realistic in the two-factor case. For example, they can have humps.

Historical attempts to add factors

- Dupire's unified theory of volatility (1996)
 - Local variances are driftless in the butterfly measure.
 - We can impose dynamics on local variances.
- Stochastic implied volatility (1998)
 - The implied volatility surface is allowed to move.
 - Under diffusion, complex no-arbitrage condition, impossible to work with in practice.
- Variance curve models (1993-2005)
 - Variances are tradable!
 - Simple no-arbitrage condition.

Dupire's unified theory of volatility

- The price of the calendar spread $\partial_T C(K, T)$ expressed in terms of the butterfly $\partial_{K,K} C(K, T)$ is a martingale under the measure $Q_{K,T}$ associated with the butterfly.
- Local variance $v_L(K, T)$ is given by (twice) the current ratio of the calendar spread to the butterfly.
- We may impose any dynamics such that the above holds and local variance stays non-negative.
- For example, with one-factor lognormal dynamics, we may write:

$$v(S, t) = v_L(S, t) \frac{\exp \{-b^2/2 t - b W_t\}}{\mathbb{E}[\exp \{-b^2/2 t - b W_t\} \mid S_t = S]}$$

where it is understood that $v_L(\cdot)$ is computed at time $t = 0$.
Note that the denominator is hard to compute!

Stochastic implied volatility

- The evolution of implied volatilities is modeled directly as in $\sigma_{BS}(k, T, t) = G(\mathbf{z}; k, T - t)$ with $\mathbf{z} = \{z_1, z_2, \dots, z_n\}$ for some factors z_i .
 - For example, the stochastic factors z_i could represent level, term structure and skew.
- The form of $G(\cdot)$ is highly constrained by no-arbitrage conditions
 - An option is valued as the risk-neutral expectation of future cashflows – it must therefore be a martingale.
 - Even under diffusion assumptions, the resulting no-arbitrage condition is very complicated.
- Nobody has yet written down an arbitrage-free solution to a stochastic implied volatility model that wasn't generated from a conventional stochastic volatility model.
 - SABR is a stochastic implied volatility model, albeit without mean reversion, but it's not arbitrage-free.
- *Stochastic implied volatility is a dead end!*

Why model variance swaps?

- Dupire's UTV is hard to implement because local variances are not tradable.
- Stochastic implied volatility isn't practical because implied volatilities are not tradable.
- Variance swaps are tradable.
 - Variance swap prices are martingales under the *risk-neutral measure*.
 - Moreover variance swaps are now relatively liquid
 - and forward variance swaps are natural hedges for cliquets and other exotics.
- Thus, as originally suggested by Dupire in 1993, and then latterly by Duanmu, Bergomi, Buehler and others, we should impose dynamics on forward variance swaps.

Modeling forward variance

Denote the variance curve as of time t by

$\hat{W}_t(T) = \mathbb{E} \left[\int_t^T v_s ds \mid \mathcal{F}_t \right]$. The forward variance $\zeta_t(T) := \mathbb{E} [v_T \mid \mathcal{F}_t]$ is given by

$$\zeta_t(T) = \partial_T \hat{W}_t(T)$$

A natural way of satisfying the martingale constraint whilst ensuring positivity is to impose lognormal dynamics as in Dupire's (1993) example:

$$\frac{d\zeta_t(T)}{\zeta_t(T)} = \sigma(T - t) dW_t$$

for some volatility function $\sigma(\cdot)$.

Lorenzo Bergomi does this and extends the idea to n-factors.

Bergomi's model

In the 2-factor version of his model, we have

$$\frac{d\zeta_t(T)}{\zeta_t(T)} = \xi_1 e^{-\kappa(T-t)} dW_t + \xi_2 e^{-c(T-t)} dZ_t$$

This has the solution

$$\zeta_t(T) = \zeta_0(T) \exp \left\{ \xi_1 e^{-\kappa(T-t)} X_t + \xi_2 e^{-c(T-t)} Y_t + \text{drift terms} \right\}$$

with

$$X_t = \int_0^t e^{-\kappa(t-s)} dW_s; \quad Y_t = \int_0^t e^{-c(t-s)} dZ_s;$$

Thus, both X_t and Y_t are Ornstein-Ühlenbeck processes. In particular, they are easy to simulate. The Bergomi model is a market model: $\mathbb{E}[\zeta_t(T)] = \zeta_0(T)$ for any given initial forward variance curve $\zeta_0(T)$.

Variance curve models

- The idea (similar to the stochastic implied volatility idea) is to obtain a factor model for forward variance swaps. That is,

$$\zeta_t(T) = G(\mathbf{z}; T - t)$$

with $\mathbf{z} = \{z_1, z_2, \dots, z_n\}$ for some factors z_j and some *variance curve functional* $G(\cdot)$.

- Specifically, we want \mathbf{z} to be a diffusion so that

$$d\mathbf{z}_t = \mu(\mathbf{z}_t) dt + \sum_j^d \sigma^j(\mathbf{z}_t) dW_t^j \quad (1)$$

- Note that both μ and σ are n -dimensional vectors.

Buehler's consistency condition

Theorem

The variance curve functional $G(\mathbf{z}_t, \tau)$ is consistent with the dynamics (1) if and only if

$$\begin{aligned} \partial_\tau G(\mathbf{z}; \tau) &= \sum_{i=1}^n \mu_i(\mathbf{z}) \partial_{z_i} G(\mathbf{z}; \tau) \\ &\quad + \frac{1}{2} \sum_{i,k=1}^n \left(\sum_{j=1}^d \sigma_i^j(\mathbf{z}) \sigma_k^j(\mathbf{z}) \right) \partial_{z_i, z_k} G(\mathbf{z}; \tau) \end{aligned}$$

To get the idea, apply Itô's Lemma to $\zeta_t(T) = G(z, T - t)$ with $dz = \mu dt + \sigma dW$ to obtain

$$\mathbb{E} [d\zeta_t(T)] = 0 = \left\{ -\partial_\tau G(z, \tau) + \mu \partial_z G(z, \tau) + \frac{1}{2} \sigma^2 \partial_{z,z} G(z, \tau) \right\} dt$$

Examples of consistent variance curves

- In the Heston model, $G(v, \tau) = v + (v - \bar{v}) e^{-\kappa \tau}$.
 - This variance curve functional is obviously consistent with Heston dynamics with time-independent parameters κ , ρ and η .
- Bergomi's forward variance curve with

$$\mathbb{E}[\zeta_t(T)] = \zeta_0(T)$$

is explicitly a martingale and so also obviously consistent.

Buehler's affine variance curve functional

- Consider the following variance curve functional:

$$G(\mathbf{z}; \tau) = z_3 + (z_1 - z_3) e^{-\kappa \tau} + (z_2 - z_3) \frac{\kappa}{\kappa - c} (e^{-c \tau} - e^{-\kappa \tau})$$

- This looks like the Svensson parametrization of the yield curve.
- The short end of the curve is given by z_1 and the long end by z_3 .
- The middle level z_2 adds flexibility permitting for example a hump in the curve.

Double CEV dynamics

- Buehler's affine variance curve functional is consistent with double mean reverting dynamics of the form:

$$\begin{aligned}
 \frac{dS}{S} &= \sqrt{v} dW \\
 dv &= -\kappa(v - v') dt + \eta_1 v^\alpha dZ_1 \\
 dv' &= -c(v' - z_3) dt + \eta_2 v'^\beta dZ_2
 \end{aligned} \tag{2}$$

for any choice of $\alpha, \beta \in [1/2, 1]$.

- We will call the case $\alpha = \beta = 1/2$ *Double Heston*,
- the case $\alpha = \beta = 1$ *Double Lognormal*,
- and the general case *Double CEV*.
- All such models involve a short term variance level v that reverts to a moving level v' at rate κ . v' reverts to the long-term level z_3 at the slower rate $c < \kappa$.

Check of consistency condition

- Because $G(\cdot)$ is affine in z_1 and z_2 , we have that

$$\partial_{z_i, z_j} G(\{z_1, z_2\}; \tau) = 0 \quad i, j \in \{1, 2\}.$$

- Then the consistency condition reduces to

$$\begin{aligned} \partial_\tau G(\{z_1, z_2\}; \tau) &= \sum_{i=1}^2 \mu_i(\{z_1, z_2\}) \partial_{z_i} G(\{z_1, z_2\}; \tau) \\ &= -\kappa(z_1 - z_2) \partial_{z_1} G - c(z_2 - z_3) \partial_{z_2} G \end{aligned}$$

- It is easy to verify that this holds for our affine functional.
- In fact, the consistency condition looks this simple for affine variance curve functionals with any number of factors!

Double Lognormal vs Bergomi

- Recall that the Bergomi model has dynamics (with $\tau = T - t$)

$$\frac{d\zeta_t(T)}{\zeta_t(T)} = \xi_1 e^{-\kappa\tau} dZ_1 + \xi_2 e^{-c\tau} dZ_2$$

- Now in the Double Lognormal model

$$\begin{aligned} d\zeta_t(T) &= dG(v, v'; \tau) \\ &= \xi_1 v e^{-\kappa\tau} dZ_1 + \xi_2 v' \frac{\kappa}{\kappa - c} (e^{-c\tau} - e^{-\kappa\tau}) dZ_2 \end{aligned}$$

- We see that the two sets of dynamics are very similar.
- Bergomi's model is a market model and Buehler's affine model is a factor model.
 - However any variance curve model may be made to fit the initial variance curve by writing

$$\zeta_t(T) = \frac{\zeta_0(T)}{G(\mathbf{z}_0, T)} G(\mathbf{z}_t, T)$$

Bergomi's model as a factor model

Recall that

$$\zeta_t(T) = \zeta_0(T) \exp \left\{ \xi_1 e^{-\kappa(T-t)} X_t + \xi_2 e^{-c(T-t)} Y_t + \text{drift terms} \right\}$$

with

$$X_t = \int_0^t e^{-\kappa(t-s)} dW_s; \quad Y_t = \int_0^t e^{-c(t-s)} dZ_s;$$

We should then be able to write:

$$\zeta_0(T) = \zeta_{-\infty}(T) \exp \left\{ \xi_1 e^{-\kappa T} \tilde{X} + \xi_2 e^{-c T} \tilde{Y} + \text{drift terms} \right\}$$

with

$$\tilde{X} = \int_{-\infty}^0 e^{\kappa s} dW_s; \quad \tilde{Y} = \int_{-\infty}^0 e^{c s} dZ_s;$$

so \tilde{X} and \tilde{Y} are normal rv's with covariance

$$\mathbb{E} \left[\tilde{X}^2 \right] = \frac{1}{2\kappa}; \quad \mathbb{E} \left[\tilde{Y}^2 \right] = \frac{1}{2c}; \quad \mathbb{E} \left[\tilde{X} \tilde{Y} \right] = \frac{\rho}{\kappa + c};$$

Bergomi's variance curve functional

It is easy to see that $\zeta_{-\infty}(T) = \mathbb{E}[v_t] =: \bar{v}$. Then the Bergomi variance curve functional becomes

$$G(\mathbf{z}, \tau) = \zeta_t(T) = \bar{v} \exp \left\{ \xi_1 e^{-\kappa \tau} z_1 + \xi_2 e^{-c \tau} z_2 + \text{drift terms} \right\}$$

with

$$z_1 = \int_{-\infty}^t e^{-\kappa(t-s)} dW_s; \quad z_2 = \int_{-\infty}^t e^{-c(t-s)} dZ_s$$

- We see that in this factor version of Bergomi's model, $\log \zeta_t(T)/\bar{v}$ is linear in the factors z_1 and z_2 .
 - Note that we no longer match the initial curve $\zeta_0(T)$ in general.
- It turns out that Bergomi's model generates almost no skew for VIX options so for this and other reasons, we stick with Buehler's 2-factor variance curve functional.

Discriminating between stochastic volatility models using VIX options

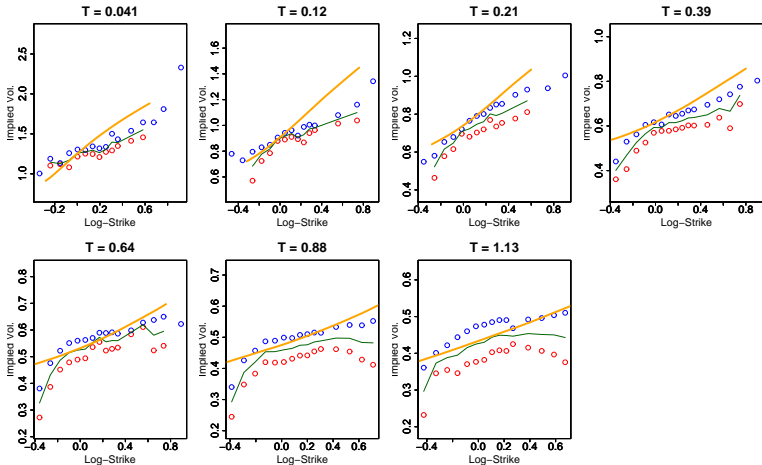
- Before fitting Double CEV with some arbitrary exponent α , let's compare Double Lognormal and Double Heston fits
 - The idea is that VIX option prices should allow us to discriminate between models.
- Specifically, we want to assess:
 - Fit quality
 - Parameter stability

Fit of Double Lognormal model to VIX options

As of 03-Apr-2007, from Monte Carlo simulation with parameters

$$z_1 = 0.0137; z_2 = 0.0208; z_3 = 0.0421; \kappa = 12; \xi_1 = 7; c = 0.34; \xi_2 = 0.94;$$

we get the following fits (orange lines):

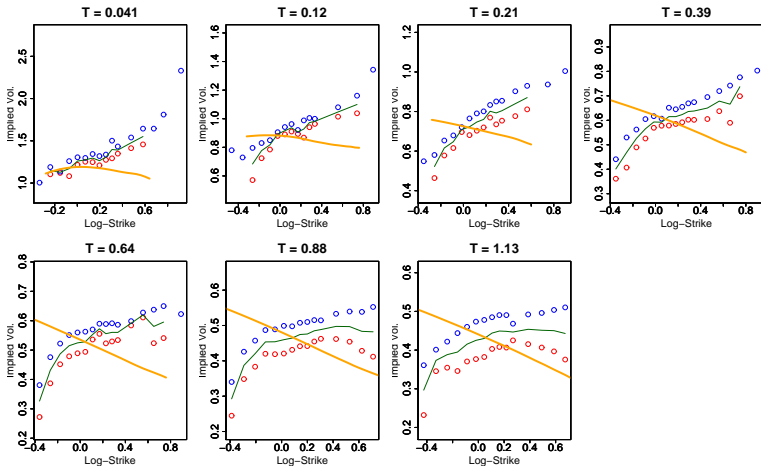


Fit of Double Heston model to VIX options

As of 03-Apr-2007, from Monte Carlo simulation with parameters

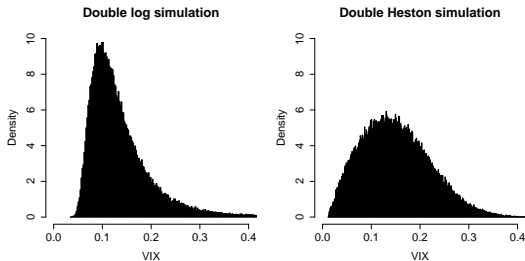
$$z_1 = 0.0137; z_2 = 0.0208; z_3 = 0.0421; \kappa = 12; \xi_1 = 0.7; c = 0.34; \xi_2 = 0.14;$$

we get the following fits (orange lines):



In terms of densities of VIX

- When we draw the densities of VIX for the last expiration ($T = 1.13$) under each of the two modeling assumptions, we see what's happening:



- In the (double) Heston model, v_t spends too much time in the neighborhood of $v = 0$ and too little time at high volatilities.

Parameter stability

- Suppose we keep all the parameters unchanged from our 03-Apr-2007 fit. How do model prices compare with market prices at some later date?
- Recall the parameters:
 - Lognormal parameters:

$$\kappa = 12; \xi_1 = 7; c = 0.34; \xi_2 = 0.94;$$

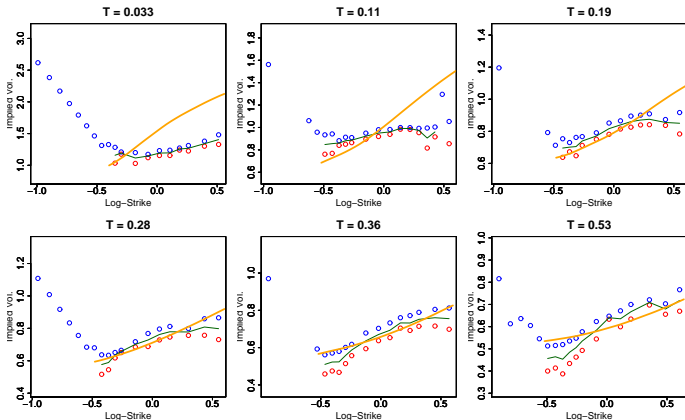
- Heston parameters:

$$\kappa = 12; \xi_1 = 0.7; c = 0.34; \xi_2 = 0.14;$$

- Specifically, consider 09-Nov-2007 when volatilities were much higher than April.
 - We re-use the parameters from our April fit, changing only the state variables z_1 and z_2 .

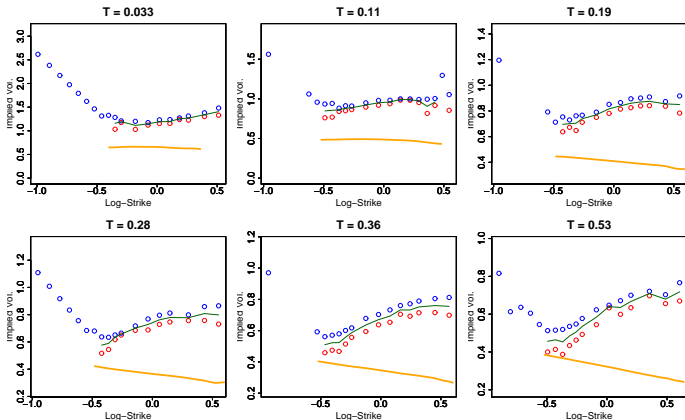
Double Lognormal fit to VIX options as of 09-Nov-2007

With $z_1 = 0.0745$, $z_2 = 0.0819$ we get the following plots (model prices in orange):



Double Heston fit to VIX options as of 09-Nov-2007

With $z_1 = 0.0745$, $z_2 = 0.0819$ we get the following plots (model prices in orange):



Observations

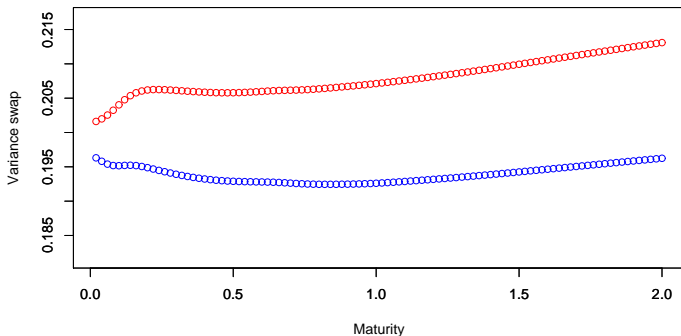
- The Double Lognormal model clearly fits the market better than Double Heston.
 - Not only does Double Lognormal fit better on a given day, but parameters are more stable over time.
 - The fair value of a put on VIX struck at 5% should be negligible but Double Heston says not.
 - VIX option prices are inconsistent with Double Heston dynamics.
- Would a more general Double CEV model with $\alpha \neq 1/2$ or 1 fit better?
 - Can we estimate the CEV exponent α ?

Double CEV parameter estimation strategy

- Variance swaps don't depend on volatility of volatility
 - Analyze time series of variance swap curves to get κ , c and z_3 .
 - Get the correlation ρ between volatility factors from time series of z_1 and z_2 .
- Fit SABR to SPX smiles to estimate the CEV exponent α .
- Fit Double CEV model to VIX options to deduce ξ_1 and ξ_2 .
- Fit Double CEV model to SPX options to deduce ρ_1 and ρ_2 .

Average variance swap curves

We proxy expected variance to each maturity with the usual strip of European options. Averaging the resulting curves over 7 years generates the following plot (SVI curve in red, ML curve in blue):



We note that the log-strip is only an approximation to the variance swap: interpolation and extrapolation methodology can make a big difference!

Double CEV variance swap curve

- Recall that in the Double CEV model, given the state $\{z_1, z_2\}$, the variance swap curve is given by

$$z_3 + (z_1 - z_3) \frac{1 - e^{-\kappa T}}{\kappa T} + (z_2 - z_3) \frac{\kappa}{\kappa - c} \left\{ \frac{1 - e^{-c T}}{c T} - \frac{1 - e^{-\kappa T}}{\kappa T} \right\}$$

Calibration of κ , c and z_3

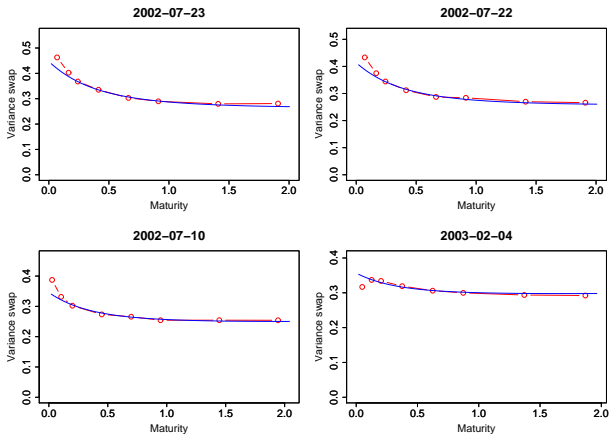
- We proceed as follows:
- For each day, and each choice of κ , c and z_3
 - 1 Impute z_1 and z_2 using linear regression
 - In the model, variance swaps are linear in z_1, z_2, z_3
 - The coefficients are functions of T and the parameters κ , c
 - 2 Compute the squared fitting error
- Iterate on κ , c and z_3 to minimize the sum of squared errors
- Optimization results are:

Parameterization	κ	c	z_3
SVI	4.874	0.110	0.082
ML	5.522	0.097	0.074

- Processes have half-lives of roughly 7 weeks and 7 years respectively. Parameters are not too different from Bergomi's.

Four worst fits

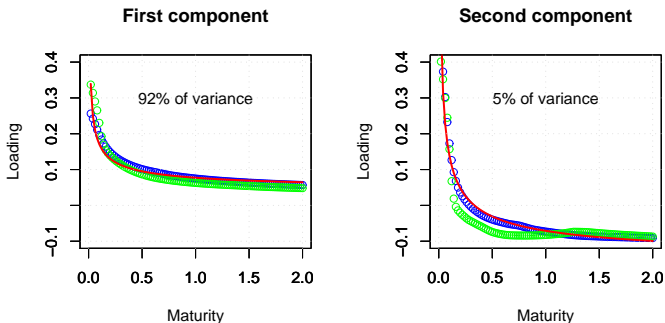
- The four worst individual SVI fits were as follows:



- We see real structure in the variance curve that the fit is not resolving.

PCA on variance swap curves

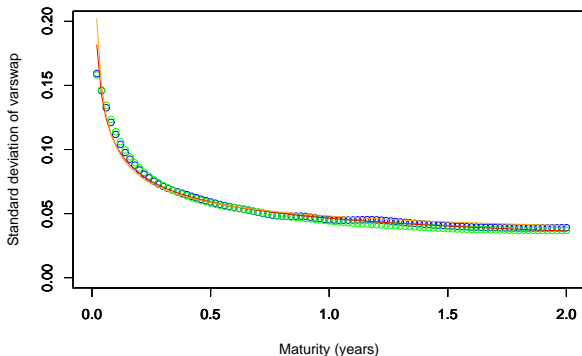
- Perform PCA on log-differences of the SVI curves to obtain the following two factors:



- The blue and green points are from conventional and robust PCA respectively. The red lines are fits of the form $a + b/\sqrt{T}$.

Volatility envelope

For each maturity, we compute the standard deviation of log-differences of the curves. ML and SVI results are green and blue respectively. The red and orange lines are proportional to $1/\sqrt{T}$ and $1/T^{0.36}$ respectively.



$\frac{1}{\sqrt{T}}$ seems to be a good approximation to the term structure of the volatility envelope!

Motivation for fitting SABR

- It seems that volatility dynamics are roughly lognormal
 - Option prices and time series analysis lead us to the same conclusion.
- SABR is the simplest possible lognormal stochastic volatility model
 - And there is an accurate closed-form approximation to implied volatility.
- The lognormal SABR process is:

$$\begin{aligned}\frac{dS}{S} &= \Sigma dZ \\ \frac{d\Sigma}{\Sigma} &= \nu dW\end{aligned}\tag{3}$$

with $\langle dZ, dW \rangle = \rho dT$.

- As suggested by Balland, fitting SABR might allow us to impute effective parameters for a more complicated model (such as Double Lognormal).

The SABR formula

As shown originally by Hagan et al., for extremely short expirations, the solution to (3) in terms of the Black-Scholes implied volatility σ_{BS} is approximated by:

$$\sigma_{BS}(k) = \sigma_0 f\left(\frac{k}{\sigma_0}\right)$$

where $k := \log(K/F)$ is the log-strike and

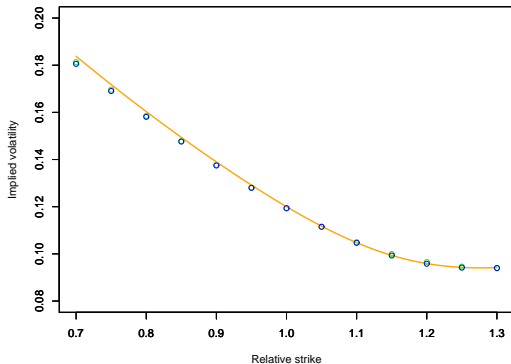
$$f(y) = -\frac{\nu y}{\log\left(\frac{\sqrt{\nu^2 y^2 + 2\rho\nu y + 1 - \nu y - \rho}}{1 - \rho}\right)}$$

It turns out that this simple formula is reasonably accurate for longer expirations too.

- Note that the formula is independent of time to expiration T

How accurate is the SABR formula?

With $T = 1$ and SABR parameters $\nu = 0.5$, $\rho = -0.7$ and $\sigma_0 = 0.12$, the following plot compares the analytical approximation with Monte Carlo and numerical PDE computations (in green and blue respectively):



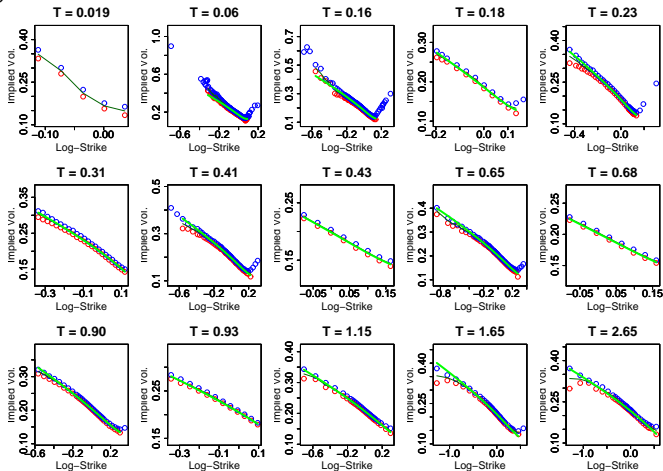
The formula looks pretty good, even for 1 year!

Fitting SABR to SPX implied volatilities

- Consider the SPX option market on a given day (25-Apr-2008 for example).
- We fit the lognormal SABR formula to each timeslice of the volatility surface.
 - Then for each expiry, we impute ν , ρ , σ_0 .

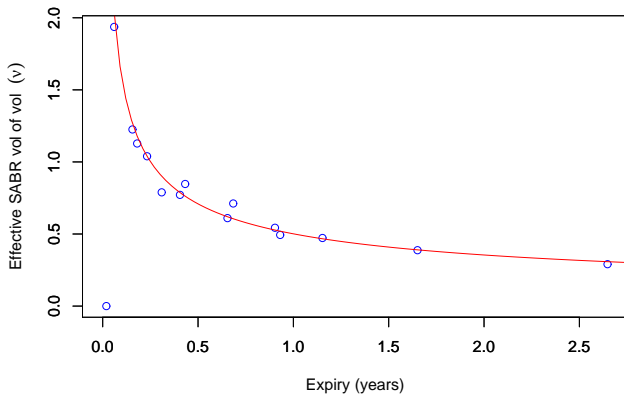
Empirical and fitted volatility smiles

As of 25-Apr-2008, we obtain the following fits (SABR fits in green):



The term structure of ν

As of 25-Apr-2008, plot ν for each slice against T_{exp} :



The red line is the function $\frac{0.501}{\sqrt{T}}$.

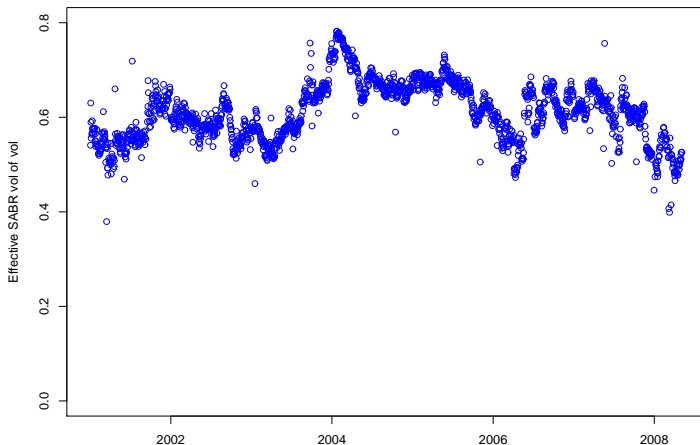
Empirical observations

- We see that the term structure of ν is almost perfectly $1/\sqrt{T}$.
 - Consistent with the empirical term structure of standard deviation of variance swaps
- This is found to hold for every day in the dataset.
- We can then parameterize the volatility of volatility on any given day by a single number: ν_{eff} such that

$$\nu(T) = \frac{\nu_{eff}}{\sqrt{T}}$$

SABR fits to SPX: ν_{eff}

Computing ν_{eff} every day for seven years gives the following time-series plot:

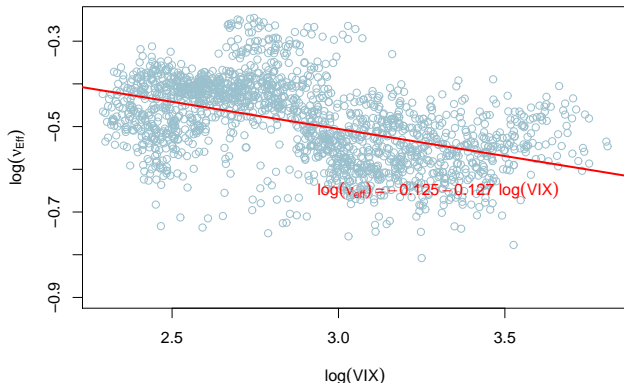


Observations from ν_{eff} time-series

- Lognormal volatility of volatility ν_{eff} is empirically rather stable
 - The dynamics of the volatility surface imply that volatility is roughly lognormal.
- Can we see any patterns in the plot?
 - For example, does ν_{eff} depend on the level of volatility?

Regression of ν_{eff} vs VIX

Regression does show a pattern!



$VIX \sim \sqrt{v}$ so we conclude that $dv \sim v^{0.94} dZ$.¹

¹The exponent of 0.94 coincides with an estimate of Aït-Sahalia and Kimmel from analysis of the VIX time series

How to price options on VIX

A VIX option expiring at time T with strike K_{VIX} is valued at time t as

$$\mathbb{E}_t \left[\left(\sqrt{\mathbb{E}_T \left[\int_T^{T+\Delta} v_s ds \right]} - K_{VIX} \right)^+ \right]$$

where Δ is around one month (we take $\Delta = 1/12$).

In the Double CEV model, the inner expectation is linear in v_T, v'_T and z_3 so that

$$VIX_T^2 = \mathbb{E}_T \left[\int_T^{T+\Delta} v_s ds \right] = a_1 v_T + a_2 v'_T + a_3 z_3$$

with some coefficients a_1, a_2 and a_3 that depend only on Δ .

Monte Carlo

- Monte Carlo simulations of stochastic volatility models suffer from bias because even if variance remains positive in the continuous process, discretized variance may be negative.
- Various schemes have been suggested increase the efficiency of simulation of such models. For example:
 - Andersen (2006)
 - Lord, Koekkoek and Van Dijk (2008) (LKV)
- Given known moments, Andersen implements an Euler scheme for a different variance process that cannot go negative and whose moments match the first few of the known moments.
- The LKV approach is to slightly amend the Euler discretization at the boundary $v = 0$.
 - Since we don't have closed-form moments in general, we adopt a bias-corrected Euler scheme of the sort described in LKV.

Monte Carlo discretization

We implement the following discretization of the Double CEV process (2):

$$\begin{aligned} v'_{t+\Delta t} &= v'_t - c(v'_t - z_3) \Delta t + \xi_2 v'^{+\beta} \sqrt{\Delta t} Z_2 \\ v_{t+\Delta t} &= v_t - \kappa(v_t - v'_t) \Delta t + \xi_1 v^{+\alpha} \sqrt{\Delta t} Z_1 \\ x_{t+\Delta t} &= -\frac{1}{4}(v_t + v_{t+\Delta t}) \Delta t + \sqrt{v^+} \sqrt{\Delta t} \{\rho_2 Z_2 + \phi_v Z_1 + \phi_x W\} \end{aligned}$$

with $\langle Z_1, Z_2 \rangle = \rho$; $\langle Z_i, W \rangle = \rho_i$ ($i = 1, 2$), $x := \log S/S_0$ and

$$\begin{aligned} \phi_v &= \frac{\rho_1 - \rho \rho_2}{\sqrt{1 - \rho^2}} \\ \phi_x &= \sqrt{1 - \rho_2^2 - \phi_v^2} \end{aligned}$$

This discretization scheme would be classified as “partial truncation” by LKV.

Numerical PDE solution

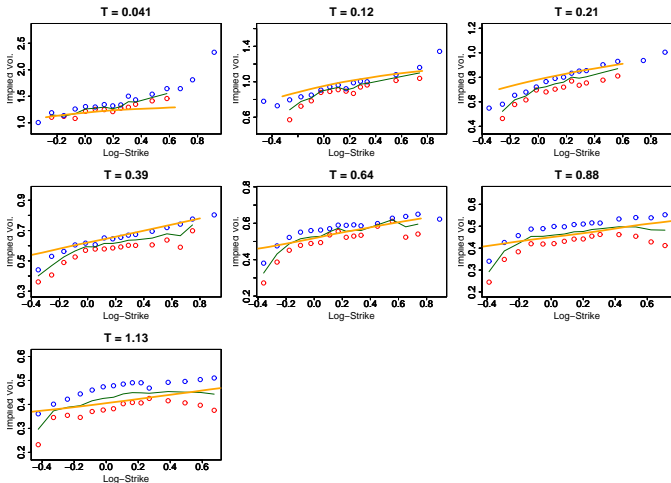
- The drift term of the Double CEV process is linear in v and v' , so VIX^2 is linear in v and v' at the option expiration.
- If we had the joint distributions of v and v' , we would also have the distributions of VIX and all the option prices.
- This suggests trying to solve the Fokker-Planck (forward equation).
- We haven't succeeded in making such a scheme work so instead, we solve the backward equation for each strike and expiration.
 - We solve this equation using an ADI scheme.

Numerical PDE vs Monte Carlo

- Although the numerical PDE solution is faster than Monte Carlo for a given accuracy, the code needs to be called once for each option.
- Monte Carlo can generate the entire joint distribution of x , v and v' for each expiration.
 - With these joint distributions, we can price any option we want, including options on SPX and exotics.
- Implementation of a 3-dimensional numerical PDE solution is hard and the resulting code would be slow.
 - It's only practical to price options on VIX with numerical PDE.
- Accordingly, we use Monte Carlo in practice.

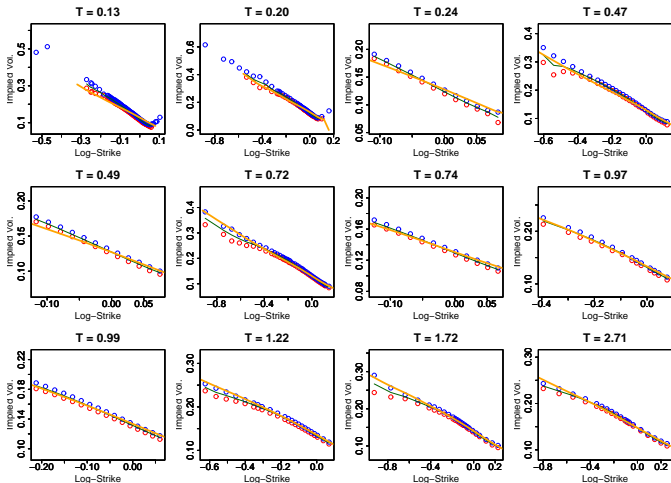
Double CEV fit to VIX options as of 03-Apr-2007

Setting the correlation ρ between volatility factors z_1 and z_2 to its historical average (see later) and iterating on the volatility of volatility parameters ξ_1 and ξ_2 to minimize the differences between model and market VIX option prices, we obtain the following fits (orange lines):



Double CEV fit to SPX options as of 03-Apr-2007

Minimizing the differences between model and market SPX option prices, we find $\rho_1 = -0.9$, $\rho_2 = -0.7$ and obtain the following fits to SPX option prices (orange lines):



Parameters

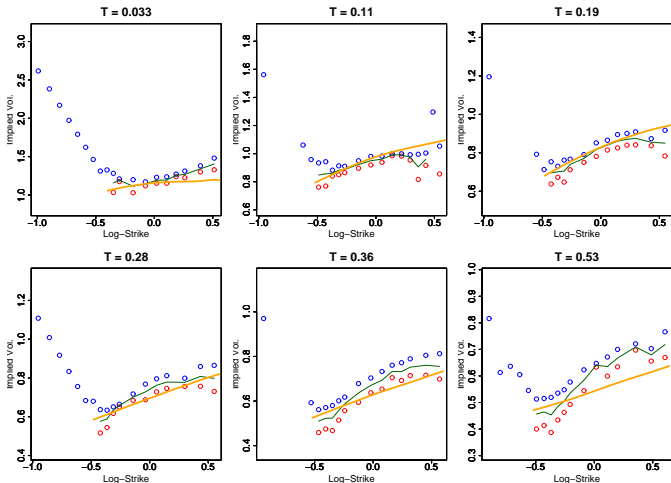
- Collating all our calibration results:

Parameter	Value
κ	5.50
c	0.10
z_3	0.078
ξ_1	2.6
α	0.94
ξ_2	0.45
β	0.94
ρ	0.59
ρ_1	-0.90
ρ_2	-0.70

- Let's call this the *April-2007 parameter set*.

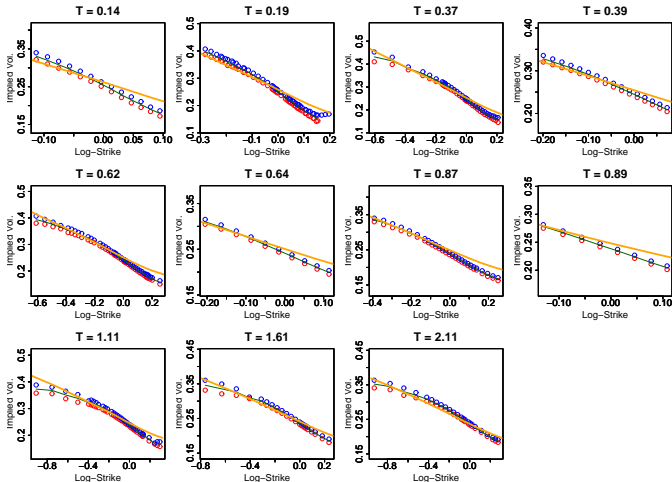
Double CEV fit to VIX options as of 09-Nov-2007

From Monte Carlo simulation with our April-2007 parameters, we get the following fits to VIX options (orange lines):



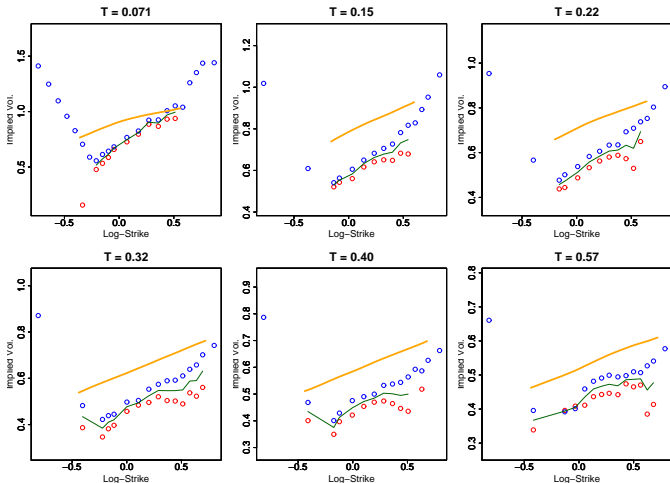
Double CEV fit to SPX options as of 09-Nov-2007

Again from Monte Carlo simulation with our April-2007 parameters, we get the following fits to SPX options (orange lines):



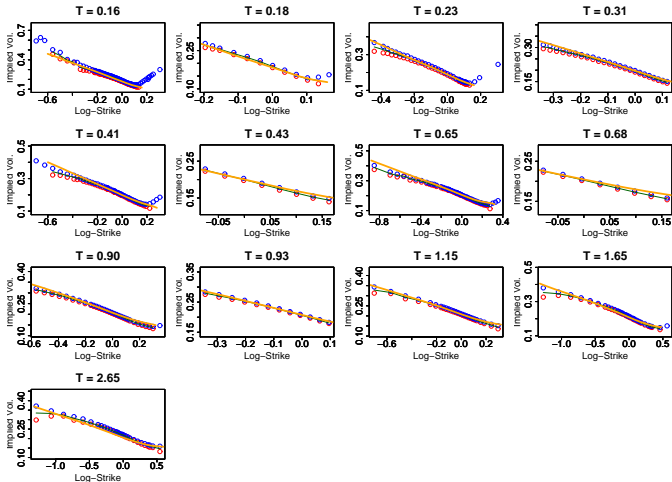
Double CEV fit to VIX options as of 25-Apr-2008

From Monte Carlo simulation with our April-2007 parameters, we get the following fits to VIX options (orange lines):



Double CEV fit to SPX options as of 25-Apr-2008

Again from Monte Carlo simulation with our April-2007 parameters, we get the following fits to SPX options (orange lines):



Is volatility of volatility stable?

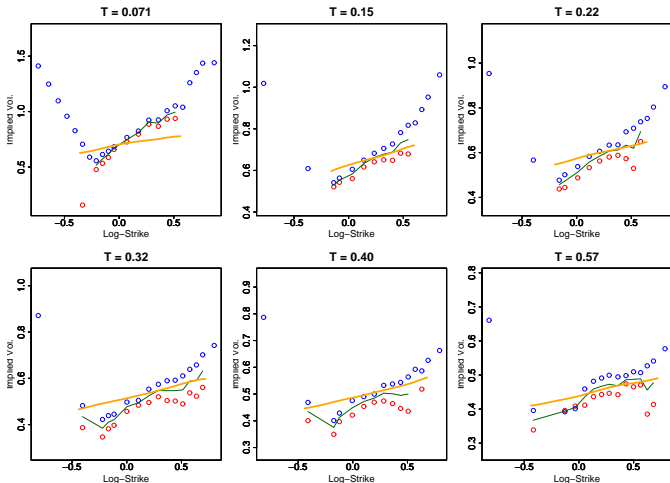
- Is volatility of volatility stable?
 - Of course not!
- Referring back to our SABR fits, we find:

Date	ν_{eff}
03-Apr-2007	0.68
09-Nov-2007	0.61
25-Apr-2008	0.50

- So volatility of volatility decreased!

Double CEV fit to VIX options as of 25-Apr-2008

From Monte Carlo simulation with proportionally lower vol-of-vol parameters, we get the following revised fits to VIX options (orange lines):



Observations so far

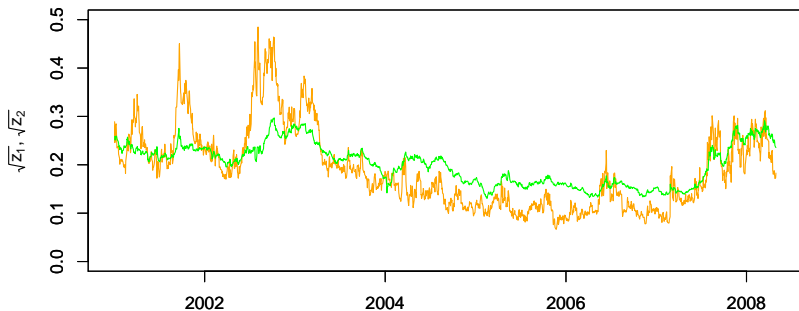
- Double Lognormal fits better than Heston with better parameter stability.
- Double CEV with $\alpha = 0.94$ fits even better with still better parameter stability.
- However, parameters are still not perfectly stable
 - In particular, volatility of volatility is not constant.
 - Implied volatilities of volatility of SPX and VIX options move together.

Implied vs Historical

- Just as option traders like to compare implied volatility with historical volatility, we would like to compare the risk-neutral volatility of volatility, mean reversion and correlation parameters that we got by fitting the Double CEV model to the VIX and SPX options markets with the historical behavior of the variance curve.

Extracting time series for z_1 and z_2

- In the Double CEV model, given estimates of κ , c and z_3 , we may estimate z_1 and z_2 using linear regression.
- From two years of SPX variance curves estimates with parameters $\kappa = 5.5$, $c = 0.1$ and $z_3 = 0.078$, we obtain the following time series for $\sqrt{z_1}$ (orange) and $\sqrt{z_2}$ (green):



Statistics of z_1 and z_2

- Let's naïvely compute the standard deviations of log-differences of z_1 and z_2 . We obtain

Factor	Historical vol.	Implied vol. (from VIX)
z_1	2.25	2.60
z_2	0.55	0.45

- Historical and implied volatilities are similar
 - in contrast to single-factor stochastic volatility models.
- Historical correlation of z_1 and z_2 is 0.59.
- In the Double CEV model, the autocorrelation functions of the factors have the form

$$acf(z_2) = e^{-c\delta}$$

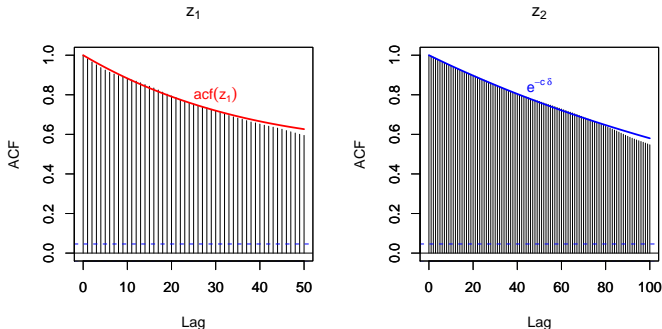
$$acf(z_1) = e^{-\kappa\delta} + \frac{\text{Cov}[z_1, z_2]}{\text{Var}[z_1]} \frac{\kappa}{\kappa - c} \left\{ e^{-c\delta} - e^{-\kappa\delta} \right\}$$

Autocorrelation plots

- Using sample estimates of $\text{Var}[z_1]$ and $\text{Cov}[z_1, z_2]$, we obtain the following estimates of decay rates:

$$\kappa = 5.12; c = 1.37;$$

and the following autocorrelation plots with model fits superimposed



How options on variance are quoted

- Define the realized variance as:

$$RV_T := \sum_i^T \Delta X_i^2$$

with $\Delta X_i = \log(S_i/S_{i-1})$.

- The price of an option on variance is quoted as

$$C = \frac{1}{2\sigma_K} \mathbb{E}[\text{payoff}]$$

where σ_K is the strike volatility.

- The price is effectively expressed in terms of volatility points on a variance swap quote. So if the quoted price of a call on variance is 2% and the strike price is 20%, the premium is $2 \times 0.2 \times 0.02 = 0.008$ and the payoff is

$$\left(\frac{RV_T}{T} - \sigma_K^2 \right)^+$$

Realized variance vs quadratic variation

- Option on variance are options on realized variance RV_T .
- In a model, we typically compute the value of an option on the quadratic variation QV_T defined as

$$QV_T := \int_0^T v_s ds$$

- Although $\mathbb{E}[RV_T] = \mathbb{E}[QV_T]$ in a diffusion model,

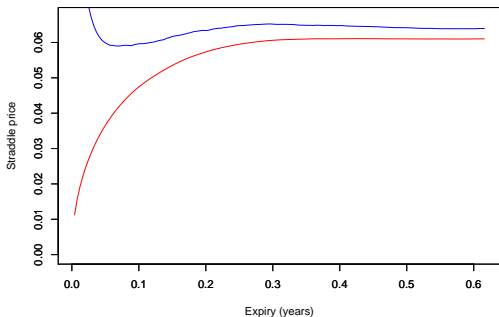
$$\text{Var}[RV_T] > \text{Var}[QV_T]$$

so options on RV are worth more than options on QV . One can think of RV_T as an approximation to QV_T with discretization error.

- It turns out that this discreteness adjustment is significant for shorter-dated options (under 3 months).

Magnitude of discreteness effect

- Double CEV model prices of QV and RV straddles (in volatility points) as a function of the number of settings. QV is in red and RV in blue. Parameters are from November 2007.



- The discreteness effect is significant!

Some broker quotes

- For the three dates for which we have computed model prices (03-Apr-2007, 09-Nov-2007 and 25-Apr-2008), we snap some broker prices of ATM variance straddles and compare our model prices.

Date	Expiry	Bid	Ask	Model	Model Adj.
03-Apr-2007	Jun-07	4.35	4.55	3.59	3.59
05-Apr-2007	Sep-07	3.90	4.70	3.74	3.74
07-Nov-2007	Jan-08	7.20	8.20	6.93	6.34
07-Nov-2007	Mar-08	4.35	7.20	7.08	6.49
13-Nov-2007	Jun-08	7.00	9.00	6.93	6.39
25-Apr-2008	Sep-08	5.60	6.10	5.25	4.20

Summary I

- It makes sense to model tradables such as variance swaps rather than non-tradables such as implied volatilities.
- Buehler's variance curve functional is particularly attractive
 - It is consistent with many dynamics of interest.
- Double Lognormal agrees much better with the market than Double Heston.
 - Whilst the rough levels of VIX option implied volatilities are determined by SPX option prices, VIX option skews are seen to be very sensitive to dynamical assumptions.
 - Double Lognormal not only fits better but its parameters are more stable over time.
- We confirmed using PCA of variance curves that two factors are necessary.
- By fitting the SABR model to SPX implied volatilities, we were able to estimate the exponent $\alpha \approx 0.94$ in the Double CEV model, close to lognormal.

Summary II

- By adding a second volatility factor, we achieved the following:
 - The term structure of SPX skew seems right even for short expirations with no need for jumps.
 - We are able to fit VIX options with time-homogeneous parameters.
 - Historical and risk-neutral estimates of the volatilities of the factors are similar
 - Recall that implied and historical vol. of vol. are very different in single-factor volatility models.
 - The volatility autocorrelation function looks more like the power law that we estimate from time series data.

Lingering Double CEV concerns

- Computation is too slow for effective calibration.
- It's not clear how to estimate mean reversion and volatility of volatility parameters independently.
- We suspect that there may be a better two-factor volatility model with power-law decay of volatility autocorrelation coefficients.
 - The timescales we settled on approximate a $1/\sqrt{T}$ volatility envelope.
 - Other parameter choices that approximate this $1/\sqrt{T}$ pattern seem to work just as well.
- In short, although it seems to work, the Double CEV model is ugly!

Current and future research

- Develop more efficient calibration techniques.
- Investigate alternative dynamics with power-law decay of volatility autocorrelations.
- Add more tradable factors (allow the skew to vary for example)

References



Yacine Aït-Sahalia and Robert Kimmel.

Maximum likelihood estimation of stochastic volatility models.
Journal of Financial Economics, 83:413–452, 2007.



Leif B.G. Andersen.

Efficient simulation of the Heston stochastic volatility model.
Technical report, Banc of America Securities, December 2006.



Philippe Balland.

Forward smile.
Presentation at Global Derivatives, Paris, 2006.



Lorenzo Bergomi.

Smile dynamics II.
Risk, 18:67–73, October 2005.



Hans Buehler.

Consistent variance curve models.
Finance and Stochastics, 10:178–203, 2006.

More references



Bruno Dupire.

Model art.

Risk, 6(9):118–124, September 1993.



Bruno Dupire.

A unified theory of volatility.

In Peter Carr, editor, *Derivatives Pricing: The Classic Collection*, pages 185–196. Risk Books, 2004.



Jim Gatheral.

The Volatility Surface: A Practitioner's Guide.

John Wiley and Sons, Hoboken, NJ, 2006.



P Hagan, D Kumar, A Lesniewski, and D Woodward.

Managing smile risk.

Wilmott Magazine September, pages 84–108, 2002.



Roger Lord, Remmert Koekoek, and Dick van Dijk.

A comparison of biased simulation schemes for stochastic volatility models.

Technical report, Erasmus University Rotterdam, February 2008.