Lecture 8: Local Volatility Models: Implications

- Practical calibration of local volatility models
- Implied trinomial trees
- Implications:

  The deltas of standard options.

  The values of exotic options: barriers, lookbacks, etc.

8.1 Practical Calibration of Local Vol Models

In practice we are given implied volatilities $\Sigma(K_i, T_i)$ over a range of discrete strikes and expirations, and must calibrate a smooth local volatility function to these discretely specified values. Earlier we mentioned that this is an ill-posed problem, and finding methods to solve it are critically important to the practical use of local volatility models. To use Dupire's equation, we need a smooth implied volatility surface that is at least twice differentiable. We must therefore create a smooth implied volatility surface.

The most straightforward way to do this is to write down a smooth parametric form for the implied volatilities, and then compute the parameters that minimize the distance between computed and observed standard options prices. One can then calculate the local volatilities by taking the appropriate derivatives of the implieds. One difficulty with this method is how to determine a realistic form of the parametrization, particularly on the wings where prices are hard to obtain. Wilmott's book has one parametrization.There are a variety of other papers on this topic, some of them mentioned in Chapter 4 of Fengler's book.

The method illustrated here will be semi parametric. The idea is to smooth the variations in market implied volatilities by averaging the data in a series of small contiguous and overlapping regions using a parametric smoothing function. One can again determine the resultant local volatilities from the theoretical relation between smooth differentiable implieds and their derivatives. Here is an example.

Let $\{x_i, y_i\}^n_{i=1}$ represent the discrete implied volatility data for a given expiration, where x_i is the moneyness, i.e. strike/spot for each option, and y_i is the corresponding implied volatility. The aim is to find a smoothed regression

$$y_i = m(x_i) + \varepsilon_i \qquad \text{Eq.8.1}$$

where $m(x)$ is a smooth function with second derivatives. and ε_i is the error.

We then estimate $m(x)$ by writing

$$m(x) = \sum_{i=1}^{n} w_{i,n}(x) y_i \qquad \text{Eq.8.2}$$

where $w_{i,n}$ are n weight functions that sum to 1, and each $w_{i,n}$ peaks around the corresponding moneyness x_i so as to give higher weight to volatilities y_i

closer to the moneyness x_i that corresponds to a particular y_i. Any argument x in the function $m(x)$ gets a contribution from all $\{x_i, y_i\}^n_{i=1}$, with the greatest contribution coming from those x_i closest to x.

As an example, we can choose

$$w_{i,n}(x) = \frac{K_h(x - x_i)}{\sum_{i=1}^{n} K_h(x - x_i)} \qquad \text{Eq.8.3}$$

where $K_h(u)$ is a function that peaks around zero with a degree of peaking determined by h. One example is the Gaussian with standard deviation h,

$$K_h(u) = \frac{1}{h}\frac{1}{\sqrt{2\pi}}e^{-u^2/2h^2} \qquad \text{Eq.8.4}$$

a function which integrates to 1. Small h produces greater localization of the smoothing, As $h \to 0$, all smoothing vanishes and the function is defined only at the observed moneyness values. The greater the number of observed implied volatilities n, the greater the density of information, and so the more information there is in a small region around the moneyness x, and so one can choose a smaller h and still obtain smoothing.

One can show that this *Nadaraya-Watson* estimator for $m(x)$ converges to the true regression function as $h \to 0$ and $n \to \infty$ with their product kept finite. One can also show that minimizing the weighted squares of the differences between the observed volatilities and the estimated volatilities, where the weights are given by Equation 8.4, leads to the solution Equation 8.3. Fengler discuss how to choose the $K_h(u)$ so as to minimize the bias between the true regression and the smoothed estimator while avoiding the oversmoothing that makes the estimator function follow every wiggle in the data.

Fengler's Chapter 4 provides much more information on this method.

8.2 Trinomial Trees with Constant Volatility

Trinomial trees provide another discrete representation of stock price movement, analogous to binomial trees[1]. Their advantage is a greater flexibility in the description of the implied stochastic process for the stock price in discrete steps, so that one can avoid arbitrage violations more easily.

Both trinomial and binomial trees are simple discrete methods of solving the partial differential equation for the options valuation model. An initial reference on trinomial trees is the paper by Derman, Kani and Chriss, *Implied Trinomial Trees of the Volatility Smile*, Journal of Derivatives, 3(4) (1996) pp 7-22; a version of this is on my web site, and the appendix of that paper has describes the construction and calibration of trinomial trees. Some of the notes below are taken from there. Other references are the book by Clewlow and Strickland, and the book by Espen Haug. Rebonato's book also has some material on this.

Binomial and trinomial trees are merely special instances of more general methods of solving partial differential equations, some of which may be much more efficient. Wilmott has a thorough and more general discussion of these methods.

We want to model the risk-neutral process

$$\frac{dS}{S} = rdt + \sigma dZ \quad \text{or} \quad d\ln S = \left(r - \frac{\sigma^2}{2}\right)dt + \sigma dZ.$$

Figure 8.1 below illustrates a single time step in a trinomial tree.

The stock price at the beginning of the time step is S. During this time step the stock price can move to one of three nodes: with probability p to the *up* node, value S_u; with probability q to the *down* node, value S_d; and with probability $1 - p - q$ to the middle node, value S_m. At the end of the time step, there are five unknown parameters: the two probabilities p and q, and the three node prices S_u, S_m and S_d.

There are two conditions – on the mean and the variance of the process – that must be satisfied in order for the tree to represent geometric Brownian motion in the continuum limit. First, for a *risk-neutral* trinomial tree, as in the binomial case, the expected value of the stock at the end of the period must be its forward price $F = Se^{(r-\delta)\Delta t}$, where δ is the dividend yield. Therefore:

1.Both trinomial and binomial trees approach the same continuous time theory as the number of periods in each is allowed to grow without limit. Nevertheless, one kind of tree may sometimes be more *convenient* than another when you are working in discrete time, before you reach the continuous limit.

FIGURE 8.1. In a single time step of a trinomial tree the stock price can move to one of three possible future values, each with its respective probability. The three transition probabilities sum to one.

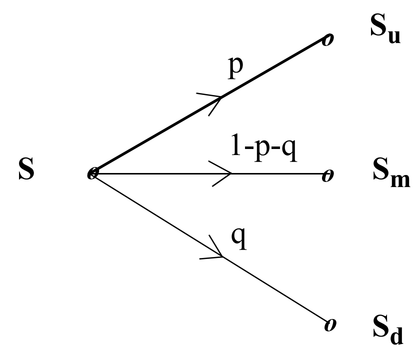

$$pS_u + qS_d + (1-p-q)S_m = F \qquad \text{Eq.8.5}$$

Second, if the stock price volatility during this time period is σ, then the node prices and transition probabilities must produce the appropriate variance, so that

$$p(S_u - F)^2 + q(S_d - F)^2 + (1-p-q)(S_m - F)^2 = S^2\sigma^2\Delta t + O(\Delta t^2) \qquad \text{Eq.8.6}$$

where $O(\Delta t^2)$ denotes terms of higher order than Δt which vanish more rapidly as we approach the continuum limit. Different discretizations of risk-neutral trinomial trees have different higher order terms in Equation 8.6. They all become negligible in the continuum limit.

Because there are two constraints on five parameters in the tree, one has much more flexibility in building the tree. In contrast, in the binomial case, the mean and variance conditions determined the location of the nodes and the risk-neutral probability with no flexibility in avoiding arbitrage violations.

Figure 8.2 below illustrates two methods of combining binomial trees to produce a trinomial tree.

Because trinomial trees are more general there are more ways to build them. Figure 8.3 illustrates a trinomial tree for the $\ln S$ that's chosen to be more symmetric. Because of the symmetry, we have to solve only for ε and q in order to match the mean and variance of $\ln S/S_0$ over time Δt. To make the tree even simpler, we choose $m = \left(r - \frac{\sigma^2}{2}\right)\Delta t$ so that the central node always coincides

FIGURE 8.2. Two equivalent methods for building constant volatility trinomial trees with spacing $\Delta\tau$. (a) Combining two steps of a CRR binomial tree with a spacing of $\Delta\tau/2$. (b) Combining two steps of a JR binomial tree with spacing $\Delta\tau/2$.

(a) Combining two steps of a Cox-Ross-Rubinstein binomial tree

$$S_u = Se^{\sigma\sqrt{2\Delta t}}$$

$$S_m = S$$

$$Sd = Se^{-\sigma\sqrt{2\Delta t}}$$

$$p = \left(\frac{e^{r\Delta t/2} - e^{-\sigma\sqrt{\Delta t/2}}}{e^{\sigma\sqrt{\Delta t/2}} - e^{-\sigma\sqrt{\Delta t/2}}}\right)^2$$

$$q = \left(\frac{e^{\sigma\sqrt{\Delta t/2}} - e^{r\Delta t/2}}{e^{\sigma\sqrt{\Delta t/2}} - e^{-\sigma\sqrt{\Delta t/2}}}\right)^2$$

(b) Combining two steps of a Jarrow-Rudd binomial tree

$$S_u = Se^{(r - \sigma^2/2)\Delta t + \sigma\sqrt{2\Delta t}}$$

$$S_m = Se^{(r - \sigma^2/2)\Delta t}$$

$$S_d = Se^{(r - \sigma^2/2)\Delta t - \sigma\sqrt{2\Delta t}}$$

$$p = 1/4$$

$$q = 1/4$$

FIGURE 8.3. In a single time step of a trinomial tree the stock price can move to one of three possible future values, each with its respective probability. The three transition probabilities sum to one. We draw the log of the stock price here.

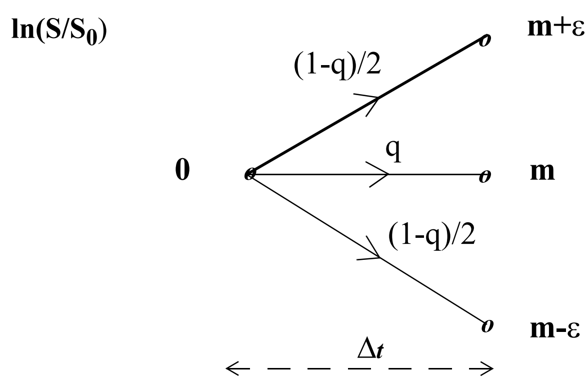

with the expected value of $\ln S/S_0$ and we also choose the probabilities to be symmetric about the center.

The expected value of the log term is then exactly *m*, since the probabilities are symmetric. To get the variance of returns right we must have

$$(1-q)\varepsilon^2 \approx \sigma^2 \Delta t$$

or

$$\varepsilon = \sigma\sqrt{\frac{\Delta t}{1-q}} \qquad \text{Eq.8.7}$$

It's often convenient to choose q = 2/3. Then the multiplicative factors for the stock become

$$\begin{aligned} M &= e^{\left(r - \frac{\sigma^2}{2}\right)\Delta t} \\ U &= Me^{\sigma\sqrt{3\Delta t}} \\ D &= Me^{-\sigma\sqrt{3\Delta t}} \end{aligned} \qquad \text{Eq.8.8}$$

This is accurate only to O(Δt), but in the limit as the spacing goes to zero, higher order terms become negligible.

Figure 8.4 illustrates a risk-neutral trinomial tree with constant volatility.

FIGURE 8.4. Example of a risk-neutral trinomial tree with constant volatility

Risk-neutral trinomial tree with constant volatility

r continuous	0.1	0.1	0.1	0.1	0.1	0.1	
f	1.0101	1.0101	1.0101	1.0101	1.0101	1.0101	
dt	0.1000	0.1000	0.1000	0.1000	0.1000	0.1000	
sig	0.2000	0.2000	0.2000	0.2000	0.2000	0.2000	
m	1.0080	1.0080	1.0080	1.0080	1.0080	1.0080	**exp(r-sig^2/2)dt with prob 2/3**
u	1.1247	1.1247	1.1247	1.1247	1.1247	1.1247	**m*exp(sig*sqrt(3dt)) with prob 1/6**
d	0.9034	0.9034	0.9034	0.9034	0.9034	0.9034	**m*exp(-sig*sqrt(3dt)) with prob 1/6**

stock

				160.0279
			142.2810	143.4238
		126.5021	127.5182	128.5425
	112.4732	113.3766	114.2872	115.2052
100.0000	100.8032	101.6129	102.4290	103.2518
	90.3441	91.0697	91.8012	92.5386
		81.6206	82.2761	82.9370
			73.7394	74.3316
				66.6192

pv of stock

			142.2810
		126.5021	127.5182
	112.4732	113.3766	114.2872
100.0000	100.8032	101.6129	102.4290
	90.3441	91.0697	91.8012
		81.6206	82.2761
			73.7394

strike					
100.0000				60.0279	
call option			43.2760	43.4238	
		28.4823	28.5132	28.5425	**C=[1/6(up)+2/3(middle)+1/6(dn)]/f**
	15.9075	15.5599	15.2822	15.2052	
7.1968	6.5036	5.6828	4.6552	3.2518	
	1.6931	1.1223	0.5366	0.0000	
		0.0885	0.0000	0.0000	
			0.0000	0.0000	
				0.0000	

8.3 Trinomial Trees with Local Volatility σ(S,t)

In dealing with binomial local volatility trees, we discovered that for finite Δt calibrating a binomial tree to a variable local volatility sometimes lead to negative probabilities or violations of the no-arbitrage principle. For trinomial trees, we will show that the calibration to local volatilities can be done by adjusting the probabilities after the stock price nodes are chosen independently, thereby more easily avoiding negative probabilities.

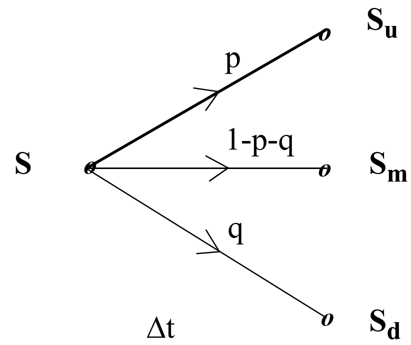

In the figure at right, the conditions to satisfy are

$$pS_u + (1-p-q)S_m + qS_d = F$$
$$p(S_u - F)^2 + (1-p-q)(S_m - F)^2 + q(S_d - F)^2 \approx S^2\sigma^2\Delta t \qquad \text{Eq.8.9}$$

To make life easier, we choose $S_m \equiv F$, so the middle node always coincides with the forward. Then the equations above simplifies to

$$pS_u + qS_d = (p+q)F$$
$$p(S_u - F)^2 + q(S_d - F)^2 \approx S^2\sigma^2\Delta t \qquad \text{Eq.8.10}$$

Given the nodes S_u and S_d, we can solve for p and q:

$$p = \frac{S^2\sigma^2\Delta t}{(S_u - F)(S_u - S_d)}$$
$$q = \frac{S^2\sigma^2\Delta t}{(F - S_d)(S_u - S_d)} \qquad \text{Eq.8.11}$$

We can therefore choose a grid of stock prices in the future that allows us to determine p's and q's that lie strictly between 0 and 1 and still match the correct forward and variance.

Below are two examples of trees built with different grids and that lead to different probabilities p and q on the tree, but will nevertheless produce the same options prices in the limit as $\Delta t \to 0$. We can first choose the grid and then determine the probabilities. In the example below we choose stock prices that lie on an initial grid formed simply by using a CRR stock price generators.

We choose $U = \exp(\sigma_g \sqrt{2\Delta t})$ and $D = \exp(-\sigma_g \sqrt{2\Delta t})$, Note that volatility σ_g (the *generator* volatility) used to generate the grid is not the true local volatility, but just some constant (convenient, fake, approximately representative local) volatility used to generate the lattice of prices.

Here below is a risk-neutral trinomial tree with local volatility $\sigma(S) = 0.1 + (S/100 - 1)$ built on a lattice generated with a 15% CRR volatility of prices.

Risk-neutral trinomial tree with constant volatility

r continuous	0			
f	1.0000			
dt	0.0100			
local sig	a+b(S/100 -	a	0.1000 b	1.0000
For generation of initial lattice				
vol generator	0.1500			
U	1.0214 UD=M^2 to close tree			
D	0.9790			
M	1.0000			

$$\sigma(S) = 0.1 + (S/100 - 1)$$
$$U = \exp(\sigma\sqrt{2\Delta t}) = \exp(0.15\sqrt{0.02}) = 1.0214$$

stock state space

				108.8557
			106.5708	106.5708
		104.3339	104.3339	104.3339
	102.1440	102.1440	102.1440	102.1440
100.0000	100.0000	100.0000	100.0000	100.0000
	97.9010	97.9010	97.9010	97.9010
		95.8461	95.8461	95.8461
			93.8343	93.8343
				91.8647

Jarrow-Rudd generated lattice with 15% volatility

p_up S^2*sig^2*dt/((S_u - F)(S_u - S_d))

			0.3019
		0.2259	0.2259
	0.1621	0.1621	0.1621
0.1099	0.1099	0.1099	0.1099
	0.0686	0.0686	0.0686
		0.0376	0.0376
			0.0162

1-p-q

			0.3898
		0.5434	0.5434
	0.6723	0.6723	0.6723
0.7778	0.7778	0.7778	0.7778
	0.8613	0.8613	0.8613
		0.9241	0.9241
			0.9673

q_dn S^2*sig^2*dt/((F-S_d)(S_u - S_d))

			0.3083
		0.2307	0.2307
	0.1656	0.1656	0.1656
0.1123	0.1123	0.1123	0.1123
	0.0701	0.0701	0.0701
		0.0384	0.0384
			0.0165

strike
102.0000
call option

				6.8557
			4.5708	4.5708
		2.4103	2.3339	2.3339
	0.8723	0.7004	0.4752	0.1440
0.1955	0.1273	0.0645	0.0158	0.0000
	0.0054	0.0011	0.0000	0.0000
		0.0000	0.0000	0.0000
			0.0000	0.0000
				0.0000

discounted call value for strike 102

Below is another risk-neutral trinomial tree built on 20% vol-generating lattice with local volatility = 0.15 - 0.1(S/100 - 1) and time steps of one year, just as an example. Of course, for accurate convergence to the continuous time limit, one needs much smaller time steps.

Risk-neutral trinomial tree

r continuous		0.04879			
f		1.0500			
dt		1.0000			
local sig	a+b(S/100 -	a	0.1500	b	-0.1000
For generation of initial lattice					
vol generator		**0.2000**			
U		1.3932	UD=M^2 to close tree		
D		0.7913			
M		1.0500			

$$\sigma(S) = 0.15 - 0.1(S/100 - 1)$$

stock state space

				376.7949	104.999983
			270.4449	283.9671	
		194.1121	203.8176	214.0085	
	139.3241	146.2903	153.6048	161.2850	
100.0000	105.0000	110.2500	115.7624	121.5505	
	79.1320	83.0886	87.2430	91.6051	
		62.6188	65.7497	69.0371	
			49.5515	52.0290	
				39.2111	

Jarrow-Rudd generated lattice with 20% volatility

p_up S^2*sig^2*dt/((S_u - F)(S_u - S_d))

			0.0020
		0.0151	0.0103
	0.0593	0.0521	0.0450
0.1089	0.1018	0.0945	0.0872
	0.1413	0.1348	0.1282
		0.1699	0.1643
			0.1945

1-p-q

			0.9953
		0.9648	0.9760
	0.8620	0.8789	0.8953
0.7466	0.7632	0.7800	0.7971
	0.6712	0.6862	0.7017
		0.6046	0.6177
			0.5475

q_dn S^2*sig^2*dt/((F-S_d)(S_u - S_d))

			0.0027
		0.0201	0.0137
	0.0787	0.0691	0.0597
0.1445	0.1350	0.1254	0.1157
	0.1875	0.1789	0.1701
		0.2255	0.2180
			0.2581

adjust probabilities

strike

102.0000				274.7949
call option			173.3020	181.9671
		101.5950	106.6748	112.0085
	51.4482	53.8486	56.4619	59.2850
20.5266	20.3209	20.0519	19.7653	19.5505
	5.3875	4.0986	2.3873	0.0000
		0.3864	0.0000	0.0000
			0.0000	0.0000
				0.0000

discounted call value for strike 102 is 20.5266

Here is one more risk-neutral trinomial tree built on a 13% vol-generating lattice: stock prices are different, probabilities are different, but options prices are about the same, and will be identical as $\Delta t \to 0$

Risk-neutral trinomial tree

r continuous		0.04879			
f		1.0500			
dt		1.0000			
local sig	a+b(S/100 -	a	0.1500	b	-0.1000
For generation of initial lattice					
vol generator		**0.1300**			
U		1.2619	UD=M^2 to close tree		
D		0.8737			
M		1.0500			

$$\sigma(S) = 0.15 - 0.1(S/100 - 1)$$

stock state space

				253.5906	104.999983
			200.9555	211.0032	
		159.2453	167.2075	175.5679	
	126.1924	132.5020	139.1271	146.0834	
100.0000	105.0000	110.2500	115.7624	121.5505	
	87.3665	91.7349	96.3216	101.1376	
		76.3291	80.1456	84.1528	
			66.6861	70.0204	
				58.2614	

Jarrow-Rudd generated lattice with 13% volatility

p_up S^2*sig^2*dt/((S_u - F)(S_u - S_d))

			0.0292
		0.1001	0.0833
	0.1863	0.1678	0.1494
0.2735	0.2555	0.2374	0.2190
	0.3215	0.3044	0.2870
		0.3666	0.3506
			0.4084

1-p-q

			0.9356
		0.7796	0.8166
	0.5898	0.6306	0.6710
0.3979	0.4374	0.4774	0.5178
	0.2922	0.3297	0.3680
		0.1929	0.2280
			0.1008

q_dn S^2*sig^2*dt/((F-S_d)(S_u - S_d))

			0.0351
		0.1203	0.1001
	0.2239	0.2017	0.1796
0.3286	0.3071	0.2853	0.2632
	0.3863	0.3659	0.3450
		0.4406	0.4214
			0.4908

adjust probabilities

strike

102.0000				151.5906
call option			103.8126	109.0032
		66.7283	70.0647	73.5679
	38.4824	40.0265	41.9842	44.0834
20.2896	19.9541	19.5063	18.8357	19.5505
	8.6451	7.1391	5.3443	0.0000
		1.8658	0.0000	0.0000
			0.0000	0.0000
				0.0000

discounted call value for strike is 20.2896

Finally, here is a trinomial tree built on a 5% volatility-generating lattice This generating volatility is too small to properly match or represent the much larger local volatilities generated by the formula, and so the nodes are not far enough apart to give probabilities that lie between 0 and 1.This is an illustration of a lattice that doesn't work. But, because of the greater number of degrees of freedom in building trinomial trees, one can always find a lattice that doesn't violate the no-arbitrage principle.

Risk-neutral trinomial tree

r continuous	0.04879				
f	1.0500				
dt	1.0000				
local sig	a+b(S/100 -	a	0.1500	b	-0.1000
For generation of initial lattice					
vol generator	**0.0500**				
U	1.1269	UD=M^2 to close tree			
D	0.9783				
M	1.0500				

$$\sigma(S) = 0.15 - 0.1(S/100$$

stock state space

				161.2850	104.999983
			143.1184	150.2743	
		126.9980	133.3479	140.0153	
	112.6934	118.3281	124.2444	130.4566	
100.0000	105.0000	110.2500	115.7624	121.5505	
	97.8318	102.7234	107.8595	113.2525	
		95.7106	100.4961	105.5209	
			93.6354	98.3171	
				91.6051	

Jarrow-Rudd generatec with 5% volatility

p_up S^2*sig^2*dt/((S_u - F)(S_u - S_d))

			0.9991
		1.3232	1.1901
	1.6489	1.5164	1.3831
1.9679	1.8389	1.7081	1.5760
	2.0252	1.8971	1.7671
		2.0820	1.9549
			2.1384

1-p-q

		-1.74
	-2.4186	-2.14
-3.0799	-2.8125	-2.54
	-3.1987	-2.93
		-3.31

q_dn S^2*sig^2*dt/((F-S_d)(S_u - S_d))

			1.0723
		1.4202	1.2773
	1.7697	1.6275	1.4845
2.1121	1.9736	1.8333	1.6915
	2.1736	2.0361	1.8965
		2.2346	2.0981
			2.2951

adjust probabilities:
ARBITRAGE
VIOLATIONS

strike
102.0000
call option

				59.2850
			45.9755	48.2743
		34.4810	36.2050	38.0153
	24.5819	25.8110	27.1016	28.4566
-152.2057	43.7111	17.7329	18.6196	19.5505
	-34.8299	24.4765	10.7166	11.2525
		2.6732	10.7124	3.5209
			7.1706	0.0000
				0.0000

discounted call valu
is -152.2057

NONSENSE

Thus, we have more flexibility in building trinomial trees; we can choose a lattices of stock prices that don't violate arbitrage, and then adjust the probabilities to match the stochastic process, provided the lattice was reasonable. In contrast, with binomial trees, we were forced to a definite lattice which sometimes violated the no-arbitrage conditions.

8.4 Deltas and Exotics in Local Volatility Models

8.4.1 Four rules of thumb for local volatilities in the small slope at-the-money approximation:

Rule of Thumb 1: The Rule of 2: Local volatility varies with market level about twice as rapidly as implied volatility varies with strike.

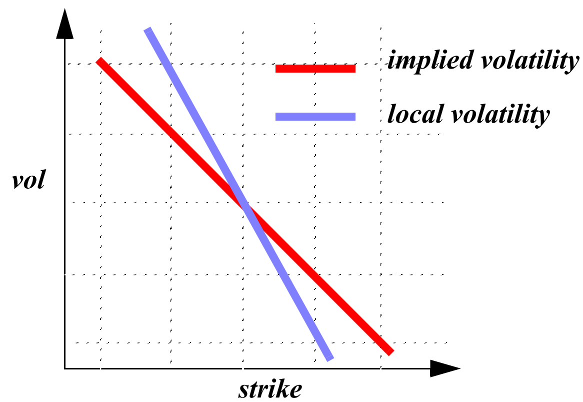

Comment: In equity markets with negative skew, the implied volatility for all strikes and maturities decrease as the market level increases.

Rule of Thumb 2: Relation between sensitivity of implied volatility to spot and strike. The change in implied volatility of a given option for a change in market level is about the same as the change in implied volatility for a change in strike level.

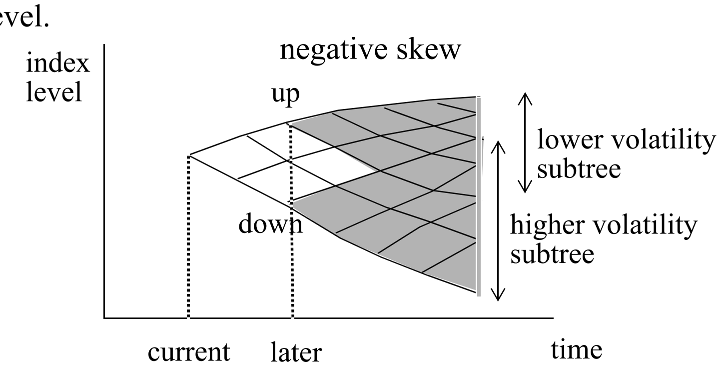

$$\Sigma(S, K) \approx \sigma_0 - \beta(S + K) + 2\beta S_0$$

Hedge ratios of standard options in the presence of (negative) skew are therefore smaller than Black-Scholes hedge ratios.

Rule of Thumb 3: The correct exposure Δ of an option is approximately given by the chain rule formula

$$\Delta = \Delta_{BS} + V_{BS} \times \beta \qquad \text{Eq.8.12}$$

For example, a one-year S&P option with a B-S hedge ratio of 60% probably has a true hedge ratio of 50%, because volatility moves down as the market moves up. Suppose S = 1000.

V_{BS} = 400 dollars; β = -0.0002 vol point per strike pt.: $V_{BS}\beta \sim 0.1$

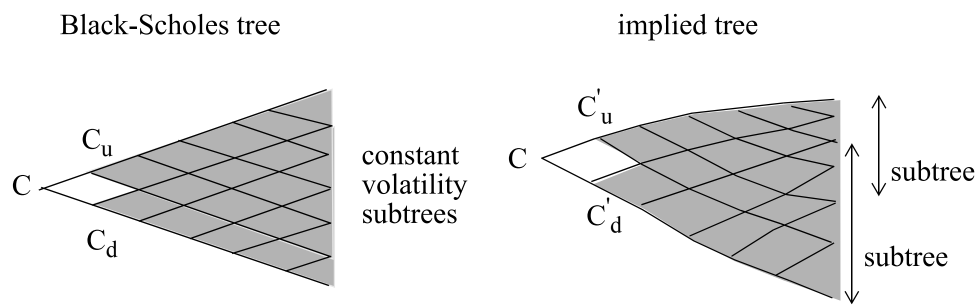

Rule 4. For short times to expiration, the inverse of the implied volatility for a given strike is the harmonic average of the local volatilities across ln(S) from spot to strike.

8.4.2 Theoretical Value of Barrier Options in Local Vol Models

In this section we illustrate the effect of local volatility models on exotic options, taking barrier options as an example. Barrier options values are especially sensitive to the risk-neutral probability of index remaining in the region between the strike and the barrier, and hence to the local volatility in this region. strike and barrier, which depends on the skew:. Here we are going to calculate their value in local volatility models and try to gain some intuition about them.

Example 1: An Up-and-Out Call. with Strike 100 and Barrier 110

In the lecture on static hedging, we showed that you can approximately replicate a down and out call by means of a long position in the call itself combined with a short position in a put whose strike is (logarithmically) reflected through the barrier. In a flat-volatility world, the value of both of these calls is determined by the constant Black-Scholes volatility. In a skewed world, however, each call has an implied volatility which is approximately the average of the local volatilities between spot and strike.

For an option with strike at 100 and barrier at 110, the reflected strike is approximately at 120. Thus, in a local volatility model, the approximate value of the Black-Scholes implied volatility for the up-and-out call is the average of the local volatilities between 100 and 120.In the figure below, the local volatility varies between 0.1 and 0.07 in this range, with an average of a about 0.085. The value of the down and out option in the local volatility model is about 1.1, which corresponds to a Black-Scholes implied volatility of about 0.09, so this intuition about averaging works reasonably.

.

LOCAL VOLATILITY AS A FUNCTION OF SPOT

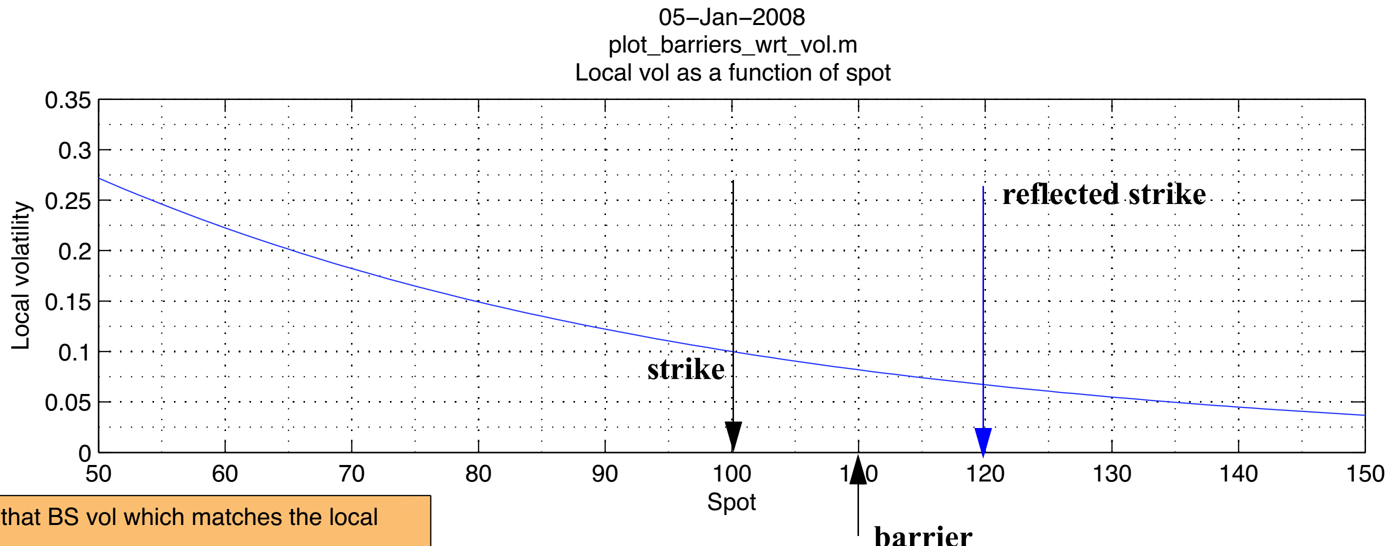

Note that BS vol which matches the local vol
price is about 0.09 which si the average of the local vols between 100 and 120, where 120 is the reflection of K in B.
(0.0 +0.10)/2 = 0.09

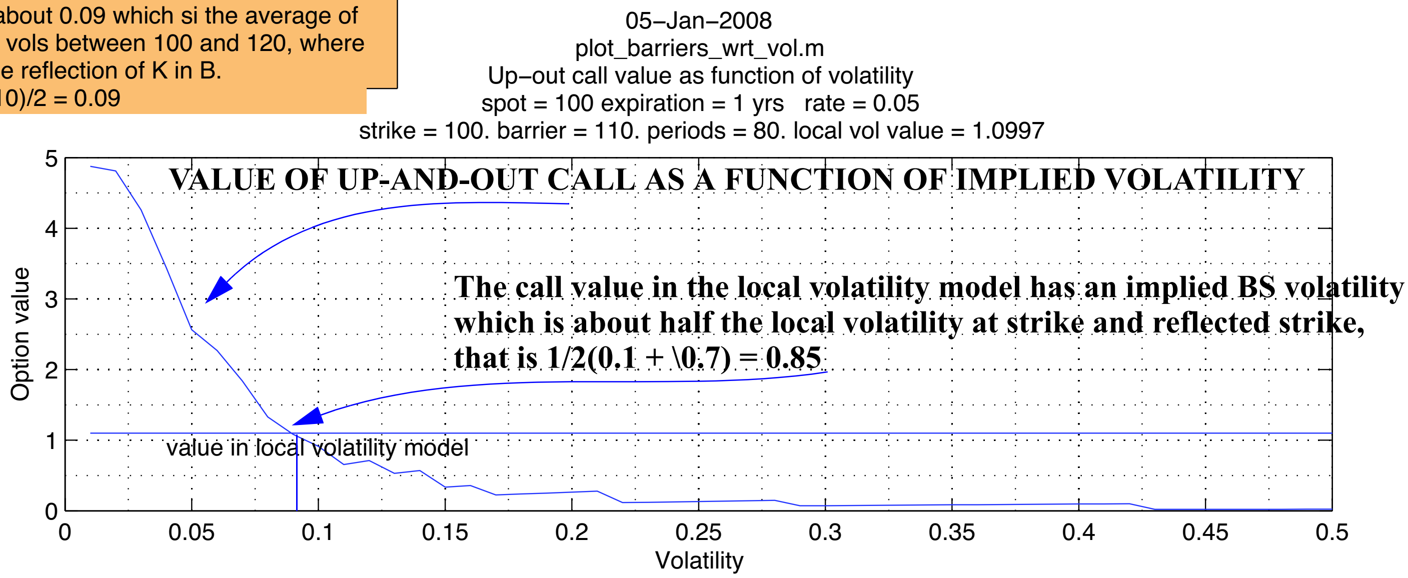

Example 1. An Up-and Out Call that has no Black-Scholes Implied Volatility

In some cases, the local volatilities can produce options values that cannot be matched by *any* Black-Scholes implied volatility. No amount of intuition can get you the exactly correct value. Consider the case below, with a spot and strike at 100, and the barrier at 130, and the skew as shown in *Figure* 8.5.

FIGURE 8.5. A hypothetical volatility skew for options of any expiration. We assume a constant riskless discount rate of 5% and a zero dividend yield. The arrows show the strike (100) and barrier (130) level of the up-and-out option under consideration.

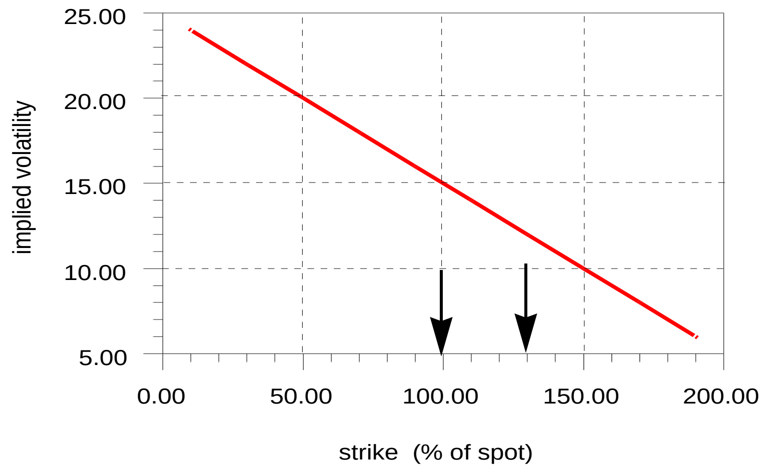

We can value the Up-and-Out Call by building an implied tree calibrated to this skew. The resultant value of the barrier option in this local volatility model is 6.46. What Black-Scholes volatility does this call price correspond to?

No skew: Up-and-Out call value as a function of Black-Scholes Implied Volatility

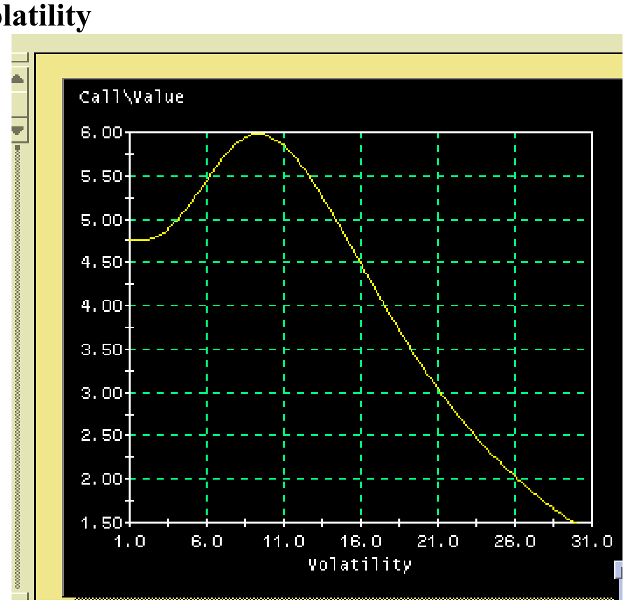

The maximum Black-Scholes value in a no-skew world is 6.00 corresponding to a 9.5% implied volatility. This value is smaller than the "correct" value in the local volatility model. There is NO Black-Scholes implied volatility which gives the local-volatility "correct" option value.

The implied volatility that comes closest to it is about 10%. We can understand this as follows.The slope of the skew is 1 vol pt. per 10 strike points. The rule of 2 then indicates that the slope of the local volatilities will be about 1 vol pt. per 5 strike points.

Now, we showed in the previous lecture that you can think of an up-and-out option with strike 100 and barrier 130 as being replicated by an ordinary call with strike 100 and a reflected call with strike 160. Therefore, the local volatility that is relevant to valuation ranges between spot prices 100 and 160 with a slope of approximately 1 vol pt. per 5 strike points, that is from values of 15% to $15-(60/5) = 3\%$. The average local volatility in this range is about 9%, which substantiates the approximate claim the implied volatility is the average of the local volatilities between spot and strike.

Local volatility models have analogous effects on the values of other exotic options, moving their values away from Black-Scholes values. Lookback calls (that pay out the final value of the index less the minimum value of the index between inception and expiration), for example, have higher deltas in a local volatility model than they do in Black-Scholes.[1]

1. Derman, Kamal, Zou: The Local Volatility Surface.