

Practicalities of Pricing Exotic Derivatives

John Crosby

Glasgow University / UBS

My website is: <http://www.john-crosby.co.uk>

If you spot any typos or errors, please email me.

My email address is on my website.

Lecture given 1st July 2011 for Module 4
of the M.Sc. course in Mathematical Finance
at the Mathematical Institute at Oxford University.

File date 2nd June 2011 at 08.15

- Your course of studies has now shown you how to price several types of exotics analytically. You have studied generic pricing techniques such as Monte Carlo simulation and PDE lattice methods which can, potentially, price many different types of exotic options under many different types of stochastic processes.
- So pricing exotic options is easy, right?
- We are given the payoff of the exotic option by our trader. We choose our favourite stochastic process. Naturally, to show how clever we are, it has stochastic vol, stochastic skew, jumps, local volatility, stochastic interest-rates,...
- We calibrate our model to the market prices of vanilla (standard European) options. Naturally, the fit to vanillas is wonderful which just reinforces how great our model must be.
- We price the exotic option in question. The trader now has an arbitrage-free price - so nothing could possibly be wrong with it (!). The trader puts a bid-offer spread around our price (usually the spread is 10 per cent or more of the price). The sales person may add on a further sales margin (which may be another 10 per cent or more).

- Everyone is happy.
- Nothing could go wrong, could it?

- Everyone is happy.
- Nothing could go wrong, could it?
- Well, actually lots of things can go wrong.
- Pricing a book of exotics is not easy. It can sometimes be an art as well as a science.
- I will attempt to highlight a few of the problems - and a few possible solutions.

- Talk outline:
- Robust pricing of exotics.
- Choice of model.
- Vanilla and barrier fx options.
- Model calibration.



- Because the price of exotic options can often be highly sensitive to the chosen model, *robust* pricing methodologies are very attractive. A *robust* pricing methodology can, loosely, be defined as one which works for two or more (classes of) models. Because it works for two or more models, there is robustness against model error (either misspecification of the type of stochastic process and/or misestimation of the parameters of the process).
- An example of this is as follows:

- Consider a binary cash-or-nothing call option (sometimes called a European digital option) which pays one dollar if the stock price S_T at expiry T is greater than or equal to the strike K and it pays zero otherwise.
- Suppose we form a portfolio of long $1/(2\Delta K)$ vanilla call options with strike $K - \Delta K$ and short $1/(2\Delta K)$ vanilla call options with strike $K + \Delta K$, where $\Delta K > 0$.
- The payoff of the portfolio at expiry is:
 $1/(2\Delta K)[S_T - (K - \Delta K)] - 1/(2\Delta K)[S_T - (K + \Delta K)] = 1$ if $S_T \geq K + \Delta K$,
 0 if $S_T < K - \Delta K$,
 $0 < 1/(2\Delta K)[S_T - (K - \Delta K)] < 1$ if $K - \Delta K \leq S_T < K + \Delta K$.
- If we consider the limit as $\Delta K \rightarrow 0$, then we see that the payoff of the portfolio replicates that of the binary cash-or-nothing call option.

- Let us denote by $V(K, \sigma(K))$ the price of a vanilla call option with strike K and with implied Black and Scholes (1973) volatility $\sigma(K)$.
- Then, in the absence of arbitrage, the price of the binary cash-or-nothing call option is:

$$\begin{aligned} \lim_{\Delta K \downarrow 0} \frac{V(K - \Delta K, \sigma(K - \Delta K)) - V(K + \Delta K, \sigma(K + \Delta K))}{2\Delta K} &= -\frac{dV}{dK} \\ &= -\frac{\partial V}{\partial K} - \frac{\partial V}{\partial \sigma} \frac{\partial \sigma}{\partial K} \end{aligned}$$

- Note that we can compute the term $-\frac{\partial V}{\partial K}$ by noting that it is the discount factor $\exp(-rT)$ multiplied by the $N(d_2)$ term in the Black and Scholes (1973) formula, evaluated with strike K and volatility $\sigma(K)$:

$$-\frac{\partial V}{\partial K} = \exp(-rT)N(d_2)$$

- We can compute the term $\frac{\partial V}{\partial \sigma} \frac{\partial \sigma}{\partial K}$ by noting it is the “vega” multiplied by the “volatility skew”.

- This formula gives us a non-parametric methodology for pricing binary cash-or-nothing call options in the presence of a volatility skew or smile.
- If the volatility surface is negatively-skewed as it is for equities, then the term $\frac{\partial V}{\partial \sigma} \frac{\partial \sigma}{\partial K}$ is negative and so the presence of the skew increases the price of a binary cash-or-nothing call option relative to its price in a (constant volatility) Black and Scholes (1973) world.
- The vanilla call spread gives us the hedge. In practice, there is a problem in that, in the limit $\Delta K \rightarrow 0$, we have to simultaneously both buy and sell an infinite quantity of vanilla options (!). Recognising that this is not possible, a safer strategy is sub- or super-replication.

- Suppose we write a binary cash-or-nothing call option with strike K (an analogous strategy can be used if we are buying the option) and there are vanilla options, with strikes K_1 and K_2 , which we can respectively buy and sell in the market at prices $V(K_1, \sigma(K_1))$ and $V(K_2, \sigma(K_2))$, where $K_1 < K_2 < K$.
- Then if we buy $1/(K_2 - K_1)$ of the options with strike K_1 and we sell $1/(K_2 - K_1)$ of the options with strike K_2 , we can super-replicate the binary cash-or-nothing call option payoff (ie we never lose money). The closer to K that K_2 is, the cheaper this super-replicating portfolio is. Note that in the event that it turns out that the stock price is between K_1 and K_2 at expiry, then we always make money on our strategy.
- On the other hand, if we sold at the initial cost of the super-replicating portfolio, the option buyer has paid over the odds. A proportion of the business of investment banks in the area of exotic options appears to consist of trying to find poorly-informed or less technical investors who make this mistake.

- There are potentially economies of scale related to the analysis above in the sense that if we already have vanilla options with strike K on our books, we may be able to use these as part of our hedge.
- This illustrates a general principle of hedging a book of exotic options which is that it is much easier if you already have a large vanilla options book.
- Note that we can get the price of a binary cash-or-nothing put option by a “put-call parity” relation (namely, that a binary cash-or-nothing call option plus a binary cash-or-nothing put option equals the discount factor to expiry).



- Two papers by Schoutens, Simons and Tistaert illustrate a general feature of pricing exotic options.
- Specifically, in the first paper, they calibrated seven different stochastic models to the market prices of 144 different vanilla options on the Eurostoxx 50 index with maturities ranging from less than a month to just over five years. The models were sophisticated models (HEST, HESTJ, BNS, VGCIR, VGOU, NIGCIR, NIGOU) with between five and eight free parameters.
- Any model with this number of parameters should be able to give a decent fit to the market prices of vanilla options and, indeed, the paper demonstrates this. We leave aside for now important issues such as what happens if the error function has many local minima, potential parameter instability arising and the effect of parameter instability through time on hedging performance (see later).
- Using Monte Carlo simulation, a number of different exotic options were priced.

- A brief selection of prices are:
Lookback call options.

Model	<i>HEST</i>	<i>HESJ</i>	<i>BNS</i>	<i>VGCI</i>	<i>VGOU</i>	<i>NIGCI</i>	<i>NIGOU</i>
Price	838.48	845.19	771.28	724.80	713.49	730.84	722.34

Cliquet options with local floor at -0.03, local cap at 0.05 and global floor at -0.05.

Model	<i>HEST</i>	<i>HESJ</i>	<i>BNS</i>	<i>VGCI</i>	<i>VGOU</i>	<i>NIGCI</i>	<i>NIGOU</i>
Price	0.0729	0.0724	0.0788	0.1001	0.1131	0.0990	0.1092

All prices are quoted from Schoutens, Simons and Tistaert (“A Perfect Calibration! Now What?”, Wilmott magazine, 2004 or see Wim Schouten’s website).

- We can see that there is a large variation in the prices of these exotic options.
- In the case of the lookback options, the variation is more than 18 per cent.
- In the case of the cliquet options, the variation is more than 56 per cent.
- In the case of barrier options, results in Schoutens, Simons and Tistaert show that the variation can be more than 200 per cent (especially if the spot price is close to the barrier level).
- Now it is worth remembering that the bid-offer spreads will typically be greater than ten per cent (even much more for barrier options when the spot price is close to the barrier level).
- However, the variation in the exotic option prices over the seven different models is still such that the highest bid will be greater than the lowest offer.

- This means that, for example, a hedge fund could, in principle, make money by buying from the lowest offer and simultaneously selling at the highest bid if it knew that these models were being used by particular market-making investment banks.
- This is potentially a very dangerous situation for either market-making bank.
- This is called “adverse price discovery risk”.
- It is the risk you do a lot of business because your prices are out of line with other banks.
- Totem prices from MarkIT can help identify this risk.
- A number of banks (for example, UBS, Barclays Capital and others) provide “live” screen-based prices for first and second generation exotics. These prices can also be used to identify this risk.

- It is important to see where your prices lie in relation to other banks' prices.
- This is a variation on the saying of John Maynard Keynes: “It is better to be wrong conventionally than right unconventionally”.
- This, in turn, means that models have to be chosen not only so as to match the market prices of vanilla options but also, where possible, to match the observable prices of exotic options. This may rule out the use of some models and/or require the use of different models for different types of exotic options and/or require the development of new types of models with the built-in flexibility to fit the market prices of both vanilla and exotic options.



- A paper by Lipton and McGhee illustrates some of these features.
- In the foreign exchange (fx) options markets, barrier options are very actively traded. Many different types of barrier options (single barrier, double barrier, with or without rebates, partial, window, etc) trade but, by far, the most actively traded barrier options are double-no-touch (DNT) options. These pay one unit of domestic currency at expiry if the spot fx rate (quoted as the number of units of domestic currency per unit of foreign currency) never trades equal to or outside a lower barrier level nor an upper barrier level. If either the lower barrier level or the upper barrier level are touched or breached prior to expiry, the option expires worthless.
- The market prices of DNT options on major currency pairs are widely available to market-making banks - either through inter-dealer brokers or through “live” screen prices. Hence, it is desirable that any model can fit these market prices.
- The market-standard working assumption made in pricing these DNT options is that the spot fx rate is monitored continuously to see if the barrier levels have been hit (in practice, monitoring is continuous from Monday morning in Sydney until Friday close in New York).

- Lipton and McGhee make a number of observations:
- If a Dupire (1994) local volatility model is calibrated to the market prices of vanilla options, and then used to price DNT options, the resulting model DNT option prices are less than the market prices i.e. a local volatility model tends to under-price DNT options relative to the market prices. For example, the model price may be around 50 to 75 per cent of the market price.
- If a stochastic volatility (for example, Heston (1993)) model is calibrated to the market prices of vanilla options, and then used to price DNT options, the resulting model DNT option prices are greater than the market prices i.e. a stochastic volatility model tends to over-price DNT options relative to the market prices. For example, the model price may be around 125 to 150 per cent of the market price.
- In either case, the degree of mis-pricing is well in excess of the bid-offer spread (which is usually around 2 to 4 per cent of notional - eg. if the mid-market price is 0.48, then the bid-offer spread might be 0.465 to 0.495).

- The above observations are fairly robust to which type of stochastic volatility model is chosen and which currency pair is chosen (at least for major currency pairs).
- Adding a compound Poisson jump process may give a more realistic fit to the short-term skew/smile but it doesn't change the broad conclusion regarding DNT option prices.
- Lipton and McGhee did not consider pure jump Levy processes such as CGMY but my experience is that, if calibrated to vanilla options, the CGMY model over-prices DNT options relative to market prices.
- The above observations lead Lipton and McGhee to propose their “Universal Volatility” model.

- Let us denote the spot fx rate, at time t , by $S(t) \equiv S$. Domestic (respectively, foreign) interest-rates are denoted by r_d (respectively, r_f). Lipton and McGhee proposed the following risk-neutral dynamics:

$$\begin{aligned}\frac{dS}{S} &= (r_d - r_f - \lambda m)dt + \sigma_L(t, S(t))\sqrt{V(t)}dz_S(t) + (\exp(J) - 1)dN(t), \\ dV(t) &= \kappa(\theta - V(t))dt + \epsilon\sqrt{V(t)}dz_V(t), \quad V(0) \equiv V_0,\end{aligned}$$

where λ is the intensity rate of the Poisson process $N(t)$, the Brownian motions $z_S(t)$ and $z_V(t)$ have constant correlation ρ , $m = E[\exp(J) - 1]$ and where $V(t)$ is a stochastic volatility (or stochastic variance to be more exact) and $\sigma_L(t, S(t))$ is a local volatility function.

- If $\sigma_L(t, S(t)) \equiv 1$, then the model is the same as Heston (1993) (plus jumps).
- If $\epsilon \equiv 0$, then the model is the same as Dupire (1994) (plus jumps).

- Either Heston (1993) (plus jumps or without jumps) or Dupire (1994) (plus jumps or without jumps) could (on their own) be calibrated to the market prices of vanilla options. The key issue is that we can “mix” the local volatility and the stochastic volatility. Essentially, by “tweaking” the vol-of-vol parameter ϵ to some value which is between that value it would take if $\sigma_L(t, S(t)) \equiv 1$ (the value it would take in a wholly stochastic volatility model) and zero (the value it would take in a wholly local volatility model), we find the value which can approximately match the market prices of DNT options.
- To put it another way, a local volatility model tends to under-price DNT options relative to the market prices, a stochastic volatility model tends to over-price DNT options relative to the market prices, so we should be able to “mix” the two models by “tweaking” the vol-of-vol parameter ϵ in such a way that we match the market prices of DNT options. (In practice, “tweaking” ϵ will likely cause other parameters to change somewhat).

- A few comments are in order:
- Lipton and McGhee was published in 2002. Since publication, this model has become a kind of market standard model and is used by many banks. It is quite common - in fact, nearly standard - to NOT have a jump term as it complicates the solution but this is to the detriment of fitting short-dated options.
- Parameterizations of the local volatility function of the form $\sigma_L(t, S(t))S(t) = a + bS(t)$ i.e. displaced diffusion do NOT seem to fit the market at all. One typically parameterises the local volatility function $\sigma_L(t, S(t)) = \alpha + \beta(S(t) - S(0)) + \gamma(S(t) - S(0))^2$, where α , β and γ are independent of $S(t)$.
- No useful (by which we mean except in limited special cases) analytic formulae exist for either vanilla option prices or barrier option prices. Therefore, all options are priced by numerically solving the relevant two-factor (plus time) PDE.
- One invariably uses term structures of interest-rates in both currencies which poses no extra difficulty in solving the PDE numerically.

- With an efficient ADI implementation, a whole grid of options can typically be priced in around 2 seconds.
- The model has a number of attractive features. However, it is not perfect. Bid-offer spreads have declined and volatility skews have become more severe (especially dollar/yen and even before the credit crunch) since the model was first introduced. Now, in order to closely match both vanilla and DNT (or other barrier option) prices, one needs to introduce time-dependent parameters.
- This involves making the parameters α , β and γ of the local volatility function as well as, at least, the vol-of-vol parameter ϵ time-dependent.
- Often, calibrations to real market data show these parameters varying in a way that is not credible or believable from a financial perspective. By the time one has finished, one might ask whether one has a model or a giant interpolation machine.

- Much more recently, two other alternative models have been proposed. The first was proposed by Dherminder Kainth and the second by Peter Carr and John Crosby. Both models produce excellent (in fact, astonishingly good) fits without the need for time-dependent parameters and both models give very realistic (risk-neutral) dynamics. We will now describe both models in more detail.

- The Kainth model utilises what he calls a “double Heston” process. Let us denote the spot fx rate, at time t , by $S(t) \equiv S$. Domestic (respectively, foreign) interest-rates are denoted by r_d (respectively, r_f). The risk-neutral dynamics are:

$$\begin{aligned}\frac{dS}{S} &= (r_d - r_f)dt + \sum_{i=1}^2 \sqrt{V_i(t)}dW_i(t), \\ dV_1(t) &= \kappa_1(\theta_1 - V_1(t))dt + \epsilon_1\sqrt{V_1(t)}dz_1(t), \quad V_1(0) \equiv V_{0,1}, \\ dV_2(t) &= \kappa_2(\theta_2 - V_2(t))dt + \epsilon_2\sqrt{V_2(t)}dz_2(t), \quad V_2(0) \equiv V_{0,2}.\end{aligned}$$

The parameters $\kappa_1, \theta_1, \epsilon_1, V_{0,1}, \kappa_2, \theta_2, \epsilon_2, V_{0,2}$ are positive constants. There are two stochastic variance processes $V_1(t)$ and $V_2(t)$, with initial (time zero) values $V_{0,1}$ and $V_{0,2}$ respectively. Both stochastic variance processes drive the process for the spot fx rate. The processes $W_1(t)$, $W_2(t)$, $z_1(t)$ and $z_2(t)$ are standard Brownian motions with the following correlation structure:

$$E[dz_1(t)dW_1(t)] = \rho_1, \quad E[dz_2(t)dW_2(t)] = \rho_2.$$

All other correlations are identically equal to zero. In particular,

$$E[dz_1(t)dz_2(t)] = 0, \quad E[dW_1(t)dW_2(t)] = 0, \quad E[dz_1(t)dW_2(t)] = 0, \quad E[dz_2(t)dW_1(t)] = 0.$$

- The way the correlation structure is set up is important. In particular, it is part of the model design that it is intended that ρ_1 and ρ_2 will have opposite signs (eg. ρ_1 is negative and ρ_2 is positive).
- If we define $\Sigma^2 \equiv V_1(t) + V_2(t)$, then the dynamics of $S(t)$ can be written in the form:

$$d(\log S) = (r_d - r_f - \frac{1}{2}\Sigma^2)dt + \Sigma dW(t).$$

This shows that what we have is a three-factor model.

- Both vanilla options and DNT options (or other types of barrier options or exotic options) can be priced by numerically solving a three-factor PDE (i.e. three factors plus time).
- However, for vanilla options, (much faster) Fourier inversion methods can also be used because the characteristic function can be computed in closed-form.

- Exotic options can also be priced by Monte Carlo simulation.
- Using term structures of interest-rates in both currencies poses no extra difficulty.
- One could also easily allow some of the parameters (eg. ϵ_1 and ϵ_2) to be time-dependent.
- However, that seems to be unnecessary. Kainth shows that, at least, for major currency pairs and the data-set considered, this double Heston model can give an excellent fit to the market prices of both vanilla and DNT fx options with all parameters constant.
- The double Heston model also has two additional attractive features:

- Kainth (2007) (talk given at the ICBI Global Derivatives conference, Paris, May 2007) gives an example of parameter values obtained from a calibration:

Process	First	Second
i	1	2
$\sqrt{V_{0,i}}$	0.09135	0.06599
$\sqrt{\theta_i}$	0.08583	0.05800
κ_i	2.132	5.493
ρ_i	-0.8646	0.9612

- Notice how this seems to suggest two distinct processes, the first with a low mean reversion rate and the second with a high mean reversion. The idea of two characteristic time-scales for mean reversion is certainly appealing.
- Furthermore, the first correlation is large and negative and the second correlation is large and positive.

- If we consider the process for $\Sigma^2 = V_1(t) + V_2(t)$, we find (from Ito's lemma) that:

$$d\Sigma = [\dots]dt + \frac{\epsilon_1 \sqrt{V_1(t)}}{2\sqrt{V_1(t) + V_2(t)}} dz_1(t) + \frac{\epsilon_2 \sqrt{V_2(t)}}{2\sqrt{V_1(t) + V_2(t)}} dz_2(t). \quad (1)$$

- Furthermore, the instantaneous correlation between $d(\log S)$ and $d\Sigma$ is:

$$\frac{\epsilon_1 \rho_1 V_1(t) + \epsilon_2 \rho_2 V_2(t)}{2((V_1(t) + V_2(t))(\epsilon_1^2 V_1(t) + \epsilon_2^2 V_2(t)))^{1/2}}, \quad (2)$$

which (loosely speaking) is proportional to $\epsilon_1 \rho_1 V_1(t) + \epsilon_2 \rho_2 V_2(t)$. Clearly, this instantaneous correlation is stochastic since it changes as $V_1(t)$ and $V_2(t)$ change.

- We have noted that we expect ρ_1 and ρ_2 to have opposite signs (borne out by the calibration above). Also, $\epsilon_1, V_1(t), \epsilon_2, V_2(t)$ are all non-negative.
- We see that the instantaneous correlation between $d(\log S)$ and $d\Sigma$ is not only stochastic but can be negative or positive and can actually switch sign through time (provided, to repeat, ρ_1 and ρ_2 are of opposite sign).

- This is linked to the idea of stochastic skew.
- Stochastic skew was a term first coined in a paper by Carr and Wu in connection with foreign exchange (fx) options. Consider the implied volatility of a call option with a delta of 0.25 (and a given maturity) and subtract from it the implied volatility of a put option with a delta of -0.25 (and the same maturity). This is called the 25-delta risk-reversal and is very actively traded in the fx options markets.
- Carr and Wu document that, not only the magnitude, but, also, the sign of the 25-delta risk-reversal changes through time for major currency pairs. This is in contrast to the equity options markets where risk-reversals are always negative. Since the risk-reversal is essentially a measure of the skewness of the (risk-neutral) distribution, we see that in order to capture this feature, we need the instantaneous correlation between $d(\log S)$ and $d\Sigma$ to be able to change sign.

- Most models cannot capture this. For example, in the Heston (1993) model (i.e. single Heston), the instantaneous correlation between $d(\log S)$ and $d\Sigma$ always has the sign of the ρ parameter and cannot change sign (since ρ is a constant).
- In their paper, Carr and Wu propose a class of stochastic processes which can capture stochastic skew of which the double Heston model is a special case.
- The ability of the double Heston model to capture stochastic skew is a very attractive feature.
- Taking everything into account, the double Heston model is very appealing.
- Its main drawback is that one is very reliant on solving a three-factor (plus time) PDE numerically or Monte Carlo simulation for the pricing of exotic options (both of which will be relatively slow).

- The models of Lipton and McGhee and of Kainth could broadly be described as: “Make the stochastic process richer, more complicated with more parameters and hope that, eventually, you have enough degrees of freedom to match the market prices of both vanilla and barrier fx options”. In fact, this is broadly the approach of nearly all models in the financial literature.
- By contrast, the paper by Carr and Crosby tries a rather different line-of-attack.
- Firstly, we define some notation. We denote the initial time (today) by t_0 and denote calendar time by t , $t \geq t_0$. We denote the spot fx rate, at time t , by $S(t)$. Domestic (respectively, foreign) interest-rates are denoted by r_d (respectively, r_f).
- We introduce levels L and U (with $L < S(t_0) < U$). Intuitively, these correspond to the lower and upper barrier levels of DNT options to which the model will be calibrated (that statement provides the intuition - the model is actually more flexible than that).
- We refer to the region (L, U) as the “corridor”.

- We denote the first exit time of the spot FX rate from the corridor by τ i.e. we define:
$$\tau = \inf \{t : S(t) \leq L \text{ or } S(t) \geq U, t \geq t_0\}.$$

If the spot FX rate $S(t)$, $t \geq t_0$, has always been strictly between L and U i.e. the spot FX rate has never exited from the corridor, then the convention is that we set $\tau = \infty$.
- We specify the dynamics of the spot FX rate $S(t)$ by specifying two auxiliary stochastic processes $S_1(t)$ and $S_2(t)$.
- We assume that the dynamics, under the risk-neutral measure, of the spot FX rate $S(t)$, at time t , for $t \in (t_0, \tau)$ are: $S(t) = S_1(t)$, where $S_1(t)$, for all $t \in [t_0, \infty)$, follows an arbitrage-free stochastic process with $S_1(t_0) = S(t_0)$. Note that $S(t_0)$ and $S_1(t_0)$ are equal and known at time t_0 and that, for all $t \in [t_0, \tau)$, the spot FX rate satisfies:

$$E_{t_0}[S(t)] = E_{t_0}[S_1(t)] = S_1(t_0) \exp((r_d - r_f)(t - t_0)) = S(t_0) \exp((r_d - r_f)(t - t_0)).$$

- The dynamics of the spot FX rate, for all $t \in [t_0, \infty)$, are constructed as follows: We have already defined the stochastic process $S_1(t)$, for all $t \in [t_0, \infty)$. Now we define $S_2(t)$ to be an arbitrage-free stochastic process. In particular, for all $t \geq \tau$,

$$E_\tau[S_2(t)] = S_2(\tau) \exp((r_d - r_f)(t - \tau)).$$

Furthermore, we require:

$$S_2(\tau) = S_1(\tau).$$

We assume that the dynamics, under the risk-neutral measure, of the spot FX rate $S(t)$, at time t , for $t \in [\tau, \infty)$, are:

$$S(t) = S_2(t).$$

- Hence, it is clear that the dynamics of the spot FX rate $S(t)$, for all $t \in [t_0, \infty)$, are of the form:

$$S(t) = S_1(t)I(t < \tau) + S_2(t)I(t \geq \tau),$$

where $I(\bullet)$ denotes the indicator function. The dynamics are those of a mixture model but with random weights governed by the stopping time τ . Note, it can easily be verified using the optional stopping theorem that for all $t \in [t_0, \infty)$,

$$E_{t_0}[S(t)] = S(t_0) \exp((r_d - r_f)(t - t_0)),$$

which is required in the absence of arbitrage.

- In words, the key idea of Carr and Crosby is that the spot fx rate follows one (arbitrage-free) stochastic process while the spot fx rate has never exited from the corridor. From the instant, $t = \tau$, at which the spot fx rate touches or breaches either the lower or the upper barrier level for the first time, the spot fx rate follows a different (arbitrage-free) stochastic process.

- The change in dynamics of the spot fx rate is a one-off change in the sense that if the spot fx rate later returns to the corridor (L, U) , the dynamics of the spot fx rate do NOT return to those of the first stochastic process $S_1(t)$.
- The condition $S_2(\tau) = S_1(\tau)$ means, in words, that “the second stochastic process $S_2(t)$ starts off from where the first stochastic process $S_1(t)$ left it”. In particular, it rules out ridiculous behaviour such as the spot fx rate breaches the upper barrier level U and then the second stochastic process starts off from below the lower barrier level L .
- The definition of the spot fx rate in terms of the auxiliary stochastic processes $S_1(t)$ and $S_2(t)$ just formalises the key idea that there is a one-off change in the dynamics of the spot fx rate at the first exit time from the corridor.

- The reason for defining the auxiliary stochastic processes $S_1(t)$ and $S_2(t)$ is to make more formal the definition of the model. In particular, Carr and Crosby allow the stochastic processes $S_1(t)$ and $S_2(t)$ to be jump processes. Specifically, this means that although $S_2(\tau) = S_1(\tau)$ is certainly a requirement, it is not necessarily the case that $S_1(\tau) = L$ or $S_1(\tau) = U$. This is because it is possible for the stochastic process $S_1(t)$, and hence the spot fx rate, to jump straight through (“overshoot”) either the lower barrier level L or the upper barrier level U .
- Hence, it is possible that $S_1(\tau) < L$ or $S_1(\tau) > U$.
- The possibility of “overshoot” complicates the model and means that analytical solutions are only available for certain jump processes (hyperexponential jump-diffusions (possibly time-changed)).
- Nonetheless, the model framework allows for a rich variety of dynamics for $S_1(t)$ and $S_2(t)$, including multiple jump processes and stochastic volatility (or stochastic time-changes).

- A key point is that the prices of barrier options (including DNT options and double barrier knockout options) which knockout at the barrier levels L and U are unaffected by the change in dynamics at the first exit time from the corridor (L, U) because if the spot fx rate exits from the corridor, these barrier options expire worthless.
- The model framework of Carr and Crosby allows the prices of barrier options (including DNT options and double barrier knockout options), with barriers at L and U to be determined analytically (upto Laplace inversion). Hence, we can estimate the model parameters for the stochastic process $S_1(t)$ by calibrating the model to the market prices of DNT options (provided the barrier levels are L and U or are inside (L, U)). This can be done independently of any assumptions about the nature of or the parameters of the stochastic process $S_2(t)$.

- In order to price vanilla options, note that the price of a vanilla option is the price of a double barrier knockout option plus a double barrier knockin option (“in-out parity”). The model framework allows double barrier knockout options to be priced analytically (upto Laplace inversion). Again, the prices of double barrier knockout options are unaffected by the change in dynamics at the first exit time from the corridor (L, U) because if the spot fx rate exits from the corridor, double barrier knockout options expire worthless.
- Carr and Crosby show that double barrier knockin options can also be priced analytically (for the class of stochastic processes they consider). Hence, vanilla option prices can be determined from “in-out parity”.
- Calibration of the Carr and Crosby model is in two stages:

- In the first stage, they estimate the model parameters of the stochastic process $S_1(t)$ by calibrating to the market prices of DNT options.
- They then regard these model parameters as fixed.
- They then price double barrier knockout options with the same strikes and maturities as the vanilla options to which they will calibrate in the second stage (using the model parameters just estimated).
- They subtract these double barrier knockout option prices from the market prices of the corresponding vanilla options to get the implied market prices of double barrier knockin options which they will use in the second stage.
- In the second stage, they estimate the model parameters of the stochastic process $S_2(t)$ by calibrating to the market prices of vanilla options using as an intermediate step the implied market prices of double barrier knockin options which were obtained in the first stage.

- The benefit of the methodology is that, instead of one large calibration, the problem is broken down into two independent smaller calibrations.
- Carr and Crosby demonstrate impressive fits to the market prices of both barrier and vanilla fx options.
- Calibration of the Carr and Crosby model is made faster by virtue of the fact that both vanilla and barrier options can be priced analytically (for the types of stochastic processes they considered).
- Additional positive features of the model are that it can allow for a rich variety of stochastic processes, including allowing for stochastic volatility and stochastic skew.
- Once you introduce jumps into a model, one is required to solve a PIDE (partial integro-differential equation) which is inevitably much harder than solving a PDE. The PIDE in this model would be particularly challenging. Hence, a drawback of the Carr and Crosby model is that almost all other types of exotic options (i.e. ones for which no analytical results exist) require Monte Carlo simulation.

- The key assumption in the Carr and Crosby model is that the dynamics of the spot fx rate can change at the first exit time from the corridor. With this assumption, it is possible to get excellent fits (by design) to the market prices of both vanilla and barrier fx options.
- In the paper, they provide some possible explanations and tentative reasons why the assumption of the change of dynamics at the first exit time may be justified or plausible - theoretically and empirically.



- The typical calibration problem can be described as follows:
- We are given a set of N market option prices (usually, they are all vanilla options) $C_i(K_i, T_i)$ of different strikes K_i and maturities T_i , for $i = 1, \dots, N$.
- We are given a stochastic processes with M parameters labelled $\bar{\theta} = \{\theta_1, \theta_2, \dots, \theta_M\}$ which generates N model prices $G_i(\bar{\theta}, K_i, T_i)$, for $i = 1, \dots, N$. Note $M \leq N$.
- Then we use a “solver” algorithm (simplex, simulated annealing, Newton-Raphson, Broyden’s method, etc) to find the values of the parameters $\bar{\theta}$ which minimise:

$$\sum_{i=1}^N \omega_i (G_i(\bar{\theta}, K_i, T_i) - C_i(K_i, T_i))^2$$

(One does not have to use the L^2 -norm but it is usual).

- This inverse problem is not well-posed. In other words, there may be many local minima, one of which is located by the “solver”, as opposed to the global minimum, the results may be highly sensitive to the starting point of the “solver” algorithm and small changes in $C_i(K_i, T_i)$, for some $i = 1, \dots, N$ may result in large changes in $\bar{\theta}$.
- There are surprisingly few papers or books which even acknowledge the problem here. Exceptions include chapter 13 of Cont and Tankov (2003) and Cont and Tankov (2004). Cont and Tankov use the concept of relative entropy, which, in fullness, is beyond the scope of this lecture, and they focus mostly on jump-diffusion models.
- I will be less formal than Cont and Tankov and try to suggest some “rules-of-thumb”.

- Changing the target function (i.e. the function to be minimised) may sometimes help:
- Choose the vanilla options to be puts or calls according to which is out of the money.
- One could choose the weights ω_i such that they are inversely proportional to the square of the bid-offer spreads. This gives more weight to the most liquid options. Alternatively, choose the weights ω_i such that they are inversely proportional to the square of the Black and Scholes (1973) “vegas”. This is roughly equivalent to minimising the sum of squares of differences in implied volatilities (which is more financially relevant and interesting).

- Cont and Tankov's idea of relative entropy could be summarised non-mathematically: "Use what you know from historical data eg. real-world physical measure parameter estimates or parameter estimates from the previous day's calibration".
- Suppose we are calibrating a jump-diffusion model with one Poisson process with either Merton (1976) lognormal jumps or with fixed jump amplitudes. Cont and Tankov explain how calibrating such models is particularly prone to parameter instability because as one increases the intensity rate and decreases the mean jump size, the target function can move along a flat valley.

- We expect jumps to be rare, extreme events. Therefore, we do not expect there to be 100 jumps per year on the average. We would expect a few (eg less than 5) jumps per year on the average. This gives a very rough estimate of the real-world physical measure jump intensity rate. More refined estimates could be obtained from time-series data (eg recursive filtering - which is roughly counting the number of jumps per year bigger than x standard deviations for some choice of x).
- Real-world and risk-neutral intensity rates can be completely different in theory and will typically not be the same in practice. However, with very rough estimates of risk-aversion parameters, we can roughly go from one to the other. For equities, this implies risk-neutral intensity rates will typically be higher than real-world intensity rates. For commodities, empirical evidence suggests that risk-neutral intensity rates and real-world intensity rates are typically quite close to one another.

- Therefore, it should be possible to estimate a value λ_{est} for the risk-neutral intensity rate λ , before starting the calibration. Then one can incorporate this estimate into the calibration by use of a penalty function. Specifically, we minimise:

$$(w_\lambda(\lambda_{\text{est}} - \lambda)^2) + \sum_{i=1}^N \omega_i (G_i(\bar{\theta}, K_i, T_i) - C_i(K_i, T_i))^2,$$

for some positive weight w_λ (which can be obtained with the help of a little trial-and-error based on how “tightly” we wish to keep the risk-neutral intensity rate to λ_{est}).

- One could also use time-series data to estimate the mean jump size (which if the jumps are random in size may, of course, be different in the real-world and risk-neutral measures). Indeed, in a similar vein to above, we would probably expect the estimates of the mean jump size to be equivalent (in magnitude) to, say, at least 2 standard deviation daily moves (the sign can be determined from the slope of the implied volatility as a function of strike).
- By incorporating historical data into the calibration, we increase the curvature of the target function to be minimised and so, intuitively, increase the chance of a more stable calibration.

- If we are recalibrating a model daily (as is usual), in some cases, we might use the estimate of a parameter from the previous day in a penalty function. This can be useful if we are worried about parameter instability - though at risk of deteriorating calibrations if repeated over long time periods.
- We might expect mean-reversion rates to reflect characteristic time-scales. So for example, we might expect characteristic time-scales for fx option models with mean-reverting stochastic volatility to be 3 months to 3 years while, for interest-rate derivatives models, we might expect characteristic time-scales to be 3 years to 40 years. This might translate into mean-reversion rates (at least for single-factor models with a single mean-reversion rate - we might have to modify our reasoning if there are two mean-reversion rates) of the order of magnitude of 0.33 to 4.0 for fx options models with mean-reverting stochastic volatility and of the order of magnitude of 0.025 to 0.33 for interest-rate derivatives.
- Bear in mind that strictly speaking (depending on the specifications of the market price(s) of risk), mean-reversion rates can be different in the real-world measure compared to the risk-neutral measure but given the approximations inherent in our heuristics, this seems to be the least of our worries.

- With some models, it is not uncommon, in practice, in the markets for quants or traders to simply fix mean-reversion rates to plausible levels (which is just a special case of a penalty function with an infinite weight).
- Historical time-series data may be able to give rough estimates of “vol-of-vol” parameters.

- Such parameter estimates are very rough and approximate but, nonetheless, it should be possible to “guesstimate” many (if not all) parameter values (at least to the right order of magnitude) before the calibration takes place. We can then incorporate our “guesstimates” into penalty functions.
- The penalty function need not necessarily be of the form above. For example, if we believed that a parameter θ was very unlikely to be above some value θ_{high} , we could try a penalty function of the form: $(wI_{\theta > \theta_{\text{high}}}(\theta - \theta_{\text{high}})^2)$, where I denotes the indicator function, for some choice of the positive weight w .
- Similarly, we can add penalty functions if we believe that a parameter θ is very unlikely to be below some value θ_{low} .
- All penalty functions should be smooth.

- Consider the G2++ Gaussian HJM interest-rate model in which continuously-compounded instantaneous forward-rates $f(t, T)$, at time t , to tenor T follow the SDE:

$$df(t, T) = \dots dt + \sigma_1 \exp(-\alpha_1(T - t))dz_1(t) + \sigma_2 \exp(-\alpha_2(T - t))dz_2(t),$$

where the Brownian increments $dz_1(t)$ and $dz_2(t)$ have constant correlation ρ and where σ_1 , α_1 , σ_2 , α_2 are positive constants (and for brevity we don't write down the drift term - it is easily derived from the HJM no-arbitrage condition).

- In the special case that $\alpha_1 = \alpha_2$, the model is degenerate and it would make financial sense to include a penalty function which penalises if α_1 and α_2 are too close together. In a similar vein, it may be helpful to include a penalty function which forces the choice of which of α_1 and α_2 is the larger.
- Humped volatility curves (which, in practice, are very common in the caps and swaptions market) are only possible if the parameter ρ is negative. So, if this is the case, we can penalise positive values of ρ . (Clearly, we can always penalise values of ρ outside of, say, $[-0.995, 0.995]$).

- Historical data tells us that instantaneous forward interest-rates $f(t, T_1)$ and $f(t, T_2)$ of different tenors T_1 and T_2 should be positively, but not perfectly, correlated. We can, from historical data (probably using forward LIBOR rates as proxies for continuously-compounded instantaneous forward interest-rates), estimate $\text{correl}(df(t, T_1), df(t, T_2))$ for, say, $T_1 = 1$ and $T_2 = 10$ (typical values will be around 0.4).
- We can compute $\text{correl}(df(t, T_1), df(t, T_2))$ analytically within this model and use this as a penalty function within the calibration.
- We can then use all this information when we calibrate our model parameters to the market prices of caps or European swaptions.
- The “information” contained within knowledge of $\text{correl}(df(t, T_1), df(t, T_2))$ is different from that contained within knowledge of the implied Black (1976) volatilities (which are essentially equivalent to variances). This, intuitively, should help the calibration.

- If we were very sure about the accuracy of our estimate of $\text{correl}(df(t, T_1), df(t, T_2))$, we could even use the available analytical results to eliminate one of the parameters from the calibration and in doing so fix $\text{correl}(df(t, T_1), df(t, T_2))$.
- Together with intuition on the mean-reversion rates, it is clear that we can and should give the calibration a significant amount of information in the hope of making it more stable.

- One often extends the G2++ Gaussian HJM model to allow for time-dependent (usually piecewise constant) volatilities so as to be able to, for example, exactly match the market prices of at-the-money-forward swaptions of different expiries.
- The best way to do this is to calibrate the model with constant parameters $\rho, \sigma_1, \alpha_1, \sigma_2, \alpha_2$. Then set $\beta \equiv \sigma_2/\sigma_1$. Then hold $\rho, \beta, \alpha_1, \alpha_2$ fixed. Then define the model:

$$df(t, T) = \dots dt + \eta(t)(\exp(-\alpha_1(T - t))dz_1(t) + \beta \exp(-\alpha_2(T - t))dz_2(t)),$$

where $\eta(t)$ is a time-dependent function of t .

- This parameterization preserves the stationary correlation structure between forward rates of different tenors, whereas a parameterization, for example, of the form:

$$df(t, T) = \dots dt + \sigma_1(t) \exp(-\alpha_1(T - t))dz_1(t) + \sigma_2 \exp(-\alpha_2(T - t))dz_2(t),$$

with $\sigma_1(t)$ time-dependent and σ_2 constant, does not.

- Then (keeping $\rho, \beta, \alpha_1, \alpha_2$ fixed), we solve for $\eta(t)$ by a series of one-dimensional searches matching market prices of increasing expiries sequentially.

- Consider the Heston (1993) model:

$$\begin{aligned}\frac{dS}{S} &= (r_d - r_f)dt + \sqrt{V(t)}dz_S(t), \\ dV(t) &= \kappa(\theta - V(t))dt + \epsilon\sqrt{V(t)}dz_V(t), \quad V(0) \equiv V_0, \\ \text{correl}(dz_S(t), dz_V(t)) &= \rho.\end{aligned}$$

It is not difficult to see that the value of the initial variance V_0 should be approximately equal to the implied volatility of a very short-dated at-the-money-forward option. Hence, we can use a penalty function which penalises a value of V_0 a long way from this.

- Now $V(t)$ follows a mean-reverting diffusion process with long-run reversion level θ . If we calibrate the model and find that the value of θ is very significantly different from V_0 , it would beg the question why? After all, $V(t)$ is a mean-reverting process which reverts to θ and so it should not move too far away from θ at any time, including at time 0.

- Hence, it is tempting to also use a penalty function which penalises if θ is too far from V_0 .
- However, it is a little harder to define what “too far” might mean.
- It can be shown that if $2\kappa\theta < \epsilon^2$ (usually called the Feller condition) then the process for $V(t)$ can reach zero (and given that the long run mean level is θ , it must presumably be able to move significantly above θ also). Of course, we will not know what κ and ϵ are until after we have completed the calibration.

- The last slide demands a more considered approach. Suppose we have the last D consecutive working days data for the market prices of N options. We denote the market prices by $C_i(d)$, for each $d = 1, \dots, D$ and each $i = 1, \dots, N$.
- Now the parameter V_0 is different from the other parameters κ, θ, ϵ and ρ . Whereas the latter are supposed to stay constant, the parameter V_0 represents the current value of the instantaneous variance and it is supposed to vary through time. Suppose it has the values $V_0^{(1)}, V_0^{(2)}, \dots, V_0^{(D)}$ on days $d = 1, \dots, D$ respectively.
- Let us denote by $P(V(u), u, V(s), s)$ the transition probability density function for the probability of $V(s)$, at time s , transitioning to $V(u)$, at time u , for $u \geq s$. Denote by $G_i(\kappa, \theta, \epsilon, \rho, V_0^{(d)})$ the model price of the option on day d when the variance is $V_0^{(d)}$, for each $i = 1, \dots, N$. Our calibration now minimises, by choice of the $D + 4$ parameters $\kappa, \theta, \epsilon, \rho, V_0^{(1)}, V_0^{(2)}, \dots, V_0^{(D)}$:

$$\sum_{d=1}^D \sum_{i=1}^N \omega_i (G_i(\kappa, \theta, \epsilon, \rho, V_0^{(d)}) - C_i(d))^2 - \sum_{d=2}^D P(V_0^{(d)}, d, V_0^{(d-1)}, d-1).$$

- The rationale for the minus sign in front of the second summation is that we want to maximise the $P(V_0^{(d)}, d, V_0^{(d-1)}, d-1)$ terms (because that will mean that the successive values of the variance $V_0^{(d)}$ are more likely to be consistent with the underlying stochastic process for $V(t)$ which means minimising the negative of $P(V_0^{(d)}, d, V_0^{(d-1)}, d-1)$).
- A typical value of D might be between 2 and 10.
- What has this achieved?
- On the downside, we now have to value ND options rather than N and we have to minimize by choice of $D+4$ parameters rather than by choice of 5 parameters so we have made the calibration larger.
- On the upside, we can use data for D days and try to make sure that successive days calibrations are internally consistent with the stochastic process that we are trying to fit.
- Note that the final calibrated value of V_0 that we are interested in is simply $V_0^{(D)}$.

- The transition probability density function $P(V(u), u, V(s), s)$ is known in closed-form (in terms of Bessel functions) for the Heston (1993) process (see p607, chapter 13 of Lipton (2001)).
- We have illustrated this approach for the Heston (1993) model but the same approach could, for example, be used for the double Heston model or indeed any stochastic volatility model.
- It is better for this approach if the transition probability density function is known, in closed-form, for the model under consideration (as it is for Heston (1993) model) but, even if it is not known in closed-form, if the first few moments are known, then the transition probability density function could probably be satisfactorily approximated by a Edgeworth series expansion.

- Most of the possible methods described above for potentially improving the quality of calibrations could be described as ad-hoc or as fudges. Nevertheless, it is worth reflecting that the whole concept of calibrating a model by solving an ill-posed and unstable inverse problem is also unsatisfactory.
- Taken in this context, and given the importance of calibration, it seems reasonable to use every possible piece of information at our disposal.

- Pricing a book of exotic options is not easy. The perfect model does not exist - even for fx which I would say is the simplest-to-model asset class.
- A lot of lessons have been learnt by quants the hard way and a lot of models have been arrived at by trial and error.

- Carr P. and J. Crosby (2010) “A class of Levy process models with almost exact calibration to both barrier and vanilla fx options” *Quantitative Finance*, Vol. 10, No. 10, p1115-1136. Also available at <http://www.john-crosby.co.uk>.
- Carr P. and L. Wu (2007) “Stochastic skew in currency options” *Journal of Financial Economics* Vol. 86 (October 2007) p213-247.
- Cont R. and P. Tankov (2003) “Financial modelling with jump processes” Chapman and Hall, CRC Press, 2003.
- Cont R. and P. Tankov (2004) “Nonparametric calibration of jump-diffusion option pricing models” *Journal of Computational Finance*, Vol. 7 p1-49.
- Kainth D. (2007) “Pricing FX options under multi-factor Heston” Presentation given at the ICBI Global Derivatives conference, Paris, 23rd May 2007.

- Lipton A. (2001) “Mathematical methods for foreign exchange” World Scientific
- Lipton A. and W. McGhee (2002) “Universal barriers” Risk Vol. 15 (May 2002) p81-85.
- Schoutens W., E. Simons and J. Tistaert (2004) “A Perfect Calibration! Now What?”, Wilmott magazine, 2004 or see Wim Schouten’s website
- Schoutens W., E. Simons and J. Tistaert (2005) “Model risk for exotic and moment derivatives” in “Exotic options and advanced Levy models” A. Kyprianou, W. Schoutens and P. Wilmott (eds) Wiley