IEOR E4602: Quantitative Risk Management Spring 2016

Model Risk

We discuss model risk in these notes, mainly by way of example. We emphasize (i) the importance of understanding the physical dynamics and properties of a model and (ii) the dangers of price *extrapolation* whereby a model that has been calibrated to the market prices of a certain class of securities is then used to price a very different class of securities. Our examples will draw on equity, credit and fixed income applications as well as some structured products within these markets. That said, our conclusions will apply to all markets and models. We also discuss the very important role of model calibration and recognize that it is an integral part of the pricing process.

1 An Introduction to Model Risk

Our introduction to model risk will proceed by discussing several examples which emphasize different aspects of model risk. Other related examples will be discussed in later sections. We begin with the pricing of barrier options using the Black-Scholes / geometric Brownian motion (GBM) model.

Example 1 (Pricing Barrier Options Assuming GBM Dynamics)
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 (i) a put or a call (ii) a knock-in or knock-out and (iii) digital[1] or vanilla resulting in eight different payoff combinations. Of course other combinations are possible if you also allow for early exercise features as well as multiple barriers. For example, a knockout put option with strike K, barrier B and maturity T has a payoff given by

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

Similarly a digital down-and-in call option with strike K, barrier B and maturity T has a payoff given by

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

Knock-in options can be priced from knock-out options and vice versa since a knock-in plus a knock-out 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 GBM dynamics. We can use these analytic solutions to develop some intuition regarding barrier options and for this we will concentrate on knock-out put options which are traded quite frequently in practice. Figure 1(a) shows the Black-Scholes price of a knock-out put option as a function of the volatility parameter, σ. The barrier is 85, the strike is 100, current stock price is 105, time-to-maturity is six months and $r = q = 0$ where r and q are the risk-free rate and dividend yield, respectively. The price of the same knock-out option is plotted (on a different scale) alongside the price of a vanilla put option in Figure 1(b).

It is clear that the knock-out put option is always cheaper than the corresponding vanilla put option. For low values of σ, however, the prices almost coincide. This is to be expected as decreasing σ decreases the chances of the knock-out option actually being knocked out. While the vanilla option is unambiguously increasing in σ the same is not true for the knock-out option. Beyond a certain level, which depends on the strike and other parameters, the value of the knock-out put option decreases as σ increases.

[1] A digital call option pays \$1 if the underlying security expires above the strike and 0 otherwise. A digital put pays \$1 if the security expire below the strike. By "vanilla" option we mean the same payoff as a European call or put option.

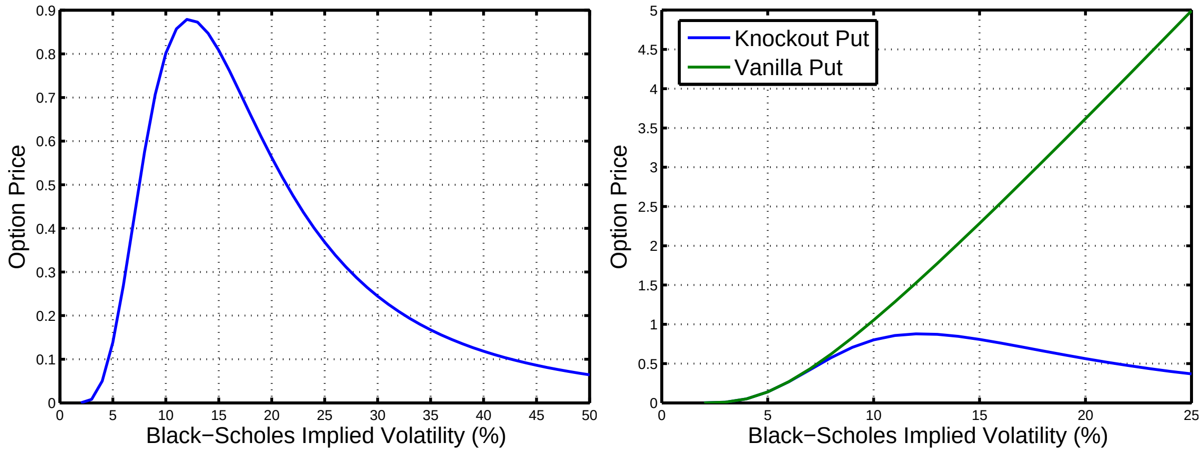

(a) Price of Knockout Put (b) Knockout Put v. Vanilla Put

Figure 1: Knockout Put Options in Black-Scholes: $S_0 = 105$, $K = 100$, Barrier $= 85$, $T = 6$ months, $r = q = 0$.

Exercise 1 *What do you think would happen to the value of the knock-out put option as $\sigma \to \infty$?*

There are many reasons why the Black-Scholes-GBM model is **so inadequate** when it comes to pricing barrier options and other exotic securities. These reasons include

1. There is only one free parameter, σ, in the GBM model and of course the model's implied volatility surface must always be constant as a function of time-to-maturity and strike. Clearly then the model cannot replicate the shape of the market implied volatility surface. In particular, this means that the Black-Schole model cannot be calibrated to the market prices of even simple securities like vanilla calls and puts.

2. Even if we were willing to overlook this calibration problem, what value of σ should be used when pricing a barrier option: $\sigma(T, K)$, $\sigma(T, B)$, some function of the two or some other value entirely? And if some "rule" is used for determining what value of σ to use it is very possible that the market would learn this and "game" you.

3. Suppose now there was no market for vanilla European options (so that the calibration problem discussed above did not exist) but that there was a market for the barrier option in question so that its price could be observed. Then how would you compute the implied volatility for this option using the Black-Scholes model. Figure 1(a) suggests that you would have either two solutions, $\sigma_{1,b} < \sigma_{2,b}$, or no solution.

   (a) If there are two solutions, which would you choose?

   (b) If there is no solution, what would that tell you about your model? In particular, would you interpret this situation as representing an arbitrage opportunity?

4. Assuming the same setup as the previous point, how would you compute the Greeks for your barrier option? If we took $\sigma_{1,b}$ as the implied volatility then we would see a positive vega but if we took $\sigma_{2,b}$ then our vega would be negative! Similarly, we might see a delta that has the opposite sign to the true delta. Hedging using these Greeks might therefore only serve to *increase* the riskiness of the position.

Of course all of the above problems arise because the GBM model is a very poor model for security price dynamics. Indeed, if security prices really did behave as a GBM then the implied volatility surface would be flat and we could easily identify the correct value of σ which we could then use to compute the correct Greeks etc. However, we know GBM is not a plausible model. For example, it is well known that when a security falls in value then this fall is usually accompanied by an increase in volatility. Similarly, a rise in security prices is often accompanied by a decrease in volatility. The inability of the GBM model to capture this behavior is reflected in the difficulties we encounter when trying to use GBM to price barrier options.

Note that depending on the security we are trying to price, these kinds of problems can occur with any model in any market. It seems to be the case that no matter how a good a model is, there will always be some security class for which it is particularly unsuited. For this reason, it is important to have available a range of plausible models that can be easily calibrated and then used to price other more exotic securities. The properties of these models should also be well understood by their users so that they are not used in a black-box fashion. These are the central messages of Section 2. ∎

Example 2 (Parameter Uncertainty and Hedging)
Our second example again concerns the use of the Black-Scholes model but this time with a view to hedging a vanilla European call option in the model. Moreover, we will assume that the assumptions of Black-Scholes are correct so that the security price has GBM dynamics, it is possible to trade continuously at no cost and borrowing and lending at the risk-free rate are also possible. It is then possible to dynamically *replicate* the payoff of the call option using a self-financing (s.f.) trading strategy. The initial value of this s.f. strategy is the famous Black-Scholes arbitrage-free price of the option. The s.f. replication strategy requires the continuous *delta-hedging* of the option but of course it is not practical to do this and so instead we hedge periodically. (Periodic or discrete hedging then results in some *replication error* but this error goes to 0 as the interval between rebalancing goes to 0.)

Towards this end, let P_t denote the time t value of the discrete-time self-financing strategy that attempts to replicate the option payoff and let C_0 denote the initial value of the option. The replicating strategy is then given by

$$P_0 \;:=\; C_0 \tag{1}$$
$$P_{t_{i+1}} \;=\; P_{t_i} \;+\; \left(P_{t_i} - \delta_{t_i} S_{t_i}\right) r\Delta t \;+\; \delta_{t_i}\left(S_{t_{i+1}} - S_{t_i} + q S_{t_i}\Delta t\right) \tag{2}$$

where $\Delta t := t_{i+1} - t_i$ is the length of time between re-balancing (assumed constant for all i), r is the annual risk-free interest rate (assuming per-period compounding), q is the dividend yield and δ_{t_i} is the Black-Scholes delta at time t_i. This delta is a function of S_{t_i} and some assumed implied volatility, σ_{imp} say. Note that (1) and (2) respect the self-financing condition. 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 will be 0 along every price path in the limit as $\Delta t \to 0$.

In practice, however, we do not know σ and so the market (and hence the option hedger) has no way to ensure a value of σ_{imp} such that $\sigma = \sigma_{imp}$. This has interesting implications for the trading P&L and it means in particular that we cannot exactly replicate the option even if all of the assumptions of Black-Scholes are correct. In Figure 2 we display histograms of the P&L in (3) that results from simulating 100'000 sample paths of the underlying price process with $S_0 = K = \$100$. (Other parameters and details are given below the figure.) In the case of the first histogram the true volatility was $\sigma = 30\%$ with $\sigma_{imp} = 20\%$ and the option hedger makes (why?) substantial loses. In the case of the second histogram the true volatility was $\sigma = 30\%$ with $\sigma_{imp} = 40\%$ and the option hedger makes (why?) substantial gains.

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[2] to highlight the

[2] We do acknowledge that this example is somewhat contrived in that if the true price dynamics really were GBM dynamics then we could estimate σ_{imp} perfectly and therefore exactly replicate the option payoff (in the limit of continuous trading). That said, it can also be argued that this example is not at all contrived: options traders in practice know the Black-Scholes model is incorrect but still use the model to hedge and face the question of what is the appropriate value of σ_{imp} to use. If they use a value that doesn't match (in a general sense) some true "average" level of volatility then they will experience P&L profiles of the form displayed in Figure 2.

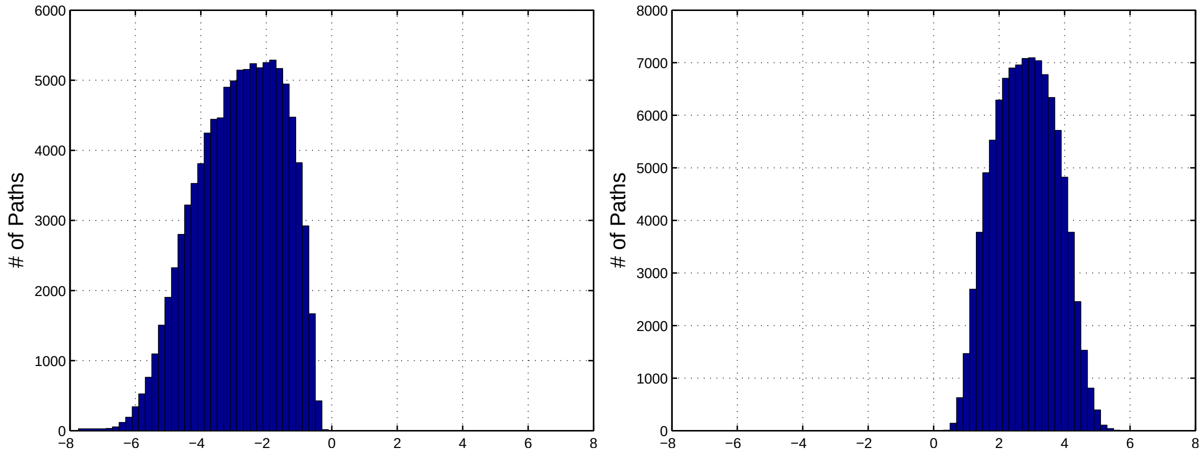

(a) Delta-hedging P&L: true vol. = 30%, imp. vol. = 20%. (b) Delta-hedging P&L: true vol. = 30%, imp. vol. = 40%.

Figure 2: Histogram of P&L from simulating 100k paths where we hedge a short call position with $S_0 = K = \$100$, $T = 6$ months, true volatility $\sigma = 30\%$, and $r = q = 1\%$. A time step of $dt = 1/2,000$ was used so hedging P&L due to discretization error is negligible. The hedge ratio, i.e. delta, was calculated using the implied volatility that was used to calculate the initial option price.

importance of not just having a good model but also having the correct model parameters. In general, however, it can be very difficult to estimate model parameters and we will discuss this further in Section 3.

Note that the payoff from delta-hedging an option is in general *path-dependent*, i.e. it depends on the price path taken by the stock over the entire time interval. In fact, it can be shown that the payoff from continuously delta-hedging an option satisfies

$$\text{P\&L} \;=\; \int_0^T \frac{S_t^2}{2}\frac{\partial^2 V_t}{\partial S^2}\;\left(\sigma_{imp}^2 - \sigma_t^2\right)\;dt \tag{4}$$

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*. It is always positive for a call or put option, but it goes to zero as the option moves significantly into or out of the money.

Returning to self-financing trading strategy of (1) and (2), note that we can choose any model we like for the security price dynamics. In particular, we are not restricted to choosing GBM and other diffusion or jump-diffusion models could be used instead. It is interesting to simulate these alternative models and to then observe what happens to the replication error in (4) where the δ_{t_i}'s are computed assuming (incorrectly) GBM price dynamics. Note that it is common to perform numerical experiments like this when using a model to price and hedge a particular security. The 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. ∎

Example 3 (Calibration and Extrapolation)
Binomial models for the short-rate were particularly popular[3] for pricing fixed income derivatives in the 1990's and into the early 2000's. The lattice below shows a generic binomial model for the short-rate, r_t which is a 1-period risk-free rate. The risk-neutral probabilities of up- and down-moves in any period are given by q_u and $q_d = 1 - q_u$, respectively. Securities are then priced in this lattice using risk-neutral pricing in the usual backwards evaluation manner.

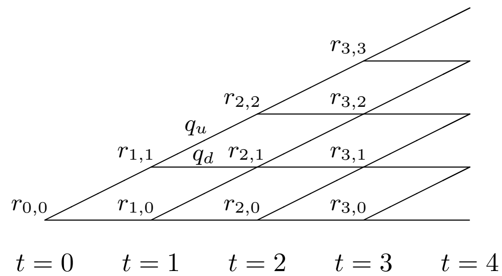

For example, if we wish to find the time t price, Z_t^T, of a zero-coupon bond (ZCB) maturing at time T then we can do so by setting $Z_T^T \equiv 1$ and then calculating

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

for $t = T-1, \ldots, 0$ and $j = 0, \ldots, t$ and where $B_t = \prod_{i=0}^{t-1}(1 + r_i)$ is the the time t value of the cash-account. More generally, risk-neutral pricing for a "coupon" paying security takes the form

$$\begin{aligned} Z_{t,j} &= \mathrm{E}_t^{\mathbb{Q}}\left[\frac{Z_{t+1} + C_{t+1}}{1 + r_{t,j}}\right] \\ &= \frac{1}{1 + r_{t,j}}\left[q_u\left(Z_{t+1,j+1} + C_{t+1,j+1}\right) \;+\; q_d\left(Z_{t+1,j} + C_{t+1,j}\right)\right] \end{aligned} \tag{5}$$

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 then. We can iterate (5) to obtain

$$\frac{Z_t}{B_t} = \mathrm{E}_t^{\mathbb{Q}}\left[\sum_{j=t+1}^{t+s}\frac{C_j}{B_j} + \frac{Z_{t+s}}{B_{t+s}}\right] \tag{6}$$

Other securities including caps, floors, swaps and swaptions can all be priced using (5) and appropriate boundary conditions. Moreover, when we price securities like this the model is guaranteed (why?) to be arbitrage-free.

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

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

where $\log(a_i)$ is a drift parameter for $\log(r)$ and b_i is a volatility parameter for $\log(r)$. If we want to use such a model in practice then we need to calibrate the model to the observed term-structure in the market and, ideally, other liquid fixed-income security prices. We can do this by choosing the a_i's and b_i's to match market prices[4]. Once the parameters have been calibrated we can now *consider* using the model to price less liquid or more exotic securities.

Consider now, for example, pricing a $2-8$ *payer swaption* with fixed rate $= 11.65\%$. This is:

[3]Binomial models are still used today in many applications. Moreover short-rate models are still used extensively in fixed income models, particularly in risk management applications where closed-form solutions are often required due to the need for (re-)valuing portfolios in large numbers of scenarios. For the same reason they are often used in CVA, i.e. *credit valuation adjustment*, calculations.

[4]For relatively few periods, we can do this easily in `Excel` using the `Solver` Add-In. For more realistic applications with many time periods we would use more efficient algorithms and software.

- 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. (The owner of a receiver swaption would "receive fixed and pay floating".)

In order to simplify matters we use a 10-period BDT lattice so that 1 period corresponds to 1 year. We will calibrate the lattice to the term structure of interest rates. There are therefore 20 parameters, a_i and b_i for $i = 0, \ldots, 9$, to choose in order to match 10 spot interest rates, s_t for $t = 1, \ldots, 10$, in the market place. We 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 now leaves 10 unknown parameters that we can use to match 10 spot interest rates.

We will assume a notional principal for the underlying swaps of $\$1$m and let S_2 denote the value of swap at $t = 2$. We can compute S_2 by discounting[5] the swap's cash-flows back from $t = 10$ to $t = 2$ using (of course) the risk-neutral probabilities which we take to be $q_d = q_u = 0.5$ here. The option is then exercised at time $t = 2$ if and only if $S_2 > 0$ so the value of the swaption at that time is $\max(0, S_2)$. The time $t = 0$ value of the swaption can be computed using backwards evaluation in the lattice.

After calibrating to the term structure of interest rates in the market 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 $\$1,962$, which is a price increase of approximately 50%! This price discrepancy should not be at all surprising: swaption prices clearly depend on the model's volatility and in the BDT model volatility is controlled (why?) by the b_i parameters. Moreover, ZCB prices, i.e. the term structure of interest rates, do not depend at all on volatility.

There are several important lessons regarding model risk and calibration that we can learn from this simple example:

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. If we had recognized in advance that the b_i's are volatility parameters then it would have been obvious that the price of a swaption would be very sensitive to their values and that just blindly setting them to a value of $.005$ was extremely inappropriate. Model transparency is the subject of Section 2.
2. It 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. We will discuss calibration in further detail in Section 3.
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 above. We will see other similar examples later in these notes.

∎

[5] It's easier to record time t cash flows at their predecessor nodes, discounting appropriately.

Example 4 (Regime Shift in Market Dynamics: LIBOR Pre- and Post-Crisis)
LIBOR is calculated on a daily basis when the *British Bankers' Association* (BBA) polls a pre-defined list of banks with strong credit ratings for their interest rates of various maturities. The highest and lowest responses are dropped and then the average of the remainder is taken to be the LIBOR rate. It was always understood there was some (very small) credit risk associated with these banks and that LIBOR would therefore be higher than the corresponding rates on government treasuries. However, because the banks that are polled always had strong credit ratings (prior to the crisis) the spread between LIBOR and treasury rates was generally quite small. Moreover, the pre-defined list of banks is regularly updated so that banks whose credit ratings have deteriorated are replaced on the list 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.

LIBOR is extremely important because it is a benchmark interest rate so many of the most liquid fixed-income derivative securities, e.g. floating rate notes, forward-rate agreements (FRAs), swaps, swaptions, caps and floors, are based upon it. In particular, the cash-flows associated with these securities are determined by LIBOR rates of different tenors. Before the 2008 financial crisis, however, these LIBOR rates were viewed as being essentially (default) risk-free and this led to many simplifications when it came to the pricing of the aforementioned securities. Note that LIBOR rates are quoted as **simply-compounded** interest rates, and are quoted on an annual basis. The accrual period or *tenor*, $T_2 - T_1$, is usually fixed at $\delta = 1/4$ or $\delta = 1/2$ corresponding to 3 months and 6 months, respectively.

Derivatives Pricing Pre-Crisis

We consider three simple examples related to spot and forward LIBOR that were all assumed to hold 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. The note expires at time $T = T_M$ and pays 100 (in addition to the interest) at that time.

A well known and important result was that the fair value of the FRN at any reset point, i.e. at any T_i, just after the interest has been paid, is 100. This follows by a simple induction argument. The cash-flow at time T_M is $100 \times (1 + L(T_{M-1}, T_M))$ and so its value, V_{M-1}, at time T_{M-1} (after the interest has been paid at T_{M-1}) is

$$\begin{aligned} V_{M-1} &= \frac{1}{1 + L(T_{M-1}, T_M)} \, 100 \times (1 + L(T_{M-1}, T_M)) \\ &= 100. \end{aligned} \tag{7}$$

The same argument can be applied to all earlier times to obtain an initial value of $V_0 = 100$ for the FRN. Note the key implicit assumption we made in (7) was that we could also use the appropriate LIBOR rate to discount the time T_M cash-flow $1 + L(T_{M-1}, T_M$.

E.G. 2. The **forward LIBOR rate** at time t based on simple interest for lending in the interval $[T_1, T_2]$ was given by

$$L(t, T_1, T_2) = \frac{1}{T_2 - T_1} \left(\frac{P(t, T_1)}{P(t, T_2)} - 1 \right) \tag{8}$$

where $P(t, T)$ is the time t price of a deposit maturing at time T. Note also that if we measure time in years, then (8) is consistent with $L(t, T_1, T_2)$ being quoted as an annual rate. Spot LIBOR rates were obtained by setting $T_1 = t$. With a fixed tenor, δ, in mind we could define the δ-year forward LIBOR rate at time t with maturity T as

$$L(t, T, T + \delta) = \frac{1}{\delta} \left(\frac{P(t, T)}{P(t, T + \delta)} - 1 \right). \tag{9}$$

Note that the δ-year spot LIBOR rate at time t is then given by $L(t, t + \delta) := L(t, t, t + \delta)$. We also note that $L(t, T, T + \delta)$ is the time t FRA rate for the swap(let) maturing at time $T + \delta$. That is, $L(t, T, T + \delta)$ is the 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 $t < T$.

E.G. 3. We can **compound** two consecutive 3-month forward LIBOR rates to obtain the corresponding 6-month forward LIBOR rate. In particular, we have

$$\left(1 + \frac{F_1^{3m}}{4}\right)\left(1 + \frac{F_2^{3m}}{4}\right) = 1 + \frac{F^{6m}}{2} \tag{10}$$

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

During the Crisis

All three of these pricing results broke down during the financial crisis and because these results were also required for the pricing of swaps and swaptions, caps and floors, etc., the entire approach to the pricing of fixed income derivative securities needed to be updated. The cause of this breakdown was the loss of trust in the banking system and the loss of trust between banks so that LIBOR rates were no longer viewed as being risk-free. The easiest way to demonstrate this loss of trust is via the spread between LIBOR and OIS rates.

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) such as the Fed Funds rate in the U.S., and EONIA and SONIA in the Eurozone and U.K., respectively. The fixed rate in the swap is the fixed rate the makes the swap worth zero at inception. We note there is essentially no credit / default risk premium included in OIS rates due to the fact that the floating payments in the swap are based on *overnight* lending rates. In Figure 3 we see a plot of the spread (in percentage points) between LIBOR and OIS rates for 1 month and 3-month maturities leading up to and during the financial crisis.

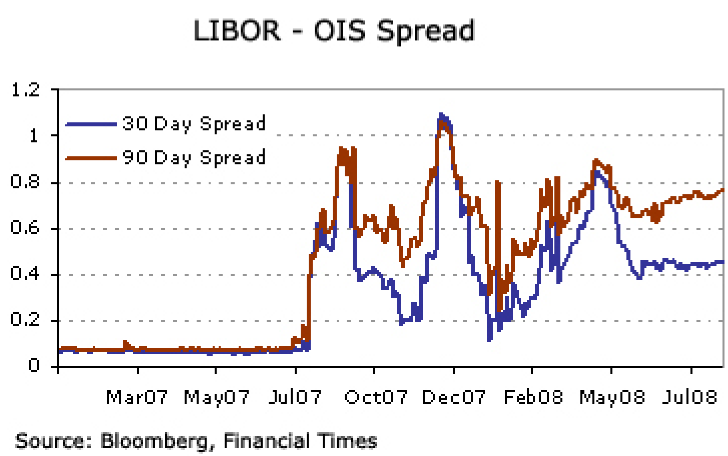

Figure 3: LIBOR-OIS spreads before and during the 2008 financial crisis.

We see that the LIBOR-OIS spreads were essentially zero leading up to the financial crisis. This is because (as mentioned earlier) the market viewed LIBOR rates as being essentially risk-free with no associated credit risk. This changed drastically during the crisis when the entire banking system nearly collapsed and market participants realized there were substantial credit risks in the interbank lending market. In the years 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 (in this case regarding the credit risk embedded in LIBOR rates) that resulted in all of the pricing models being hopelessly inadequate for the task of pricing and hedging.

Derivatives Pricing Post-Crisis

Since the financial crisis we no longer take LIBOR rates to be risk-free and indeed the risk-free curve is now computed from OIS rates. This means for example, that the discount factor in (7) should be based off the OIS curve and as a result we no longer obtain the result that an FRN is always worth par at maturity. OIS forward rates (but not forward LIBOR rates!) are calculated as in (8) or (9) 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. Forward LIBOR rates are now defined as risk-neutral expectations conditional on time $t < T$ (under the so-called *forward* measure) of the spot LIBOR rate, $L(T, T + \delta)$. Similarly, relationships such as (10) no longer hold so that we now live in a *multi-curve* world with a different LIBOR curve for different tenors. For example, we now talk about the 3-month LIBOR curve, the 6-month LIBOR curve etc.

Exercise 2 *Which if any of the 3-month and 6-month LIBOR curves will be lower? Why?*

With these new definitions, straightforward extensions[6] of the traditional pricing formulas (that held pre-crisis) can be used to price swaps, caps, floors, swaptions etc. in the post-crisis world. ▮

We could of course present many other examples of model risk. In addition to other examples we will discuss later in these notes, we briefly identify two more important examples.

- **Model approximations**: It is often the case that (derivative) security prices cannot be computed explicitly with a given model. While Monte-Carlo is often a good alternative, it can be very slow (depending on the model and complexity of the security) and analytic approximations might be preferable. In such situations it is very important to properly test the quality of these approximations. For example, it's quite possible they will work very well in typical market conditions but work poorly in extreme market conditions. It is important to test for and beware of such limitations so that we can respond appropriately in the event that extreme market conditions prevail.

- **Model extrapolation**: We have already identified price extrapolation as an example of model risk but it's worth noting there are many ways a model can be "extrapolated". In Example 3 extrapolation took the form of using market prices of zero-coupon bonds to estimate swaption prices. This of course was not a good idea. Another form of extrapolation would be to use market prices of a given derivative to estimate the price of the same derivative but with a different underlying security or portfolio of securities. This was exactly the problem faced by traders of *bespoke* CDO tranches who used the Gaussian copula model for their pricing. A bespoke tranche is a CDO tranche written on a portfolio of underlying credits that differed from the credits underlying the liquid tranches based on the CDX and ITRAXX indices. Given that CDS rates on the credits underlying the bespoke portfolio were generally observable, the extrapolation problem (assuming the Gaussian copula model) then amounted to choosing the right correlation parameters. There were many approaches[7] to this problem and some of them resulted in very poor prices for the bespoke trance.

---

2 Model Transparency

We begin by discussing some well known models in the equity derivatives space. These models or slight variations of them can also be used in the foreign exchange and commodity derivatives markets. They can be useful when trading or investing in exotic derivative securities and not just the "vanilla" securities for which prices are readily available in the market-place. If these models can be calibrated to the observable vanilla

[6] In order to evaluate these new pricing formulae, dynamics for the spread between the discount curve (calculated from OIS rates) and the LIBOR curves must be specified.

[7] These approaches are discussed in Section 3.5 of *Understanding and Managing Model Risk* (2011) by Massimo Morini. This book discusses extensively the many weaknesses of the Gaussian copula model and model risk more generally.

security prices, then they can be used to construct a range of plausible prices for the more exotic securities. As such they can be used to counter extrapolation risk as well as to provide alternative estimates of the Greeks or hedge ratios. We will not focus on stochastic calculus or the various numerical pricing techniques that can be used when working with these models; they are not the primary consideration here. Instead we simply want to emphasize that a broad class of tractable models exist and that they should be employed when necessary. In general, "quants" should be used to do the more technical work with these models but serious investors and risk managers who rely on these models should fully understand their dynamics as well as the model assumptions. This is true irrespective of whatever market – equity, credit, fixed income, commodity etc.– that you are working in.

It is 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. In Section 3.1 we 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[8] that can be called upon as needed. In this section we will simply introduce various classes of models and discuss some of the main features of these models.

2.1 Local Volatility Models

The GBM model for stock prices assumes that $dS_t = \mu S_t\, dt + \sigma S_t\, dW_t$ where μ and σ are constants. Moreover, when we use *risk-neutral pricing* to price derivative securities we know that $\mu = r - q$ where r is the risk-free interest rate and q is the dividend yield. This means that we have a single free parameter, σ, which we can fit to option prices or, equivalently, the volatility surface. It is not all surprising then that this exercise fails. As we know, 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 4.

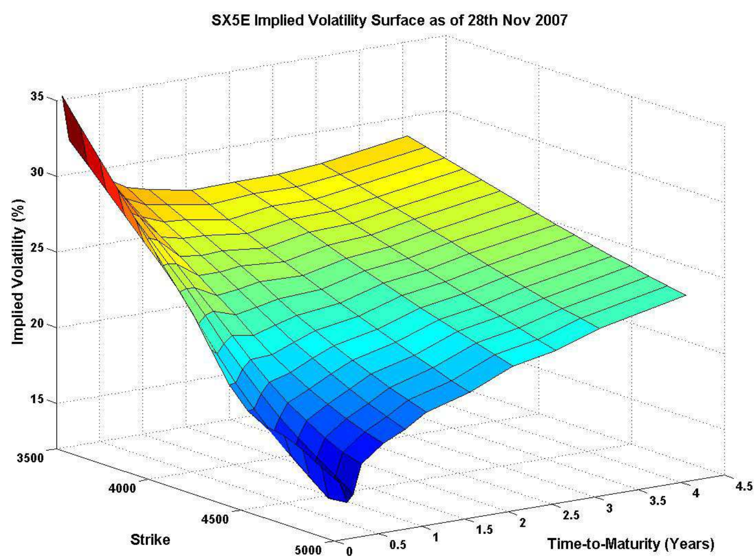

Figure 4: The Eurostoxx 50 Volatility Surface.

[8]See for example "A Stochastic Processes Toolkit for Risk Management" by Brigo, Dalessandro, Neugebauer and Triki (2007). They provide `Matlab` code to calibrate various models that can be used in many contexts. As is the case with these notes, they only focus on single-name securities and do not discuss models for pricing derivatives that have multiple underlying securities.

After this crash, researchers developed alternative models in an attempt to model the skew and correctly price vanilla European calls and puts in a model-consistent manner. While not the first[9] such model, the local volatility model is probability the simplest extension of Black-Scholes. It assumes that the stock's risk-neutral dynamics satisfy

$$dS_t \;=\; (r-q)S_t\,dt \;+\; \sigma_l(t,S_t)S_t\,dW_t \tag{11}$$

so that the instantaneous volatility, $\sigma_l(t,S_t)$, is now a function of time and stock price. The key result of the local volatility framework is the Dupire formula that links the local volatilities, $\sigma_l(t,S_t)$, to the implied volatility surface. This results states

$$\sigma_l^2(T,K) \;=\; \frac{\frac{\partial C}{\partial T} \;+\; (r-q)K\,\frac{\partial C}{\partial K} \;+\; qC}{\frac{K^2}{2}\,\frac{\partial^2 C}{\partial K^2}} \tag{12}$$

where $C = C(K,T)$ is the price of a call option as a function of strike and time-to-maturity. Calculating the local volatilities from (12) is difficult and numerically unstable as taking (mathematical) derivatives is a numerically unstable operation. As a result, it is necessary to use a sufficiently smooth Black-Scholes implied volatility surface when calculating local volatilities. It is worth emphasizing that any vanilla European option that is priced using (11) and (12) will exactly match the prices seen in the market-place that were used to construct the implied volatility surface. This local volatility model then is cable of exactly replicating the market's implied volatility surface.

Nonetheless, local volatility is known to suffer from several weaknesses. 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, the 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*.

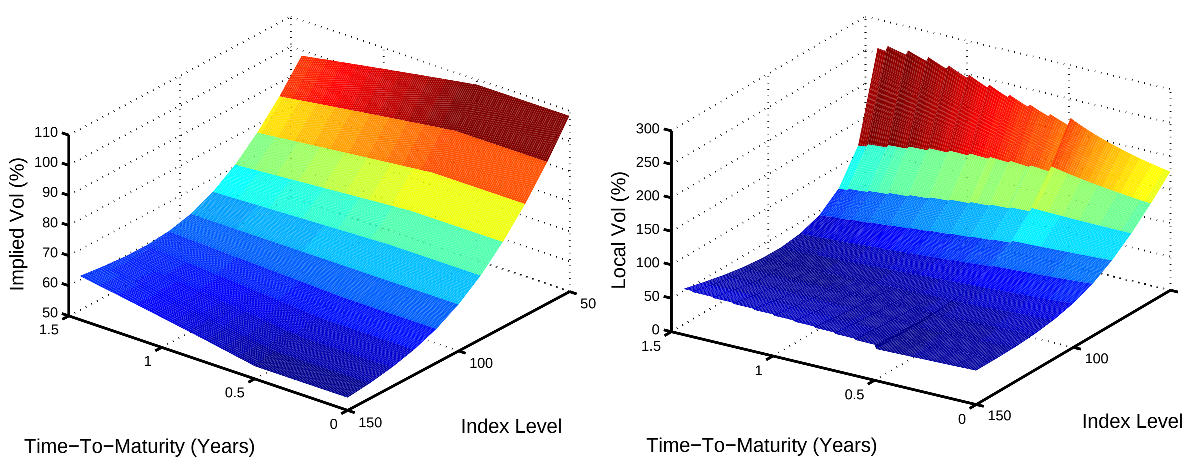

(a) Implied Volatility Surface (b) Local Volatility Surface

Figure 5: Implied and local volatility surfaces. The local volatility surface is constructed from the implied volatility surface using Dupire's formula.

In Figure 5(a) we have plotted a sample implied volatility surface while in Figure 5(b) we have plotted the corresponding *local volatility surface*. This latter surface was calculated by using the implied volatility surface to compute call option prices and then using Dupire's formula in (12) to estimate the local volatilities. The

[9]Earlier models included Merton's jump-diffusion model, the CEV model and Heston's stochastic volatility model. Indeed the first two of these models date from the 1970's. The local volatility framework was developed by Derman and Kani (1994) and in continuous time by Dupire (1994).

(mathematical) derivatives in (12) were estimated numerically. Perhaps the main feature to note is that the local volatilities also display a noticeable skew with lower values of the underlying security experiencing higher levels of (local) volatility.

2.2 Stochastic Volatility Models

The most well-known stochastic volatility model is due to Heston (1989). It is a two-factor model and assumes separate dynamics for both the stock price and instantaneous volatility. In particular, it assumes

$$dS_t = (r-q)S_t\,dt + \sqrt{\sigma_t}S_t\,dW_t^{(s)} \tag{13}$$

$$d\sigma_t = \kappa\,(\theta-\sigma_t)\,dt + \gamma\sqrt{\sigma_t}\,dW_t^{(vol)} \tag{14}$$

where $W_t^{(s)}$ and $W_t^{(vol)}$ are standard $\mathcal{Q}$-Brownian motions with constant correlation coefficient, ρ. (Hereafter we will use $\mathcal{Q}$ to denote an *equivalent martingale measure* (EMM), i.e. a risk-neutral probability measure.)

Exercise 3 *What sign you would expect ρ to have?*

Whereas the local volatility model is a *complete* model, Heston's stochastic volatility model is an *incomplete* model. This should not be too surprising as there are two sources of uncertainty in the Heston model, $W_t^{(s)}$ and $W_t^{(vol)}$, but only one risky security and so not every security is replicable by trading in just the cash account and the underlying security. The volatility process in (14) is commonly used in interest rate modeling where it is known as the CIR[10] model. It has the property that the process will remain non-negative with probability one. When $2\kappa\theta > \gamma^2$ it will always be strictly positive with probability one.

The pricing PDE that the price, $C(t, S_t, \sigma_t)$, of any derivative security must satisfy in Heston's model is given by

$$\frac{\partial C}{\partial t} + \frac{1}{2}\sigma S^2\frac{\partial^2 C}{\partial S^2} + \rho\sigma\gamma S\frac{\partial^2 C}{\partial S\partial\sigma} + \frac{1}{2}\gamma^2\sigma\frac{\partial^2 C}{\partial\sigma^2} + (r-q)S\frac{\partial C}{\partial S} + \kappa\,(\theta-\sigma)\frac{\partial C}{\partial\sigma} = rC. \tag{15}$$

The security price is then obtained by solving (15) subject to the relevant boundary conditions. Alternatively we can write

$$C(t, S_t, \sigma_t) = \mathrm{E}_t^{\mathcal{Q}}\left[e^{-r(T-t)}C_T\right] \tag{16}$$

where C_T is the value of the security at expiration, T, and then use Monte-Carlo[11] methods to estimate (16). Heston succeeded in using transform techniques to solve (15) in the case of European call and put options.

How Well Does Heston Capture the Skew?

An interesting question that arises is whether or not Heston's model accurately represents the dynamics of stock prices. This question is often reduced in practice to the less demanding question of how well the Heston model captures the volatility skew. By "capturing" the skew we have in mind the following: once the Heston model has been calibrated, then European option prices can be computed using numerical techniques such as Monte-Carlo, PDE or transform methods. The resulting option prices can then be used to determine the corresponding Black-Scholes implied volatilities. These volatilities can then be graphed to create the *model's* implied volatility surface which can then be compared to the *market's* implied volatility surface.

Figure 6 displays the implied volatility surface for the following choice of parameters: $r = .03$, $q = 0$, $\sigma_0 = \sqrt{.0654}$, $\gamma = .2928$, $\rho = -.7571$, $\kappa = .6067$ and $\theta = .0707$. Perhaps the most noticeable feature of this surface is the persistence of the skew for long-dated options. Indeed the Heston model generally captures longer-dated skew quite well but it typically struggles to capture the near term skew, particularly when the latter is very steep. The problem with a steep short-term skew is that any diffusion model will struggle to capture it as

[10] After Cox, Ingersoll and Ross who used this model for modeling the dynamics of the short interest rate.

[11] Special care must be taken when simulating the dynamics in (13) and (14) as standard Euler-type simulation schemes do not always converge well when applied to (14).

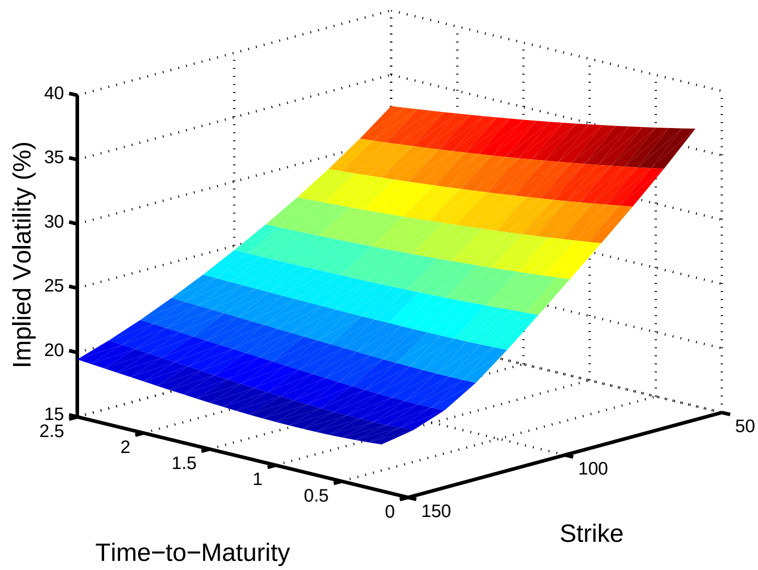

Figure 6: An Implied Volatility Surface under Heston's Stochastic Volatility Model

there is not enough time available for the stock price to diffuse sufficiently far from its current level. In order to solve this problem jumps are needed.

We now discuss the application of local and stochastic volatility models to the pricing of so-called *cliquet* options or forward-start options. We will conclude that local volatility is an inappropriate model or framework for pricing securities of this type. One of the principal reasons for this is that local volatility generates *forward skews* that are too *flat*. (Much of our discussion borrows from Chapter 10 of the *The Volatility Surface: A Practitioner's Guide* (2006) by Jim Gatheral. This chapter can be consulted for further details.)

The Forward Skew

By forward skew we mean 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. Consider Figure 4, for example, where we have plotted the implied volatility surface of the *Eurostoxx 50* as of the 28^{th} November, 2007. The skew on that date is also clear from this figure. Suppose now that we simulate some model that has been calibrated to this volatility surface forward to some date $T > 0$. On any simulated path we can then compute the implied volatility surface as of that date, T. This is the forward implied volatility surface and the forward skew simply refers to the general shape of the skew in this forward volatility surface. Note that this forward skew is model *dependent*.

A well known fact regarding local volatility models (that have been calibrated to today's volatility surface) is that forward skews tend to be very flat. This is not true of stochastic volatility or many jump diffusion models where the forward skew is generally similar in shape to the spot skew that is observed in the market-place today. This is a significant feature of local volatility models that is not at all obvious until one explores and works with the model in some detail. Yet this feature has very important implications for pricing forward-start and cliquet-style securities that we will now describe.

Example 5 (Forward-Start Options)
A forward-start call option is a call option on a security that commences at some future date and expires at a date further in the future. While the premium is paid initially at $t = 0$ the strike of the option is set on the date that the option commences. For example an at-the-money forward start call option has payoff $\max(0,\, S_{T_2} - S_{T_1})$ at date $T_2 > T_1$. The strike, S_{T_1}, is only determined at the commencement date of the option, T_1, but the option premium is paid at date $t = 0$. A cliquet option is then a series of forward-start

options, with the strikes usually set at the money. The premium is paid in advance at time $t = 0$.

It is well known (but should be clear from the form of the payoff anyway) that the model price of a forward-starting call option has a strong dependence on the volatility skew that prevails in the model at time T_1. But the security must be priced at time $t = 0$ and so the model can only be calibrated to the available market prices at time $t = 0$. It is therefore imperative that the pricing model is capable of producing a forward volatility skew at date T_1. Unfortunately local volatility models are generally not capable of producing these forward volatility skews, a fact that is easy to confirm numerically. Stochastic volatility models on the other hand are much better at producing forward volatility skews and should therefore be capable of producing better prices.

As will see in the next example, many structured products are sensitive to assumptions regarding the forward skew and so it is vital when pricing these instruments to understand whether or not the pricing model is suitable. It is not surprising that the failure of market participants to understand the properties of their models has often led to substantial trading losses. ∎

Example 6 (The Locally Capped, Globally Floored Cliquet and Other Structured Notes)
Consider the case of a *locally capped, globally floored cliquet*. The security is structured like a bond so that the investor pays the bond price upfront at $t = 0$ and in return, she is guaranteed to receive the principal at maturity as well as an annual coupon. This coupon is based on the monthly returns of the underlying security over the previous year. It is calculated as

$$\text{Payoff} \;=\; \max\left\{\sum_{t=1}^{12} \min\left(\max(r_t,\ -.01),\ .01\right),\ \text{MinCoupon}\right\} \tag{17}$$

where MinCoupon $= .02$ and each monthly return, r_t, is defined as

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

The annual coupon is therefore capped at 12% and floored at 2%. Ignoring the global floor, this coupon behaves like a strip of forward-starting monthly *call spreads*. Recall that 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 $k_2 > k_1$. Clearly a call spread is sensitive to the implied volatility skew. We would expect the value of this coupon to be very sensitive to the forward skew properties of our model. In particular, we would expect the value of the coupon to increase as the skew becomes more negative. As a result, we 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 . . ."*

And indeed this intuition can be confirmed. In fact numerical experiments suggest that a stochastic volatility model will place a considerably higher value on this bond than a comparable[12] local volatility model. Our intuition is not always[13] correct, however, and it is important to evaluate exotic securities with which we are not familiar under different modeling assumptions. It is also worth keeping in mind that any mistakes are likely to be costly. In the words of Gatheral 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!"*

Of course if the traders priced the deals too highly, they wouldn't have completed any business and while they wouldn't have lost any money, they wouldn't have earned any either. An interesting article discussing cliquets and related securities was published by *RISK magazine* and can be found at

> *http://www.risk.net/risk/feature/1506188/reverse-cliquets-road*

[12]By comparable, we mean that both models have been calibrated to the same initial implied volatility surface.

[13]Gatheral provides an example where you might expect local volatility to *underprice* a security relative to a stochastic volatility model but where the opposite is the case: local volatility ends up overpricing the security.

It is clear from this article how important it is to completely understand the products being traded as well as the suitability of the models that are used to price these products. More specifically, the article discusses reverse cliquets and Napoleon options.

A reverse cliquet is a capital-guaranteed structured product with monthly strike resets through which an investor can hope to earn a very high coupon and as much as 40% over two years if the underlying equity moves up every month. The potential 40% coupon is lowered, however, by each monthly negative performance of the underlying. A Napoleon option is also a capital guaranteed structured product through which an investor can earn a relatively high coupon each year, for example 8%, plus the worst monthly performance of the underlying. There is a high probability that this worst performance turns out to be negative in which case the coupon will be less than 8%.

The key hedging issue with both products is their exposure to *volatility gamma*. This means sellers of these products need to buy option volatility when volatility increases and sell it when volatility decreases so that they buy (volatility) high and sell low. For example, if volatility is extremely high with typical monthly moves on the order of 50% say, then a 1 percentage point change in volatility will not impact the price so the seller has no volatility exposure, or vega, at that point. This is because the high level of volatility is almost certain to consume the 40% coupon for a reverse cliquet or the 8% annual coupon for the Napoleon in one negative month and a one percentage point change will not alter this fact. In contrast, if volatility is currently very low, the product's price will have a high vega. This means with low volatility the seller is long vega and that with high volatility this vega tends to zero. The product seller therefore needs to sell volatility when volatility decreases and buy it back when volatility increases to neutralise the trading book's vega. With big swings in volatility this hedging can be very expensive, and this effect needs to be accounted for in the price. ∎

Example 7 (Simulating Heston's Stochastic Volatility Model)
A common approach for pricing derivative securities is Monte-Carlo simulation which for the Heston model means that we need to simulate the underlying stochastic differential equation (SDE), i.e. (13) and (14). The default simulation scheme for SDE's is the so-called Euler scheme which amounts to simply simulating a discretized version of (13) and (14). While the Euler scheme does converge in theory, depending on the model parameters, κ, θ and γ, it can perform poorly in practice when applied to the volatility process in (14).

For example, Andersen (2007) considers the problem of pricing an at-the-money 10-year call option when $r = q = 0$. He takes $\kappa = .5$, $\sigma_0 = \theta = .04$, $\gamma = 1$, $S_0 = 100$ and $\rho := \mathrm{Corr}(W_t^{(s)}, W_t^{(vol)}) = -0.9$. The true value of the option (which can be determined via numerical transform methods) is 13.085. Using one million sample paths and a "sticky zero" or "reflection" assumption[14], he obtains the estimates displayed in Table 1 for the option price as a function of m, the number of discretization points.

Table 1: Call Option Price Estimates Using an Euler Scheme for Heston's Stochastic Volatility Model. The true option price is 13.085.

Time Steps	Sticky Zero	Reflection
100	28.3	45.1
200	27.1	41.3
500	25.6	37.1
1000	24.8	34.6

One therefore needs to be very careful when applying an Euler scheme to this SDE. In general, it is a good idea from the perspective of implementation risk to price vanilla securities in the Monte-Carlo alongside the more exotic securities that are of direct interest. If the estimated vanilla security prices are not comparable to their analytic[15] prices, then we know the scheme has not converged. That said, superior simulation schemes may also

[14]The sticky zero assumption simply means that anytime the variance process, σ_t, goes negative in the Monte-Carlo it is replaced by 0. The reflection assumption replaces V_t with $|\sigma_t|$. In the limit as $m \to \infty$, the variance will stay non-negative with probability 1 so both assumptions are unnecessary in the limit and the resulting option price estimates should be identical.

[15]Or semi-analytic prices that may be computable via numerical transform inversion methods.

be available and of course one should use them if possible. (In his paper Andersen (2007) proposed a superior scheme for (14) which converged much more quickly.) ∎

2.3 Jump-Diffusions and Levy Processes

We now briefly describe some jump-diffusion models and other more general models such as Levy processes. It is worth emphasizing that jump-models are almost invariably incomplete models with infinitely many EMM's. A particular EMM, $\mathcal{Q}$, is typically chosen (possibly from a parameterized subset of the possible EMMs) by some calibration algorithm. The earliest *jump diffusion* model that was proposed as a model of stock price dynamics was Merton's jump diffusion. This model assumes that the time t stock price satisfies

$$S_t \;=\; S_0\, e^{(\mu-\sigma^2/2)t+\sigma W_t} \prod_{i=1}^{N_t} Y_i \tag{18}$$

where N_t is a Poisson process with mean arrival rate λ, and the Y_i's are IID log-normal random variables. Each Y_i represents the magnitude of the i^{th} jump. In between jumps, the stock price behaves like a regular GBM; hence the name "jump-diffusion" that is given to models of this type. It can be shown that in Merton's model European option prices can be expressed as an infinitely weighted sum of Black-Scholes options prices. If the dynamics in (18) are under an EMM, $\mathcal{Q}$, then μ, λ and the mean jump size are constrained in such a way that the $\mathcal{Q}$-expected rate of return must equal $r-q$. In particular it must be the case that

$$\mu+\lambda(\mu_y-1)=r-q \tag{19}$$

which is an equation in three unknowns. It therefore has infinitely many solutions and so we can conclude from the *Second Fundamental Theorem of Asset Pricing* that Merton's model is incomplete.

Exercise 4 *Show that if the dynamics in (18) are under an EMM, $\mathcal{Q}$, with the cash account as the numeraire then (19) must hold.*

Exercise 5 *Can you see how using (18) to price European options might lead to an infinitely weighted sum of Black-Scholes options prices?* Hint*: Use a conditioning argument.*

An interesting question to consider is how well Merton's jump-diffusion model can replicate the implied volatility surfaces that are typically observed in the market. Note that in contrast to the GBM model, we have five parameters[16], σ, λ, μ, μ_y and σ_y^2 and just one equation to satisfy, namely (19). We therefore have much more flexibility than GBM which can only achieve constant implied volatility surfaces.

Figure 7 displays the implied volatility surface under a Merton jump-diffusion model when $\sigma=20\%$, $r=2\%$, $q=1\%$, $\lambda=10\%$, $\mu_z=-.05$ and $\sigma_z=\sqrt{0.1}$. While other shapes are also possible by varying these parameters, Figure 7 demonstrates one of the principal weaknesses of Merton's jump-diffusion model, namely the rapid flattening of the volatility surface as time-to-maturity increases. For very short time-to-maturities, however, the model has no difficulty with producing a steep volatility skew. This is in contrast to stochastic volatility models which do not allow jumps.

More recently, Kou (2002) developed the double-exponential jump-diffusion model where the jump-sizes have a double exponential distribution. Duffie, Pan and Singleton (1998) introduced the class of *affine jump-diffusion* models and many other jump-diffusion models have been introduced by other researchers. While introducing a new model is a simple task, showing that the model is *tractable* and of practical use is the real challenge.

Another class of models that has become more popular is the class of *exponential Levy processes* and *exponential Levy processes with a stochastic clock*.

Definition 1 *A Levy process is any continuous-time process with stationary*[17] *and independent increments.*

[16] We could use σ_z and σ_z^2 in place of μ_y and σ_y^2.

[17] X_t is a stationary process if the distribution of $X_{t+h}-X_t$ does not depend on t.

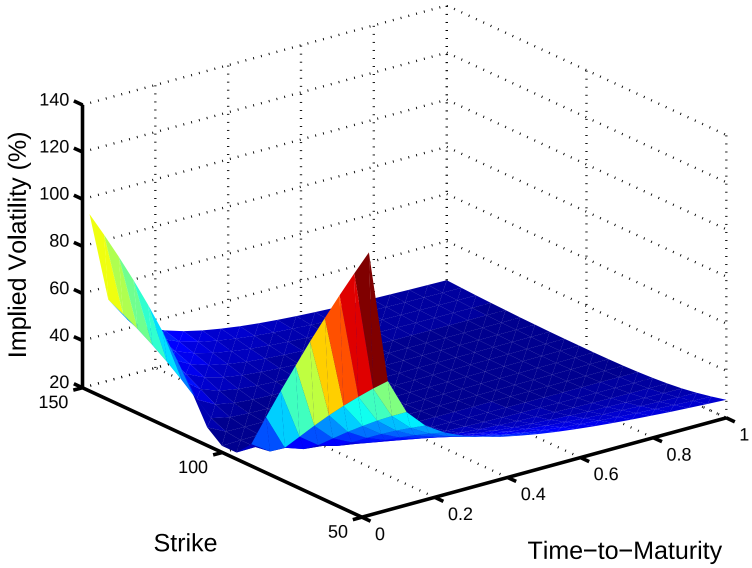

Figure 7: An Implied Volatility Surface under Merton's Jump Diffusion Model

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 (in which case most of the jumps will be very small), or of *finite activity* so that it makes only finitely many jumps in any finite time interval. The most important result regarding Levy processes is the famous *Levy-Khintchine formula* which gives the functional form that the characteristic function of any Levy Process must possess.

Definition 2 *An exponential Levy process, S_t, satisfies $S_t = \exp(X_t)$ where X_t is a Levy process.*

The Merton jump-diffusion model is an example of an exponential Levy process of finite activity. One of the weaknesses of an exponential Levy process is that it does not capture "volatility clustering", i.e. 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 a *stochastic clock*, however, can capture volatility clustering. Such a process takes teh form $S_t = \exp(X_{Y_t})$ where Y_t is a strictly increasing stochastic process or *subordinator*.

Figures 8(a) and 8(b) display three samples paths each for two time-changed exponential Levy processes. Figure 8(a) plots paths of a *normal-inverse Gaussian* process with a *CIR* subordinator. Figure 8(b) plots paths of a *variance-gamma* process[18] with a *Gamma - OU* subordinator. Note how these paths differ from those generated by a GBM model, for example. GBM paths of course cannot jump and do not display volatility clustering. Despite the complexity of these processes they are quite tractable in that vanilla option prices can be calculated via numerical transform inversion. They are also reasonably straightforward to simulate.

---

3 More on Model Risk and Calibration

As discussed in Example 3 when we used the Black-Derman-Toy (BDT) model to price swaptions, the calibration process in an integral component of the pricing process. Indeed, in the early 2000's there was debate about precisely this issue among some well known academic and quant researchers. The context for the debate was the general approach in the market at the time to use simple one-factor models (such as BDT and other

[18]See Cont and Tankov's book or the article "A Perfect Calibration! Now What?" by Schoutens, Simons and Tistaert (Wilmott, 2007) for further details on these processes.

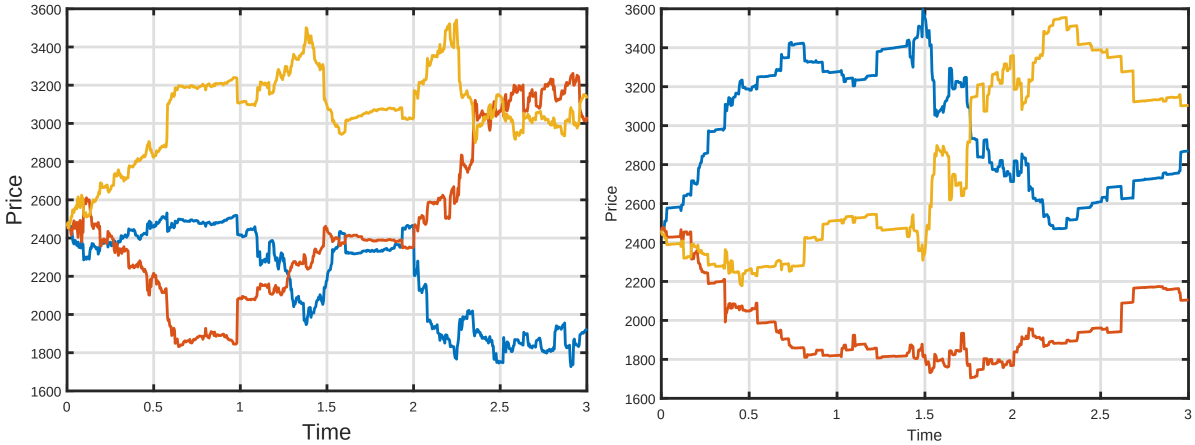

(a) NIG Process with CIR Clock (b) Variance-Gamma Process with OU-Gamma Clock

Figure 8: Sample Paths of Some Time-Changed Exponential Levy Processes

binomial lattice models of the short-rate) to price Bermudan swaptions. In a 2001 paper[19] Longstaff, Santa-Clara and Schwartz argued that since the term-structure of interest rates was driven by several factors (at least three in general given our earlier PCA analysis), using a one-factor model to price Bermudan swaptions was fundamentally flawed and would result in erroneous prices. On the face of it, this argument was sound but the authors neglected to account for the calibration process. In response Andersen and Andreasen (2001) argued that the Bermudan swaption prices were actually quite accurate even though they were indeed typically priced at the time using simple one-factor models. 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 that were observable in the market. 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 the dynamics of a model may not be important at all if the exotic security being priced (the Bermudan swaption in this case) 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. Another source of calibration risk is the topic of the following exercise.

Exercise 6 *We have a model that is used to price (liquid) securities but the model does require a numerical integration. We want to use these liquid securities as calibration securities in order to calibrate the model. What sort of algorithm would you use for your numerical integration (which will be called many times by the calibration routine)?*

3.1 A More Robust Approach to Managing Extrapolation Risk

We continue our discussion of extrapolation risk by returning again to the problem of pricing barrier options. It is well known that the price of a barrier option is *not* solely determined by the *marginal distributions*[20] of the underlying stock price. And yet the implied volatility surface only provides information about the marginal risk-neutral distributions. Therefore when pricing a barrier option with a model that has been calibrated to vanilla European options only, it is clear that we are extrapolating and probably extrapolating a long way! After all, knowing only the marginal risk distributions amounts to having *no information* regarding the copula of the joint distribution of the stock prices at various distinct times. It should be no surprise then that different models

[19] "Throwing Away a Billion Dollars: the Cost of Sub-optimal Exercise Strategies in the Swaptions Market" (2001), Journal of Financial Economics. The response, "Factor Dependence of Bermudan Swaptions: Fact or Fiction?", by Andersen and Andreasen (2001) was also published in the Journal of Financial Economics.

[20] By marginal distribution we mean the distribution of S_t for a fixed time, t. By joint distribution we mean the joint distribution of $(S_{t_1}, \ldots, S_{t_n})$ for fixed times $t_1, \ldots, t_n$.

that agree on the marginal distributions, might still extrapolate to very different exotic prices. Figure 9 in Example 8 below[21] Simons and Tistaert (2007). emphasizes this point.

Example 8 (Pricing Barrier Options)
Figure 9 displays the prices of down-and-out (DOB) and up-and-out (UOB) barrier call option prices for several different models, as a function of the barrier level. These models, which include the Heston model as well as a jump-diffusion version of the Heston model and several time-changed exponential Levy process models from Section 2.3, have all been calibrated to the same implied volatility surface.

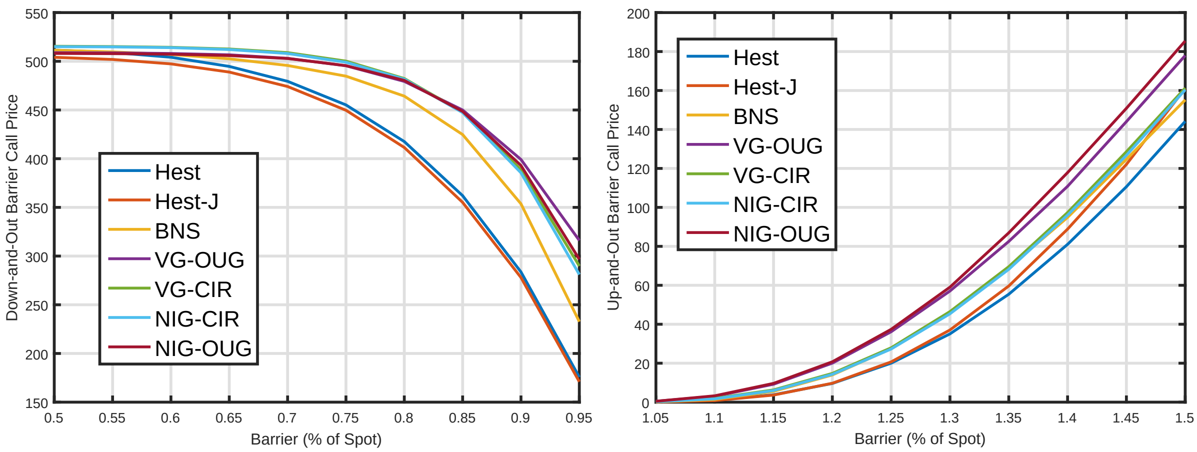

(a) DOB Call Prices (b) UOB Call Prices

Figure 9: Down-and-out (DOB) and 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.

The payoffs of the options are:

$$\begin{aligned} \text{Payoff}_{\text{DOB}} &:= (S_T - K)^+ \, 1_{\{m_T > B\}} \\ \text{Payoff}_{\text{UOB}} &:= (S_T - K)^+ \, 1_{\{M_T < B\}} \end{aligned}$$

where S_T is the terminal value of the underlying security, K is the strike, B is the barrier, and m_T and M_T are the minimum and maximum values of the underlying price, respectively, over the interval $[0, T]$. The strike K for both the DOB and UOB options was set equal to the initial level, S_0, of the underlying security. The parameter values for the various models were calibrated to the implied volatility surface of the Eurostoxx 50 Index on October 7^{th}, 2003. As a result, their marginal distributions coincide, at least to the extent that the calibration errors are small. Yet it is clear from Figures 9(a) and 9(b) that the different models result in very different barrier prices. This example again serves to emphasize the danger of extrapolation and the importance of understanding the calibration process and the selection of the calibration securities. ∎

Exercise 7 *The DOB prices in Figure 9(a) seem to coincide for all models when the barrier level is low relative to S_0. Why is this?*

How to Manage Extrapolation Risk

The question therefore arises as to what model one should use in practice when pricing (and risk managing) the barrier option. This is a difficult question to answer but perhaps the 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

[21] This example is taken from the article "A Perfect Calibration! Now What?" by Schoutens, Simons and Tistaert (Wilmott, 2007).

reasonable dynamics that make sense from a modeling viewpoint. For example, GBM could be ruled out immediately on the basis of both (a) and (b). The minimum and maximum of these prices could 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.

In the case of barrier options in the equity derivatives market, many participants appear to use local volatility models but with various ad-hoc adjustments included. These adjustments are intended to account for the possibility of jumps and other hedging difficulties that might arise. Barrier options are considerably more liquid in the foreign exchange markets, however, and so there has been a greater emphasis there on developing models that can accurately reproduce the market prices of vanilla and digital options. This has led, for example, to the development of stochastic-local volatility models and it is still a topic of ongoing research.

Sometimes model risk is (unfairly) confused with personnel risk! Apparently the following situation did actually occur (probably more than once) in the early days of barrier option trading: when asked to quote a *bid-ask* price for a particular barrier option, a trader decided to quote in terms of Black-Scholes implied volatilities. Knowing that a European option price is an increasing function of implied volatility, the trader quoted σ_1/σ_2 where $\sigma_1 < \sigma_2$. What he meant by this quote was that he was willing to buy the barrier option at a Black-Scholes volatility of σ_1 and that he was willing to sell it at a Black-Scholes volatility of σ_2. Unfortunately it turned out that the Black-Scholes barrier price evaluated at σ_1 was *greater* than the Black-Scholes barrier price evaluated at σ_2. Presumably the trader either quickly understood the importance of model transparency or quickly lost his job.

Exercise 8 *Is is possible to encounter extrapolation risk when pricing a derivative security whose price only depends on the marginal distribution of the underlying stock? If so, how might you handle this risk?*

Example 9 (Gap Risk in Pricing Structured Credit Notes)
Chapter 2 of Morini[22] discusses an interesting example regarding the pricing of a structured note whose payoff at maturity depends on the (leveraged) performance of an underlying CDS. The structured note is such that the seller of the note is very exposed to a so-called *gap risk* where the seller is on the hook for credit losses on the underlying CDS beyond some threshold. Understanding how to evaluate this gap risk is a key concern in the pricing and it turns out that the value of this gap risk is very sensitive to the assumed model dynamics. In his analysis Morini considers structural as well as reduced-form models where the intensity process driving the default event of the underlying is a CIR model or a CIR model with jumps. ▮

3.2 Other Problems With Model Calibration

One has 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 that we have a volatility surface available[23] to us then, at the very least, we can certainly calibrate our model to this surface. An obvious but potentially hazardous solution 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 the 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:

[22] In particular, see Sections 2.3 to 2.7 of *Understanding and Managing Model Risk* (2011) by Massimo Morini.

[23] Since there will only be options with finitely many strikes and maturities, it will be necessary to use interpolation and extrapolation methods to construct the volatility surface from the implied volatilities of these options. In fact it is probably a good idea to also use variance-swaps to construct the volatility surface since the prices of variance swaps will depend in part on the prices of deep out-of-the money options. Market prices will generally not exist for these options and variance-swap prices can therefore help us to determine implied volatility levels in these (deep OTM) regions.

1. In general (20) is a non-linear and non-convex optimization problem. As such it may be difficult to solve as there may be many local minima. It is possible that the optimization routine could get "stuck" in one of these local minima and not discover other local minima that could result in a much smaller objective function.

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. It is 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 but it becomes unacceptable, for example, in the following circumstances:

   Day 1: Model is calibrated to option prices and the 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. (We are ignoring the fact that there is 1 day less to maturity which we assume has an insignificant impact on prices.) In particular, the volatility surface has not changed from Day 1. 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?

3. This is related to the previous point and explains why P_2 can be so different from P_1. 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 very similar marginal distributions, but they will often result in very different joint distributions. This was demonstrated in Figure 9 when we discussed the pricing of barrier options. (In the latter case, however, different models rather than different parameter values within the same model were responsible for generating different joint distributions.)

There are several methods that can be used to resolve the problems listed above. If possible, it is a good idea to include in the calibration exercise some path-dependent securities. This will decrease the likelihood of the troublesome valleys existing. Unfortunately, liquid, path-dependent option prices may not be available in the market place and so this solution may not be feasible.

Regardless of whether or not we can include path-dependent options in the calibration, it is generally a good idea to try and "convexify" the optimization problem in (20). This can be done by adding an additional (convex) term to the objective function in (20). This term would penalize you the further the selected process is from some pre-specified "base" process. There are many ways in which this could be implemented.

Another possibility is to include a penalty term in the objective function that penalizes deviations in the parameter values from the calibrated parameter values of the previous day. By making this penalty term large, the calibration algorithm will only choose a new set of parameters when they perform significantly better than the previous day's parameters.