Stochastic Volatility Models and Products

Hans Buehler <hans.buehler@db.com>

Risk Training Course Hong Kong July 7/8 2004

A Passion to Perform.

Deutsche Bank

Contents

Section

Volatility products?

An introduction in products which trade volatility

A Passion to Perform.

Deutsche Bank

Introduction

Liquid Options and Market data

Option prices are quoted in implied volatility

- On a typical stock S, we also have a range of liquid options. By Call/Put parity, we can focus only on the Call prices.

$$C(T,K) := \mathrm{E}[DF_T(S_T - K)^+]$$

- These prices are typically quoted in "volatility" Σ by means of the Black&Scholes-formula.

$$\begin{aligned} C(T,K) &= DF_T\left(F_T\,\mathrm{N}(d^+) - K\,\mathrm{N}(d^-)\right) \\ d^\pm &:= (\ln{}^{F_T}\!/_K \pm {}^1\!/_2\Sigma^2\mathrm{T})/\Sigma\sqrt{\mathrm{T}} \end{aligned}$$

  - Here, F denotes the forward of the stock and r deterministic interest rates.

- This yields a typical "implied volatility surface".

Implied Volatility

Using the B&S formula to quote option prices.

Interpolated Implied Volatility .STOXX50E 30/06/2004 @ 2811.08

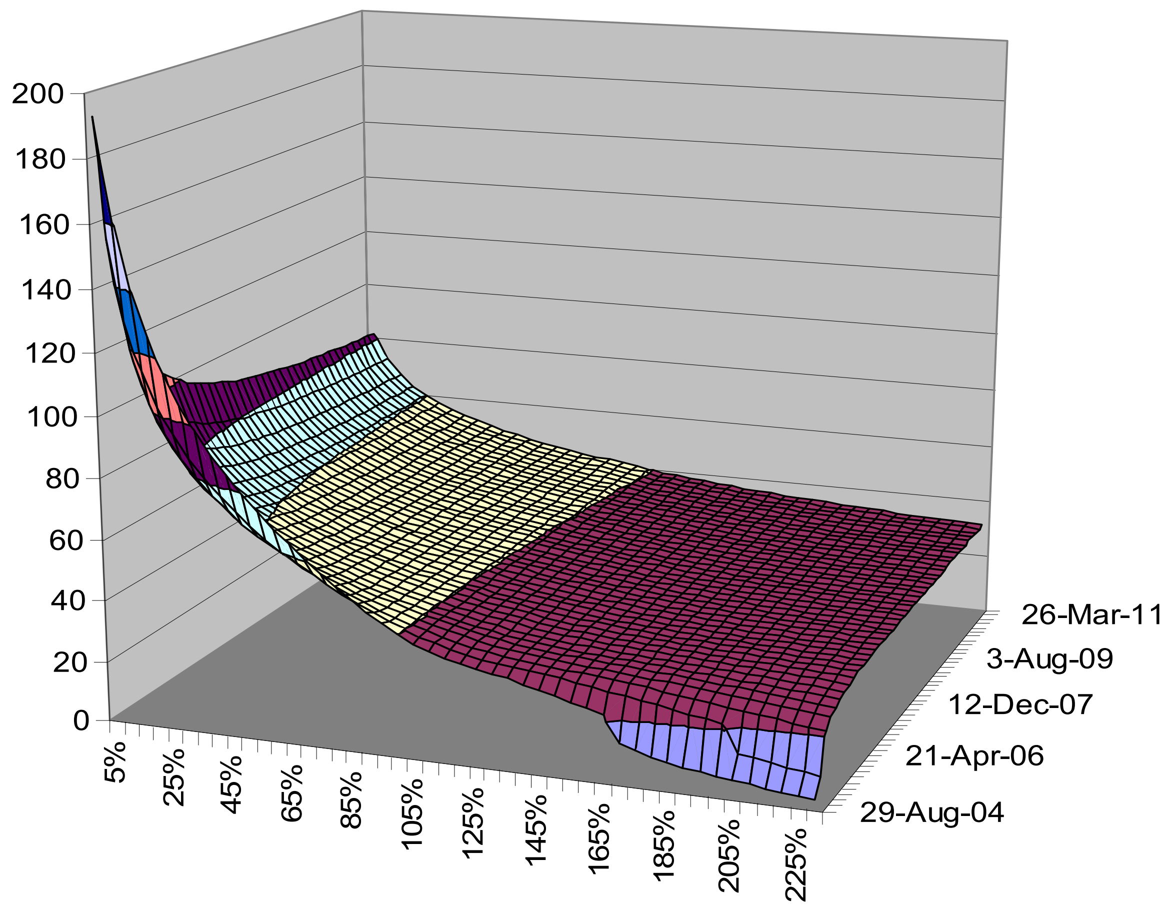

Pricing with the Smile

European options

European options can be priced without model assumptions.

- Now, we want to price options other than Vanilla Europeans.
  - Europeans: If the option is a (nice) function of the terminal stock price, we can decompose it into the difference of two positive convex functions.
  - Each of those convex functions can be approximated by a sequence of linear functions.
  - Simple examples are

Simple European payoffs

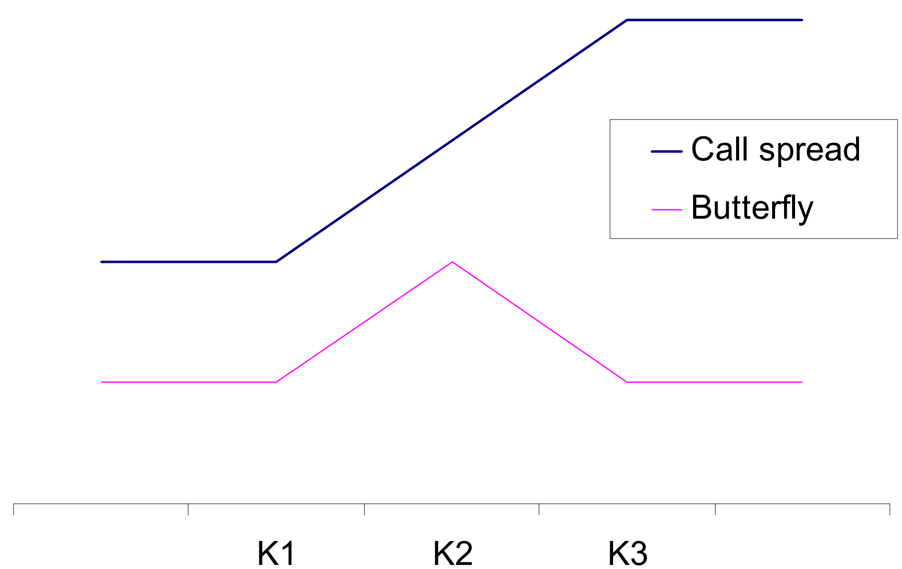

Pricing with the Smile

European options

- … more interesting indeed
  - Log-Contract

$$\ln(S_T / F_T)$$

  - Note the following small computation for positive *a* and *b*:

$$\begin{aligned}
\ln(b/a) &= -(\ln(b)-\ln(a)) - b(\tfrac{1}{a}-\tfrac{1}{b}) + (\tfrac{b}{a}-1) \\
&= -\int_a^b \tfrac{1}{x}\,dx + \int_a^b b(-\tfrac{1}{x^2})\,dx + (\tfrac{b}{a}-1) \\
&= -\int_a^b (b-x)\tfrac{1}{x^2}\,dx + (\tfrac{b}{a}-1) \\
&= -1_{b>a}\left\{\int_a^b (b-x)\tfrac{1}{x^2}\,dx\right\} - 1_{b<a}\left\{\int_b^a (x-b)\tfrac{1}{x^2}\,dx\right\} + (\tfrac{b}{a}-1) \\
&= -\int_a^\infty (b-x)^+\tfrac{1}{x^2}\,dx - \int_0^a (x-b)^+\tfrac{1}{x^2}\,dx + (\tfrac{b}{a}-1)
\end{aligned}$$

Apply this fact to $b=S_T$ and $a=F_T$. This yields

$$\tfrac{S_T}{F_T} - 1 - \ln(S_T/F_T) = +\int_{F_T}^\infty (S_T-K)^+\tfrac{1}{K^2}\,dK + \int_0^{F_T} (K-S_T)^+\tfrac{1}{K^2}\,dK$$

Pricing with the Smile
European options

- Hence we find the pricing formula

$$DF_T\,\mathrm{E}[S_T/F_T-1-\ln(S_T/F_T)] = -DF_T\,\mathrm{E}[\ln(S_T/F_T)] = \int_{F_T}^{\infty} C(T,K)\frac{1}{K^2}\,dK + \int_{0}^{F_T} P(T,K)\frac{1}{K^2}\,dK$$

- To actually price the variance swap, we construct a dominating descretization of the portfolios of the options.

$$w_i^{Call} = 1 - \frac{\ln(K_i)+\ln(K_{i-1})}{K_i - K_{i-1}} - \sum_{j=1,\dots,i-1} w_j^{Call}$$

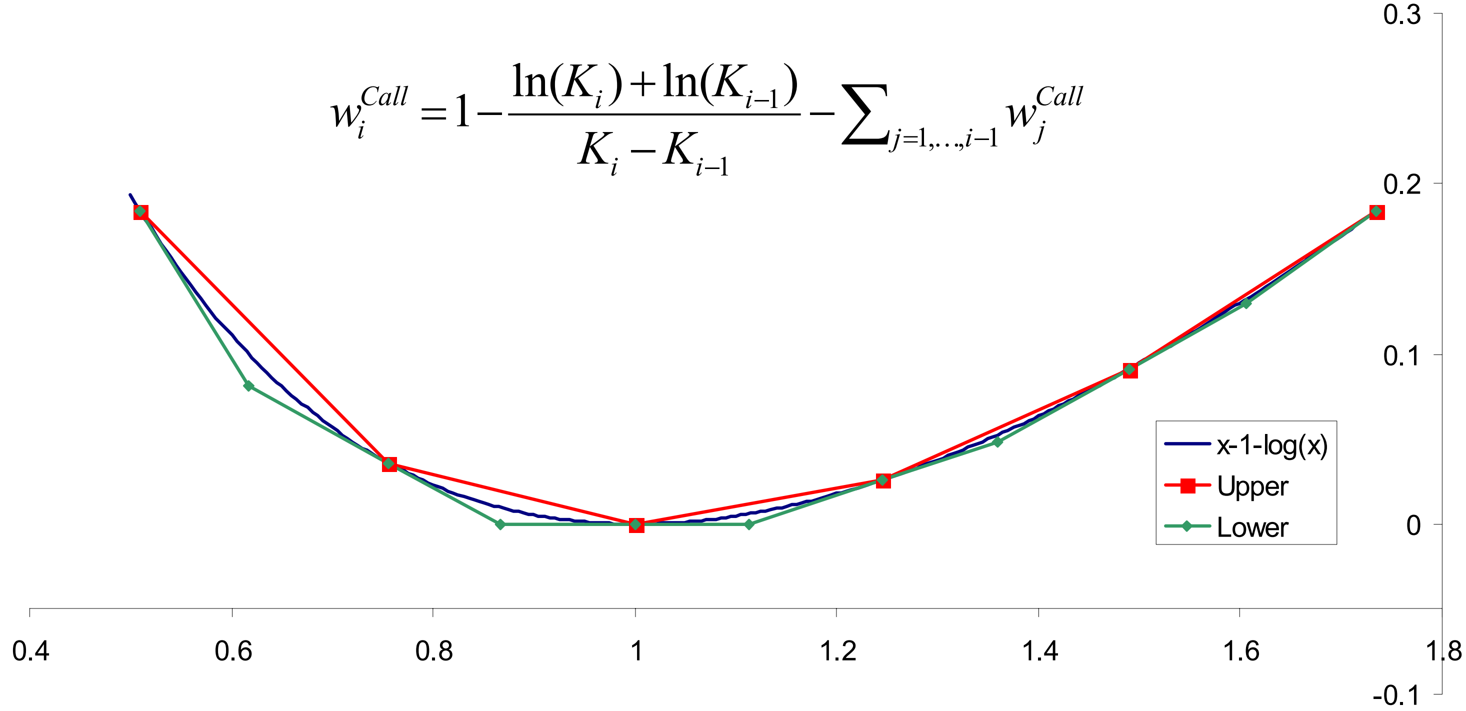

Pricing with the Smile

Path-dependency

Path-dependency may pose a problem for B&S pricing

- However, what about path-dependent options?
  - For simple products which depend only “one strike” we can use a term structure of volatility. Given the implied volatilities $(\Sigma_t)_t$ define the “short volatility”

$$\sigma_t^2(K) := \partial_t(\Sigma_t^2(K)t) \qquad \Sigma_t^2(K) =: \int_0^t \sigma_s^2(K)ds$$

  - This is quite robust, but what about KnockIn-Call where barrier B and strike K<B are different?

$$(S_T - K)^+ 1_{S_T^* > B} \quad S_T^* := \sup_{t \le T} S_t$$

Obviously, there is no clear reference strike to choose. So we have to use a model that does produce a skewed implied volatility surface.

Products with Volatility Sensitivity

Forward Started Options

- There are certain contracts which are particularly sensible to some aspects of the volatility surface. A prime example:

  Forward started Call. This contract exists as fixed and variable notional:

$$(S_T / S_t - k)^+ \qquad\qquad (S_T - kS_t)^+$$

  - The forward started call has exposure to “forward volatility”.
  - Given a *term* structure of implied volatilities $(\sigma_t)_t$ we obtain define the forward implied volatility via

$$\Sigma^2(T,t) := \frac{\Sigma_T^2 T - \Sigma_t^2 t}{T - t}$$

    because Black&Scholes is linear in variance.
  - However, this would only work if the strike of the option would be known in advance (or, rather then then skew until maturity).
  - Note that the forward started option turns into the normal option.
  - This is a consistency problem.

Variance Swaps

Introduction

- There are also products which are directly linked to the volatility resp. variance of the *returns* of the asset.
- A *variance swap* pays the *realized variance* of the returns in exchange for a previously agreed fixed variance.

$$\mathrm{var}_{realized} - \mathrm{var}_{set}$$

  – Using fixing dates $t_0, t_1, \dots, t_n$, let the returns be

$$X_i := \ln(S_{t_i} / S_{t_{i-1}})$$

  – The realized variance "per X" is usually measured using one of the two estimators (and scaled to a year with 252 business days)

    – actual mean

$$\mathrm{var}_{realized} := \tfrac{252}{n-1} \sum\nolimits_{i=1,\dots,n} (X_i - \overline{X})^2$$

    – zero mean

$$\mathrm{var}_{realized} := \tfrac{252}{n} \sum\nolimits_{i=1,\dots,n} X_i^2$$

Variance Swaps
Link to stochastic processes

- Note that the estimators do *not* estimate the variance of the terminal X_T, but the variance of the path during the life of the option.
  The variance estimator is per se a *random variable*.

- However, if the stock follows a diffusion with *arbitrary* volatility process σ,

$$dS_t = S_t \mu_t dt + S_t \sigma_t dW_t$$

we get

$$Z_t := \ln S_t = s_0 + \int_0^T (\mu_t - \tfrac{1}{2}\sigma_t^2) dt + \int_0^T \sigma_t dW_t$$

and by general properties of the quadratic variation, we have

$$\lim_{n \uparrow \infty} \sum_{i=1,\ldots,n} X_i^2 = \langle Z \rangle_T = \lim_{n \uparrow \infty} \sum_{i=1,\ldots,n} \left(Z_{i/n\,T} - Z_{i-1/n\,T} \right)^2 = \langle Z \rangle_T = \int_0^T \sigma_s^2 ds$$

Variance Swaps

Model-Independent Pricing

A variance swap can be priced with relatively mild model assumptions.

- Now recall our decomposition of the log payoff,

$$\ln(x) = \sum_i \alpha_i (x - k_i)^+ + \sum_j b_j (k_j - x)^+$$

  - We saw before, that we can price this option from the market without the need for advanced processes.
  - Hence, this is a “traded instrument”.
  - We now have

$$\begin{aligned} \mathrm{E}[\ln S_t] = \mathrm{E}[Z_t] &= \ln S_0 + \int_0^T (\mu_t - \tfrac{1}{2}\sigma_t^2) dt + 0 \\ &= \ln F_t - \tfrac{1}{2} \int_0^T \sigma_t^2 dt \end{aligned}$$

- If we use this as an approximation for the variance, we can price the variance swap using the market instruments directly. *In other words, we do not a model assumption on σ.*
  - Very nice and surprising property of this product.
  - However, it is subject to volatility interpolation and extrapolation (we can never trade *all* strikes).
  - You can not *really* hedge it statically.
  - What about jumps ?

Volatility Swaps and Options on Variance

Realized variance products

- We have seen that a variance swap essentially pays out realized variance $var_{realized}$.
- But what if the buyer wants to buy volatility

$$\mathrm{vol}_{realized} := \sqrt{\mathrm{var}_{realized}}$$

  - This is a European option on realized variance.

- Apart from the volatility swap, we also have versions of the variance swap which are subject to caps and floors:

$$\min(\mathrm{var}_{realized}, \mathrm{var}_{cap}) - \mathrm{var}_{set}$$

  - Using Put/Call parity (with the variance swap itself), it is sufficient to price

$$(\mathrm{var}_{realized} - K)^{+}$$

Options on Variance

Stochastic Volatility

- These payoffs are *options on realized variance*.
  - As with all options on an underlying, this means that the prices depend on the *volatility of the variance.*
  - There is normally no model-independent pricing scheme known, so we have to make a model assumption.
  - The model must have a "Volatility of Volatility", ie Black&Scholes is not suitable (the volatility swap in B&S has a deterministic payoff).
- This brings us to "stochastic volatility" where we model the process σ.

- Other approaches were trying to develop a model-independent pricing.
  - If correlation between stock and volatility is zero, we can indeed recover the density η of the variance from the market. Indeed, we have

$$\begin{aligned} DF_T \, \mathrm{E}[(S_T - K)^+] &= DF_T \, \mathrm{E}[(F_T e^{-\frac{1}{2}\langle Z\rangle_T + \sqrt{\langle Z\rangle_T} W_T^X} - K)^+] \\ &= \mathrm{E}[BS(DF_T, F_T, K, T, \langle Z\rangle_T)] \\ &= \int_0^\infty BS(DF_T, F_T, K, T, x)\eta(x)dx \end{aligned}$$

    hence the problem reduces to a Fredholm-type integral equation in *(K,x)*.
  - However, zero correlation assumption highly questionable.

Products
Summary

- European payoffs can always be priced consistently by decomposing them into vanilla European options.
- Quite some path-dependent options can also be priced in B&S.
- However, when it comes to strong path-dependency, B&S is no longer an option.

Stochastic Volatility and Local Volatility

Key concepts and examples

A Passion to Perform.

Deutsche Bank

Dupire's Model

Implied Local Volatility

- Dupire (1993) has shown that we can in principle construct a function s such that the SDE

$$dS_t = S_t \mu_t dt + \sigma(t, S_t) dW_t$$

has a unique solution and such that *all* market prices can be recovered, ie

$$C(T,K) = DF_T \, \mathrm{E}[(S_T - K)^+] \qquad DF_T = e^{-\int_0^T r_s ds}$$

– The idea is to use Ito's formula on the payoff of the call, ie

$$\mathrm{E}[d(S_T - K)^+] = \mathrm{E}[S_T \mu_T 1_{S_T > K} dT] + 0 + \tfrac{1}{2}\mathrm{E}[\sigma^2(T, S_T)\delta_K(S_{Tt})dT]$$

– By exchanging integration and deriving towards *T*, we obtain

$$DF_T^{-1}(\partial_T C + r_T C) = \mu_T \left(\mathrm{E}[K 1_{S_T > K}] + \mathrm{E}[(S_T - K)^+]\right) + \tfrac{1}{2}\sigma^2(T,K)\mathrm{E}[\delta_K(S_T)]$$

a side-computation yields

$$\mathrm{E}[1_{S_T > K}] = -\partial_K \, \mathrm{E}[(S_T\text{-}K)^+] \qquad \mathrm{E}[\delta_K(S_T)] = \partial^2_{KK} \, \mathrm{E}[(S_T\text{-}K)^+]$$

such that

$$\partial_T C = (\mu_T - r_T)C - \mu_T K \partial_K C + \tfrac{1}{2}\sigma^2(T,K)\partial^2_{KK} C$$

Dupire's Model

Theory ...

Local volatility is theoretically very appealing - all market prices are recovered.

- $$\sigma^2(T,K) = 2\frac{\partial_T C + \mu_T K \partial_K C - (\mu_T - r_T)C}{\partial^2_{KK} C}$$

Local Volatility (based on the interpolated implied vol)

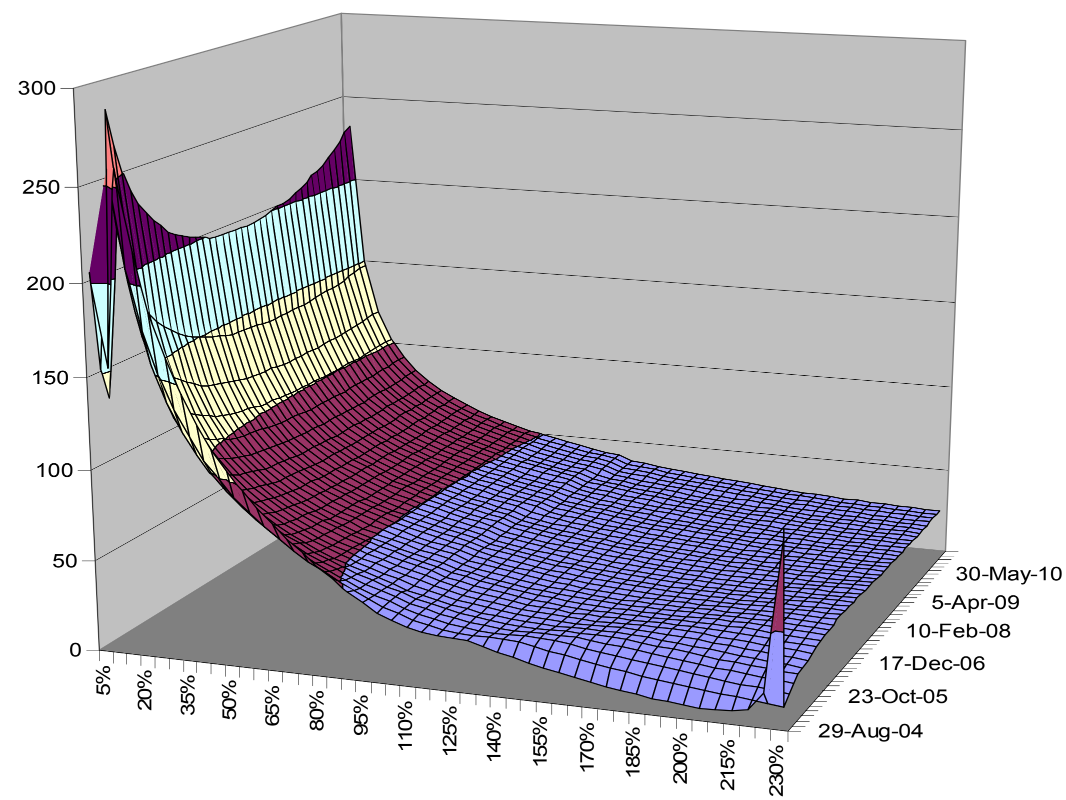

Dupire's Model

Theory ...

Local volatility is theoretically very appealing - all market prices are recovered.

- *Fits them all.*
- Complete Model - ie, Delta-Hedging is theoretically sufficient.
- One-factor model with clear concept of multi-dimensionality.
- Can be used in Monte-Carlo and Finite Difference without problems (compare with Jump-Diffusions or Levy processes)

Dupire’s Model

… and practise - implementation

Local volatility has numerical problems.

- In reality, we do not have a continuum of call prices.

- (A) Interpolation of call prices
  - Interpolation method not clear (splines are a very bad idea). Numerically complicated.

- (B) Interpolation of implied volatility
  - Arbitrage-freeness complicated (and that means local volatility explodes)
    Using the B&S - formula we can express the local volatility in terms of Sqrt of Variance, Σ:

$$\sigma^2(T,K) = \frac{2\partial_T \Sigma\, K^2 \Sigma}{(\partial_K \Sigma d^+ K + 1)^2 + (\partial_{KK}^2 \Sigma - (\partial_K \Sigma)^2 d^+) K^2 \Sigma}$$

- (C) Interpolation of Local Volatility via numerical calibration
  - Typically FD forward scheme (see also Derman & Kani 1993)
  - Accuracy error when applying to different schemes (Monte-Carlo for example)
  - Extrapolation?
  - But it guarantees a smooth function, hence SDE convergence is good.

Dupire's Model

… and practise - application

… and local volatility has the wrong dynamics.

- The implied movements of implied volatility are not correct
  - It predicts a certain implied volatility over time - forward skews wrong (cf. Cliquets).
  - Skew massively stock-level dependent (or fixed, if local volatility is modelled onto the ATM level).

- Perfect correlation between Volatility and Stock (that's a very strong assumption).

- Summary
  - Robust concept
  - Nice theoretically
  - Numerically not trivial to calibrate, but easy to use afterwards.
  - Wrong dynamics of implied surface (that makes it somewhat problematic for hedging).
  - However market prices of Barriers, Worst Ofs, Baskets seem close to LV price.

Stochastic Volatility
Ideas

- After Dupire, people started to develop “stochastic volatility models”.
- This term generally also includes models which are in fact jump models.

- Some popular models
  - Jump-Diffusion (Merton 1976)
  - Heston (1993)
  - Heston with Jumps (Bates 1996)
  - Duffie, Pan, Singleton (2000)
  - Levy processes (for example CGMY 2002)
  - SABR Mixed local volatility and stochastic volatility (2002)
- Extensive range of models.

Stochastic Volatility

Concept - Heston

Heston's model is robust and has nice properties.

- The first idea was to use, once more

$$dS_t = S_t \mu_t dt + \sigma_t S_t dW_t$$

but this time with a stochastic process σ in some form

$$\sigma_t = V(v_t) \quad dv_t = a_t(v_t)dt + b_t(v_t)dW_t^v \quad \langle W, W^v \rangle = \rho t$$

- If *a* and *b* are time-independent, this is a homogeneous diffusion.
- A prototype for a continuous Markov-process.
- Numerically easy to simulate in both forward and backward schemes.
- If ρ is zero, pricing of derivatives is a matter of weighting European prices. "Unfortunately", correlation is *strongly* not zero, but very negative.

Heston:

$$\sigma_t = \sqrt{v_t} \quad dv_t = \kappa(\theta - v_t)dt + s\sqrt{v_t}\,dW_t^v \quad (2\kappa\theta / s^2 > 1)$$

Stochastic Volatility

Heston - volatility sample paths

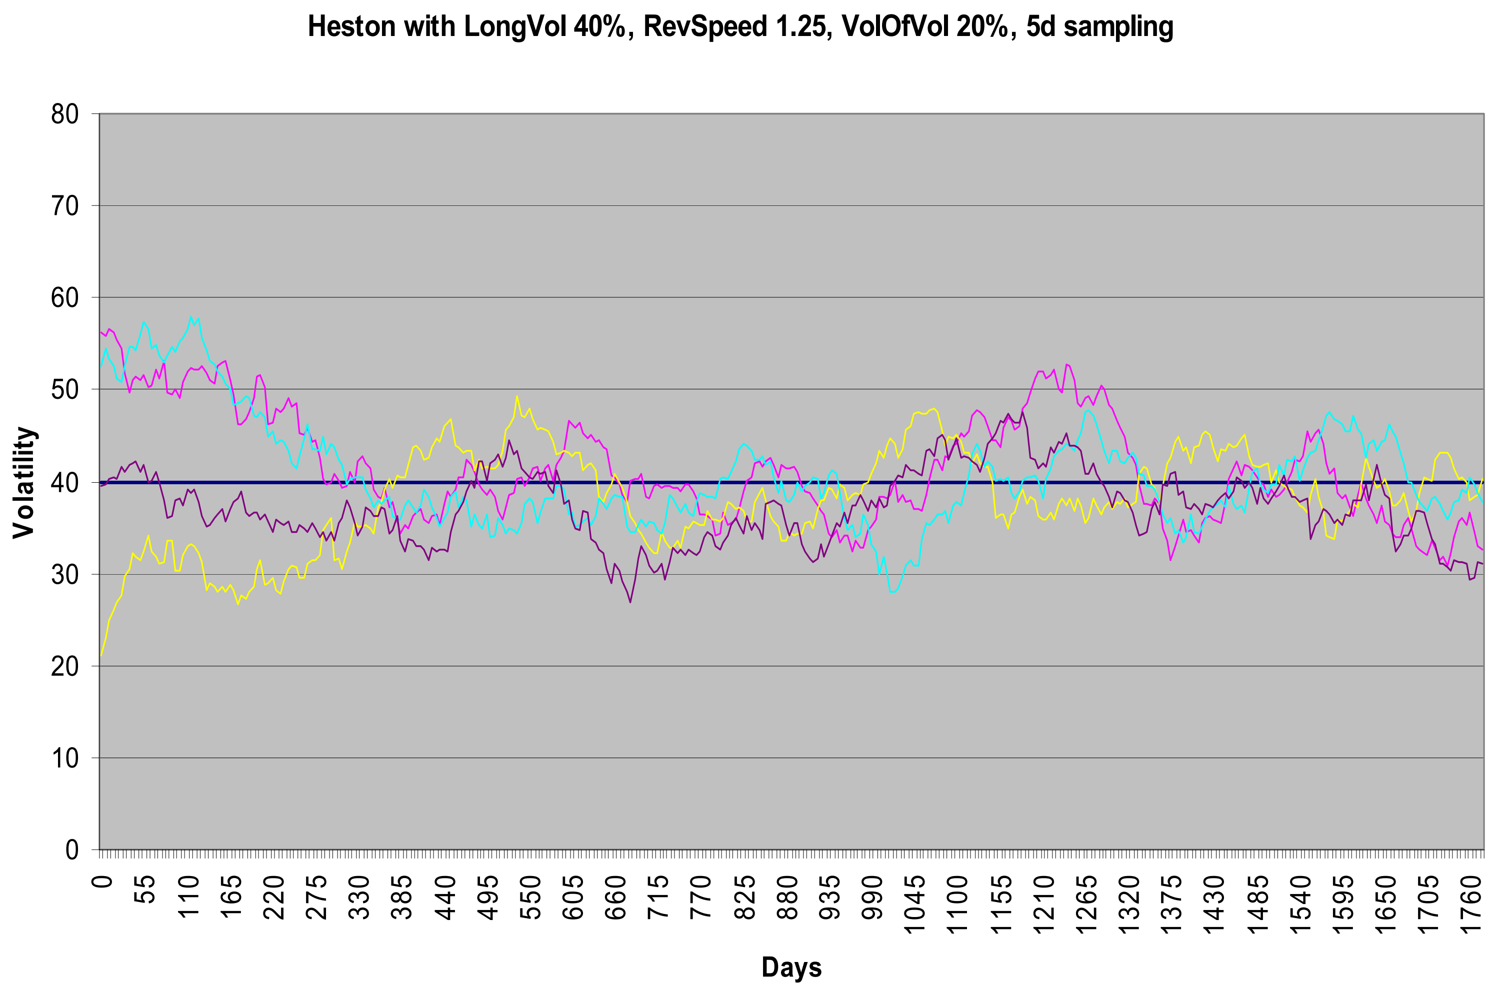

Stochastic Volatility
Heston - volatility sample paths

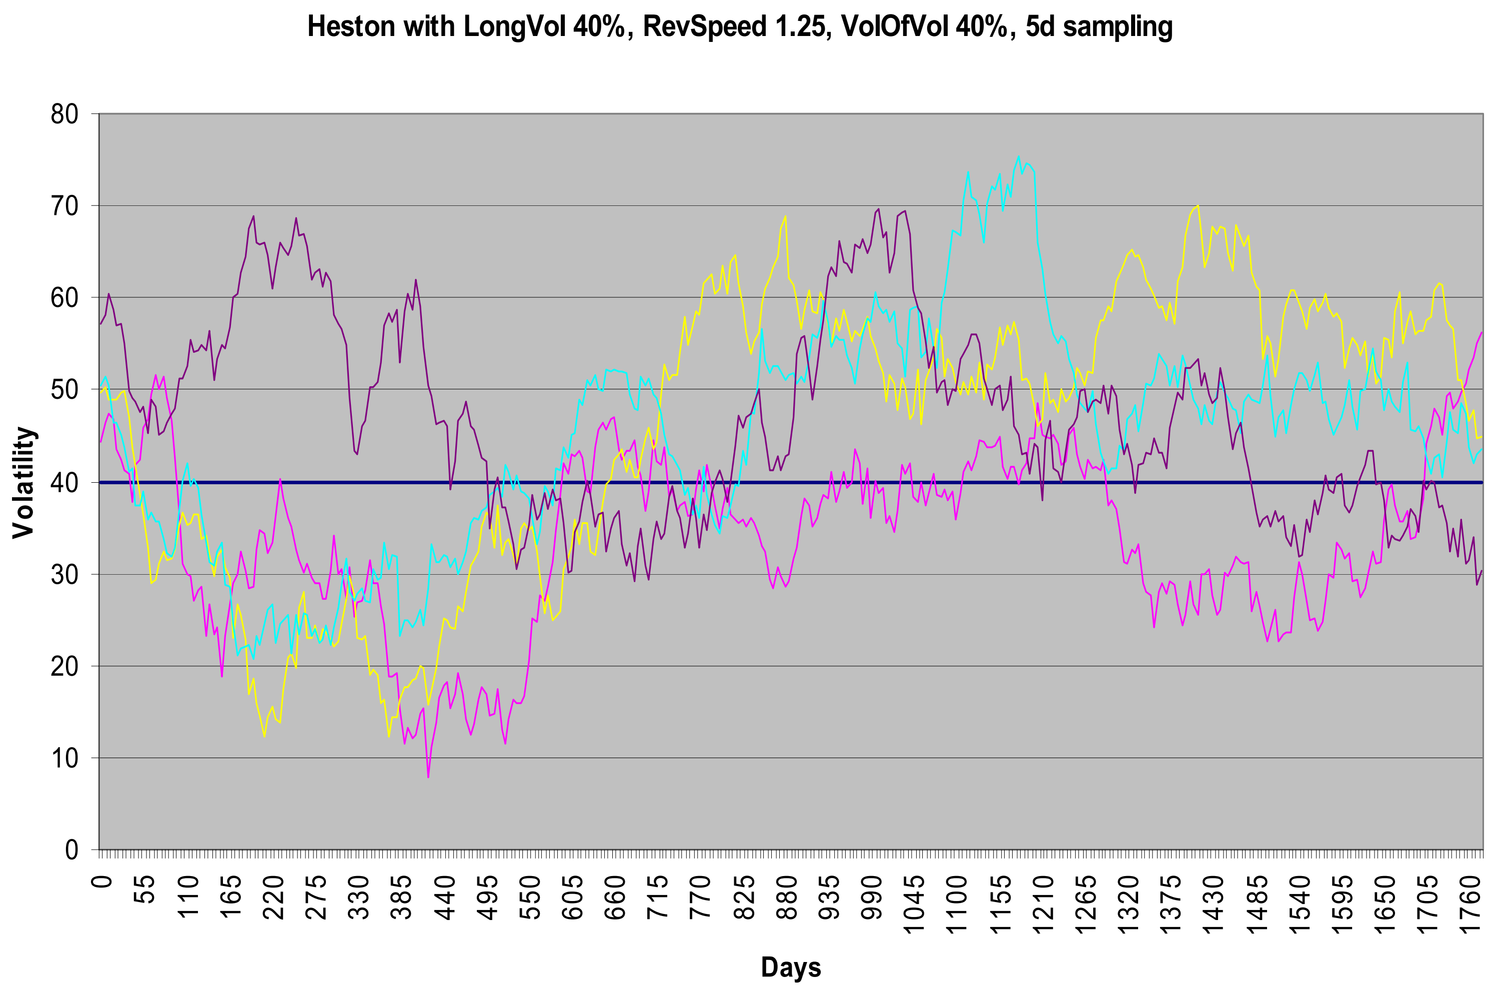

Stochastic Volatility

Heston - volatility sample paths

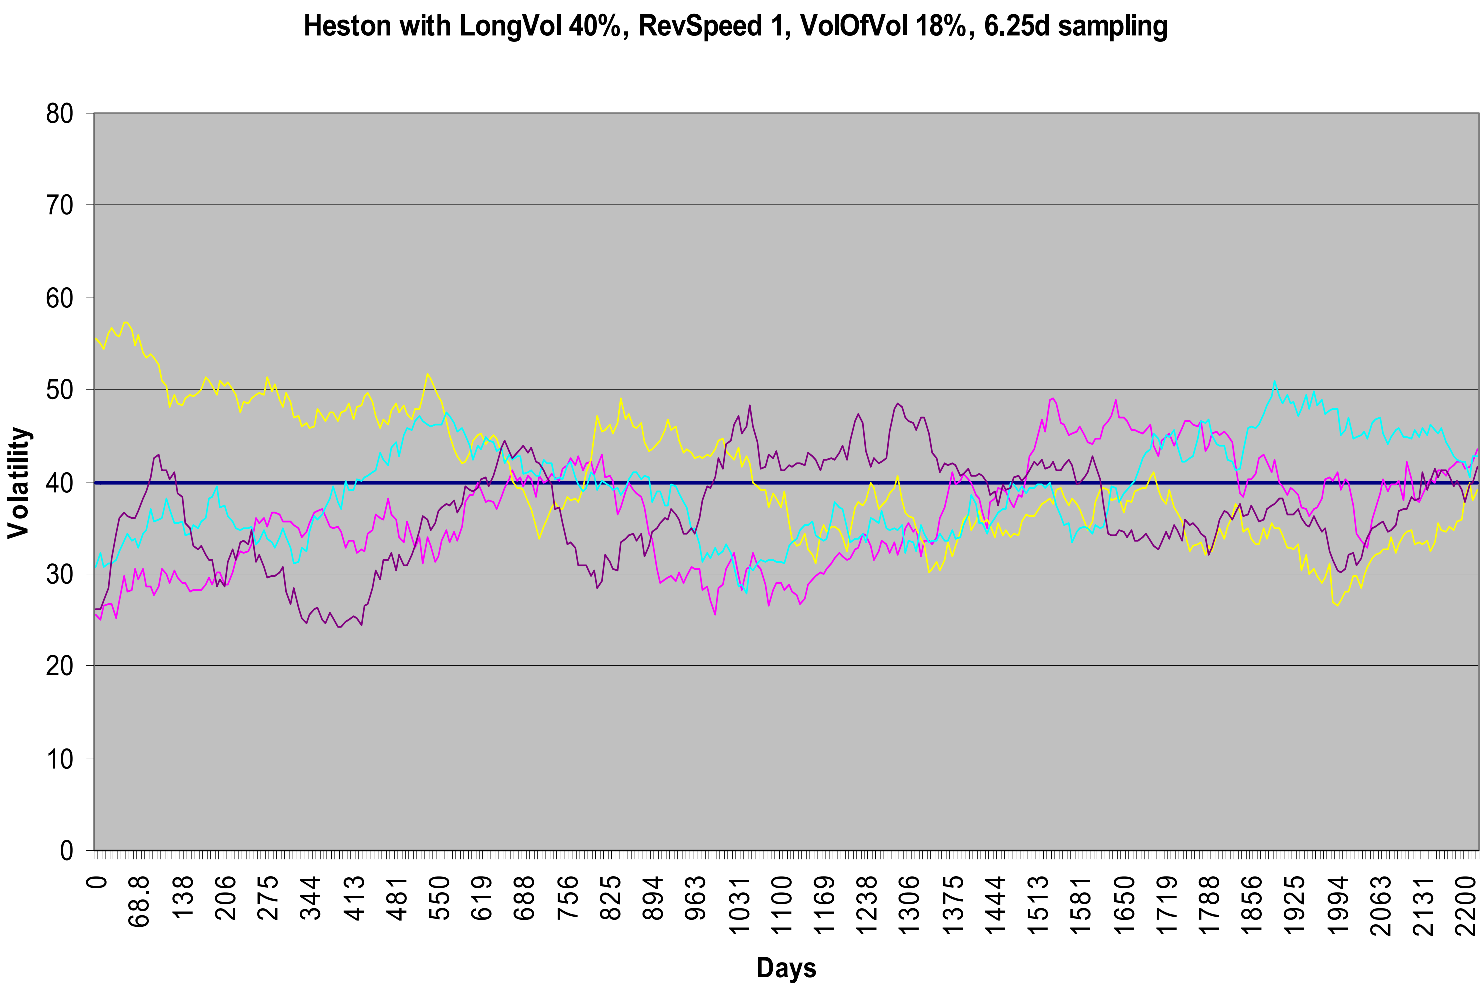

Deutsche Bank

Stochastic Volatility

Hedging with Heston

- Note that Heston’s model is not complete using the stock alone.
- However, with one more option it is in theory complete.
- In reality, we cannot rely on one option, but we also have to hedge the exposure to our model parameters.
- Hence, we need to “Vega”-hedge our parameters just like volatility in B&S.

$$\sigma_0, \kappa, \tau, s, \rho$$

- The question remains which options to choose.
  - In principle, five options are sufficient.
  - A selection criterion is needed.
- If possible, statically hedge parts of the product.

Stochastic Volatility

Heston - Implied Volatilities

Fit for short maturities is not that great.

- Heston allows to price European Options with Fourier-Inversion (more on this later).
- It has a nice smile for maturities more than 6m, but on the short end it is a bit weak.

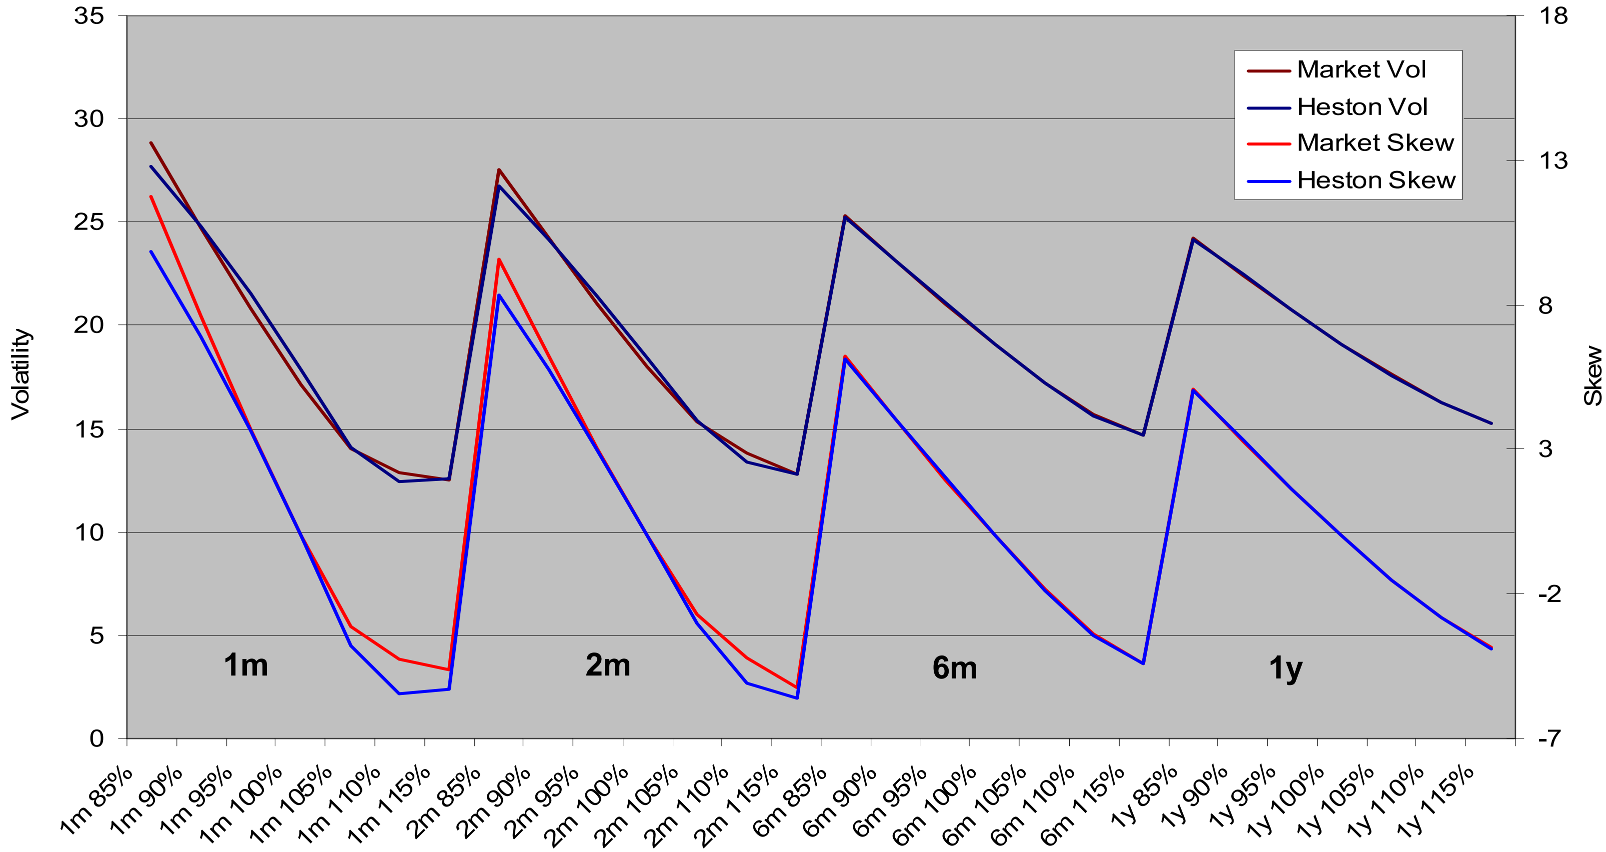

Stochastic Volatility

Heston and OU-processes

- A close relative to Heston's model which also allows for non-zero correlation has been proposed by Schoebel/Zhu (1999) in form of an Ornstein-Uhlenbeck process

$$d\sigma_t = \kappa(\theta - \sigma_t)dt + s\,dW_t^v$$

  - This process can (and will) assume negative values.
  - This is unsatisfactory from a design point of view and is a real drawback if more than one underlying is modelled.
  - But much easier to handle mathematically.

Stochastic Volatility

Merton

- Merton (1976) used Jump-Diffusion models.
  - Assume *N* is a Poisson-process with accumulated intensity Λ, and that $(\xi_i)_i$ is a sequence of iid variables which represent the jumps of the log of the process.
  - For deterministic volatility σ, assume the SDE

$$dS_t = S_t \tilde{\mu}_t dt + \sigma_t S_t dW_t + S_t d \sum_{i=1,\dots,N_t} (e^{\xi_i} - 1))$$

which is solved by a product of a standard B&S model and a jump part,

$$S_t = S_0 F_t \exp\left\{ \int_0^t \sigma_r dW_r - \tfrac{1}{2} \int_0^t \sigma_r^2 dr - \Lambda_t (\mathrm{E}[e^{\xi_1}] - 1) \right\} \times \prod_{i=1,\dots,N_t} e^{\xi_i}$$

- Note that
  - *Jumps are real.*
  - Model is incomplete if jumps are not constant (that means there is not even a perfect hedge in a theoretical sense). You cannot hedge yourself against the jumps.
  - Implied volatility flattens out very quickly (additive process).
  - Jumps are difficult to use in PDEs (but no problem in Monte-Carlo).

Stochastic Volatility

Merton - sample stock paths

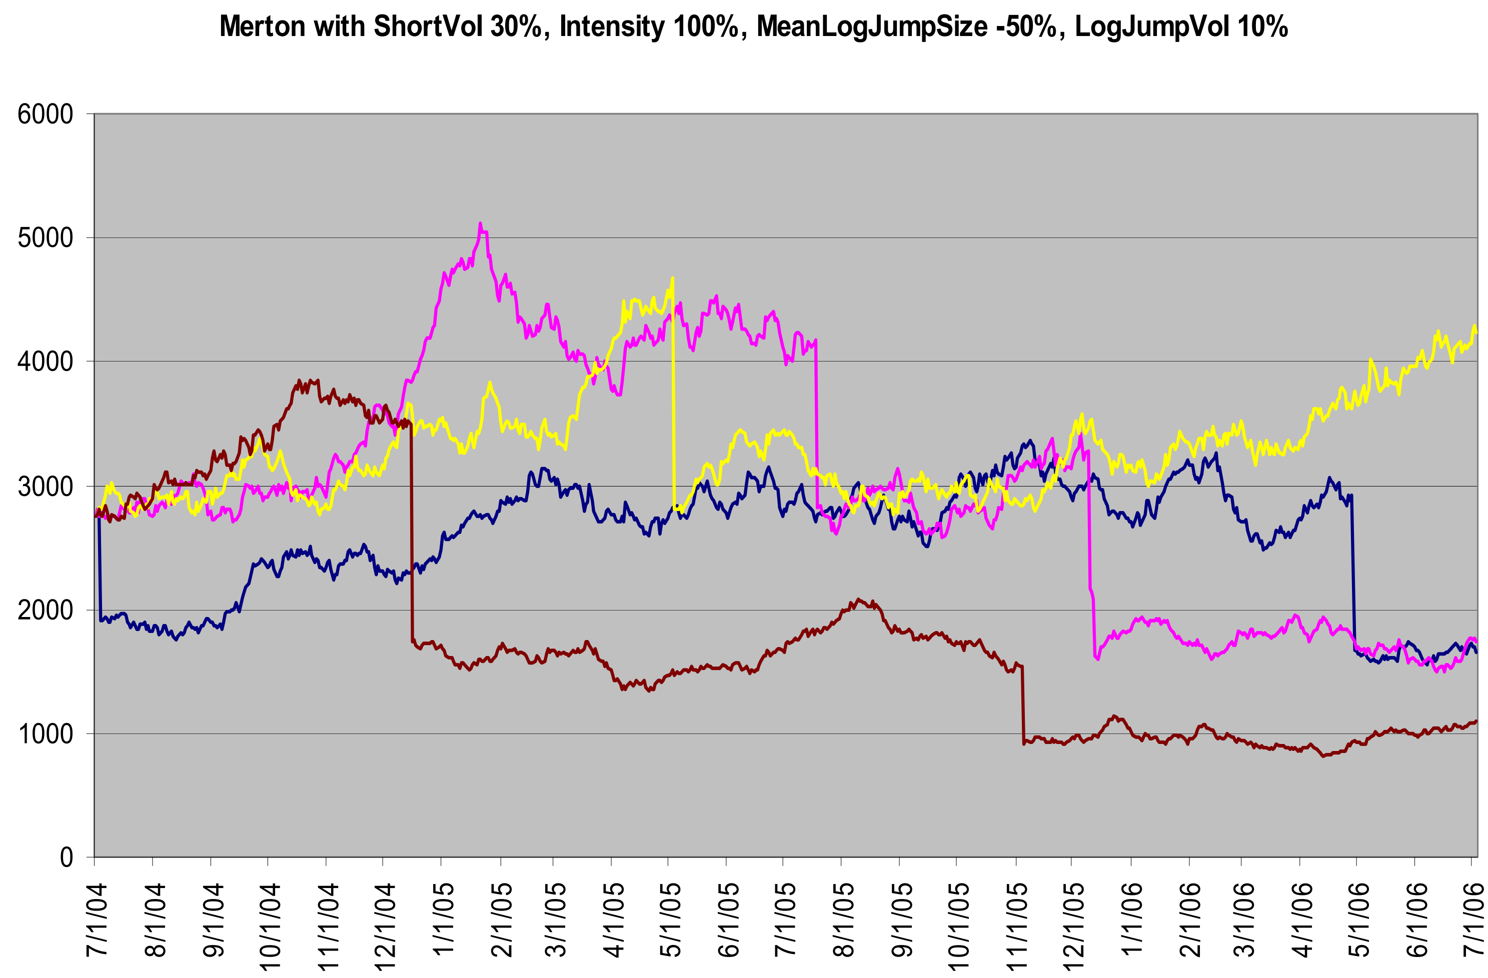

Stochastic Volatility

Merton - what about Variance Swaps

- Recall that we developed the formula for the variance swap assuming that the stock is a diffusion process.
- We used that for a diffusion

$$Z_T = \ln S_T = \ln F_T + \int_0^T \sigma_s dW_s - \tfrac{1}{2}\int_0^T \sigma_s^2 ds$$

and

$$\sum_{i=1,\dots,m} X_i^2 \to \langle Z \rangle_T = \int_0^T \sigma_s^2 ds$$

- In the presence of jumps with finite activity (ie, Jump-Diffusion), we have instead

$$\sum_{i=1,\dots,m} X_i^2 \to \langle Z \rangle_T^C + \sum_{s \le T} \Delta_s Z^2 = \int_0^T \sigma_s^2 ds + \sum_{i=1,\dots,N_T} \xi_i^2$$

while

$$Z_T = \ln S_T = \ln F_T + \int_0^T \sigma_s dW_s - \tfrac{1}{2}\int_0^T \sigma_s^2 ds + \Lambda_T (1 - \mathrm{E}[e^{\xi}]) + \sum_{i=1,\dots,N_t} \xi_i$$

Stochastic Volatility

Merton - what about Variance Swaps

- We get

$$\begin{aligned} \mathrm{E}[Z_T] &= \ln F_T - \tfrac{1}{2}\mathrm{E}[\langle Z\rangle_T] + \Lambda_T(\mathrm{E}[\xi] + 1 - \mathrm{E}[e^{\xi}]) \\ &= \ln F_T - \tfrac{1}{2}\left(\mathrm{E}[\langle Z\rangle_T] + \Lambda_T \mathrm{E}[\tfrac{\xi^2}{1} + \tfrac{\xi^3}{3} + \tfrac{\xi^4}{3\cdot 4} + \cdots]\right) \\ &= \ln F_T - \tfrac{1}{2}\left(\mathrm{E}[\langle Z\rangle_T] + \Lambda_T \xi^2 - o\right) \end{aligned}$$

  and

$$\mathrm{E}\left[\langle Z\rangle_T^C + \sum_{s\le T}\Delta_s Z^2\right] = E\left[\int_0^T \sigma_s^2 ds + \Lambda_T \xi^2\right]$$

- Hence, the variance swap formula is only approximately right.

Stochastic Volatility

Bates

Joining stochastic volatility with jumps is quite a promising concept..

- Let us combine the two

$$dS_t = S_t \tilde{\mu}_t dt + \sigma_t S_t dW_t + S_t (e^{\xi_i} - 1) dN_t$$

$$\sigma_t = \sqrt{v_t} \qquad dv_t = \kappa(\theta - v_t)dt + s\sqrt{v_t}\,dW_t^v \qquad (2\kappa\theta / s^2 > 0)$$

- Other ideas: Jumps in volatility (the jumps themselves correlated with the jumps of the stock) and so forth.

Stochastic Volatility

Duffie, Pan, Singleton 2000

- Duffie, Pan and Singleton have taken the next step and introduced jumps in both volatility and the stock price, potentially correlated (in fact, their paper covers n-dimensional affine models).

  For example, one may choose

$$dS_t = S_t \tilde{\mu}_t dt + \sigma_t S_t dW_t + S_t (e^{\xi_i} - 1) dN_t$$

$$\sigma_t = \sqrt{v_t} \qquad dv_t = \kappa(\theta - v_t)dt + s\sqrt{v_t}\, dW_t^v + h_i dN_t$$

  with log normal jumps ξ in the log of the stock price and with exponentially distributed jumps *h* in the volatility (downward jumps are tricky if variance is supposed to stay positive). Note that the jumps in this example appear in both volatility and stock at the same time.

- Various extensions are possible, see the original paper.
- But improvement on fitting the market not as good as expected.

Stochastic Volatility

Bates - sample paths

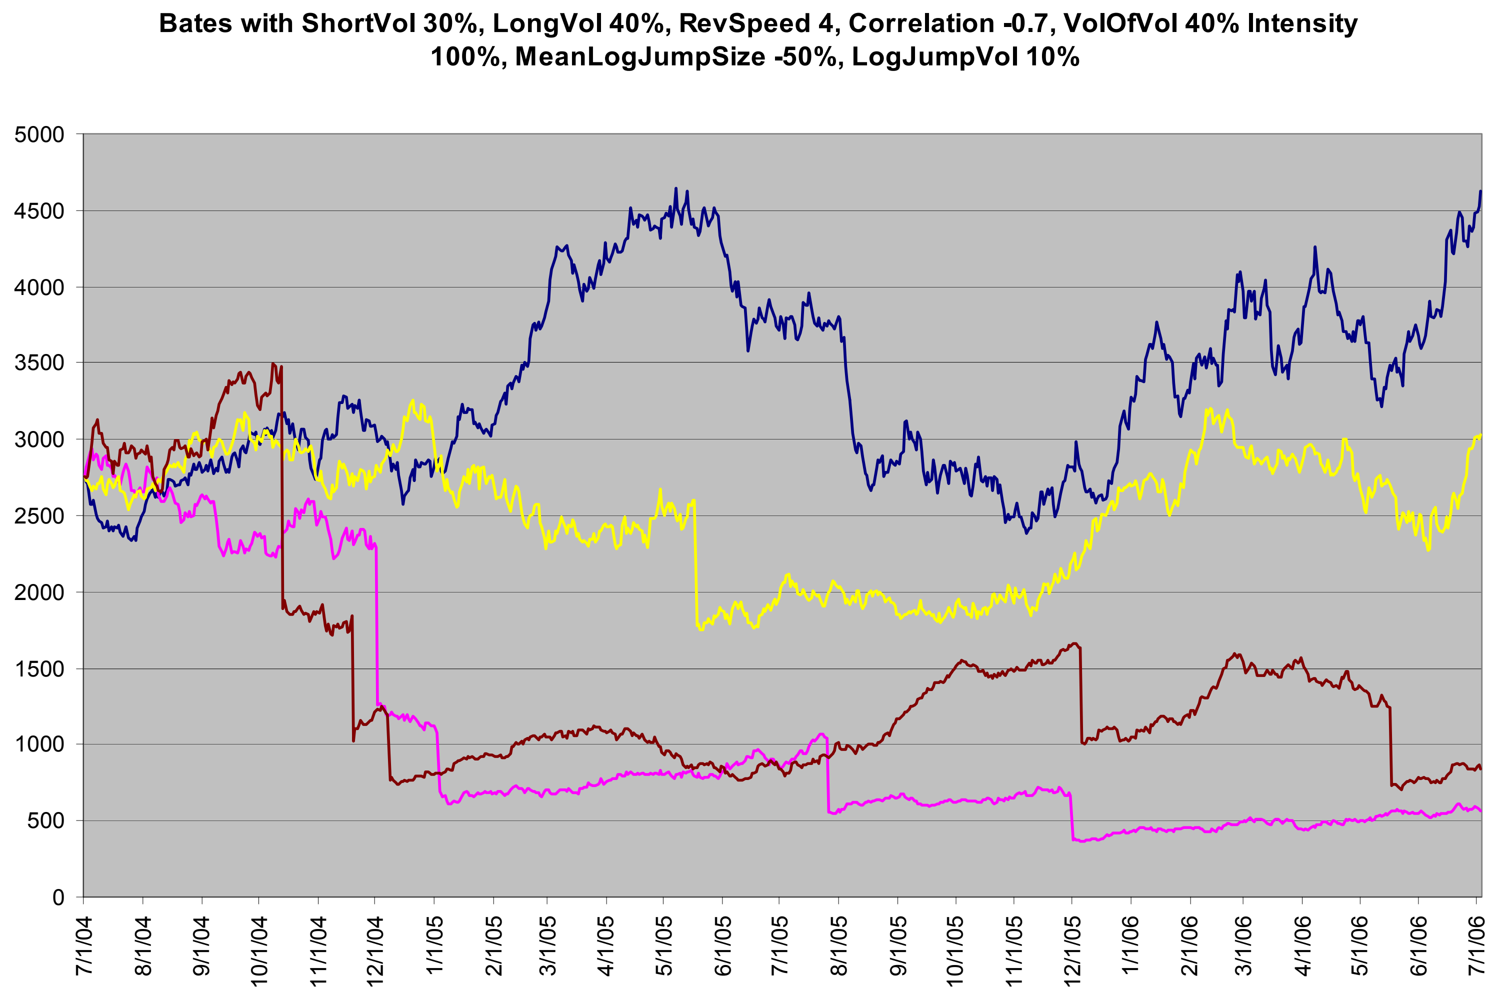

Stochastic Volatility

Additive processes

- A wide class of models has independent increments.
- As a generalisation of (stationary) Levy processes, additive processes can be used.
- These are Jump-Processes plus Diffusion component plus infinite activity.
- Widely known example: CGMY - a process with independent increments (generalisation of Variance Gamma, ie a Brownian motion time-changed with a stable tempered subordinator).
- Other examples are Barndorff-Nielsen (1997): Normal Inverse Gaussian.

- Advantages
  - Single source of randomness.
  - Nice theoretical features.
  - Many small jumps describe more realistically reality than continuous trading.
- Disadvantages
  - Hard/very hard to simulate in both FD and MC (in fact, one has to reduce to Jump-Diffusion case).
  - Incomplete markets.

Stochastic Volatility

The next generation

- None of the models of the previous category is a true market-model:
  Their parameters are not directly linked to market-observables.
- The resulting shape and dynamics of the implied volatility surface must be assessed after the model is been constructed - instead of the reverse approach to construct a model which has the desired properties per se.
- Starting with Schonbucher (1999) and Brace *et al* (2001), so-called *market models of implied volatility* or *stochastic implied volatility models* have been developed. Cont *et al* performed a series of statistical tests to investigate the dynamical behaviour of the full surface.
  - The mathematical framework behind this initially attractive approach is very difficult.
  - No publicly known application has been successful yet.
  - However, research is moving in this direction.
- Other ideas are
  - Mixed stochastic and local volatility models.
  - Distribution-based pricing.

Stochastic Volatility

Summary

- Various stochastic volatility models have been proposed
- Local Volatility captures the full smile surface and can be used to price a wide range of options.
- It has, however, the wrong dynamics in time and does not price strongly volatility-dependent products correctly.

- Most widely known stochastic volatility models are probably Heston and Merton.
- Such models are needed to overcome incorrect behaviour of the local volatility model.

Implementation

Calibration and Pricing

A Passion to Perform.

Deutsche Bank

Implementation
Naïve calibration

- Since we model the process under the martingale measure and if we assume that interest rates and forwards are known, we "only" have to ensure that the liquid options are replicated.
- The basic idea is to see the European price as a function of the model parameters Θ,

$$C(T,K) \equiv C(T,K;\Theta) = DF_T \, \mathrm{E}[(F_T Z_T - K)^+]$$

- Given the market prices $C_M(T,K)$ we therefore need to find a parameter set Θ such that

$$\left\| C(\cdot;\Theta) - C_M(\cdot) \right\|_d$$

  is minimised with respect to some metric *d*. Normally, one uses the L^2 norm, but L^1 might also sometimes be interesting (however, note that it is not smooth around zero).
  Normally weights are assigned to the various options.

- The problem of finding a minimum of a function is well-studied, and there are specialised algorithms for L^2 minimisation. They usually perform better if derivatives are known.

Implementation
Pricing European options

To calibrate, we need to price European options quickly.

Carr/Madan's algorithms helps if we know the CF of the log of the stock price.

- For calibration, we hence need to compute plenty of option prices.
- Using Carr&Madan (1998), we can price European calls if we know the characteristic function of the log of the stock price using the FFT.
- The idea is that once we know the characteristic function of the *call price* as a function of its log-strike *k*, we can invert the Fourier-transformation to compute a full sequence of prices in one run.
- However, note that the call price for the log-strike is not L^2 (which is required for FFT).
- We solve that problem using the "dampened call price" (*Z* is the log of the stock),

$$C_t^{\alpha}(k) := e^{-\alpha k}\,\mathrm{E}[(e^{Z_t} - e^{k})^{+}]$$

- A few simple computations yield

$$\mathrm{F}[C_t^{\alpha}](z) = \frac{\mathrm{F}[Z_t]\big(z + i(\alpha - 1)\big)}{(iz - \alpha)(iz - \alpha + 1)}$$

- FFT can be used and the original call price are available by inverting the dampening.

Implementation

Calibration - problems

- This approach works, but has some drawbacks
  - Routine is blind to true nature of the subject.
  - The call price function for most models is not bijective, hence there are many local minima. Sometimes these minima may even be connected (for example, the zero jump case of a Merton model can be achieved using both zero intensity or zero mean and zero volatility for the jumps).
  - Moreover, the parameters are not independent. Changing the reversion speed in Heston has a similar effect than changing the long volatility.
  - The option pricer might be very sensitive to changes in one parameter.
  - Once calibrated, the parameters may fluctuate too much on a daily basis.

- Improvement
  - Reparametrization of the model, if possible.
    For Heston, the VolOfVol *s* should also appear in the drift term , ie

$$\sigma_t = \sqrt{v_t} \qquad dv_t = \tilde{\kappa} s^2 (\theta - v_t) dt + s \sqrt{v_t} dW_t^v \qquad (2\tilde{\kappa}\theta > 1)$$

  - Penalise on past parameters.

Implementation

Calibration - penalties

Penalties stabilize the calibration.

- A refined approach on penalties (Cont/Tankov 2004)
  - Calibrate unconstrained (if possible, from several randomly chosen starting points). This yields an estimate of the “best achievable” error e_1.
  - Now calibrate with penalty on the past values under the constraint that the error is no less than, say, two times e_1.
- Example:
  - Unconstrained calibration vs fixed reversion speed and “soft penalty” on the past values.
  - We calibrate on 6m skew and ATM at 1y, 2y, 3y every day for more than one year using the previous values as starting points every day.

Implementation

Calibration

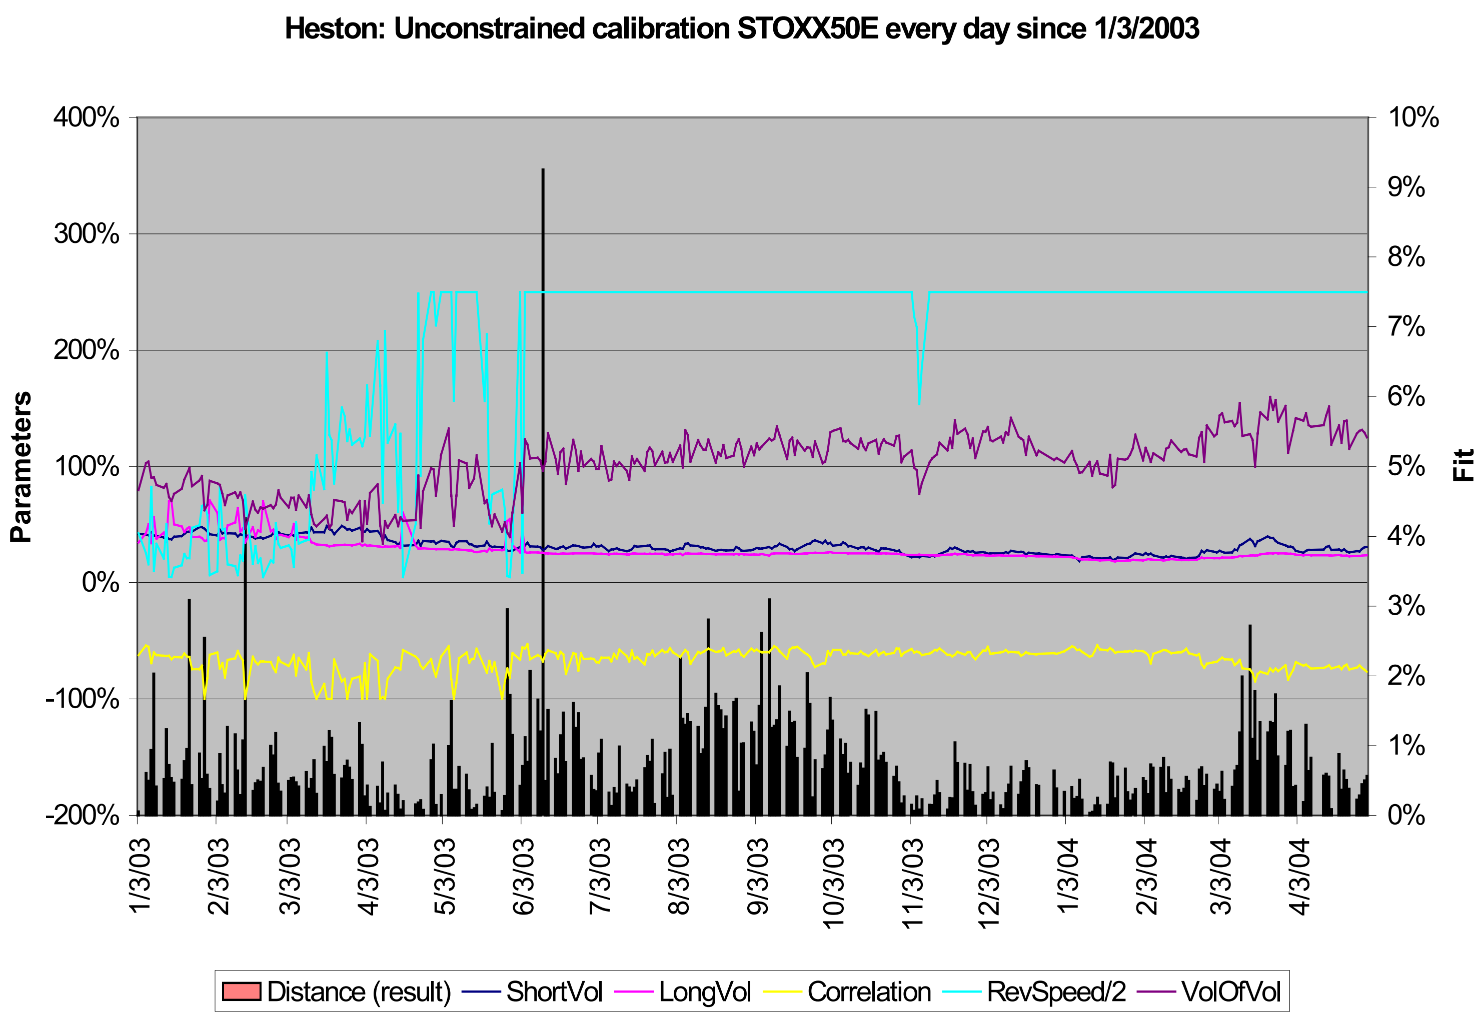

Implementation

Calibration

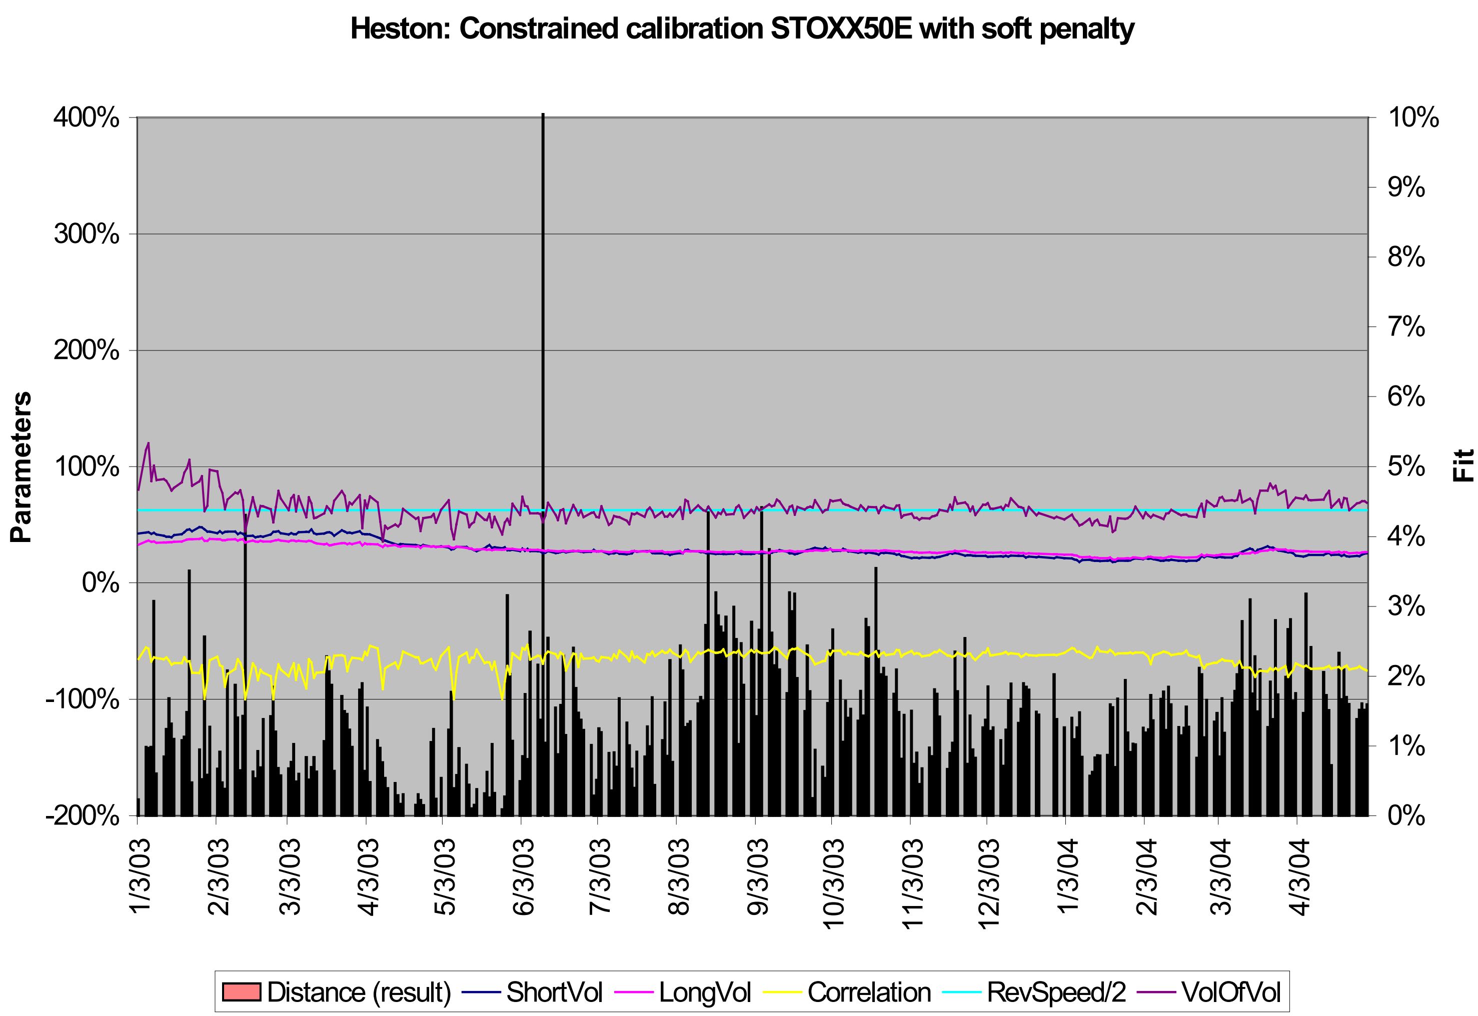

Numerical schemes

Pricing structured products - Monte-Carlo

Monte-Carlo is robust and easy to implement (at least for simple schemes)

- For stochastic volatility models with finite activity, Monte-Carlo simulation is easily available.
- Assume, for example Bates' model

$$dS_t = S_t \tilde{\mu}_t dt + \sigma_t S_t dW_t + S_t \sum_{i=1,\dots,N_t} (e^{\xi} - 1))$$

$$\sigma_t = \sqrt{v_t} \qquad dv_t = \kappa(\theta - v_t)dt + s\sqrt{v_t}\, dW_t^v \qquad (2\kappa\theta / s^2 > 0)$$

- If we have a payoff based on the stock value at some fixing dates,
  - Between two fixing dates t_j and t_{j+1}, compute the number of jumps n.
  - Conditional on the number of jumps, Bates' model is just Heston plus a sum of normals.
  - If the ξ_i are $N(o,u)$-normals, they can be simulated as one normal
  - Volatility can be simulated using Euler with m small steps of size $\delta = (t_{j+1} - t_j)/m$.

$$\sigma_t = \sqrt{v_t} \qquad v_j^{i+1} - v_j^i = \kappa(\theta - v_j^i)\delta + s\sqrt{(v_j^i)^+ \delta}\, \omega_i^{(1)} \qquad v_j^0 := v_j, v_j^m =: v_{j+1}$$

  - We also compute along the path

$$\overline{\omega}_j^{(1)} = \frac{1}{\sqrt{m}} \sum_{i=1,\dots,m} \omega_i^{(1)}, \quad \overline{v}_j = \delta \sum_{i=1,\dots,m} v_i$$

Numerical schemes

Pricing structured products - Monte-Carlo

- Hence we can compute the log of the stock as

$$
\begin{aligned}
Z_{j+1} - Z_j \quad &= \ln \tfrac{F_{j+1}}{F_j} - \tfrac{1}{2}\overline{v}_j + \sqrt{\overline{v}_j}\left(\rho \overline{\omega}_j^{(1)} + \sqrt{1-\rho_2}\,\omega_j^{(2)}\right) \\
&\quad - \lambda \Delta_j t (e^{\frac{1}{2}o^2 + u} - 1) + (\sqrt{n} o \xi + nu)
\end{aligned}
$$

- If the jump times are important (for example, when we want to price barriers)
  - Compute number *k* of jumps per interval
  - Simulate the *k* jump times (conditional on number of jumps the jump times are *k* uniform iid variables in the interval considered).
  - Simulate the jumps themselves, and perform the above Euler scheme in between.

Numerical schemes

Pricing structured products - Monte-Carlo

Control variates can massively decrease computation time.

- Performance of the scheme can be greatly improved using control-variates.
- In essence, we do not price a payoff *H* but

$$H - bG$$

where we know the price E[*G*] of G. Let us denote by *h* and *g* the estimated price of the options *H* and *G*, respectively, and use h^* as an estimator for E[*H*].

$$h^* := h - b(g - \mathrm{E}[G])$$

- What happens? Assume wlg that E[*G*] = 0.
  The rate of convergence for a product *F* is basically given by

$$\frac{\mathrm{Var}[F]}{\sqrt{N}}$$

Numerical schemes

Pricing structured products - Monte-Carlo

- In our case,

$$\mathrm{Var}[h-bg]=\mathrm{Var}[h]+b^2\,\mathrm{Var}[g]-2b\,\mathrm{Cov}[h,b]$$

this is a quadratic function in *b*. We derive it and set it to zero to find the optimal *b* as

$$b=\frac{\mathrm{Cov}[h,b]}{\mathrm{Var}[g]}$$

- Clearly, the involved quantities are now known and must by themselves by estimated.
- However, this technique is very powerful (and it can be shown that in fact most variance reduction techniques based on transforming paths or the payoff lead to expressions as above).
- The above can be extended to multiple control variates *G* (see Glassermann 04 or compute by yourself - it is just the same computation as above in more dimensions).

$$h^* := h-b(g-\mathrm{E}[G])$$

Numerical schemes

Pricing structured products - Monte-Carlo

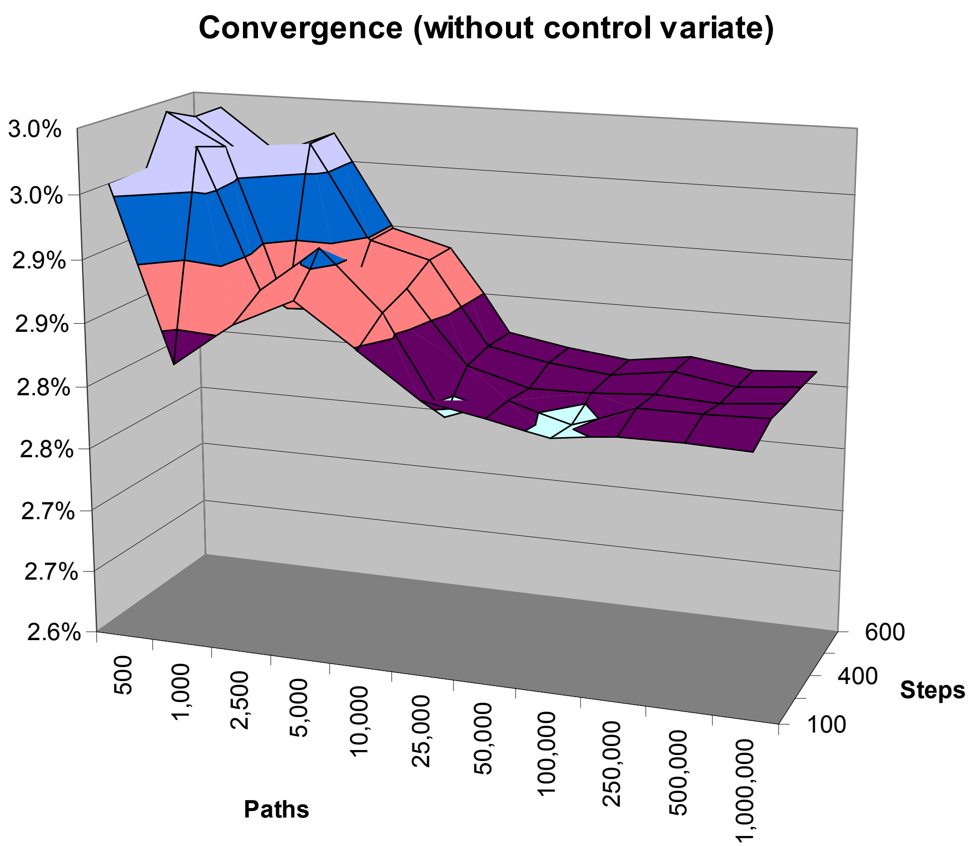

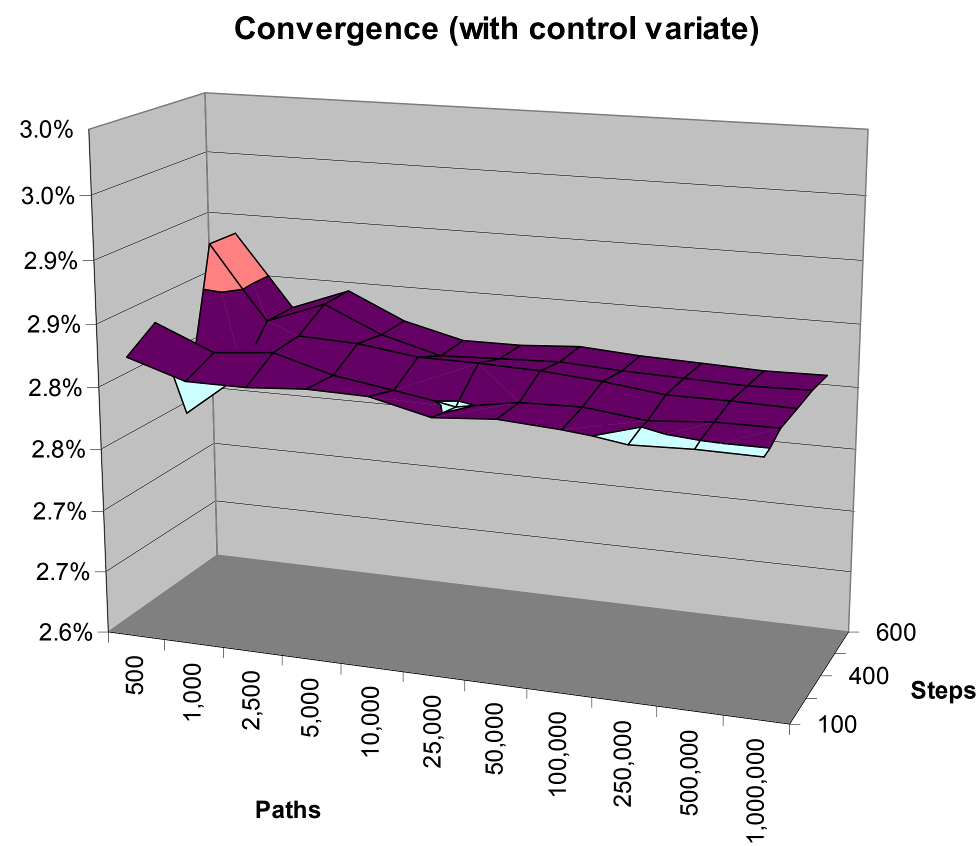

STOXX50E 1y Option On Variance, 20% Strike

Numerical schemes

Pricing structured products - Monte-Carlo

- More refined schemes than the Euler-scheme described are available
  - Broadie (Madrid 2004) - showed how to improve performance by simulating the Heston variance with its stationary (non-centred chi-square) distribution.
  - Glassermann (2000) - his book “Monte-Carlo Methods in Financial Engineering” is a very good guide through pricing with Monte-Carlo.

- If the structure in question is of barrier-type, it is more appropriate to simulate the jump-times explicitly and do refinement in between.
  - Rama/Cont (2004) wrote a good book on pricing with jump processes.

Implementation and Numerical schemes

Summary

- Calibration can be involved
  - European prices required and are available via FFT if characteristic function of the log of the stock price process can be computed.
  - Parameters might be too volatile
  - Too many local minima
- Monte-Carlo
  - Easily applicable.
- Finite Difference (see talk on Levy-processes)
  - Non-default jumps tricky.
  - Otherwise equally straight forward.

Variance products

Application of Stochastic Volatility models.

A Passion to Perform.

Deutsche Bank

Variance Swaps
Reminder

- Recall we wanted to price, more or less,

$$\mathrm{var}_{realized} := \frac{252}{n} \sum_{i=1,\ldots,n} X_i^2$$

which we approximated using

$$\langle Z \rangle_T$$

(up to scaling). In the case of a variance swap,

$$\mathrm{E}[\langle Z \rangle_T] = -2(\mathrm{E}[\ln S_T] - \ln \mathrm{F}_T)$$

which we can compute model-independently.

Variance Swaps

Calibration

- In reality, we have to limit ourselves in the sum

$$\ln(x) = \sum_i \alpha_i (x - k_i)^+ + \sum_j b_j (k_j - x)^+$$

- The market strike are between 2.5 and 4 standard-deviations.
- This is not a complete variance swap.
- When calibrating our models, we have to take this into account.
- Penalize on the value of a variance swap in the calibration of stochastic volatility models
  - We know the weights of the options from the computation above.
  - Use variance swap price as a condition in the calibration.
- The resulting model-parameters allow to recover the variance swap price.

- Note that if the characteristic function of the integrated variance is available, there is no need to run a Monte-Carlo if we are to price Europeans.

Options on Variance

Pricing using stochastic volatility

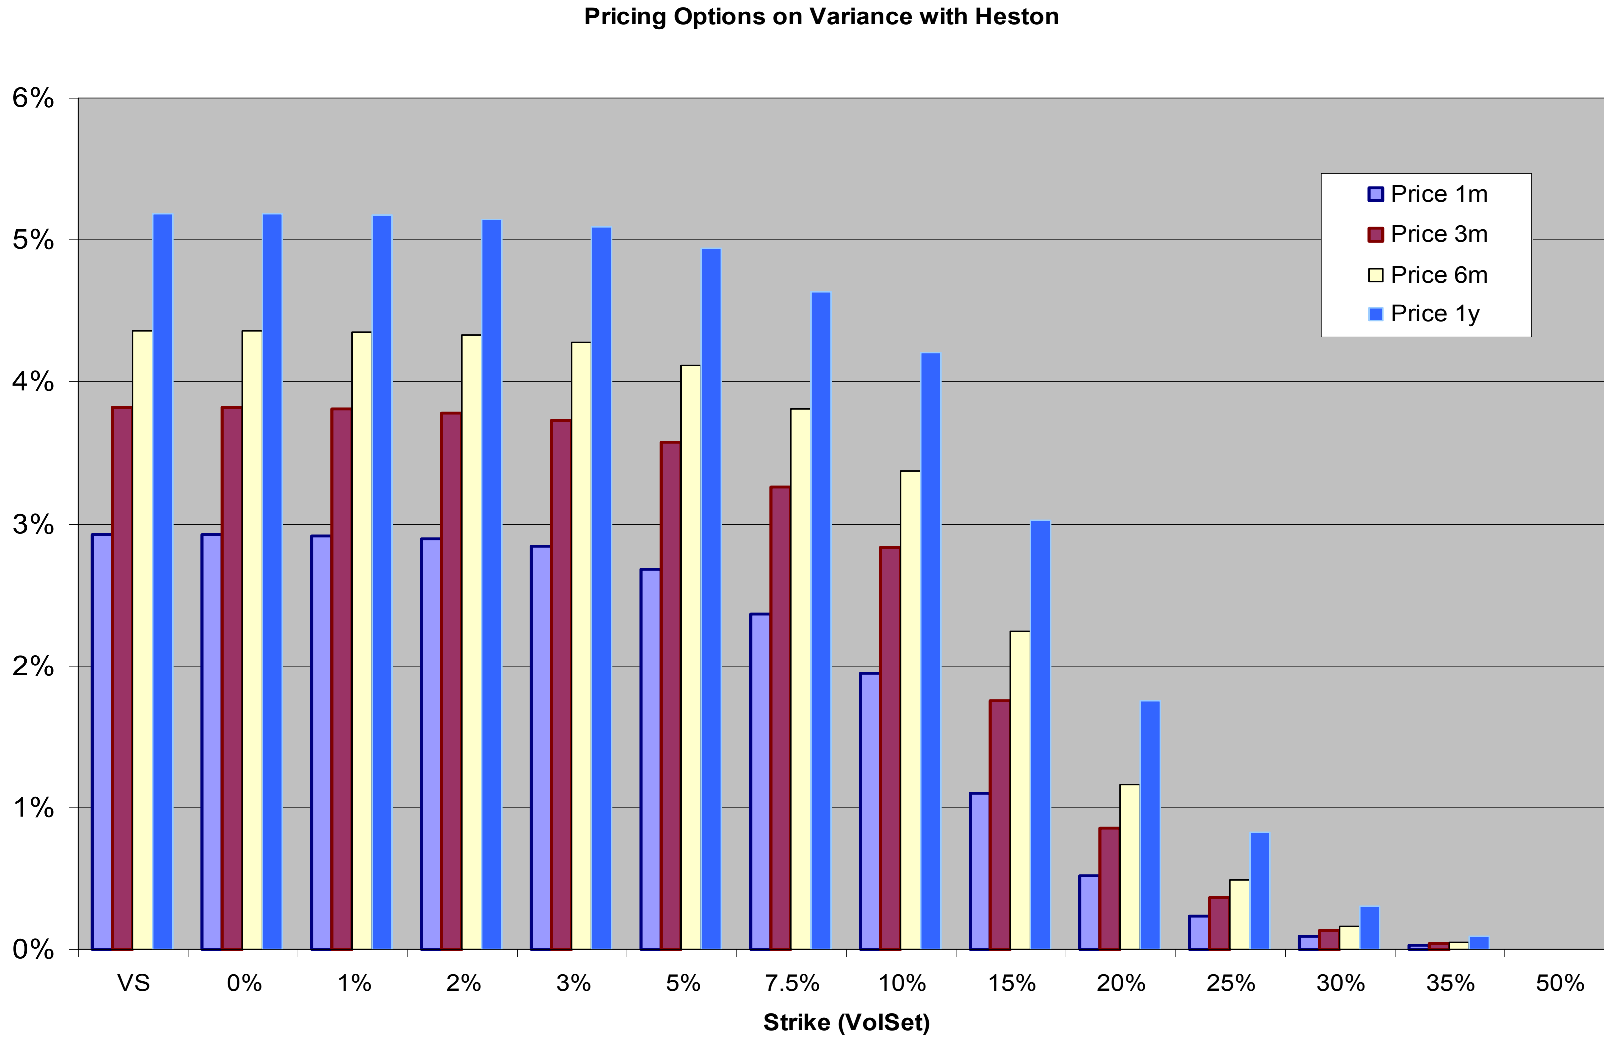

Stochastic Volatility

Pricing using stochastic Volatility

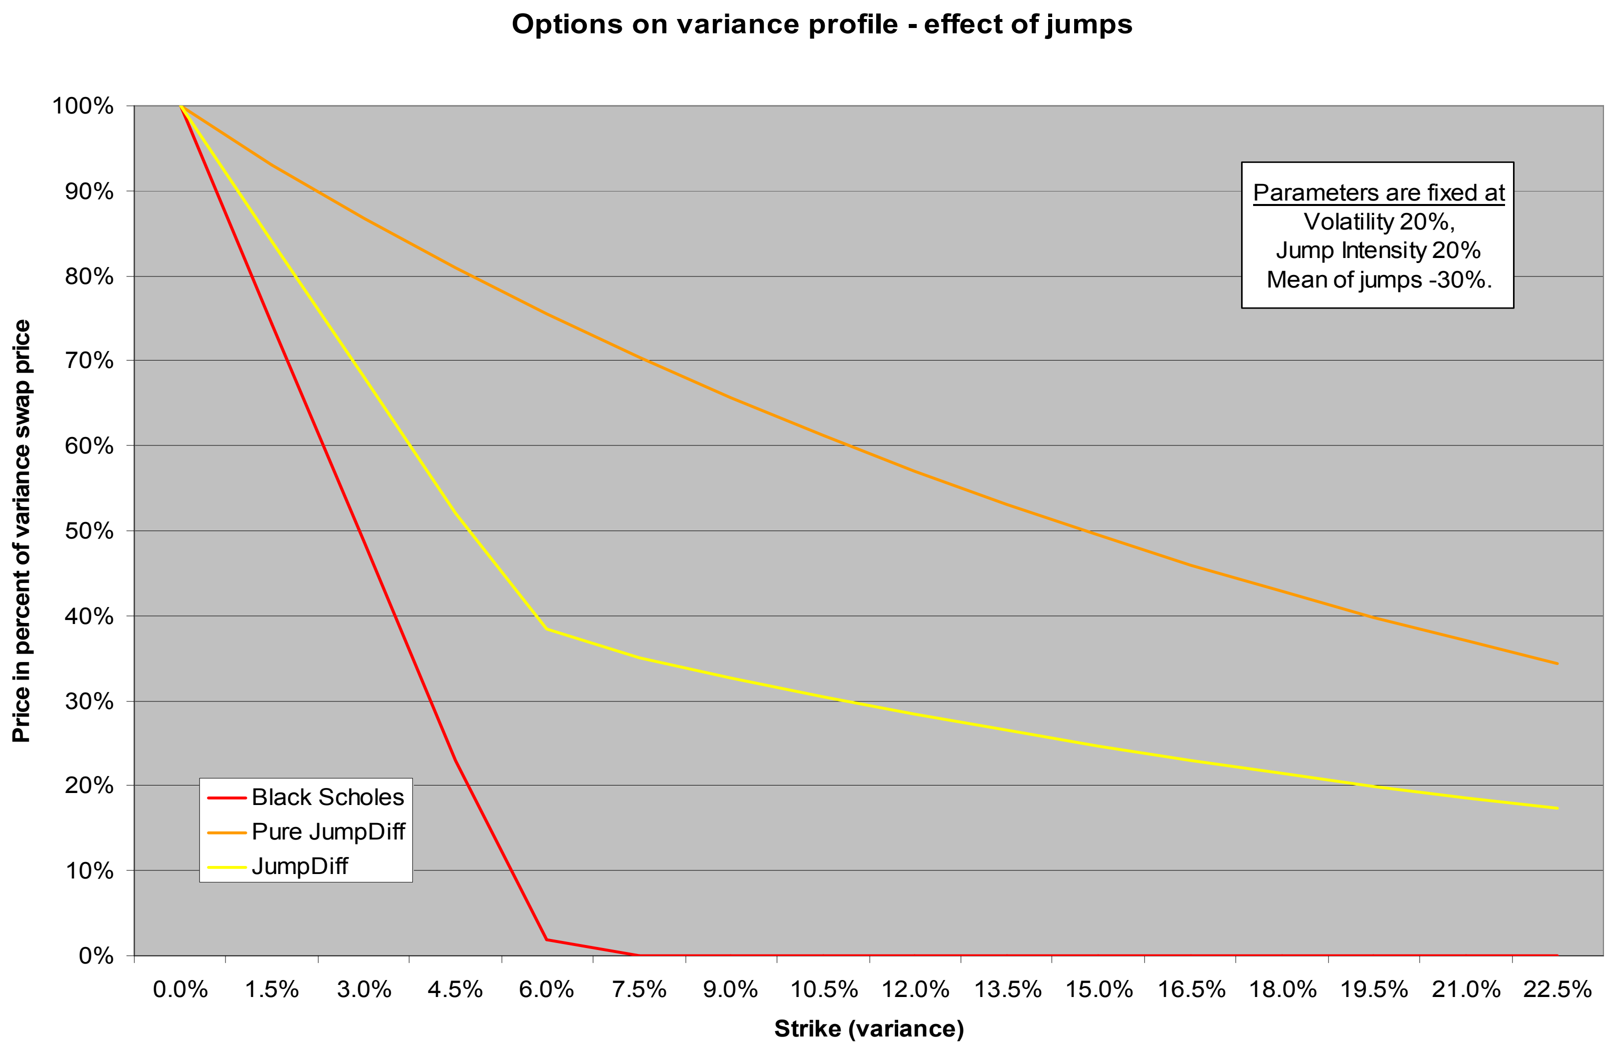

Stochastic Volatility

Pricing using stochastic Volatility

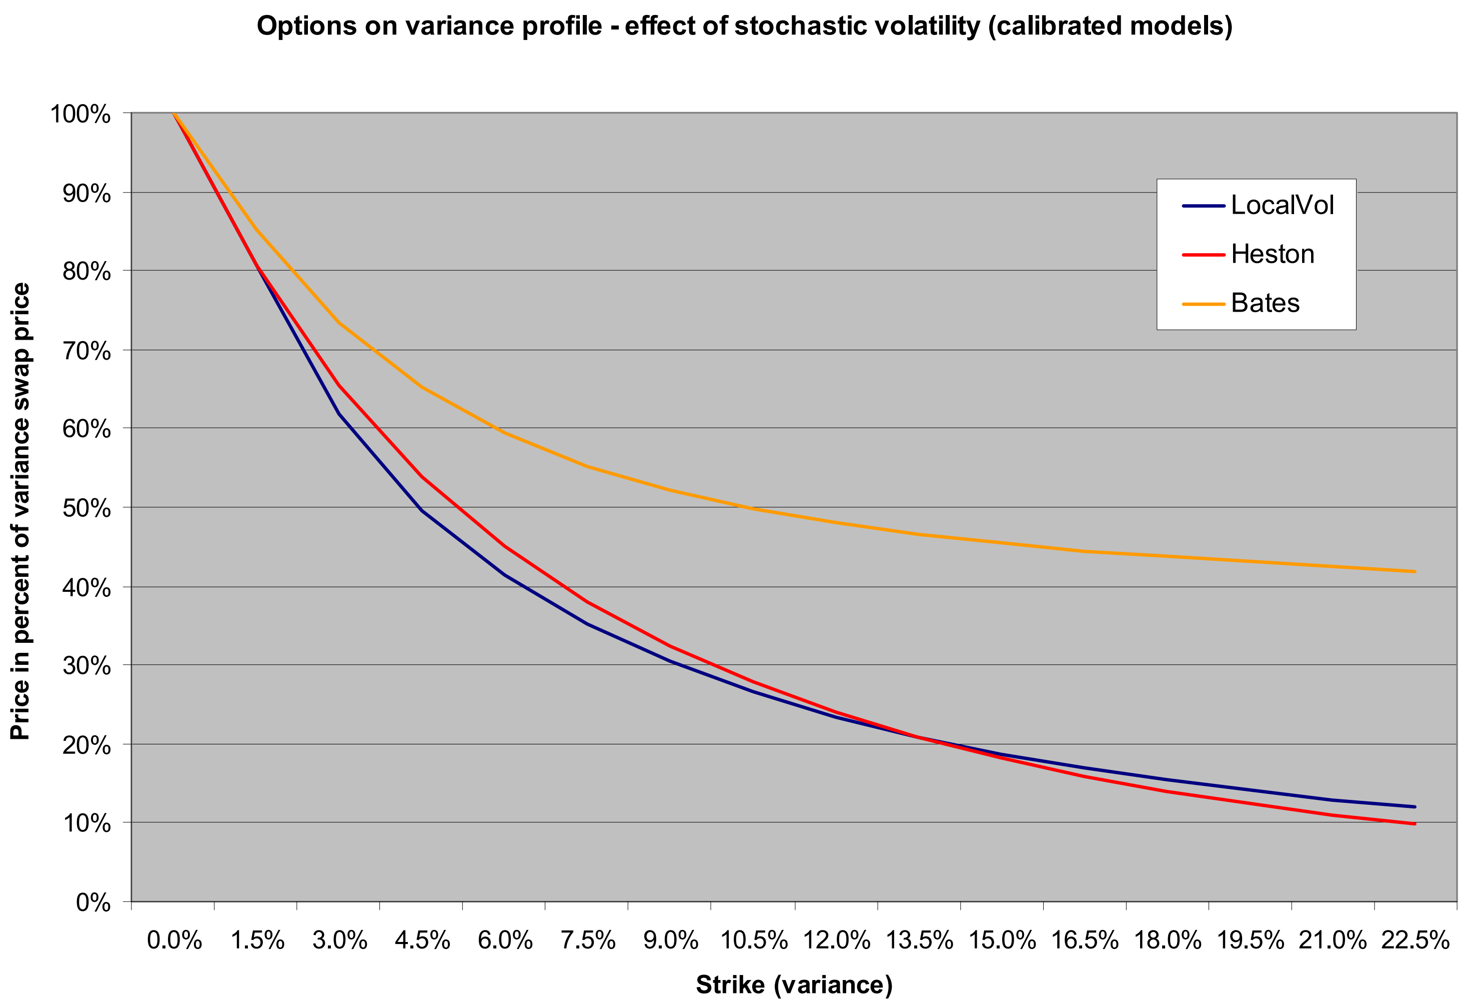

Deutsche Bank

Forward Started Options

Pricing Cliquets

- We now focus on the problem to price the *forward call spreads*

$$(S_T / S_t - k_1)^+ - (S_T / S_t - k_2)^+$$

  - These call-spreads are quite pure volatility products (note that delta is theoretically zero).
  - The price depends on the "skew" between the two strikes.
  - These options are the underlying blocks for *Cliquets*, which are of the type

$$\min\left(C, \max\left\{F, \sum_{i=1}^{n} \left(S_{t_i} / S_{t_{i-1}} - k_1\right)^+ - \left(S_{t_i} / S_{t_{i-1}} - k_2\right)^+ \right\}\right)$$

    where the fixing dates t_i are equidistant with a typical length of 1m, 3m or 6m while the whole trade lasts for 5y, for example. We will furthermore assume $k_1=90\%$ and $k_2=110\%$.
  - It is a common assumption that the *forward implied skew* between the two strikes k_1 and k_2 should be roughly the skew of today. That allows in principle to price the above structure if *cap* and *floor* are trivial.

- We search for a model which prices the skew between the strikes correctly.
  - Calibrate to a strip of forward starts with prices computed according to our forward volatility assumption.
  - In the following graphs, all forwards and discount factors are trivial.

Cliquets

Pricing using stochastic volatility

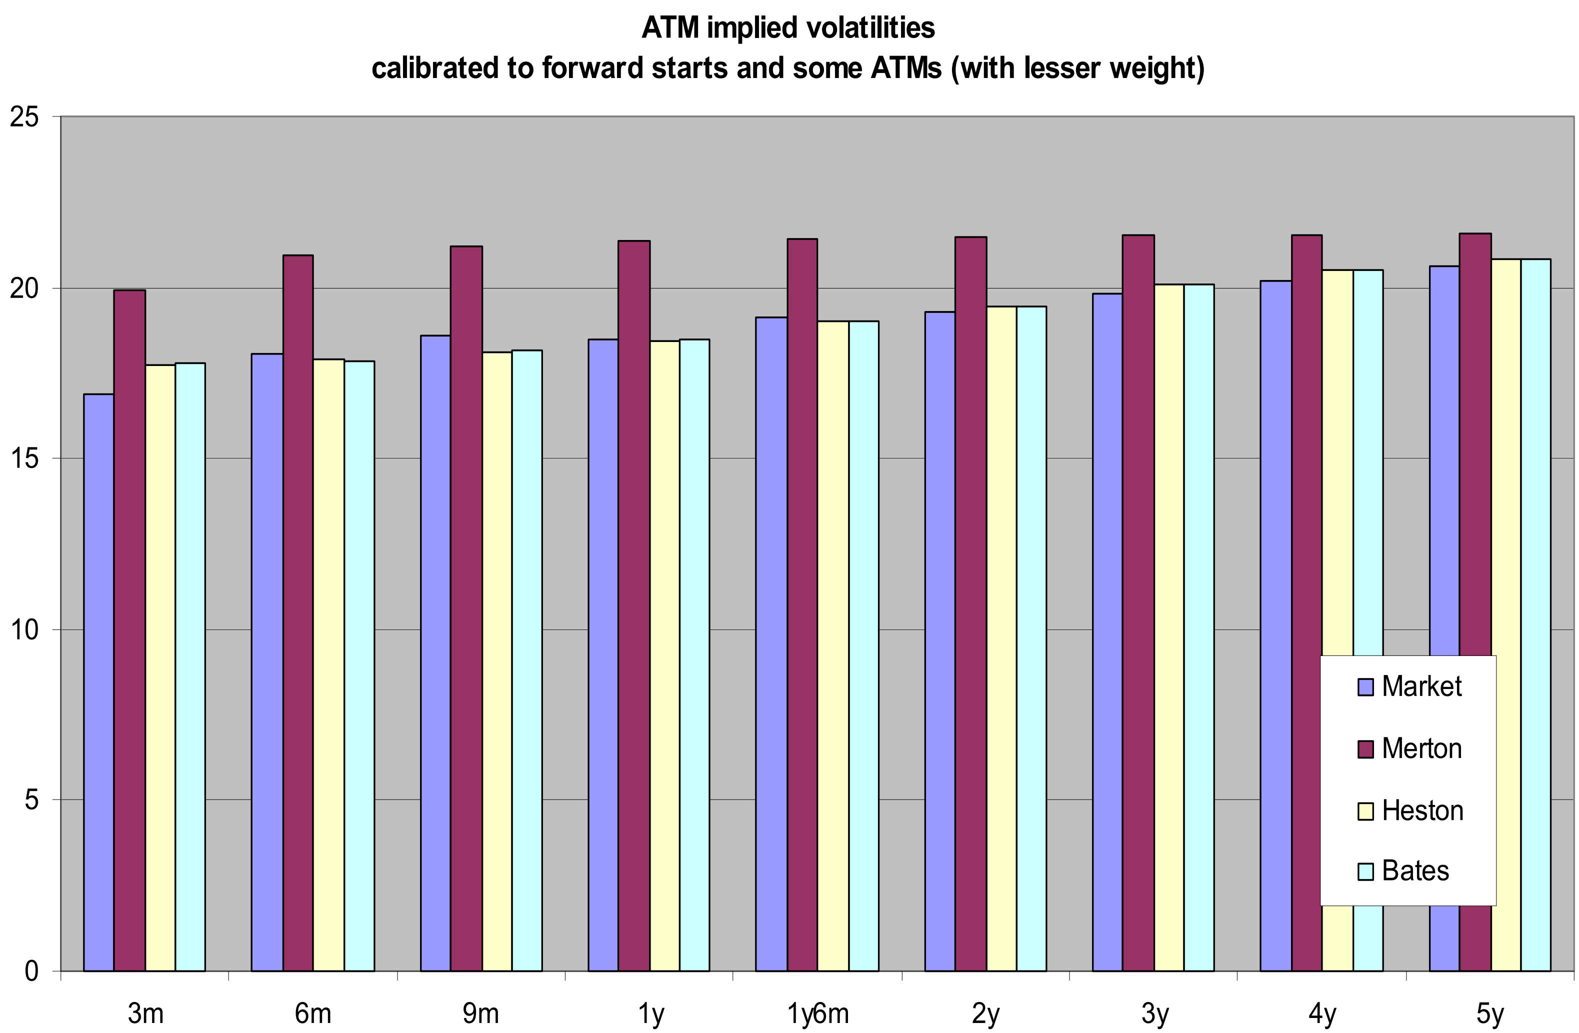

Cliquets

Pricing using stochastic volatility

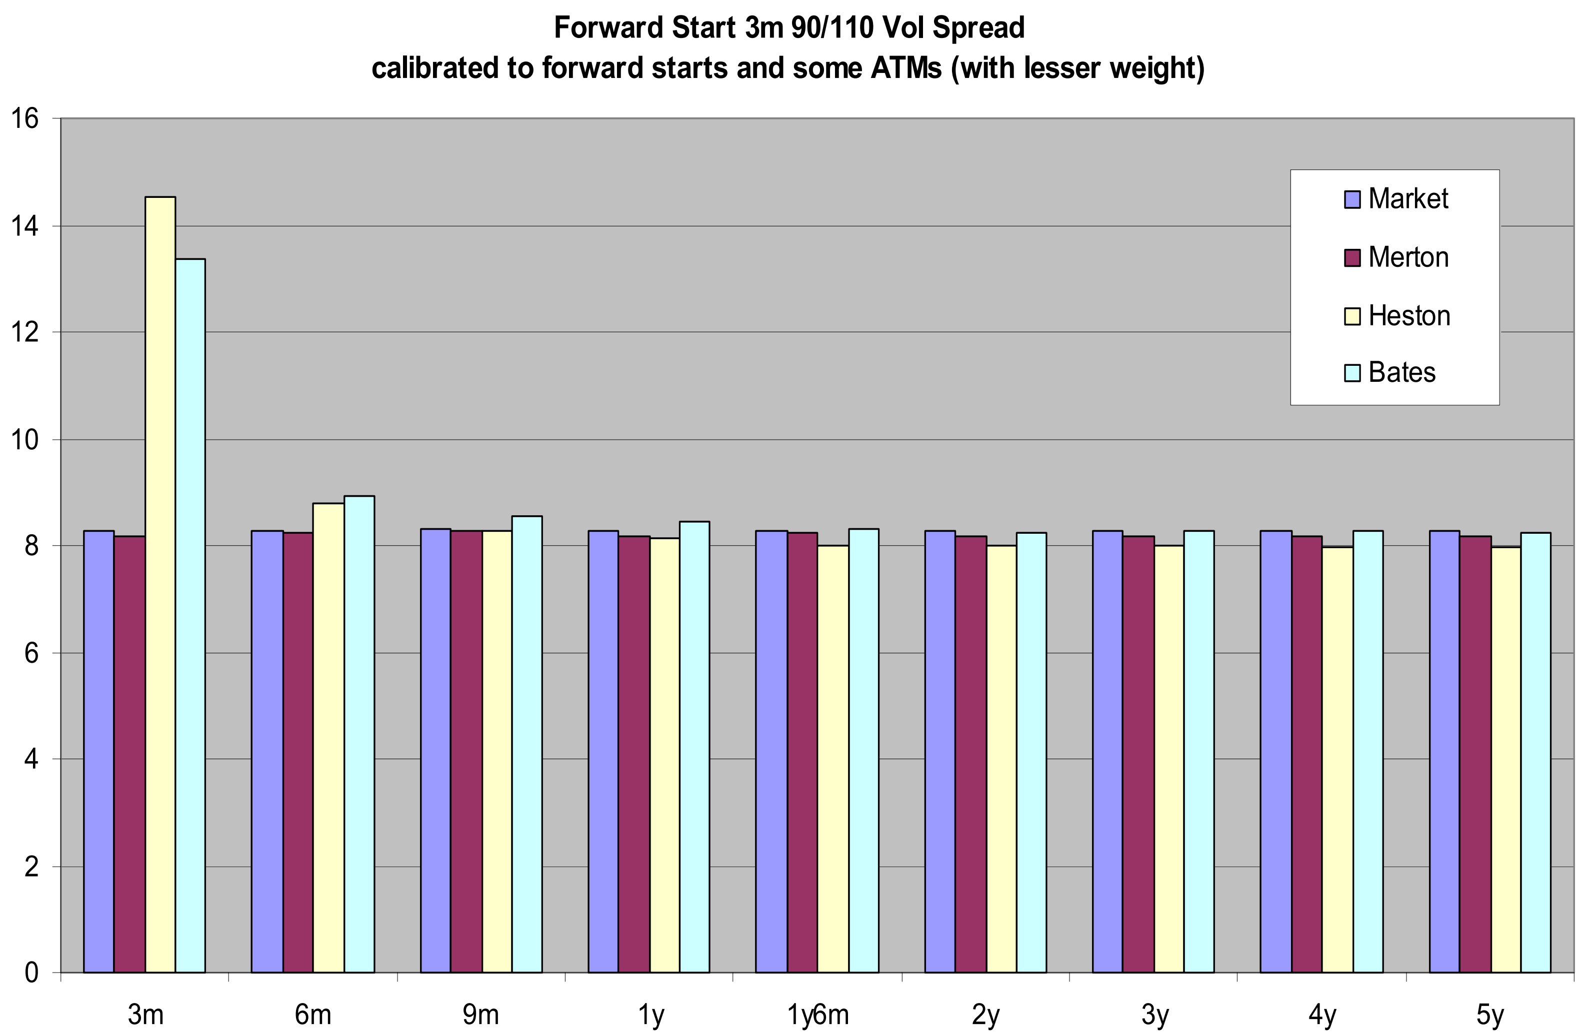

Cliquets
Calibration

- Merton's Jump-Diffusion prices the skew very well.
  - The process is a Levy-process and has stationary increments.
  - All Levy processes will have a fixed implied relative skew.
  - ATM fit quite poor.

- Heston fits better, except the short end skew.

- Bates' Heston-Jump-Diffusion behaves roughly like Heston.
  - Idea: Take the Heston-fit and modify the jump-diffusion component accordingly to achieve a better fit.
  - We take the former Heston parameters as a result, and plug them into Bates with *time-dependent* jump-diffusion parameters.

Cliquets

Adjusted calibration

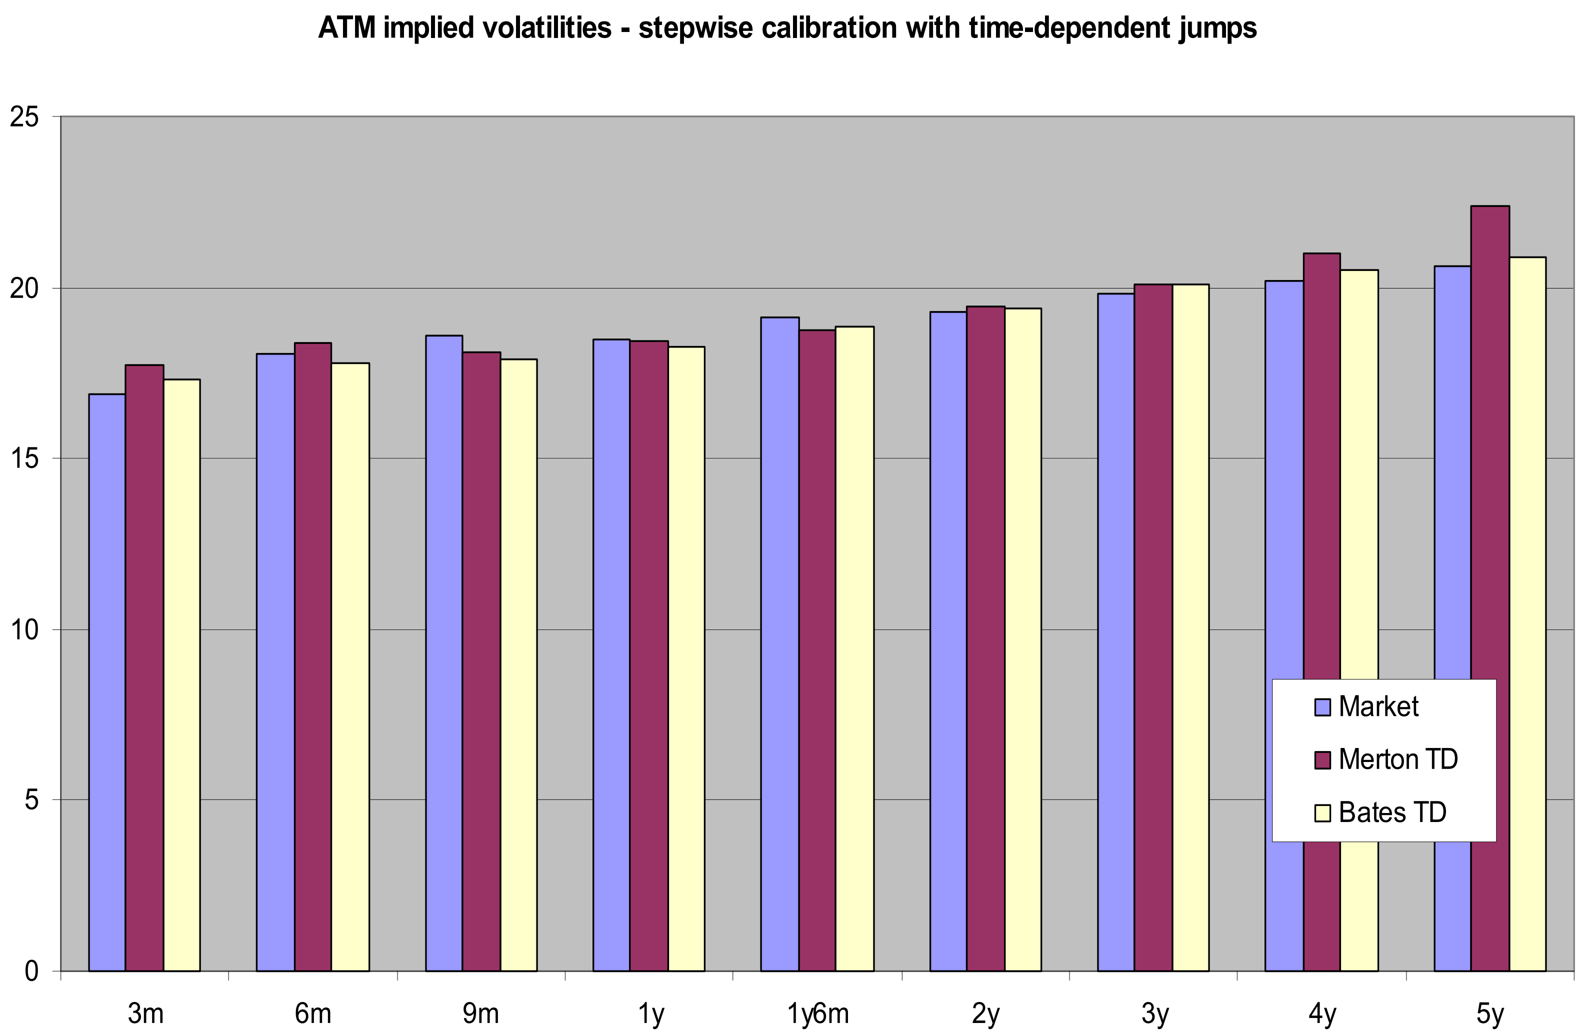

Cliquets
Adjusted calibration

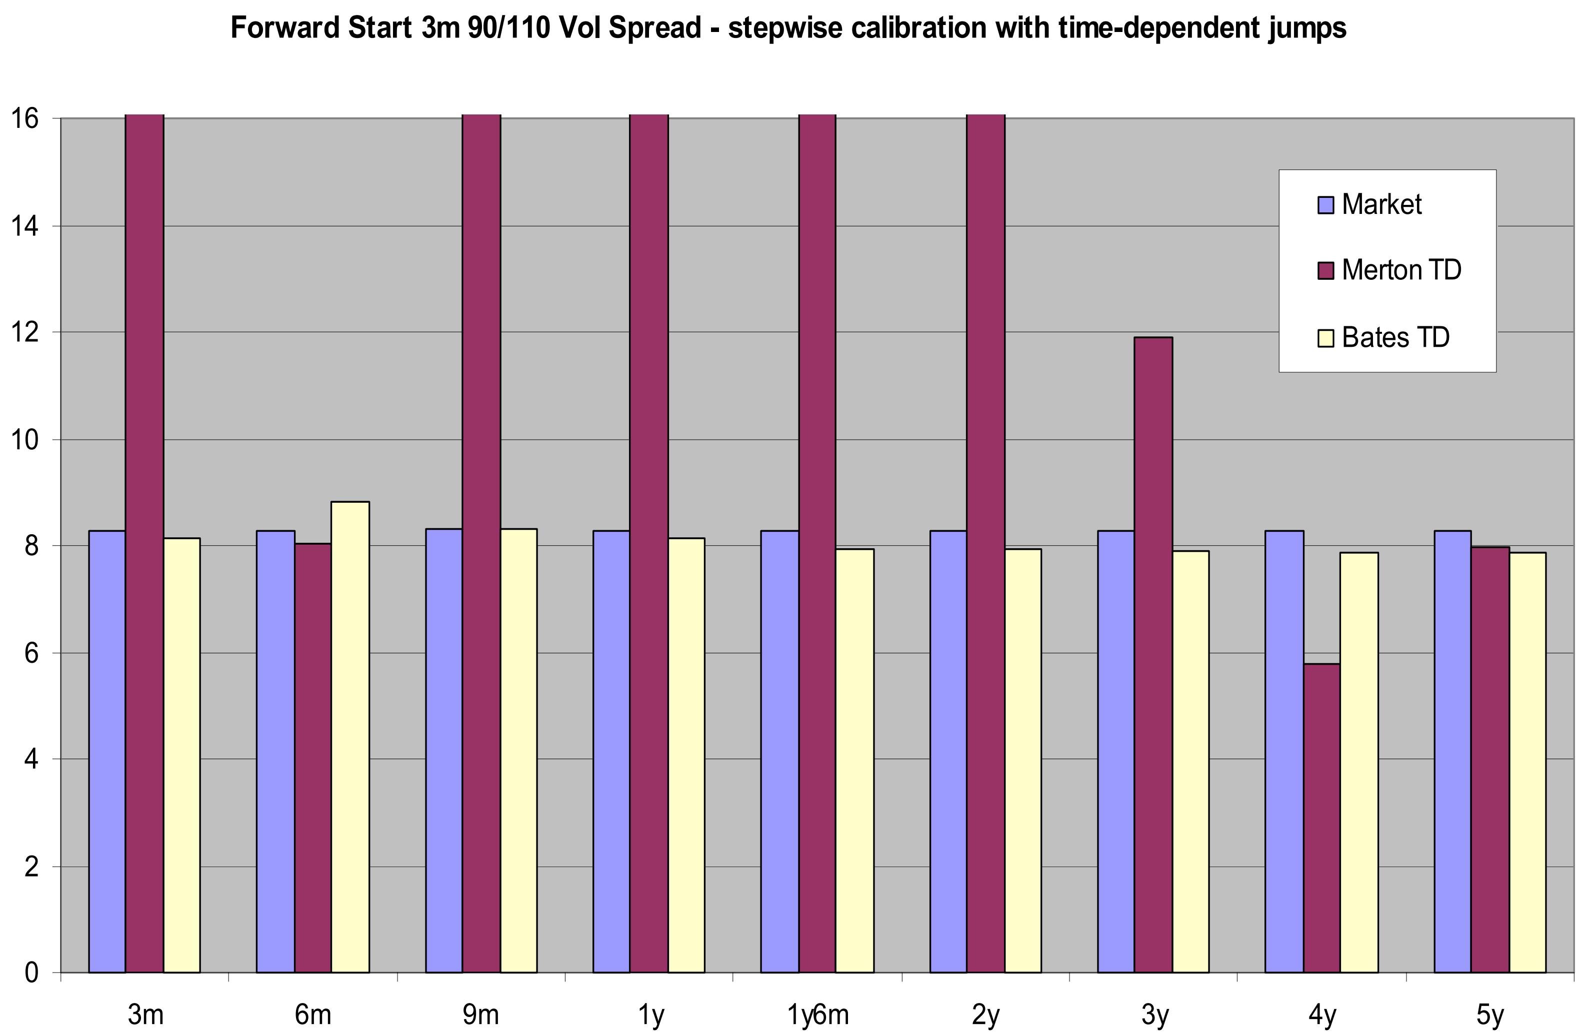

Cliquets

Prices are quite similar since they depend on forward skew.

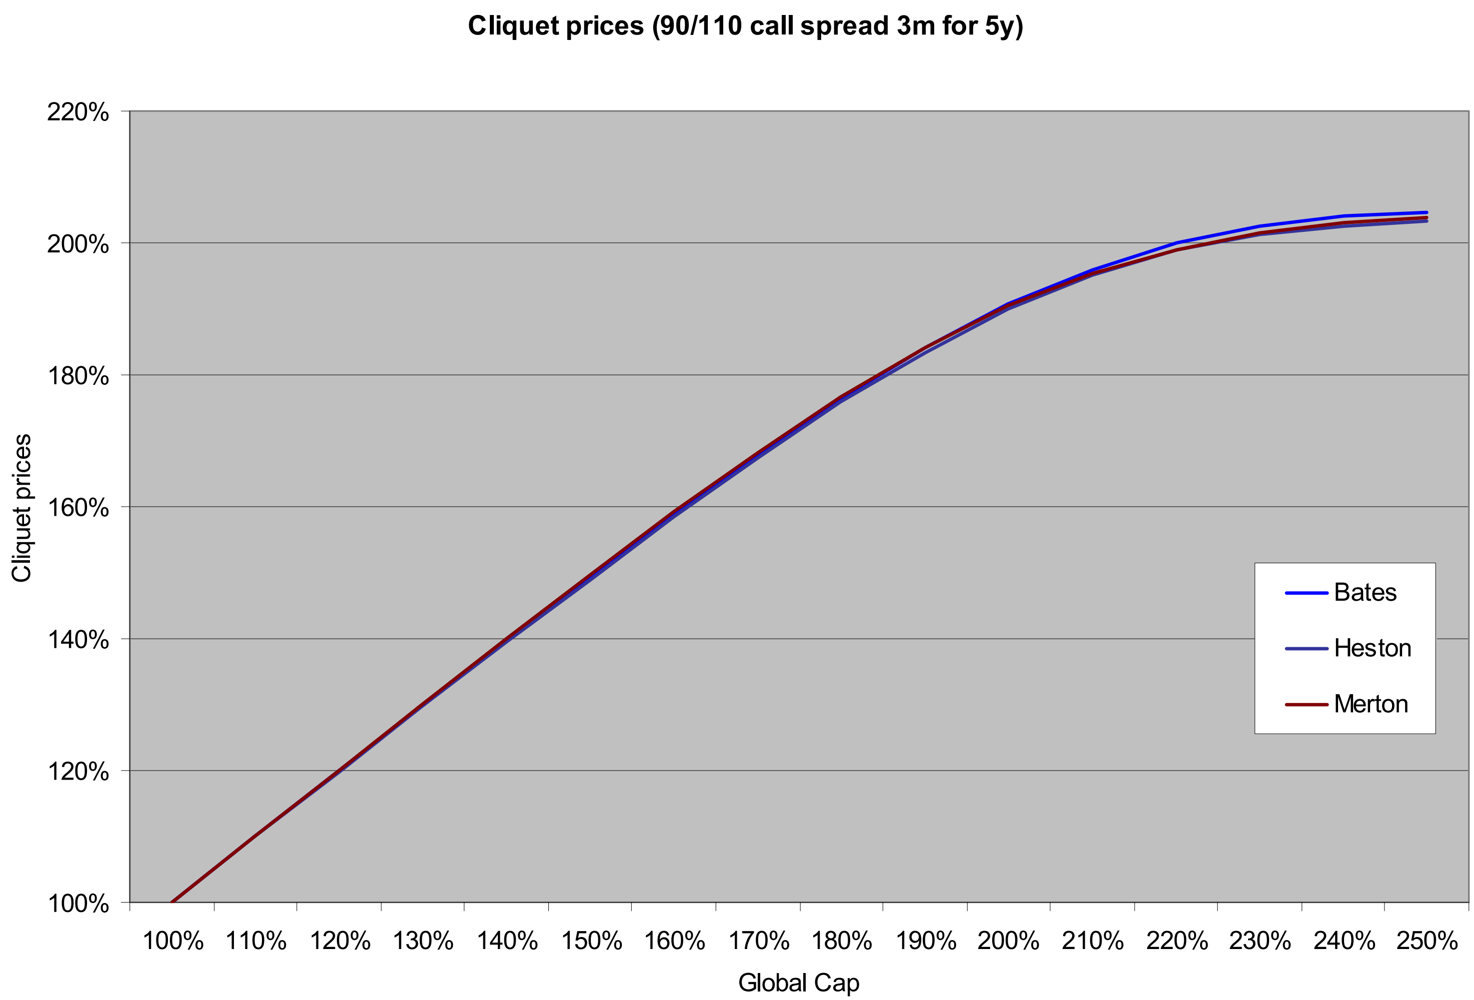

Cliquets

For other products, price can differ widely.

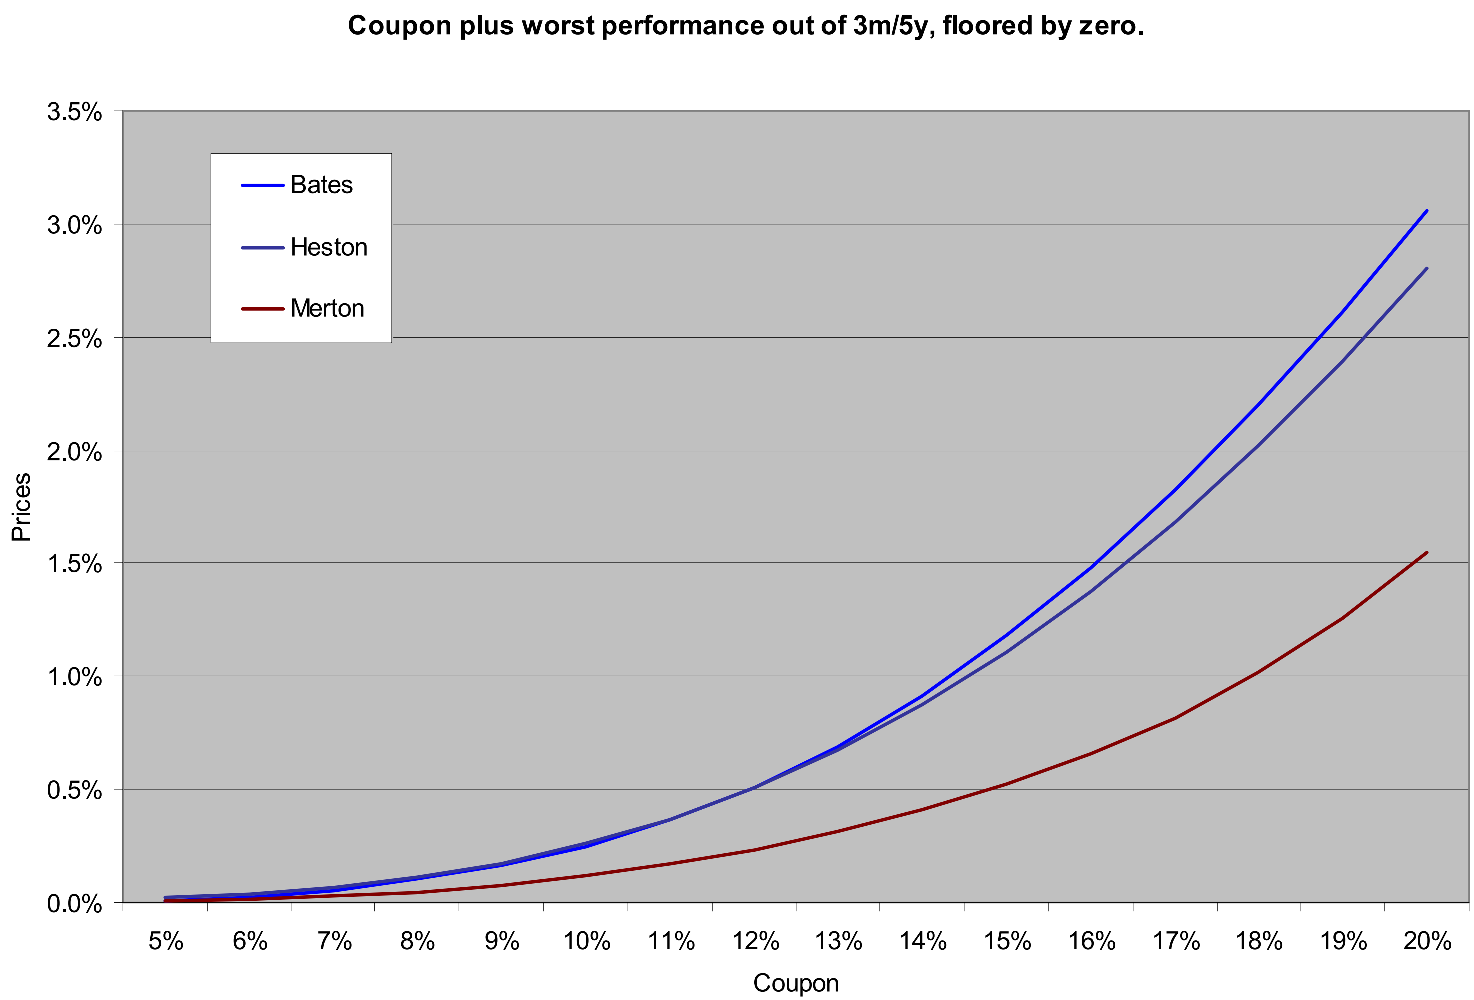

Options on Variance

Conclusion

- Of course, we have a discrepancy in our pricing between the models.
- Hence, the choice of the model depends on the risk profile and the trader’s view. Statistical investigation of historical data can also give a hint on the shape of the distribution of realized variance. A particular question is to what extend jumps have to be modelled.
- Calibration should take “related” products into account.

Thank you very much for your attention.

hans.buehler@db.com
http://www.dbquant.com

A Passion to Perform.

Deutsche Bank

Literature

Recommended reading

- Monographs
  - CT04: Cont/Tankov: “Financial Modelling with Jump Processes” (2004)
  - OV02: Overhaus et al “Equity deriviatives” (2002)
  - GL00: Glassermann “Monte-Carlo Methods in Financial Engineering” (2000)
- Papers
  - Barndorff-Nielsen, 1997: *Normal inverse Gaussian distributions and stochastic volatility modelling*. Scand J. Statist. 24, 1-13.
  - Bates, 1996: *Jumps and stochastic volatility: exchange rate process implicit in DM options*. Rev. Fin. Studies 9-1
  - Bouleau et al, 1989: *Residual risks and hedging strategies in markovian markets*. Stochastic Process. Appl., 33 (1989), pp. 131-159
  - Brace et al 2001: *Market Model of Stochastic Implied Volatility with an Application to the BGM model*, Working paper (http://www.maths.unsw.edu.au/~rsw/Finance/svol.pdf)

Literature

Recommended reading

- Papers continued

  - Carr et al, 1998: *Option Valuation Using the Fast Fourier Transform.* Journal of Computational Finance, 2, pp. 61-73
  - Carr et al, 2002: *The Fine Structure of Asset Returns: An Empirical Investigation.* Journal of Business, April 2002, Volume 75 Number 2, pp.305-32.
  - Hagan et al, 2000: *Managing Smile Risk* (http://www.math.columbia.edu/~lrb/sabrAll.pdf)
  - Heston, 1993: *A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options.*The Review of Financial Studies, 6, pp. 327-343.
  - Madan et al, 1990: *The Variance Gamma (V.G.) Model for Share Market Returns.* Journal of Business, 1990, vol. 63, issue 4, pp.511-24
  - Merton, 1976: *Option Pricing When Underlying Stock Returns are Discontinuous.* Journal of Financial Economics 3 (1976) pp. 125-144.
  - Schonbucher, 1999: *A Market Model for Stochastic Implied Volatility*. SFB 303 Working Paper No. B - 453
  - Schoebel et al, 1999: *Stochastic Volatility with an Ornstein-Uhlenbeck process: An Extension*. European Finance Review 3: 23–46.