

FX Barriers with Smile Dynamics

Glyn Baker*, Reimer Beneder[†] and Alex Zilber[‡]

December 16, 2004

Abstract

Our mandate in this work has been to isolate the features of smile consistent models that are most relevant to the pricing of barrier options. We consider the two classical approaches of stochastic and (parametric) local volatility. Although neither has been particularly successful in practice, their differing qualitative features serve our exposition. By constructing approximate static hedges we are able to closely mimic their prices. The only information we require from the models, other than the initial vanilla market to which they are calibrated, is their conditional forward smile along the barrier. This strongly supports the fact that realistic smile dynamics are of paramount importance when assessing a model to be used in pricing barrier options.

*Woolgate Exchange, 25 Basinghall Street, London EC2V 5HA, glyn.baker@westlb.co.uk

[†]Erasmus University Rotterdam, Burg. Oudlaan 50, 3062 PA Rotterdam, beneder@few.eur.nl

[‡]University of Twente, FELab, P.O.Box 217, 7500 AE Enschede, a.zilber@ewi.utwente.nl

1 Introduction

Pricing barrier options using different models can yield rather variable results. In some cases, the smile corrections to the Black-Scholes price given by a stochastic volatility and a local volatility model, both calibrated to the same vanilla prices, can even go in different directions. These “contradicting” results do not arise due to different interpolations of the implied volatility surface or some other numerical reason, but due to modelling assumptions. We will clarify this point a bit further:

Vanilla prices are fully determined by the risk-neutral transition densities from the current market state (time and spot level) to all possible market states in the future. Prices of barrier options additionally depend on the transition densities from one future state to another future state. These conditional transition densities are not observed through vanilla prices. It is straightforward to construct an infinite family of models that perfectly fit the same volatility surface initially but each with different conditional transition densities¹. Hence, by calibrating a barrier pricing model to implied volatilities, we do not make sure that some very relevant information is reflected in the model.

Of course, once a model is specified and calibrated to the volatility surface, the conditional transition densities are implicitly chosen. However, as one volatility surface admits many different solutions for the transition densities, this choice is somewhat arbitrary. In fact, as argued by Ayache et al. [2], even different calibration solutions of the same model can yield significantly different conditional density structures, and consequently, very different barrier prices.

Just as transition densities today are characterized by implied volatility smiles, transition densities at future states are characterized by conditional forward smiles. We are going to present a static hedging approach that gives a fairly easy way of quantifying the impact of different forward smiles on barrier prices. This approach enables us to take model-induced forward smiles along a trigger level as the only input and return approximate barrier prices as output. In this way we are able to establish that the key qualitative feature of a barrier pricing model is the subset of forward smiles along the trigger level. This strongly supports the fact that realistic smile dynamics are of paramount importance when assessing a model to be used in pricing barrier options.

The outline of the paper is as follows: in section 2 we introduce FX options market conventions and inspect historical smile data. As will be explained later,

¹For example, a model that combines a given stochastic volatility process with a local volatility function has in general enough flexibility to fit any implied volatility surface.

this market is well-suited for the discussion of barrier pricing and forward smiles. Some of the desired properties of a smile consistent barrier-pricing model follow from the data analysis carried out in this section. After that, in section 3, we discuss the notion of conditional forward smiles that is central to our study. Then, in the next section, we introduce a local volatility and a stochastic volatility model that are needed to illustrate our point. We show how these models can be calibrated to the vanilla market and what conditional forward smiles they imply. These forward smiles will be used later to mimic barrier prices within each respective model. In Section 5, our approximate static hedge construction is introduced. It is explained how this algorithm can take any set of forward smiles along a trigger level as input and yield an approximation to a barrier option price as output. In section 6, we use the static hedge method and conditional forward smiles implied by the stochastic and local volatility model in order to obtain approximated barrier prices in these two models. The quality of the approximations is inspected by comparing them with barrier prices obtained by solving the corresponding pricing PDEs for each model. We are able to conclude that the good quality of approximations supports our main claim. Finally, in section 7, we conclude.

2 FX Options and Empirical Smile Dynamics

The FX options market is well suited for a discussion of the smile dynamics concept, since the vanilla quotes in the over-the-counter market have clear intuitive interpretation. Moreover, it provides the combination of the most liquid vanilla markets there are (such as dollar-yen and euro-dollar options markets) with a very well developed barrier options trading. This section describes the quotation conventions and contains an empirical analysis of historical smile data.

2.1 At-The-Money Straddles, Strangles and Risk Reversals

In the FX vanilla option market the directly observable quotes consist of the at-the-money straddle (ATM), the 25-delta strangle (STR), and the 25-delta risk reversal (RR). These are very liquid in most currency pairs out to two years or so and are available for various different maturities. The at-the-money straddle comprises of a call and a put with the common strike chosen so that the Black-Scholes delta of the straddle is zero. This strike is approximately equal to the forward price. With a strangle one is long both the 25-delta call and the 25-delta put. The STR is quoted as average volatility of the call and put over the ATM. The STR quote

can be directly associated with the curvature of the smile, the excess kurtosis in the terminal distribution, or the stochasticity of the ATM. With a risk-reversal one buys the 25-delta call and sells the 25-delta put. It is quoted as the difference in the 25-delta call and the 25-delta put volatilities. The RR quote can be directly associated with the slope of the smile, the excess skew in the terminal distribution, or the correlation of the ATM with spot movements.

2.2 Empirical Smile Dynamics

For 12 major currency pairs, the British Bankers Association (BBA) puts the historical data of volatility quotes at its website². Currently, there is almost two years of data available. The quotes are provided for six maturities of ATM vols: 1 week (1W), 1 month (1M), 3 months (3M), 6 months (6M), 1 year (1Y) and 2 years (2Y), and for three maturities of RR and STR quotes: 1M, 3M and 1Y.

First inspection of the data of the last two years shows that for the major currency pairs the ATMs have typically fluctuated around levels between 5% and 10%, RRs between -2% and 2% and STRs to be reasonably stable around 0.3%. Moreover, short dated maturities have in general traded at higher levels of volatility than the longer end of the curve.

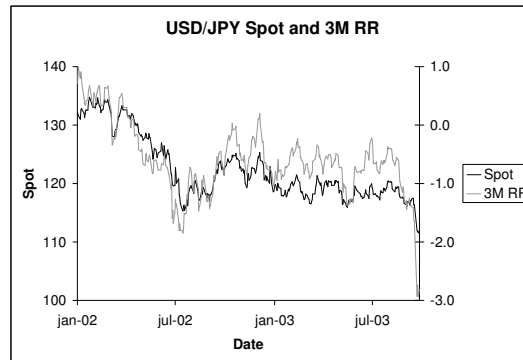


Figure 1: USD/JPY and 3M RR in 2002-2003.

Since STR quotes are very stable, our attention will be limited to the ATM and RR quotes. Figures 1 and 2 show the evolution of the USD/JPY exchange rate with the three-month RR and ATM quotes, respectively. These graphs indicate some systematic patterns. Most obvious is the correlation between changes in the

²The British Bankers Association (BBA) website can be found at <http://www.bba.org.uk/>




Figure 2: USD/JPY and 3M ATM in 2002-2003.

spot exchange rate with changes in the RR quote. The correlation is 67% here. The correlation between spot and the ATM can be significant, in this example it is -31%.

For the other currency pairs, similar analysis has been performed. It appears that RRs are always positively correlated with spot, irrespective of the currency pair. That is to say that as spot increases, the high strikes become more favorable in volatility terms. This can be observed in table 1. The table also shows that the correlation is stronger for short maturities. A recent detailed study of empirical properties of the FX options across different currency pairs can be found in Carr and Wu [6]. The authors identify the great variability of RR quotes and the strong positive correlation of changes in RR with spot returns as a unique feature of these markets.

In table 2 we provide the correlation between spot and ATM and the average RR for each currency pair. It can be observed that the average RR quote gives a good indication of the sign of the historical correlation, a negative RR (i.e. a downward sloping smile) corresponds to negative correlation and a positive RR (i.e. an upward sloping smile) corresponds to positive correlation. This result is consistent with the interpretation of the risk reversal as a measure of the correlation between spot and the ATM. If the market is expecting positive/negative correlation between the spot exchange rate and the ATM volatility, it will price this with positive/negative RR's.

Although, the empirical analysis in this section is brief, it does indicate that there are some stylized facts in the FX markets which a good model should capture. We will discuss the impact of smile dynamics on the pricing of barrier op-

	Correlation Spot/RR		
	1M	3M	1Y
AUD/USD	10%	8%	N/A
EUR/JPY	60%	57%	44%
EUR/CHF	46%	43%	33%
EUR/GBP	16%	24%	10%
EUR/SEK	38%	21%	0%
GBP/CHF	17%	7%	-1%
EUR/USD	57%	54%	40%
GBP/JPY	56%	53%	43%
GBP/USD	47%	43%	34%
USD/CAD	43%	44%	39%
USD/JPY	68%	67%	61%
USD/SEK	50%	45%	29%

Table 1: Correlation between spot and RR for major currency pairs.

	Correlation Spot/ATM			Average RR		
	1M	3M	1Y	1M	3M	1Y
AUD/USD	0%	-1%	N/A	.02	-0.03	N/A
EUR/JPY	-6%	0%	4%	-0.03	0.02	0.07
EUR/CHF	-23%	-20%	-20%	-0.33	-0.46	-0.66
EUR/GBP	12%	14%	12%	0.29	0.26	0.18
EUR/SEK	41%	41%	44%	0.83	1.08	1.48
GBP/CHF	-13%	-18%	-16%	-0.34	-0.32	-0.25
EUR/USD	34%	37%	42%	0.62	0.67	0.69
GBP/JPY	-8%	-2%	0%	-0.03	0.02	0.08
GBP/USD	30%	33%	32%	0.26	0.22	0.17
USD/CAD	-3%	-12%	-12%	-0.18	-0.17	-0.14
USD/JPY	-35%	-31%	-20%	-0.63	-0.57	-0.48
USD/SEK	-27%	-34%	-36%	-0.46	-0.55	-0.60

Table 2: Correlation between spot and ATM and average RR quote.

tions in the following sections.

3 Conditional Forward Smiles

In this section we focus on the notions of smile dynamics and conditional forward smiles. These important features strongly affect exotic options' prices in any model. However, there is no consensus in the academic literature as to these terms' exact meaning. The usual meaning of the notion "smile dynamics" is related to the way implied volatility smile evolves with spot and time. The most well-known ways of describing smile dynamics - such as sticky smile and sticky delta rules - are actually referring to the local behavior of the smile in response to the spot changes. Altogether, this leads to some confusion and therefore we find it easier to talk about possible future smiles rather than local smile dynamics.

It is well understood that the implied volatility smile observed today is nothing else than a mere reflection of the (non-lognormal) spot distribution that market participants price into vanilla options. Similarly, a smile to be observed in future will only reflect spot distributions that the market will find reasonable at this (future) moment of time. For example, a three-months-to-maturity segment of an options market is priced in accordance with the distribution of spot in three months time. Hence, asking what will the three-months volatility smile look like in half a year from now is similar to asking what the market will find (in 6 months from now) a consensus distribution for spot in even later (9 months from now) time.

Therefore, once a certain smile-consistent model is chosen, it is only natural to put oneself forward in time and, assuming a certain value for the spot process, ask what kind of volatility smile one is likely to observe. In the simplest smile model - the local volatility model - an answer to this question is straightforward, as all the conditional future smiles are determined once the model parameters are fixed. In other words, conditional on a spot level at a future time point, the volatility smile is known. However, this is not the case for stochastic volatility models, just because the conditional (on spot) future smiles will also depend on the future instantaneous volatilities. In general, conditional future smiles in non local volatility models are stochastic and are characterized by some distributions themselves. It is not clear what, say, an expectation of these future smiles is then and whether we can use it to price options. However, if we again switch to conditional densities rather than smiles, we can find a more consistent notation.

Let us introduce the concept of conditional forward (as opposed to future)

smiles by means of a simple yet meaningful example of a digital call option which knocks out if spot is below the barrier at a single intermediate observation date. This is a digital option with a European barrier feature and it is a common building block in the FX market for constructing exotic payoffs. Under the risk-neutral dynamics, the undiscounted price of this option is equal to

$$V(L, K) = E[1_{S_t > L \wedge S_T > K}] = \int_L^\infty \int_K^\infty f(S_T|S_t)f(S_t)dS_TdS_t, \quad (1)$$

where L and K denote the barrier and strike, respectively. Just as today's smile with maturity T is determined by the density $f(S_T)$, the forward smile conditional on S_t with maturity $T - t$ is determined by the conditional density $f(S_T|S_t)$. This conditional density is a “known” function in any type of model: i.e. when we know the prices for all barrier levels and strikes, then we can obtain the joint density and hence the conditional density through differentiation,

$$f(S_t, S_T) = f(S_T|S_t)f(S_t) = \frac{\partial V(L = S_t, K = S_T)}{\partial L \partial K}. \quad (2)$$

This formula can be viewed as an extension of the Breeden-Litzenberger formula that links vanilla prices to marginal densities. In this case, we link digitals with a European barrier feature to conditional densities.

Another interpretation of the forward smile is that it corresponds to the expected vanilla prices conditional on the spot level. Once a model is fully specified, these expected prices can be calculated and hence the forward smile can be backed out.

The notion of the conditional forward smile is central to this study. In the remaining sections we are going to demonstrate the importance of forward smiles in barrier pricing. For this purpose, two smile consistent models are introduced and their respective forward smiles are inspected. This will be discussed in the next section. Later on, these models and their respective forward smiles will serve our exposition.

4 Forward Smiles in a Local - and in a Stochastic Volatility Model

We consider two different smile consistent models. The first is a local volatility model (LV) with a parametric local volatility function and the second is a stochastic volatility model (SV). For both models, we will discuss the set-up, calibration, local smile dynamics and conditional forward smiles.

4.1 Local Volatility

In a local volatility model one assumes the following risk-neutral process for the exchange rate

$$dS = (r^d - r^f)Sdt + \sigma(S, t)dW, \quad (3)$$

where $\sigma(S, t)$ is a deterministic function of the spot exchange rate and calendar time t . As usual, r^d and r^f are the deterministic domestic and foreign interest rates, respectively. The local volatility function can be backed-out from a complete surface of vanilla prices as was shown by Dupire [8]. In practice however, only a limited set of option quotes will be available and, it is not trivial to construct an arbitrage-free surface of option prices from these quotes. Therefore, one approach some researchers and practitioners take is to fit an arbitrage-free model with many parameters and to use this model to generate the vanilla surface. Another approach is to specify a parametric local volatility function and to find the parameters of this function that match the option quotes. Here, we consider the second approach and we specify the following local volatility function

$$\sigma(S, t) = \sigma_0 \left(1 + \alpha \left(\frac{F_t}{F_0} - 1 \right) + \beta \left(\frac{F_t}{F_0} - 1 \right)^2 \right) S, \quad (4)$$

where $F_t = Se^{(r^d - r^f)(T-t)}$ is the forward exchange rate at time t for delivery at maturity of the underlying options. We make a simplifying assumption by calibrating to only one maturity set of vanilla quotes. It is possible to extend the model to full term structure by introducing time dependent parameters.

There are three unknown parameters in the local volatility function: σ_0 , α , and β . These parameters will correspond roughly speaking to the ATM, RR and STR quote, respectively.

A local volatility model can be calibrated by employing the Dupire forward Partial Differential Equation (PDE) as discussed in Andersen and Brotherton-Ratcliffe [1]. The Dupire result states that in a local volatility model call prices satisfy the following PDE

$$\frac{\partial C}{\partial T} + (r^d - r^f)K \frac{\partial C}{\partial K} - \frac{1}{2}\sigma(K, T)^2 K^2 \frac{\partial^2 C}{\partial K^2} + r^f C = 0, \quad (5)$$

with initial condition

$$C(T, K) = (S_0 - K)^+ \text{ at } T = 0.$$

So, when the local volatility function is known, one can solve the forward PDE to obtain call prices for all strikes and maturities.³

Table 3 contains the calibration results for three examples: it shows the correspondence between quotes and parameters.

Quotes / Parameters		Example 1		Example 2		Example 3	
ATM	σ_0	10.00%	0.097	12.00%	0.117	12.00%	0.119
RR	α	1.00%	1.73	-1.00%	-1.65	-1.00%	-1.57
STR	β	0.30%	37.9	0.30%	23.5	0.15%	12.3

Table 3: Calibration results local volatility model with $S_0 = 1.00$, $T = 0.5$, $r^d = 0.01$, $r^f = 0.03$.

In a local volatility model future smiles are deterministic conditional on the spot exchange rate. So, using the calibrated parameters, one can easily analyze the forward smiles implied by the model. E.g., in figure 3 we show the volatility smile for three-month options using the parameters corresponding to example 1.

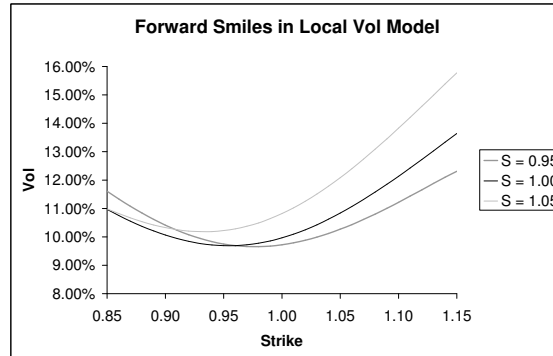


Figure 3: 3-month smile for $S_t = 0.95$, $S_t = 1.00$, $S_t = 1.05$, using parameters from example 1.

The graph shows that the volatility smile is moving in the opposite direction as the underlying exchange rate. This is a general feature of local volatility models

³In our implementation, the forward PDE is solved using a Crank-Nicholson finite difference scheme. Furthermore, we have employed Broyden's root finding method to find the three unknown parameters that match the three option quotes. This enables us to perform the calibration accurately within a second using a standard PC.

(see Hagan et al. [9]). In FX terminology, it implies that the risk reversal quote, which corresponds to the skewness of the smile, increases (decreases) when spot moves up (down). Although we observe this feature in empirical data as well, it appears that smile movements implied by a local volatility model are more extreme. Furthermore, local volatility models predict that if the spot changes significantly, smiles become less convex which contradicts empirical observations of stable strangle quotes.

4.2 Stochastic Volatility

We have chosen a simple set-up for the stochastic volatility model with instantaneous volatility following a driftless geometric Brownian motion under the risk-neutral measure. Or, formally:

$$\begin{aligned} dS &= (r^d - r^f)Sdt + \sigma SdW \\ d\sigma &= \xi\sigma dZ \end{aligned} \tag{6}$$

with correlation ρ between the two processes. The model has three unobserved parameters: initial volatility σ_0 , vol of vol ξ and correlation ρ that can be calibrated to three FX option quotes. It is possible to extend the model by including a mean-reverting drift for the volatility process or by introducing time varying parameters.

Under the specification in (6), an option price V solves the following partial differential equation:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + \rho\xi\sigma^2 S \frac{\partial^2 V}{\partial S \partial \sigma} + \frac{1}{2}\xi^2\sigma^2 \frac{\partial^2 V}{\partial \sigma^2} + (r^d - r^f)S \frac{\partial V}{\partial S} - r^d V = 0, \tag{7}$$

with appropriate boundary conditions. An analytic solution for a vanilla option price is not available for this model. However, the perturbation expansion approach by Hagan et al. [9] yields a very good quality approximation. By using these expansions one can calibrate a vanilla smile in a fraction of a second.

Nevertheless, since barriers and other exotic options are usually priced via a finite differences approach anyway, we have also developed a fairly quick numerical calibration of the model, exploiting the spot homogeneity of the model. The finite differences approach is especially useful as, unlike the perturbation expansion approach, it can also handle a model enriched with volatility mean-reversion, parameters term structure, or even a local volatility component.

By introducing an unequally spaced mesh and using a second order convergent alternating directions implicit method, we can price a single vanilla in a finite

differences grid very quickly. This is convenient, since solving the PDE once, for a vanilla with fixed strike and fixed maturity, gives vanilla prices for all strikes. This is due to the spot-homogeneity of the process, which implies that the following relation holds for vanilla option prices

$$V(K|S_0) = \frac{K}{K_{fix}} V(K_{fix}|S) \text{ with } S = \frac{K_{fix}}{K} S_0, \quad (8)$$

where $V(K_{fix}|S)$ for all S is determined by one sweep through the PDE grid.⁴

In table 4, we demonstrate some calibration examples to illustrate stochastic volatility quote-parameter correspondence. As one would expect, ATM corresponds to initial instantaneous volatility σ_0 , while RR and STR quotes are generated by ρ and ξ , respectively.

Quotes / Parameters		Example 1		Example 2		Example 3	
ATM	σ_0	10.00%	0.097	12.00%	0.117	12.00%	0.119
RR	ρ	1.00%	0.22	-1.00%	-0.22	-1.00%	-0.30
STR	ξ	0.30%	1.26	0.30%	1.15	0.15%	0.85

Table 4: Calibration results stochastic volatility model with $S_0 = 1.00, T = 0.5, r^d = 0.01, r^f = 0.03$.

Next, let us analyze possible future smiles predicted by stochastic volatility models. Stochastic volatility smiles exhibit so-called sticky delta dynamics. This means that the smile shifts in strike dimension (in response to spot changes) and/or volatility dimension (in response to instantaneous volatility changes), but does not change its shape on a moneyness scale. In the FX context it corresponds to a stochastic at-the-money volatility, but strangles and risk reversals that stay more or less constant irrespective of the spot or volatility changes. Of course, given a volatility/spot correlation one can expect a future smile conditional on spot to be higher or lower than the smile we observe today, but such an “expected smile” does not correspond to the conditional density implied by the model.

As outlined in section 3, we should rather be interested in a conditional forward smile. This smile can be backed out from conditional expectations of future vanilla option prices. To get these expected conditional option prices, we can put

⁴Like in the local volatility model, we use Broyden’s method to fit three model parameters to three option quotes exactly. In this fashion our model can be calibrated quicker than in two seconds on a standard PC.

ourselves forward in time and get vanilla prices conditional on spot for a range of possible future volatilities. Integrating over the volatility density conditional on this spot (which can be obtained efficiently by Monte-Carlo simulation using the results from Willard [13]) yields the desired expected conditional vanilla prices.

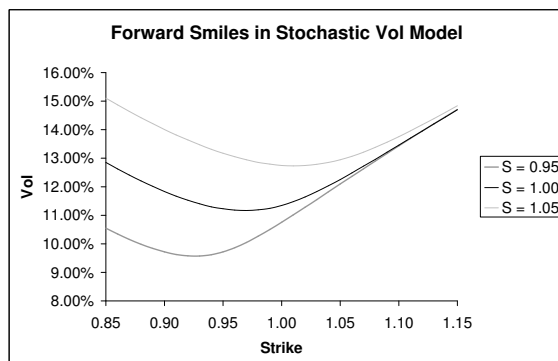


Figure 4: Forward smiles in stochastic volatility, using example 1 data.

In figure 4 we show the forward smile for different spot levels using the parameters from example 1. The SV forward smiles have roughly the same risk reversals as the initial smile, but the two other quotes change: the ATM can either increase or decline, while the strangles are ever increasing as we put ourselves farther into the future.

The change in the ATM is mainly due to the positive spot/volatility correlation. One expects the future instantaneous volatility to increase conditional on a positive spot return. Thus, we observe a higher-than-initial ATM quote for a smile conditional on 1.05 spot and a lower-than-initial ATM for a smile conditional on 0.95.

The most striking feature of the forward smile is its extra convexity compared to the initial smile. That seems to contradict the fact that SV strangles are generally quite stable. Here we provide an explanation. The main difference of the forward smile from the smile we observe today is additional uncertainty, associated with the instantaneous volatility at a future time, to which this smile corresponds. Current three months smile reflects the variation of the volatility during the next three months. However, the “three months in three months” forward smile is dependent on the instantaneous volatility in three months from now (which is yet unknown) and on its variation during the next three months. In effect, it thus depends on the volatility behaviour during the whole six months period starting

from today. The same reasoning can be made from the density point of view: conditional on any volatility level in three months from now, the transition density has roughly the same kurtosis as three months terminal density today. However, the stochasticity of this initial volatility introduces excess kurtosis. This translates into a more convex forward smile.

5 Static Hedging

This section presents our static hedge construction. It is now clear how to calculate forward smiles implied by the LV and SV models. Next, we want to analyze how forward smiles affect barrier option prices. Static hedging suits our goal, as we will only use today's smile and the forward smiles along the trigger level in order to price options. Of course, these prices will be just approximations, but their good quality (and their ability to explain the differing LV and SV prices) can confirm the importance of the forward smiles information in barrier pricing. But first we must focus on developing and explaining our static hedges.

Static hedging represents a popular alternative approach to pricing and hedging exotic options. Consider a digital call option, for example. It is basically impossible to hedge this option dynamically when approaching expiry and spot is hovering around the strike. Approximating the payoff via a call spread is thus generally preferred in practice. If one were able to put on such a spread around a single basis point the hedge would in fact be perfect. In practice however the amounts you would need to buy or sell would be extremely large and some risk still needs to be run.

Within the Black-Scholes framework there also exists a static hedge for knock-out barrier options⁵, as shown by Carr and Chou [5]. The idea is to hold a portfolio of vanilla options that produces the required vanilla payoff at maturity, but with an additional condition that its value becomes zero if spot ever trades at the trigger level. One thus buys a portfolio of vanilla options at inception and unwinds this hedge as soon as the barrier is triggered at zero cost. If one were to use this portfolio in order to hedge the barrier option in a world in which volatilities were not flat, the unwind costs will no longer be zero. One could thus think of describing barrier option prices via an approximate initial static hedge together with an estimate for the unwinding costs. Since the initial approximate static hedges just represent European style payoffs their prices are unlikely to vary too much

⁵A knock-in option can be constructed via the simple parity relationship $KI = \text{Vanilla} - KO$

from model to model. This already indicates that one could possibly explain the difference in the barrier prices emerging from different models by inspecting the conditional smiles along the barrier, since it is the volatility smile at hit that is going to characterize the unwinding costs.

Since the goal is to minimize expected unwinding costs, the algorithm should take into account expected vanilla prices at hit. From this viewpoint, using forward smiles that are derived from expected vanilla prices conditional on spot being at trigger level is more justified than taking possible future smiles at trigger⁶. However, our method will ignore the difference between the paths that have not touched the trigger prior to the conditioning time and those that have already crossed the barrier before that. Hence, a static hedge based on the forward smile rule is not a true hedge for a barrier option, when the volatility is stochastic. It is just a way to construct an approximation to the price using a minimum of model-dependent information. As such, it serves our exposition and, by virtue of using the most relevant information (a properly defined forward smile), results in a fairly good approximation to a barrier option price under the stochastic volatility model.

Throughout the rest of the section we will be using the notion “static hedge” to denote the portfolio that unwinds at zero costs under an assumption of a deterministic implied volatility smile at hit, irrespective of whether such an assumption holds for a model in question.

5.1 Static Hedge Model

In order to construct a hedge for a knock-out option, one has to build a portfolio consisting of the underlying vanilla and a collection of options that offset the vanilla’s value along the barrier and that all expire worthless if the trigger level never trades. This European style portfolio can be priced with today’s smile and hence gives us the price of the barrier option. The method will be described using a down-and-out call (DOC), but can easily be extended to other single barrier option types.

Consider a DOC struck at K , trigger $H < K$, and time to expiry T . The first instrument to include into the hedge portfolio is the underlying call itself. Now in order to synthesize the knock out feature the replicating portfolio should have a zero value whenever the spot price reaches H . It is convenient to weaken this condition by requiring the unwinding costs to be zero for a discrete subset of

⁶This distinction becomes irrelevant for local volatility models, where future and forward smiles are the same.

times. Then, with the n time points $0 < t_1 < \dots < t_n < T$, one can select n put options with maturity T and strikes $0 < K_1 < \dots < K_n < H$ and include them into the hedging portfolio with weights $w_1, \dots, w_n \in \mathbb{R}$, so that the linear set of equations

$$C(t_i, H, K, T) + \sum_{j=1}^n w_j P(t_i, H, K_j, T) = 0 \quad (9)$$

for all $i = 1, \dots, n$ is satisfied. Here $C(t, S, K, T)$ and $P(t, S, K, T)$ denote the value of call and put options at time t .

If the volatility smile at the barrier is known, one can calculate the prices of the underlying vanilla options at all time points. In that case, the system of equations can be solved to obtain the weights. Or more formally, n linear equations with unknowns w_1, \dots, w_n can be written in matrix notation (where $P(t_i, H, K_j, T)$ is denoted with $P_j(t_i)$ and $C(t_i, H, K, T)$ with $C(t_i)$) as follows,

$$\begin{bmatrix} P_1(t_1) & P_2(t_1) & \cdots & P_n(t_1) \\ P_1(t_2) & P_2(t_2) & \cdots & P_n(t_2) \\ \vdots & \vdots & \ddots & \vdots \\ P_1(t_n) & P_2(t_n) & \cdots & P_n(t_n) \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \end{bmatrix} = \begin{bmatrix} -C(t_1) \\ -C(t_2) \\ \vdots \\ -C(t_n) \end{bmatrix} \quad (10)$$

So if the matrix with put prices is invertible, the solution exists and can be easily calculated.

Note that the underlying call has time value at the barrier level, but no intrinsic value. Consequently, the portfolio of puts has zero intrinsic value and all the puts are bound to have strikes below the barrier. This technique should be modified slightly to accommodate barrier options that have positive intrinsic value just prior to triggering. Consider a DOC with $H > K$. Unlike the previous case, the underlying call has positive intrinsic value at the barrier, so it has to be offset with options that also have positive intrinsic value there. We choose to include a digital put struck at H to offset the intrinsic value of a call. If the digital put pays off one unit of base currency, then the value at the barrier level just before expiry equals one half. Therefore, we choose a weight for the digital put of $-2(H - K)$. Additionally, we use a collection of vanilla puts to offset the time value of the call.

We can now write down the set of equations defining the static hedge for the DOC in a slightly more general way, namely for all $i = 1, \dots, n$

$$C(t_i, H, K, T) + w_{dp} DP(t_i, H, H, T) + \sum_{j=1}^n w_j P(t_i, H, K_j, T) = 0 \quad (11)$$

Here $DP(t, S, K, T)$ denotes the value of a digital put at time t . Its weight w_{dp} is equal to $-2(H - K)^+$. The weights of the put options can again be found by inverting the matrix of option values.

We have not yet addressed the choice of strikes $0 < K_1 < \dots < K_n < H$. Selecting strikes that are of different orders of magnitude than spot can clearly lead to problems (we will get an ill-conditioned matrix, as the put values will be extremely close to zero). A simple yet seemingly effective rule is to take equally spaced strikes ranging from the barrier and down to two Black-Scholes standard deviations below the trigger.

In table 5 we present some errors associated with the discrete imposition of the unwind rule. In the example below, a Black-Scholes static hedge for a six-month DOC was constructed, based on eight possible hitting times and eight strikes ranging from 0.81 to 0.95. The error never exceeded 0.05 basis points of the notional. Note that a discrete barrier option with eight observations in six months would differ from the price its continuous counterpart by a far larger order of magnitude.

Strike	BS Price	Static Hedge
0.90	661.69	661.69
0.92	557.91	567.92
0.94	474.13	474.09
0.96	380.53	380.49
0.98	291.11	291.11
1.00	212.18	212.17
1.02	147.40	147.40
1.04	97.67	97.67
1.06	61.79	61.78
1.08	37.36	37.35
1.10	21.61	21.61

Table 5: Cost of setting up static hedge compared to Black-Scholes value based on the following parameters: $S_0 = 1.00$, $H = 0.95$, $T = 0.5$, $r^d = 0.01$, $r^f = 0.03$, ATM= 10.0%

We believe the approach, albeit somewhat ad hoc and far from mathematically rigorous, to be a useful tool in the pricing and risk management of barrier options. Essentially, this method provides an easy approximation of the expected hedging cost by incorporating a view of the volatility smile when the barrier is hit. Fur-

thermore, it helps to understand, that the different prices of barrier options in local and stochastic volatility models can very well be explained by their implied smile dynamics. This will be demonstrated in the next section.

6 Model Comparison

In this section we will use the static hedge technique to analyse to what extent the LV and SV barrier prices can be explained by their respective smile dynamics. The tests presented here support our claim that conditional forward smiles along trigger levels encapsulate most of the model-dependent information relevant to barrier pricing.

6.1 Test Setup

To price barrier options with the static hedge approach, one has to assume a known volatility smile at the barrier. The testing approach is as follows: we use the deterministic smiles upon hitting the barrier as predicted by our LV model, pass them to the static hedge algorithm and obtain approximations to barrier prices. Then we do the same for the SV model. These approximations are then compared to the barrier prices produced by fully-specified LV and SV models (calculated via numerically solving the corresponding pricing PDE). The forward smiles required for these tests are obtained according to the methods described in sections 4.1 and 4.2. Note that these conditional forward smiles are the only model-dependent information passed to the static hedge algorithm. Hence, the quality of the approximations is witness to the relevance of smile dynamics to barrier pricing.

6.2 Test Results

We have priced a number of down-and-out call options using LV and SV models, and using the static hedge model based on conditional forward smiles. We used eight possible hit points and hence nine vanilla options to construct our static hedges according to the method described in section 5. For both models, we analyzed three different market data scenarios: a smile skewed in favour of out-of-the money calls, a symmetric smile, and a smile skewed in favour of out-of-the money puts.

Figure 5 illustrates the down-and-out call test results. It shows the absolute values of the adjustments to the Black-Scholes price on the left (as a function of

the option strike), and the relative adjustment to the Black Scholes on the right. Each row contains a different market data scenario, positively skewed smile in row 1, symmetric smile in row 2 and negatively skewed smile in row 3.

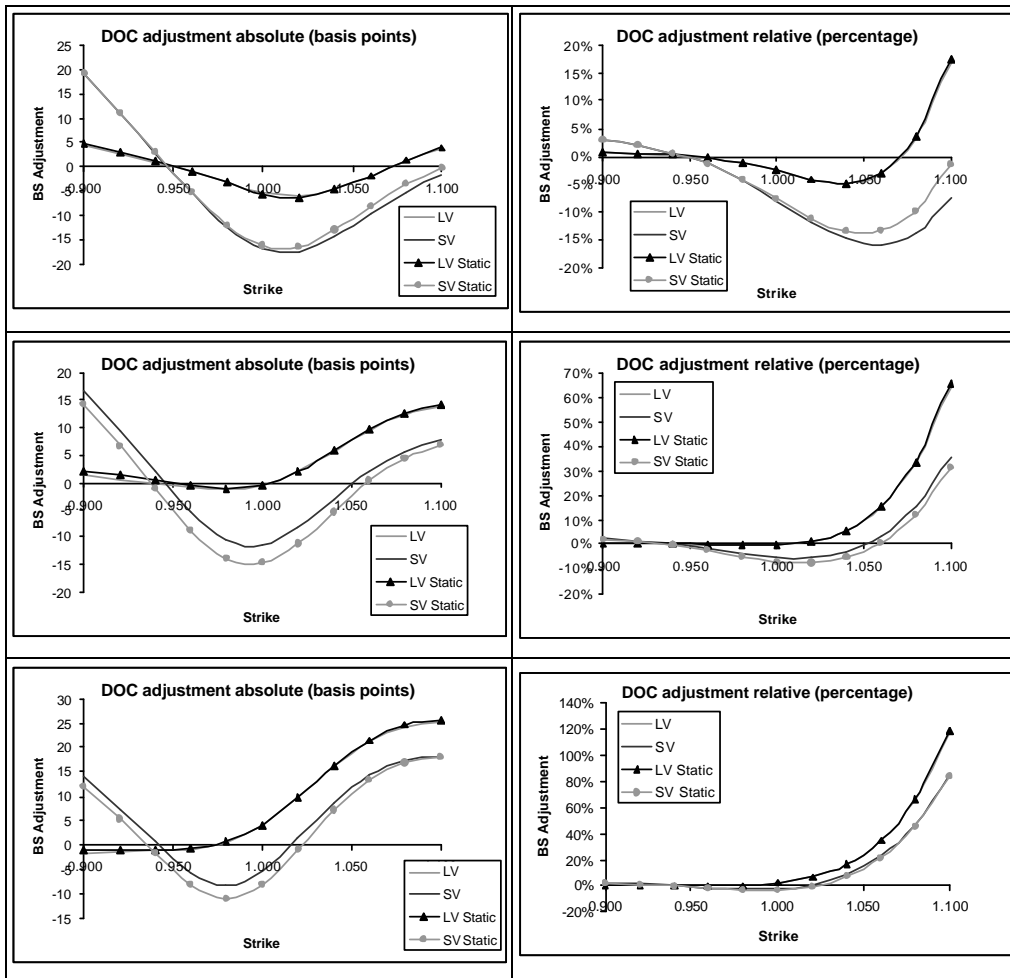


Figure 5: Down-and-out call adjustments to the Black-Scholes price as a function of strike, for three different scenarios: Top: $RR = 1.0\%$ (scenario 1), middle: $RR = 0.0\%$ (scenario 2), bottom: $RR = -1.0\%$ (scenario 3). Additional parameters: $S_0 = 1.00$, $H = 0.95$, $T = 0.5$, $r^d = 0.01$, $r^f = 0.0$, $ATM = 10.0\%$, $STR = 0.3\%$, 1 basis point = 0.01% of the notional

As the graphs demonstrate, our static hedge model is able to capture the price

patterns of SV and LV very well. Moreover, these prices are good approximations of the full models' prices, though they were based solely on conditional forward smiles information. In fact, the lines corresponding to the full LV model and its static hedge approximation are practically indistinguishable on the graphs.

7 Conclusion

As the forward smile conditional on the barrier level is the only model specific input for the static hedges, the good quality of resulting approximations illustrates the importance of smile dynamics in barrier pricing. One can say that forward smiles along trigger levels encapsulate very important information that matters for barrier prices and is not reflected in vanilla prices. Hence, a conclusion that searching for models that would be able to produce realistic conditional forward smiles is the key direction in option pricing research. At least so far as the barrier pricing is concerned, improving the fit to vanilla surface by introducing more and more sophisticated models is unlikely to yield a good model unless the forward smiles problem is addressed directly.

In light of the above conclusion studying the empirical smile dynamics becomes an essential part of our research agenda. Some observations of FX implied volatility behavior were discussed in section 2 of this paper. The most important feature we saw was the strong correlation of spot and risk reversal across all currency pairs. In other words, FX smiles exhibit certain systematic behaviour that can not be described by some simple rule like "sticky strike" and "sticky delta". This observation also makes us doubt that the classical affine jump diffusion models can produce forward smiles agreeing with the market.

There several ways to infer desired smile dynamics on the model. The first way is building models with parameters that explicitly govern smile dynamics and hence can be calibrated to some pre-specified smile sensitivities. The Reech model [4] is an example of this approach. Another example of hard-wiring certain forward smiles into a model can be found in Madan [11]. The question remains as to how exactly should we measure the market smile dynamics, and to what extent it should be enforced on the model.

An extension of this approach is based on the fact that all the relevant information should be reflected in barrier option prices. Admittedly, calibrating a pricing model to a few liquid barrier-type products will result in obtaining the model with more realistic forward smiles. This model can be then used to price other barrier type options. For examples of such approaches see Johnson and Lee [10], where a

mixture approach is used, and Ayache et al. [2], where a regime-switching model is calibrated to one-touch option quotes.

In the paper by Carr and Wu [6], two separate Levy processes are used to generate stochasticity in RR. The resulting model is able to capture the features of the currency options market better than traditional affine jump diffusion framework. To the best of our knowledge, the applicability of this approach to barrier pricing has not yet been investigated. The Reech model is another candidate for direct barrier calibration. Some research has been done in Schönbucher [12], Cont and Fonseca [7] and Balland [3] on modelling the implied volatilities directly. This approach also offers some possibilities for inferring interdependencies between, say, risk reversal and spot.

References

- [1] L.B.G. Andersen and R. Brotherton-Ratcliffe. The equity option volatility smile: an implicit-finite difference approach. *Journal of Computational Finance*, 2:5–38, 1997.
- [2] E. Ayache, P. Henrotte, S. Nassar, and X. Wang. Can anyone solve the smile problem? *Wilmott Magazine*, pages 78–95, January 2004.
- [3] P. Balland. Deterministic implied volatility models. *Quantitative Finance*, 2:31–44, 2002.
- [4] G. Blacher. A new approach for designing and calibrating stochastic volatility models for optimal delta-vega hedging of exotics. *Global Derivatives*, 2001.
- [5] P. Carr and A. Chou. Breaking barriers. *RISK Magazine*, pages 139–145, September 1997.
- [6] P. Carr and L. Wu. Stochastic skew in currency options. *Courant Institute, Working Paper*, 2004.
- [7] R. Cont and J. Fonseca. Dynamics of implied volatility surfaces. *Quantitative Finance*, 2, 2002.
- [8] B. Dupire. Pricing with a smile. *RISK Magazine*, pages 18–20, January 1994.

- [9] P. Hagan, D. Kumar, A. Lesniewski, and E. Woodward. Managing smile risk. *Wilmott Magazine*, 2002.
- [10] S. Johnson and H. Lee. Capturing the smile. *RISK Magazine*, March 2003.
- [11] D. Madan. From local volatility to local levy: Hard wiring skews. Eurandom workshop: Exotic option pricing under advanced Levy models, Eindhoven, May 2004.
- [12] P. Schonbucher. A market model for stochastic implied volatility. *University of Bonn, Working paper*, 1999.
- [13] G. Willard. Calculating prices and sensitivities for path-dependent derivatives securities in multifactor models. *Journal of Derivatives*, pages 45–61, Fall 1997.