IEOR E4602: Quantitative Risk Management

Model Risk

Martin Haugh

Department of Industrial Engineering and Operations Research
Columbia University
Email: martin.b.haugh@gmail.com

Outline

A Simple Example: Pricing Barrier Options

Barrier options are options whose payoff depends in some way on whether or not a particular barrier has been crossed before the option expires.

A barrier can be:

- A put or a call
- **A knock-in** or **knock-out**
- A **digital** or vanilla
  - so (at least) eight different payoff combinations.

e.g. A knockout put option with strike K, barrier B and maturity T has payoff

$$\text{Knock-Out Put Payoff} \;=\; \max\left(0,\; (K - S_T)\, 1_{\{S_t \geq B \text{ for all } t \in [0,T]\}}\right).$$

e.g. A digital down-and-in call option with strike K, barrier B and maturity T has payoff

$$\text{Digital Down-and-In Call} \;=\; \max\left(0,\; 1_{\{\min_{t\in[0,T]} S_t \leq B\}} \times 1_{\{S_T \geq K\}}\right).$$

Barrier Options

Knock-in options can be priced from knock-out options and vice versa since a knock-in plus a knock-out – each with the same strike – is equal to the vanilla option or digital with the same strike.

Analytic solutions can be computed for European barrier options in the Black-Scholes framework where the underlying security follows a geometric Brownian motion (GBM).

Will not bother to derive or present these solutions here, however, since they are of little use in practice

- this is because the Black-Scholes model is a terrible(!) model for pricing barrier options.

But can still use GBM to develop intuition and as an example of model risk.

Will concentrate on knock-out put option

- they are traded quite frequently in practice.

Value of a Knockout Put Option Using Black-Scholes GBM

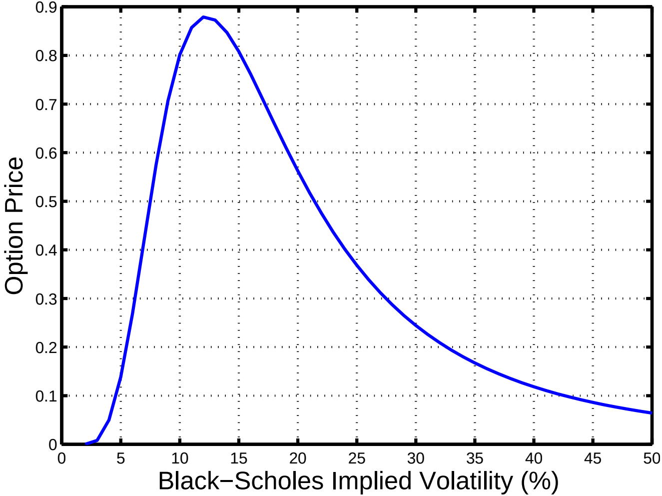

$S_0 = 105$, $K = 100$, Barrier $= 85$, $T = 6$ months, $r = q = 0$

Value of Knockout Put with Corresponding Vanilla Put

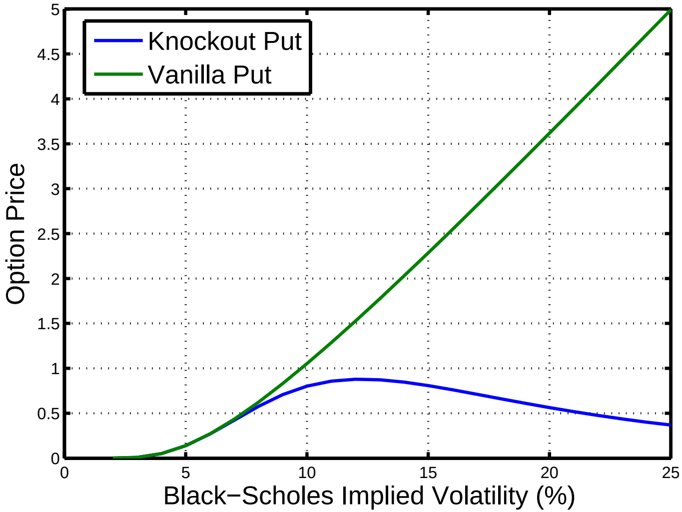

Barrier Options

So knock-out put option always cheaper than corresponding vanilla put option.

For low values of σ, however, the prices almost coincide. Why?

While the vanilla option is unambiguously increasing in σ the same is not true for the knock-out option. Why?

Question: What do you think would happen to the value of the knock-out put option as $\sigma \to \infty$?

Black-Scholes model is not a good model for pricing barrier options:

- It cannot price vanilla options correctly so certainly cannot price barriers.
- Suppose there was no market for vanillas but there was a liquid market for knockout puts. What value of σ should be used to calibrate your BS price to the market price?
- The Black-Scholes Greeks would also be very problematic: what implied volatility would you use to calculate the Greeks?

In summary, the Black-Scholes model is a disaster when it comes to pricing barrier options.

E.G: Parameter Uncertainty and Hedging in Black-Scholes

Now consider the use of the Black-Scholes model to hedge a vanilla European call option in the model.

Will assume that assumptions of Black-Scholes are correct:

- Security price has GBM dynamics
- Possible to trade continuously at no cost
- Borrowing and lending at the risk-free rate are also possible.

Then possible to dynamically replicate payoff of the call option using a self-financing (s.f.) trading strategy

- initial value of this s.f. strategy is the famous Black-Scholes arbitrage-free price of the option.

The s.f. replication strategy requires **continuous** delta-hedging of the option but of course not practical to do this.

Instead we hedge **periodically** – this results in some **replication error**

- but this error goes to 0 as the interval between rebalancing goes to 0.

E.G: Parameter Uncertainty and Hedging in Black-Scholes

P_t denotes time t value of the discrete-time s.f. strategy and C_0 denotes initial value of the option.

The replicating strategy is then satisfies

$$
\begin{aligned}
P_0 &:= C_0 && (1)\\
P_{t_{i+1}} &= P_{t_i} + \left(P_{t_i} - \delta_{t_i} S_{t_i}\right) r\Delta t + \delta_{t_i}\left(S_{t_{i+1}} - S_{t_i} + qS_{t_i}\Delta t\right) && (2)
\end{aligned}
$$

where:

- $\Delta t := t_{i+1} - t_i$
- $r =$ risk-free interest rate
- q is the dividend yield
- δ_{t_i} is the Black-Scholes delta at time t_i
  - a function of S_{t_i} and some assumed implied volatility, σ_{imp}.

Note that (1) and (2) respect the s.f. condition.

E.G: Parameter Uncertainty and Hedging in Black-Scholes

Stock prices are simulated assuming $S_t \sim \mathsf{GBM}(\mu, \sigma)$ so that

$$S_{t+\Delta t} \;=\; S_t e^{(\mu - \sigma^2/2)\Delta t + \sigma\sqrt{\Delta t} Z}$$

where $Z \sim \mathsf{N}(0, 1)$.

In the case of a short position in a call option with strike K and maturity T, the final trading P&L is then defined as

$$\text{P\&L} \;:=\; P_T - (S_T - K)^+ \tag{3}$$

where P_T is the terminal value of the replicating strategy in (2).

In the Black-Scholes world we have $\sigma = \sigma_{imp}$ and the P&L $= 0$ along every price path in the limit as $\Delta t \to 0$.

In practice, however, we cannot know σ and so the market (and hence the option hedger) has no way to ensure a value of σ_{imp} such that $\sigma = \sigma_{imp}$.

E.G: Parameter Uncertainty and Hedging in Black-Scholes

This has interesting implications for the trading P&L: it means we cannot exactly replicate the option even if all of the assumptions of Black-Scholes are correct!

In figures on next two slides we display histograms of the P&L in (3) that results from simulating 100k sample paths of the underlying price process with $S_0 = K = \$100$.

E.G: Parameter Uncertainty and Hedging in Black-Scholes

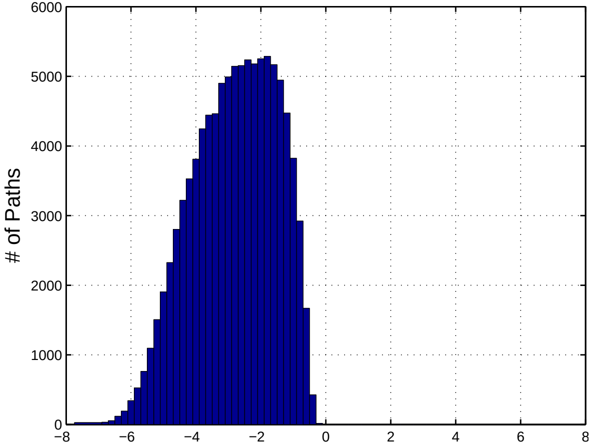

Histogram of delta-hedging P&L with true vol. = 30% and implied vol. = 20%.

Option hedger makes substantial loses. **Why?**

E.G: Parameter Uncertainty and Hedging in Black-Scholes

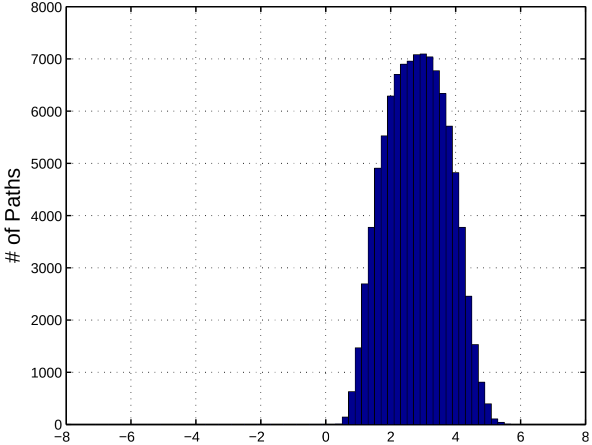

Histogram of delta-hedging P&L with true vol. = 30% and implied vol. = 40%.

Option hedger makes substantial gains. **Why?**

E.G: Parameter Uncertainty and Hedging in Black-Scholes

Clearly then this is a situation where substantial errors in the form of non-zero hedging P&L's are made

- and this can only be due to the use of incorrect model parameters.

This example is intended to highlight the importance of not just having a good model but also having the correct model parameters.

The payoff from delta-hedging an option is in general path-dependent.

Can be shown that the payoff from continuously delta-hedging an option satisfies

$$\text{P\&L} \;=\; \int_0^T \frac{S_t^2}{2} \frac{\partial^2 V_t}{\partial S^2} \;\left(\sigma_{imp}^2 - \sigma_t^2\right)\; dt$$

where V_t is the time t value of the option and σ_t is the realized instantaneous volatility at time t.

We recognize the term $\frac{S_t^2}{2} \frac{\partial^2 V_t}{\partial S^2}$ as the dollar gamma

- always positive for a vanilla call or put option.

E.G: Parameter Uncertainty and Hedging in Black-Scholes

Returning to s.f. trading strategy of (1) and (2), note that we can choose any model we like for the security price dynamics

- e.g. other diffusions or jump-diffusion models.

It is interesting to simulate these alternative models and to then observe what happens to the replication error from (1) and (2).

It is common to perform numerical experiments like this when using a model to price and hedge a particular security.

Goal then is to understand how robust the hedging strategy (based on the given model) is to alternative price dynamics that might prevail in practice.

Given the appropriate data, one can also back-test the performance of a model on realized historical price data to assess its hedging performance.

E.G: Calibration and Extrapolation in Short-Rate Models

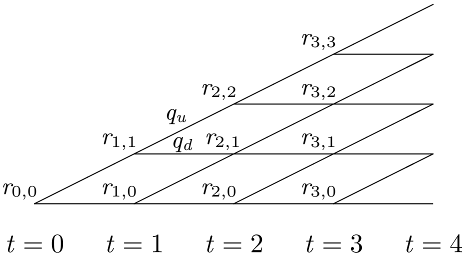

Binomial models for the short-rate were very popular for pricing fixed income derivatives in the 1990's and into the early 2000's.

Lattice above shows a generic binomial model for the short-rate, r_t which is a 1-period risk-free rate.

Risk-neutral probabilities of up- and down-moves in any period are given by q_u and $q_d = 1 - q_u$, respectively.

Securities then priced in the lattice using risk-neutral pricing in the usual backwards evaluation manner.

e.g. Can compute time t price, Z_t^T, of a zero-coupon bond (ZCB) maturing at time T by setting $Z_T^T \equiv 1$ and then calculating

$$Z_{t,j} = \mathsf{E}_t^{\mathbb{Q}}\left[\frac{B_t}{B_{t+1}}\,Z_{t+1}\right] = \frac{1}{1+r_{t,j}}\left[q_u \times Z_{t+1,j+1} + q_d \times Z_{t+1,j}\right]$$

for $t = T-1, \ldots, 0$ and $j = 0, \ldots, t$.

More generally, risk-neutral pricing for a "coupon" paying security takes the form

$$\begin{aligned} Z_{t,j} &= \mathsf{E}_t^{\mathbb{Q}}\left[\frac{Z_{t+1}+C_{t+1}}{1+r_{t,j}}\right] \\ &= \frac{1}{1+r_{t,j}}\left[q_u\left(Z_{t+1,j+1}+C_{t+1,j+1}\right) \;+\; q_d\left(Z_{t+1,j}+C_{t+1,j}\right)\right] \end{aligned} \tag{4}$$

where $C_{t,j}$ is the coupon paid at time t and state j, and $Z_{t,j}$ is the "ex-coupon" value of the security at time t and state j.

E.G: Calibration and Extrapolation in Short-Rate Models

Can iterate (4) to obtain

$$\frac{Z_t}{B_t} = \mathsf{E}_t^{\mathbb{Q}}\left[\sum_{j=t+1}^{t+s} \frac{C_j}{B_j} + \frac{Z_{t+s}}{B_{t+s}}\right] \tag{5}$$

Other securities including caps, floors, swaps and swaptions can all be priced using (4) and appropriate boundary conditions.

Moreover, when we price securities like in this manner the model is guaranteed to be arbitrage-free. **Why?**

The Black-Derman-Toy (BDT) Model

The Black-Derman-Toy (BDT) model assumes the interest rate at node $N_{i,j}$ is given by

$$r_{i,j} \;=\; a_i e^{b_i j}$$

where $\log(a_i)$ and b_i are drift and volatility parameters for $\log(r)$, respectively.

To use such a model in practice we must first **calibrate** it to the observed term-structure in the market and, ideally, other liquid fixed-income security prices

- can do this by choosing the a_i's and b_i's to match market prices.

Once the parameters have been calibrated we can now **consider** using the model to price less liquid or more exotic securities.

Using BDT to Price a $2 - 8$ Payer Swaption

Consider pricing a $2 - 8$ *payer swaption* with fixed rate $= 11.65\%$.

This is:

- An option to enter an 8-year swap in 2 years time.
- The underlying swap is settled in arrears so payments would take place in years 3 through 10.
- Each floating rate payment in the swap is based on the prevailing short-rate of the previous year.
- The “payer" feature of the option means that if the option is exercised, the exerciser “pays fixed and receives floating" in the underlying swap.

Using BDT to Price a $2-8$ Payer Swaption

We use a 10-period BDT lattice so that 1 period corresponds to 1 year.

Lattice was calibrated to the term structure of interest rates in the market

- there are therefore 20 free parameters, a_i and b_i for $i = 0, \ldots, 9$, to choose
- and want to match 10 spot interest rates, s_t for $t = 1, \ldots, 10$.

Therefore have 10 equations in 20 unknowns and so the calibration problem has too many parameters.

We can (and do) resolve this issue by simply setting $b_i = b = .005$ for all i

- which leaves 10 unknown parameters.

Will assume a notional principal for the underlying swaps of $\$1$m.

Let S_2 denote swap value at $t = 2$

- can compute S_2 by discounting the swap's cash-flows back from $t = 10$ to $t = 2$ using risk-neutral probabilities which we take to be $q_d = q_u = 0.5$.

Using BDT to Price a $2-8$ Payer Swaption

Option then exercised at time $t=2$ if and only if $S_2 > 0$

- so value of swaption at $t=2$ is $\max(0, S_2)$.

Time $t=0$ value of swaption can be computed using backwards evaluation

- after calibration we find a swaption price of $\$1,339$ (when $b = .005$).

It would be **naive**(!) in the extreme to assume that this is a fair price for the swaption

- after all, what was so special about the choice of $b = .005$?

Suppose instead we chose $b = .01$.

In that case, after recalibrating to the same term structure of interest rates, we find a swaption price of $\$1,962$

- a price increase of approximately 50%!

This price discrepancy should not be at all surprising! **Why?**

Using BDT to Price a $2-8$ Payer Swaption

Several important lessons regarding model risk and calibration here:

1. Model transparency is very important. In particular, it is very important to understand what type of dynamics are implied by a model and, more importantly in this case, what role each of the parameters plays.

2. Should also be clear that calibration is an intrinsic part of the pricing process and is most definitely not just an afterthought.

   Model selection and model calibration (including the choice of calibration instruments) should **not** be separated.

3. When calibration is recognized to be an intrinsic part of the pricing process, we begin to recognize that pricing is really no more than an exercise in interpolating / extrapolating from observed market prices to unobserved market prices.

   Since extrapolation is in general difficult, we would have much more confidence in our pricing if our calibration securities were "close" to the securities we want to price.

   This was not the case with the swaption!

E.G: Regime Shifts: LIBOR Pre- and Post-Crisis

LIBOR calculated on a daily basis:

- The *British Bankers' Association* (BBA) polls a pre-defined list of banks with strong credit ratings for their interest rates of various maturities.
- Highest and lowest responses are dropped
- Average of remainder is taken to be the LIBOR rate.

Understood there was some (very small) credit risk associated with these banks

- so LIBOR would therefore be higher than corresponding rates on government treasuries.

Since banks that are polled always had strong credit ratings (prior to the crisis) spread between LIBOR and treasury rates was generally quite small.

Moreover, the pre-defined list of banks regularly updated so that banks whose credit ratings have deteriorated are replaced with banks with superior credit ratings.

This had the practical impact of ensuring (or so market participants believed up until the crisis(!)) that forward LIBOR rates would only have a very modest degree of credit risk associated with them.

E.G: Regime Shifts: LIBOR Pre- and Post-Crisis

LIBOR extremely important since it's a benchmark interest rate and many of the most liquid fixed-income derivative securities are based upon it.

These securities include:

1. Floating rate notes (FRNs)
2. Forward-rate agreements (FRAs)
3. Swaps and (Bermudan) swaptions
4. Caps and floors.

Cash-flows associated with these securities are determined by LIBOR rates of different tenors.

Before 2008 financial crisis, however, these LIBOR rates were viewed as being essentially (default) risk-free

- led to many simplifications when it came to the pricing of the aforementioned securities.

Derivatives Pricing Pre-Crisis

e.g.1. Consider a floating rate note (FRN) with face value 100.

At each time $T_i > 0$ for $i = 1, \ldots, M$ the note pays interest equal to $100 \times \tau_i \times L(T_{i-1}, T_i)$ where $\tau_i := T_i - T_{i-1}$ and $L(T_{i-1}, T_i)$ is the LIBOR rate at time T_{i-1} for borrowing / lending until the maturity T_i.

Note expires at time $T = T_M$ and pays 100 (in addition to the interest) at that time.

A well known and important result is that the fair value of the FRN at any reset point just after the interest has been paid, is 100

- follows by a simple induction argument.

Derivatives Pricing Pre-Crisis

e.g.2. The forward LIBOR rate at time t based on simple interest for lending in the interval $[T_1, T_2]$ is given by

$$L(t, T_1, T_2) \ = \ \frac{1}{T_2 - T_1}\left(\frac{P(t, T_1)}{P(t, T_2)} - 1\right) \tag{6}$$

where $P(t, T)$ is the time t price of a deposit maturing at time T.

Spot LIBOR rates are obtained by setting $T_1 = t$.

LIBOR rates are quoted as simply-compounded interest rates, and are quoted on an annual basis.

The accrual period or tenor, $T_2 - T_1$, usually fixed at $\delta = 1/4$ or $\delta = 1/2$

- corresponding to 3 months and 6 months, respectively.

Derivatives Pricing Pre-Crisis

With a fixed value of δ in mind can define the δ-year forward rate at time t with maturity T as

$$L(t, T, T+\delta) \;=\; \frac{1}{\delta}\left(\frac{P(t,T)}{P(t,T+\delta)} - 1\right). \tag{7}$$

The δ-year spot LIBOR rate at time t then given by $L(t, t+\delta) := L(t, t, t+\delta)$.

Also note that $L(t, T, T+\delta)$ is the FRA rate at time t for the swap(let) maturing at time $T+\delta$.

That is, $L(t, T, T+\delta)$ is unique value of K for which the swaplet that pays $\pm(L(T, T+\delta) - K)$ at time $T+\delta$ is worth zero at time $t < T$.

Derivatives Pricing Pre-Crisis

e.g.3. We can **compound** two consecutive 3-month forward LIBOR rates to obtain corresponding 6-month forward LIBOR rate.

In particular, we have

$$\left(1+\frac{F_1^{3m}}{4}\right)\left(1+\frac{F_2^{3m}}{4}\right)=1+\frac{F^{6m}}{2} \tag{8}$$

where:

- $F_1^{3m} := L(0,3m,6m)$
- $F_1^{3m} := L(0,6m,9m)$
- $F^{6m} := L(0,3m,9m)$.

During the Crisis

All three pricing results broke down during the 2008 financial crisis.

Because these results were also required for pricing of swaps and swaptions, caps and floors, etc., the entire approach to the pricing of fixed income derivative securities needed broke down.

Cause of this breakdown was the loss of trust in the banking system and the loss of trust between banks.

This meant that LIBOR rates were no longer viewed as being risk-free.

Easiest way to demonstrate this loss of trust is via the spread between LIBOR and OIS rates.

Overnight Indexed Swaps (OIS)

An overnight indexed swap (OIS) is an interest-rate swap where the periodic floating payment is based on a return calculated from the daily compounding of an overnight rate (or index).

> e.g. Fed Funds rate in the U.S., EONIA in the Eurozone and SONIA in the U.K.

The fixed rate in the swap is the fixed rate the makes the swap worth zero at inception.

Note there is essentially **no** credit / default risk premium included in OIS rates

- due to fact that floating payments in the swap are based on **overnight** lending rates.

LIBOR - OIS Spread

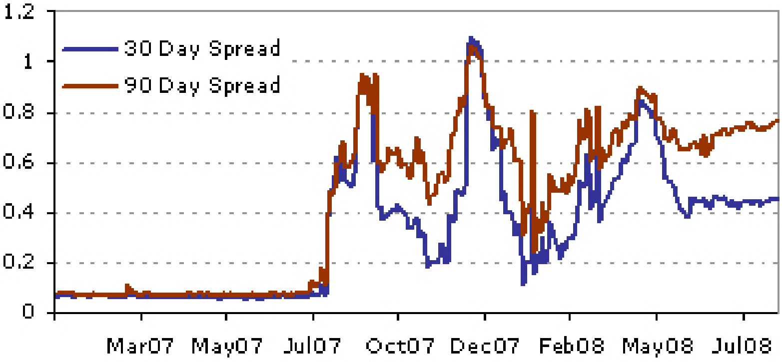

Source: Bloomberg, Financial Times

Overnight Indexed Swaps (OIS)

We see that the LIBOR-OIS spreads were essentially zero leading up to the financial crisis

- because market viewed LIBOR rates as being essentially risk-free with no associated credit risk.

This changed drastically during the crisis when entire banking system nearly collapsed and market participants realized there were substantial credit risks in the interbank lending market.

Since the crisis the spreads have not returned to zero and must now be accounted for in all fixed-income pricing models.

This regime switch constituted an extreme form of model risk where the entire market made an assumption that resulted in all pricing models being hopelessly inadequate!

Derivatives Pricing Post-Crisis

Since the financial crisis we no longer take LIBOR rates to be risk-free.

And risk-free curve is now computed from OIS rates

- so (for example) we no longer obtain the result that an FRN is always worth par at maturity.

OIS forward rates (but **not** LIBOR rates!) are calculated as in (6) or (7) so that

$$F_d(t, T_1, T_2) \;=\; \frac{1}{T_2 - T_1}\left(\frac{P_d(t, T_1)}{P_d(t, T_2)} - 1\right)$$

where P_d denotes the discount factor computed from the OIS curve and F_d denotes forward rates implied by these OIS discount factors.

Derivatives Pricing Post-Crisis

Forward LIBOR rates are now defined as risk-neutral expectations (under the forward measure) of the spot LIBOR rate, $L(T, T+\delta)$.

Relationships such as (8) no longer hold and we now live in a multi-curve world

- with a different LIBOR curve for different tenors
- e.g. we have a 3-month LIBOR curve, a 6-month LIBOR curve etc.

Question: Which if any of the 3-month and 6-month LIBOR curves will be lower? Why?

With these new definitions, straightforward extensions of the traditional pricing formulas (that held pre-crisis) can be used to price swaps, caps, floors, swaptions etc. in the post-crisis world.

Model Transparency

Will now consider some well known models in equity derivatives space

- these models (or slight variations of them) also be used in foreign exchange and commodity derivatives markets.

Can be useful when trading in exotic derivative securities and not just the “vanilla" securities for which prices are readily available.

If these models can be calibrated to observable vanilla security prices, they can then be used to construct a range of plausible prices for more exotic securities

- so they can help counter extrapolation risk
- and provide alternative estimates of the Greeks or hedge ratios.

Will not focus on stochastic calculus or various numerical pricing techniques here.

Instead simply want to emphasize that a broad class of tractable models exist and that they should be employed when necessary.

Model Transparency

Very important for the users of these models to fully understand their various strengths and weaknesses

- and the **implications** of these strengths and weaknesses when they are used to price and risk manage a given security.

Will see later how these models and others can be used *together* to infer prices of exotic securities as well as their Greeks or hedge ratios.

In particular will emphasize how they can be used to avoid the pitfalls associated with **price extrapolation**.

Any risk manager or investor in exotic securities should therefore maintain a library of such models that can be called upon as needed.

Recalling the Implied Volatility Surface

Recall the GBM model:

$$dS_t \;=\; \mu S_t \; dt + \sigma S_t \; dW_t.$$

When we use risk-neutral pricing we know that $\mu = r - q$.

Therefore have a single free parameter, σ, which we can fit to option prices or, equivalently, the volatility surface.

Not all surprising then that this exercise fails: the volatility surface is never flat so that a constant σ fails to re-produce market prices.

This became particularly apparent after the stock market crash of October 1987 when market participants began to correctly identify that lower strike options should be priced with a higher volatility, i.e. there should be a volatility skew.

The volatility surface of the Eurostoxx 50 Index on 28^{th} November 2007 is plotted in Figure 39.

The Implied Volatility Surface

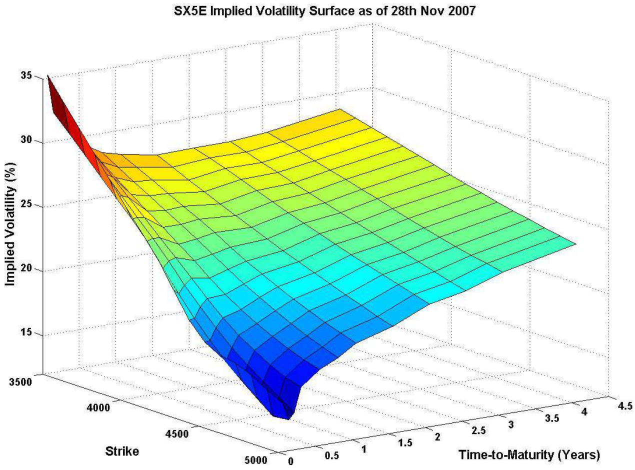

The Eurostoxx 50 Volatility Surface.

Local Volatility Models

The local volatility framework assumes risk-neutral dynamics satisfy

$$dS_t = (r-q)S_t\,dt + \sigma_l(t, S_t)S_t\,dW_t \tag{9}$$

– so $\sigma_l(t, S_t)$ is now a function of time and stock price.

Key result is the Dupire formula that links the local volatilities, $\sigma_l(t, S_t)$, to the implied volatility surface:

$$\sigma_l^2(T, K) = \frac{\frac{\partial C}{\partial T} + (r-q)K\,\frac{\partial C}{\partial K} + qC}{\frac{K^2}{2}\,\frac{\partial^2 C}{\partial K^2}} \tag{10}$$

where $C = C(K, T) =$ option price as a function of strike and time-to-maturity.

Calculating local volatilities from (10) is difficult and numerically unstable.

Local volatility is very nice and interesting

- a model that is guaranteed to replicate the implied volatility surface in the market.

But it also has weaknesses ...

Local Volatility Models

For example, it leads to unreasonable skew dynamics and underestimates the volatility of volatility or “vol-of-vol".

Moreover the Greeks that are calculated from a local volatility model are generally not consistent with what is observed empirically.

Understanding these weaknesses is essential from a risk management point of view.

Nevertheless, local volatility framework is theoretically interesting and is still often used in practice for pricing certain types of exotic options such as barrier and digital options.

They are known to be particularly **unsuitable** for pricing derivatives that depend on the forward skew such as forward-start options and cliquets.

Local Volatility

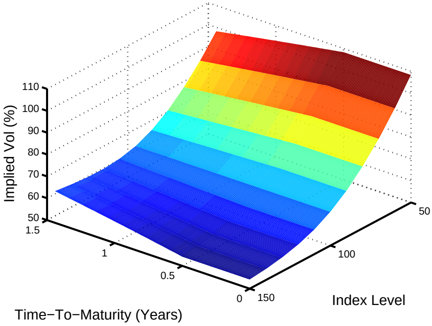

(a) Implied Volatility Surface

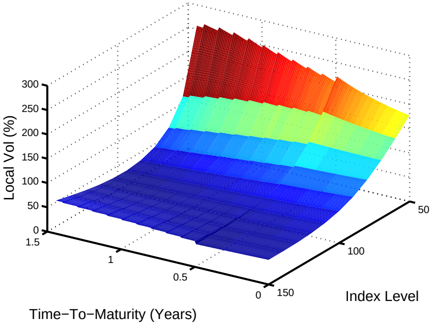

(b) Local Volatility Surface

Implied and local volatility surfaces: local volatility surface is constructed from implied volatility surface using Dupire's formula.

Stochastic Volatility Models

Most well-known stochastic volatility model is due to Heston (1989).

It's a two-factor model and assumes separate dynamics for both the stock price and instantaneous volatility so that

$$dS_t = (r-q)S_t\,dt + \sqrt{\sigma_t}S_t\,dW_t^{(s)} \tag{11}$$

$$d\sigma_t = \kappa\,(\theta - \sigma_t)\,dt + \gamma\sqrt{\sigma_t}\,dW_t^{(vol)} \tag{12}$$

where $W_t^{(s)}$ and $W_t^{(vol)}$ are standard $\mathcal{Q}$-Brownian motions with constant correlation coefficient, ρ.

Heston's stochastic volatility model is an **incomplete** model
- why?

The particular EMM that we choose to work with would be determined by some calibration algorithm
- the typical method of choosing an EMM in incomplete market models.

The volatility process in (12) is commonly used in interest rate modeling
- where it's known as the CIR model.

Stochastic Volatility Models

The price, $C(t, S_t, \sigma_t)$, of any derivative security under Heston must satisfy

$$\frac{\partial C}{\partial t} + \frac{1}{2}\sigma S^2 \frac{\partial^2 C}{\partial S^2} + \rho\sigma\gamma S \frac{\partial^2 C}{\partial S \partial \sigma} + \frac{1}{2}\gamma^2 \sigma \frac{\partial^2 C}{\partial \sigma^2} + (r-q) S \frac{\partial C}{\partial S} + \kappa\,(\theta - \sigma) \frac{\partial C}{\partial \sigma} = rC. \tag{13}$$

Price then obtained by solving (13) subject to the relevant boundary conditions.

Once parameters of the model have been identified via some calibration algorithm, pricing can be done by either solving (13) numerically or alternatively, using Monte-Carlo.

Note that some instruments can be priced analytically in Heston's model

- e.g. continuous-time version of a variance swap.

Heston generally captures long-dated skew quite well but struggles with short-dated skew, particularly when it is steep.

Stochastic Volatility Models

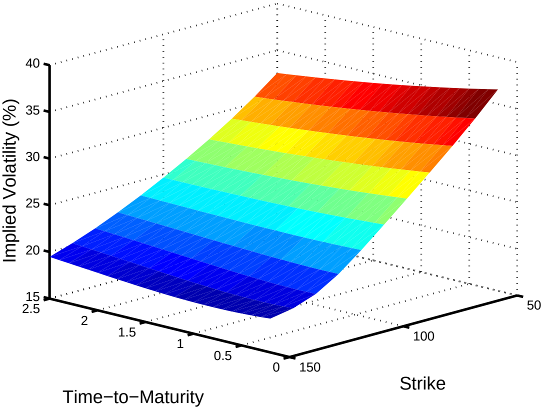

Figure displays an implied volatility surface under Heston's stochastic volatility model.

A Detour: Forward Skew

Forward skew is the implied volatility skew that prevails at some future date and that is consistent with some model that has been calibrated to **today's** market data.

e.g. Suppose we simulate some model that has been calibrated to today's volatility surface forward to some date $T > 0$.

On any simulated path can then compute the implied volatility surface as of that date T.

This is the date T forward implied volatility surface

- forward skew refers to general shape of the skew in this forward vol. surface
- note that this forward skew is **model dependent**.

Forward Skew and Local Volatility

Well known that forward skew in local volatility models tends to be very flat.

A significant feature of local volatility models that is not at all obvious until one explores and works with the model in some detail.

Not true of stochastic volatility or jump diffusion models

- where forward skew is generally similar in shape to the spot skew.

These observations have very important implications for pricing cliquet-style securities.

The failure of market participants to understand the properties of their models has often led to substantial trading losses

- such losses might have been avoided had they priced the securities in question with several different models.

Example: The Locally Capped, Globally Floored Cliquet

The locally capped, globally floored cliquet is structured like a bond:

- Investor pays the bond price upfront at $t = 0$.
- In return he is guaranteed to receive the principal at maturity as well as an annual coupon.
- The coupon is based on monthly returns of underlying security over previous year. It is calculated as

$$\text{Payoff} \;=\; \max\left\{\sum_{t=1}^{12} \min\left(\max(r_t,\ -.01),\ .01\right),\ \text{MinCoupon}\right\} \tag{14}$$

where MinCoupon $= .02$ and each monthly return satisfies

$$r_t \;=\; \frac{S_t - S_{t-1}}{S_{t-1}}.$$

- Annual coupon therefore capped at 12% and floored at 2%.

Ignoring global floor, coupon behaves like a strip of forward-starting monthly call-spreads.

Example: The Locally Capped, Globally Floored Cliquet

A call spread is a long position in a call option with strike k_1 and a short position in a call option with strike $k_2 > k_1$.

A call spread is sensitive to the implied volatility skew.

Would therefore expect coupon value to be very sensitive to the forward skew in the model.

In particular, would expect coupon value to **increase** as the skew becomes more negative.

So would expect a local volatility model to **underestimate** the price of this security.

In the words of Gatheral (2006):

"We would guess that the structure should be very sensitive to forward skew assumptions. Thus our prediction would be that a local volatility assumption would substantially underprice the deal because it generates forward skews that are too flat . . ."

Example: The Locally Capped, Globally Floored Cliquet

And indeed this intuition can be confirmed.

In fact numerical experiments suggest that a stochastic volatility model places a much higher value on this bond than a comparable local volatility model.

But our intuition is not always correct!

So important to evaluate exotic securities with which we are not familiar under different modeling assumptions.

And any mistakes are likely to be **costly**!

In the words of Gatheral (2006) again:

" . . . since the lowest price invariably gets the deal, it was precisely those traders that were using the wrong model that got the business . . .
The importance of trying out different modeling assumptions cannot be overemphasized. Intuition is always fallible!"

Jump-Diffusion Models

Merton's jump diffusion model assumes that time t stock price satisfies

$$S_t \;=\; S_0 \; e^{(\mu-\sigma^2/2)t+\sigma W_t} \prod_{i=1}^{N_t} Y_i \tag{15}$$

where $N_t \sim \text{Poisson}(\lambda t)$ and Y_i's $\sim$ log-normal and IID.

Each Y_i represents the magnitude of the i^{th} jump and stock price behaves like a regular GBM between jumps.

If the dynamics in (15) are under an EMM, $\mathcal{Q}$, then μ, λ and the mean jump size are constrained in such a way that $\mathcal{Q}$-expected rate of return must equal $r - q$.

Question: Can you see how using (15) to price European options might lead to an infinitely weighted sum of Black-Scholes options prices?

Other tractable jump-diffusion models are due to Duffie, Pan and Singleton (1998), Kou (2002), Bates (1996) etc.

Merton's Jump-Diffusion Model

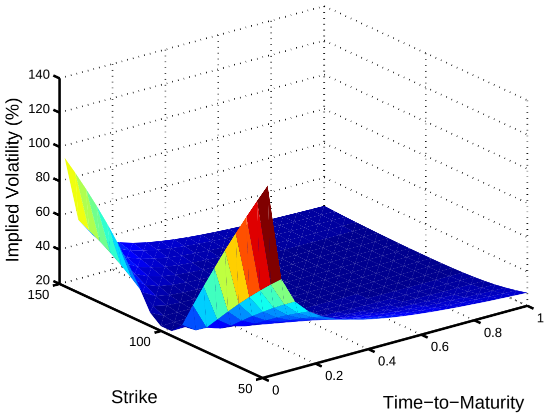

An implied volatility surface under Merton's jump diffusion model
- any observations?

Levy Processes

Definition: A Levy process is any continuous-time process with stationary and independent increments.

Definition: An exponential Levy process, S_t, satisfies $S_t = \exp(X_t)$ where X_t is a Levy process.

- The two most common examples of Levy processes are Brownian motion and the Poisson process.
- A Levy process with jumps can be of infinite activity so that it jumps infinitely often in any finite time interval, or of finite activity so that it makes only finitely many jumps in any finite time interval.
- The Merton jump-diffusion model is an example of an exponential Levy process of finite activity.
- The most important result for Levy processes is the Levy-Khintchine formula which describes the characteristic function of a Levy process.

Time-Changed Exponential Levy Processes

Levy processes cannot capture "volatility clustering", the tendency of high volatility periods to be followed by periods of high volatility, and low volatility periods to be followed by periods of low volatility.

This is due to the stationary and independent increments assumption.

Levy models with stochastic time, however, can capture volatility clustering.

Definition: A Levy subordinator is a non-negative, non-decreasing Levy process.

A subordinator can be used to change the "clock speed" or "calendar speed".

More generally, if y_t is a positive process, then we can define our stochastic clock, Y_t, as

$$Y_t := \int_0^t y_s \, ds. \tag{16}$$

Using Y_t to measure time instead of the usual t, we can then model a security price process, S_t, as an exponential time-changed Levy process.

Time-Changed Exponential Levy Processes

Can write the $\mathcal{Q}$-dynamics for S_t as

$$S_t = S_0 e^{(r-q)t} \frac{e^{X_{Y_t}}}{\mathbf{E}_0^{\mathcal{Q}}\left[e^{X_{Y_t}} | y_0\right]}. \tag{17}$$

Volatility clustering captured using the stochastic clock, Y_t.

e.g. Between t and $t + \Delta t$ time-changed process will have moved approximately $y_t \times \Delta t$ units.

So when y_t large the clock will have moved further and increments of S_t will generally be more volatile as a result.

Note that if the subordinator, Y_t, is a jump process, then S_t can jump even if the process X_t cannot jump itself.

Note that in (17) the dynamics of S_t are $\mathcal{Q}$-dynamics corresponding to the cash account as numeraire. Why is this the case?

- typical of how incomplete markets are often modeled: we directly specify $\mathcal{Q}$-dynamics so that martingale pricing holds by construction
- other free parameters are then chosen by a calibration algorithm.

Normal-Inverse-Gaussian Process with CIR Clock

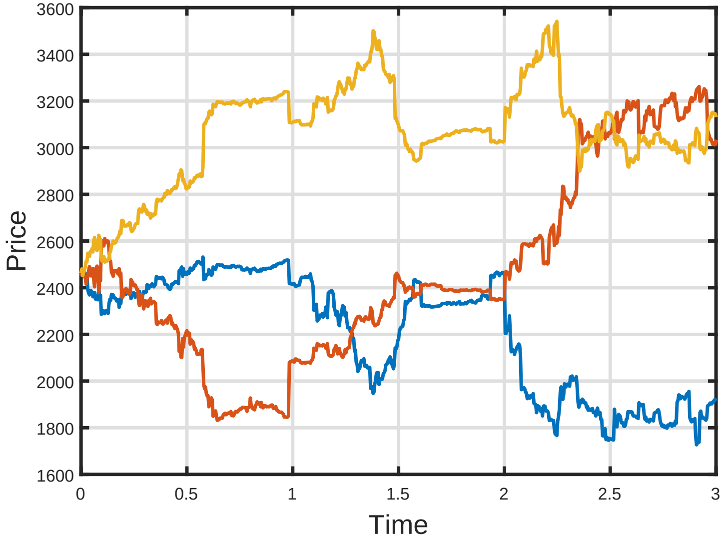

Variance-Gamma Process with OU-Gamma Clock

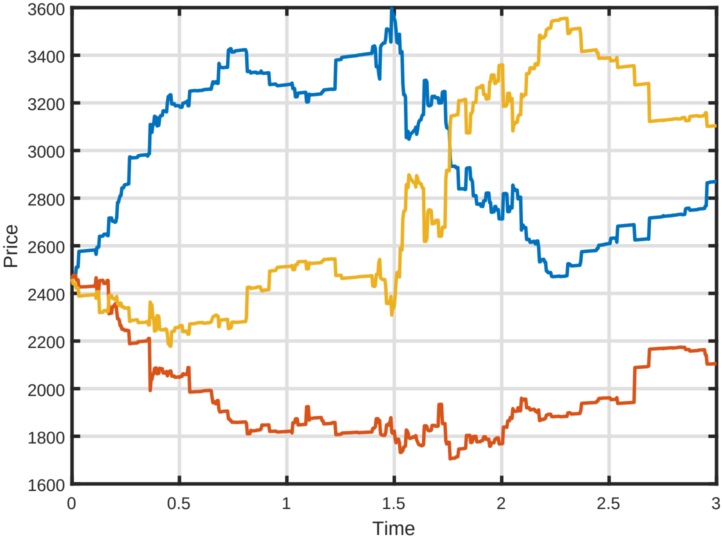

Time-Changed Exponential Levy Processes

Many Levy processes are quite straightforward to simulate.

However, the characteristic functions of the $\log$-stock price are often available in closed form, even when the Levy process has a stochastic clock.

In fact it can be shown that

$$\begin{aligned} \phi(u,t) &:= \mathsf{E}\left[e^{iu\ \log(S_t)} \mid S_0, y_0\right] \\ &= e^{iu((r-q)t+\log(S_0))}\ \frac{\varphi(-i\psi_x(u);t,y_0)}{\varphi(-i\psi_x(-i);t,y_0)^{iu}} \end{aligned} \tag{18}$$

where

$$\psi_x(u) := \log \mathsf{E}\left[e^{iuX_1}\right] \tag{19}$$

is the characteristic exponent of the Levy process and $\varphi(u;t,y_0)$ is the characteristic function of Y_t given y_0.

Therefore if we know the characteristic function of the integrated subordinator, Y_t, and the characteristic function of the Levy process, X_t, then we also know the characteristic function of the $\log$-stock price

- and (vanilla) options can be priced using numerical transform methods.

Calibration is an Integral part of the Pricing Process

When we used the BDT model to price swaptions,we saw the calibration process in an integral component of the pricing process.

Indeed, in early 2000's there was debate about precisely this issue.

Context for the debate was general approach in the market to use simple **one-factor models** to price Bermudan swaptions.

Longstaff, Santa-Clara and Schwartz (2001) argued that since the term-structure of interest rates was driven by several factors, using a one-factor model to price Bermudan swaptions was fundamentally flawed.

This argument was sound ... but they neglected to account for the **calibration process**.

In response Andersen and Andreasen (2001) argued the Bermudan swaption prices were actually quite accurate even though they were indeed typically priced at the time using simple one-factor models.

Calibration is an Integral part of the Pricing Process

Their analysis relied on the fact that when pricing Bermudan swaptions it was common to calibrate the one-factor models to the prices of **European swaptions**.

On the basis that Bermudan swaptions are actually quite "close" to European swaptions, they argued the extrapolation risk was small.

This debate clearly tackled the issue of **model transparency** and highlighted that model dynamics may not be important at all if the exotic security being priced is within or close to the span of the securities used to calibrate the model.

Essentially two wrongs, i.e. a bad model and bad parameters, can together make a right, i.e. an accurate price!

Using Many Models to Manage Model Risk

The models we have described are quite representative of the models that are used for pricing many equity and FX derivative securities in practice.

Therefore very important for users of these models to fully understand their strengths and weaknesses and the **implications** of these strengths and weaknesses when used to price and risk manage a given security.

Will see how these models and others can be used **together** to infer the prices of exotic securities as well as their Greeks or hedge ratios.

In particular we will emphasize how they can be used to avoid the pitfalls associated with price extrapolation.

Any risk manager or investor in exotic securities would therefore be well-advised to maintain a library of such models that can be called upon as needed.

Another Example: Back to Barrier Options

Well known that the price of a barrier option is **not** solely determined by the marginal distributions of the underlying stock price.

Figures on next 2 slides emphasizes this point:

- They display prices of down-and-out (DOB) and up-and-out (UOB) barrier call options as a function of the barrier level for several different models.
- In each case strike was set equal to the initial value of underlying security.
- Parameter values for all models were calibrated to implied volatility surface of the Eurostoxx 50 Index on October 7^{th}, 2003.
- So their marginal distributions coincide, at least to the extent that calibration errors are small.

Down-and-Out Barrier Call Prices Under Different Models

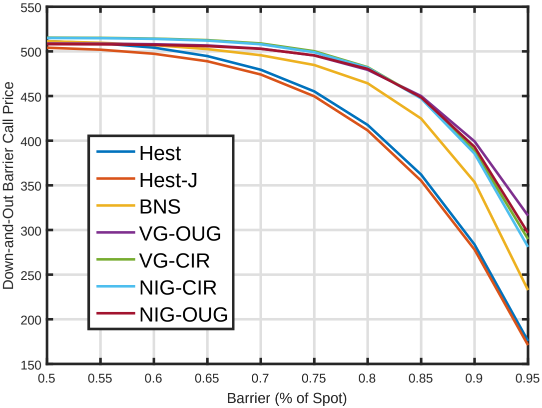

Down-and-out (DOB) barrier call option prices for different models, all of which have been calibrated to the same implied volatility surface.

See "*A Perfect Calibration! Now What?*" by Schoutens, Simons and Tistaert (2003) for further details.

Up-and-Out Barrier Call Prices Under Different Models

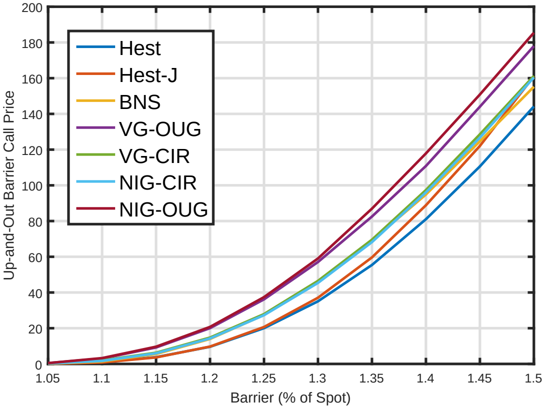

Up-and-out (UOB) barrier call option prices for different models, all of which have been calibrated to the same implied volatility surface.

See "*A Perfect Calibration! Now What?*" by Schoutens, Simons and Tistaert (2003) for further details.

Barrier Options and Extrapolation Risk

Clear that the different models result in very different barrier prices.

Question therefore arises as to what model one should use in practice?

- a difficult question to answer!

Perhaps best solution is to price the barrier using several different models that:

(a) have been calibrated to the market prices of liquid securities and

(b) have reasonable dynamics that make sense from a modeling viewpoint.

The minimum and maximum of these prices could then be taken as the **bid-offer** prices if they are not too far apart.

If they are far apart, then they simply provide guidance on where the fair price might be.

Using many plausible models is perhaps the best way to avoid extrapolation risk.

Model Calibration Risk

Have to be very careful when calibrating models to market prices!

In general (and this is certainly the case with equity derivatives markets) the most liquid instruments are vanilla call and put options.

Assuming then that we have a volatility surface available to us, we can at the very least calibrate our model to this surface.

An obvious but potentially hazardous approach would be to solve

$$\min_{\gamma} \sum_{i=1}^{N} \omega_i \ \left(\mathsf{ModelPrice}_i(\gamma) - \mathsf{MarketPrice}_i\right)^2 \tag{20}$$

where $\mathsf{ModelPrice}_i$ and $\mathsf{MarketPrice}_i$ are model and market prices, respectively, of the i^{th} option used in the calibration.

The ω_i's are fixed weights that we choose to reflect either the importance or accuracy of the i^{th} observation and γ is the vector of model parameters.

There are many problems with performing a calibration in this manner!

Problems with Calibrating Using (20)

1. In general (20) is a non-linear and non-convex optimization problem. It may therefore be difficult to solve as there may be many local minima.

2. Even if there is only one local minimum, there may be "valleys" containing the local minimum in which the objective function is more or less flat.

   Then possible that the optimization routine will terminate at different points in the valley even when given the same input, i.e. market option prices
   - this is clearly unsatisfactory!

   Consider following scenario:

   Day 1: Model is calibrated to option prices and resulting calibrated model is then used to price some path-dependent security. Let P_1 be the price of this path-dependent security.

   Day 2: Market environment is unchanged. The calibration routine is rerun, the path-dependent option is priced again and this time its price is P_2. But now P_2 is very different from P_1, despite the fact that the market hasn't changed. What has happened?

Problems with Calibrating Using (20)

Recall that the implied volatility surface only determines the marginal distributions of the stock prices at different times, T.

It tells you nothing(!) about the joint distributions of the stock price at different times.

Therefore, when we are calibrating to the volatility surface, we are only determining the model dynamics up-to the marginal distributions.

All of the parameter combinations in the “valley" might result in the very similar marginal distributions, but they will often result in very different joint distributions.

This was demonstrated when we saw how different models, that had been calibrated to the same volatility surface, gave very different prices for down-and-out call prices.

Also saw similar results when we priced interest-rate swaptions.