


Laughter in the Dark - The Problem of the Volatility Smile


Emanuel Derman
May 26, 2003

According to classic theory, the Black-Scholes implied volatility of an option should be independent of its strike and expiration. Plotted as a surface, it should be flat, as shown at right.

The volatility surface according to Black-Scholes



The volatility surface according to S&P options markets



Prior to the stock market crash of October 1987, the volatility surface of index options was indeed fairly flat.

Since the crash, the volatility surface of index options has become skewed. Referred to as the volatility smile, the surface changes over time. Its level at any instant is a varying function of strike and expiration, as shown at left.

The smile phenomenon has spread to stock options, interest-rate options, currency options, and almost ever other volatility market. Since the Black-Scholes model cannot account for the smile, trading desks have begun to use more complex models to value and hedge their options.



After 15 years, there is still no overwhelming consensus as to the correct model. Each market has its own favorite (or two). This talk covers a short history of attempts to model the smile. Despite initial optimism about finding a model to replace Black-Scholes, we are still in many ways searching in the dark.

Introduction

I'm very happy to have been invited to be here in Amsterdam, but I'm a little ashamed to say that it's the first time in my life I've ever been in the Netherlands. That's not very appropriate for someone who spent most of his life in two ex-Dutch colonies. I was born and grew up in Cape Town, South Africa, *waar ek het Afrikaans elke jaar geleer, maar dit was lank gelede en nou het ek veel vergeet. Maar ek kan nog se "Dis baie lekker" hier in Amsterdam.* It took me many years of living in a second Dutch colony, New Amsterdam, to get over the common South African habit of using the world *lekker* for anything you really liked.


In the last few years I also became a big fan of Harry Mulisch, another reason I'm glad to be here. Finally, I'm also happy to see some of the many well-known Dutch quantitative finance people, some of whom I've had the privilege of meeting in the past. I met Ton Vorst many years ago at Goldman Sachs in 1986 or 1987, and I'm also pleased to meet up again with Farshid Jamshidian whose previous talks on the volatility smile helped me a great deal in preparing this talk.

Overview


Though the smile first appeared in options markets in 1987, I had it brought to my attention in Tokyo in December 1990. Until that year I had worked exclusively in fixed-income, and wasn't really aware of the phenomenon at all. Then, in early 1990 I moved to work in the Goldman Sachs Equity Derivatives area, a business that was beginning a period of rapid growth following the collapse of the Berlin Wall and the Soviet empire. Suddenly, clients were much more interested in easy global investing, and derivatives provided some of the best methods.

That December, our trader in Tokyo showed me the implied volatilities of Nikkei options on his screen. What perturbed him was that three-month options of low strike had much higher implied volatilities than three-month options of high strike, as shown in Figure 1. This skew or asymmetry in the implied volatilities had been absent before the infamous stock market crash of Oct. 19 1987, and had begun to appear shortly afterwards, in index markets all over the world.

FIGURE 1. A typical implied volatility smile for the three-month options on the Nikkei index in late 1994. The dotted line shows the lack of skew that was common prior to the 1987 crash.



A similar but not identical pattern held for all expirations, so that it became common to speak of an implied volatility surface whose height varied as a function of strike and expiration. Figure 2a shows a typical surface for S&P 500 options. The shape is fairly typical, but the exact details change from moment to moment and day to day. The lack of flatness came to be called the volatility smile, because in the currency world it really did resemble a slight upturned curve of the lips. In fact, for indexes, the skew was described as negative, since volatilities were anti-correlated with strike prices.



These surfaces and curves were a challenge to theorists everywhere. The classic Black-Scholes model attributed a single lognormal volatility to an underlier at all times and all levels and therefore predicted a dull featureless plateau-like implied volatility surface, as shown in Figure 2b. According to Black-Scholes, options of all strikes should have the same volatility. But according to the smile, each option reported a different volatility for the same underlier.

What was wrong with Black-Scholes and what kind of new model could possibly match and explain this skewed surface? This wasn't just an intellectual challenge, but one of importance to the business too. Our equity derivatives desk made markets in index and single-stock options all over the world. Even if we knew the market price of a liquid option, we needed a model to hedge it. If the Black-Scholes model couldn't account for an implied volatility, it couldn't produce a reliable hedge. The big question was – what was the right hedge ratio for an index option?


There was a second quest of equal or greater importance. The era of exotic options and structured products – knockouts, lookbacks, averages and quantos – was just taking off, and we were part of it. If each option had a Black-Scholes implied volatility that depended on its strike price, then what implied volatility should you use to value an exotic option with, effectively, several different strikes incorporated in it? S&P implied volatilities from about 14% to 24% as a function of strike, a stupendous range. What volatility should you use? It was a serious dilemma.

The smile first appeared after the 1987 crash and was clearly connected in some way with the visceral shock of discovering, for the first time since 1929, that a giant market could drop by 20% or

more in a day. Clearly low-strike puts should be worth more than high-strike calls when you thought about the higher probabilities associated with that kind of move.

Over the next 15 years the volatility smile spread to most other options markets, but in each market it took its own idiosyncratic form. Slowly, and then more rapidly, traders and quants in every product area had to model the smile. At Goldman Sachs in the past few years, not only did each front-office trading desk have quantitative strategists working on their particular smile model, but the Firmwide Risk Management group had a group of quants to value independently many of the structured deals that depend on the details of the volatility smile. I would think it's safe to say that there is no area where model risk is more of an issue than in the modeling of the volatility smile.

A brief overview of a few other smiles. Though indexes tend to jump down more than they jump up, single stocks have more symmetric jumps up and down. Indexes tend to jump down. Perhaps in consequence, the volatility smile for single stocks often looks like Figure 3, more symmetric and smile-like.




In FX markets, the smile can be even more symmetrical, resembling a real grin, especially if the two currencies are of equal strength. The smiles are more symmetric for “equally powerful” currencies, less so for “unequal” ones, as shown in Figure 4. In contrast, the volatilities of fixed-income options tend to more skewed, with the implied volatility of low rates dominating.

From a behavioral point of view, it seems likely that implied volatilities are greatest where market movements are likely to cause the greatest shock and awe. In index markets, that's the downside; in single stock markets, both up and down jumps can occur. In the gold market, since gold is more likely to be a haven, that jumps up when stocks move down, in recent years, a positive volatility skew has occurred in that market.


FIGURE 4. FX implied volatility smiles.



ATM Strike = 0.90



9.85



123.67

If there is one safe conclusion, it is that as options markets have become more experienced about the kinds of shocks that can occur, they have started to display more sophisticated patterns of implied volatility smiles. The smile is a sign of market sophistication. You should think of Black-Scholes implied volatilities as a somewhat simplistic quoting convention for the prices of a better underlying model.

A Better Model of the Smile

The Black-Scholes model assumes for the underlying stock an idealized continuous Brownian motion with a single constant volatility at all times, as well as the ability to hedge continuously without transactions costs. Given a stock and a bond price, it gives you the price of a hybrid, the option, part stock and bond. But, its results are inconsistent with observed smiles in most markets.

Experienced people in the options business tend to think of options models as sophisticated interpolation models that take you from the transparent prices of liquid securities to the obscure and unknown prices of illiquid or exotic securities. When Black-Scholes was first invented, stocks and bonds were the liquid base from which you started. Nowadays, in many markets, vanilla options are liquid too; in equity index markets, forward starting options (cliquets) are relatively transparent. In FX markets, one-touch American-style binary options are liquid enough to serve as a calibration base. A better model of the smile should be capable of calibration to liquid stock, bond *and* options prices, and can then be used to interpolate to the untransparent hedge ratios of vanillas and to the prices of exotics.


A better model of the smile should also involve a more realistic modeling of the movement of underlier levels. There may be, indeed there are, many ways to fit the smile, and the best one is the one that simulates the market's behavior best. But each model will produce hedge ratios that differ from those of the Black-Scholes model, and different prices for exotics.

Local Volatility Models of the Smile

The earliest models of the smile were what are now called local volatility models. These models were inspired by staying as close as possible to the eminently successful world-view of Black-Scholes, loosening its framework just enough to support a smile. They were also motivated by the experience of fixed-income modeling, always a more quantitative field than equities, and in particular by the great idea of forward rates.

Look at bond yields as an analogy. Once upon a time the bond market used yield to maturity as its measure of the value of future cash flows. Nowadays bond yields are merely the simple “model” with which traders quote bond prices. To obtain the arbitrage-free value of a more complex fixed-income instrument, one discounts them using the forward rates implied by the yield curve. The yield curve provides the liquid market data; the forward rates are obtained by calibration; with options or volatility data you can build an even more realistic yield curve model that adds a range of future rates calibrated to option and bond prices, as shown in Figure 5.


FIGURE 5. Bond yields, forward rates, stochastic rate models.



Yield curves imply forward rates that can be locked in by calendar spread trades. Breeden and Litzenberger had shown in the 1970s that a stock’s option prices imply its risk-neutral distribution at expiration that can be locked in by option butterfly positions. Local volatility models went a step further. They showed that options prices (or their implied volatilities) imply future short-term local volatilities that can be locked in by butterfly and calendar spread positions in options. Once you know these future local volatilities, you can use them to value and hedge any kind of option on the same underlier.


Local volatility models achieved this by generalizing the Black-Scholes lognormal diffusion process with constant volatility to a process with a stock-price-dependent volatility. Volatility was stochastic, but stochastic as a function of stock price. Under these assumptions, you can back out the *unique* future local volatilities from the smile of options prices, as shown in Figure 6. You can think of it as an implied tree, an extension of a standard Cox-Ross-Rubinstein tree with constant volatility.

FIGURE 6. An implied binomial tree of future index moves corresponding to a particular implied volatility surfaced




Local volatility models work as illustrated in Figure 7. Roughly speaking, just as a yield to maturity is a *global average* over all *forward rates* to its maturity, so the implied volatility of an option with a definite strike and expiration is a *global average* all the *local volatilities* that the underlier will experience within the tree as it moves towards its expiration.

FIGURE 7. The implied volatility of the option expiring at the red circle is an average over the local volatilities in the red region of the tree. The implied volatility of the option expiring at the blue circle is an average over the local volatilities in the blue region of the tree.



Looked at this way, it was only one more step to realizing that by buying a long-dated option and selling two slightly shorter-dated ones at nearby strikes, you could determine and lock in the forward volatility, as shown in Figure 8.

FIGURE 8. Three different options determine the local volatility at a node. (a) Two options with adjacent strikes K_U and K_D illuminates the red interior nodes. (b) A single option with strike K_+ expiring one period later illuminates the blue interior nodes that consists of the red nodes plus one extra blue node. (c) The spread $(K_+ - K_U - K_D)$ determines the value of the single blue node.



Thus, current options prices and their spreads uniquely determined the local forward volatilities in the tree, just as bond prices determined forward rates.


The idea behind local volatility models was that you could build an implied tree calibrated to liquid index options prices and then use that tree to value exotics and hedged vanillas. Because the index skew was negative, that is, because implied volatilities fell as strike levels increased, the local volatility of the index had to decrease as index levels rose in the tree. Therefore these models predicted a hedge ratio smaller than the Black-Scholes hedge ratio.

The Problem with Local Volatility Models

It seemed great to back out the market's expectations of future volatility from current options prices. Unlike Black-Scholes, the model could accommodate the smile, was theoretically self-consistent and complete – you could hedge an option purely with stocks. You could use the local volatilities with trees or partial differential equations or Monte Carlo simulations to value almost any other volatility-dependent security, including American-style options. Since implied volatilities were the average of future local volatilities, these models also provided nice intuition.

But there were problems that arose after the initial burst of excitement. It wasn't easy to get smooth continuous local volatility surfaces from a few discrete options prices. Their details depended on how you smoothed or parameterized the few discrete implied volatilities you started with. Methods for generating smooth local volatility surfaces have multiplied, but it's still a complex and computationally intensive problem.


Setting aside the problems of computation, there were more significant semantic issues. When you looked at the future local volatilities in these models consistent with today's implied volatilities, they weren't reassuring. Local volatility models have a scale; they depended specifically on future index levels and time. Far in the future, the local volatilities were roughly constant, predicting a future smile that was unintuitively much flatter than current smiles, an uncomfortable and unrealistic forecast that contradicted the omnipresent nature of the skew.




If these models forecast unrealistic future volatilities, how much can you trust their prices and hedge ratios, self-consistent though they are? For all these reasons it became compelling to look at other models.

Stochastic Volatility Models

Stochastic volatility models take a more time-invariant view of the world; they avoid the scale built into local volatility models. They plausibly postulate, at any instant, volatility itself is volatile, fluctuating but reverting towards a long-term mean. Crudely pictured, you can think of stochastic volatility models as represented by a mixture of two (or more) stock price evolutions, each with its own volatility, as represented by the following diagram:



In these models, the mixture of high and low volatilities produces fat tails in the distribution of index returns, and, as this picture suggests, you can think of the model's option price as being an average of the Black-Scholes values for high and low volatility. The result is a smile or implied volatility smile that, in the simplest case when there is no correlation between the underlier and its volatility, is symmetric, a pure parabola, a true smile reminiscent of the equal-strength currency smile.



If you allow for a correlation between volatility and underlier, you add a linear distorting term to the parabola, producing asymmetric smiles that can match index smiles too. Instead of simply averaging over Black-Scholes prices at different volatilities, the options formula now involves averaging over shifted stock prices and volatilities, so that

$$C(S, \sigma) = E[C_{BS}(\tilde{S}\{\rho, \sigma\}, \sigma(\rho))]$$


Stochastic volatility models have attractive features. Their smile is stable, unchanging over time, and in that sense more like real-world smiles. They therefore produce more realistic future smiles.

But there are disadvantages too. If volatility is stochastic you have to hedge it to replicate and price the option. What is its current value and what do you hedge it with? How does it evolve? Unlike a stock or a currency, volatility is not a traded variable with an observable price. In practise you must hedge one option with another, and calibrate the evolution of future volatility in the model to fit current options prices in order to get going.

So, a stochastic volatility model is potentially more realistic but also filled with more complex unknowns about the evolution of volatility. And hedging options with options, which are less liquid than stocks or futures, is harder in theory and in practise. Stochastic volatility models have much in common with stochastic interest rate models – there are even models in which local volatility is allowed to vary stochastically – and stochastic volatility models have a long way to go before they become easily usable.

Jump Diffusion Models

In the interests of greater realism, practitioners have added jumps and crashes to the standard Black-Scholes diffusion of stock prices. More recent models have layered jumps on top of local volatility models. The model is represented by Figure 12.




The advantages of jump diffusion models are that they are realistic – they take account of the sorts of jumps that occurred in October 1987 and 1997. They can fit the observed smile reasonably well in a stationary way, with jumps accounting for the steep short end of the smile and local volatility for the long end.

The disadvantage is that, unlike Black-Scholes, local volatility or even stochastic volatility, these models are incomplete from the point of view of replication – you cannot hedge jumps with the stock alone or even, as in stochastic volatility models, with one additional option. For each possible jump size, you need one more underlying option to hedge with, and there could be all sorts of jumps. The idea of a delta-hedge which hedges you against small moves when in fact jumps are possible is a bit of a contradiction in terms. But maybe that is the way the world is.

A Really Ambitious Model

The most sophisticated and perhaps the most realistic model I know of has been developed by ito33.com. They correctly regard the essence of a hedging as the minimization of the variance or risk of the ultimate P&L of your hedged portfolio. If the world behaves as Black and Scholes assumed, and you can hedge continuously without transactions costs, then you can replicate an option exactly and its value is independent of any stock move, with zero variance. However, if stocks jumps and if you hedge at discrete rather continuous intervals, you cannot replicate perfectly, and the P&L of your hedged option portfolio at expiration has a finite variance.



The uncertainty of the P&L increases with the discontinuities in the stock process. Ito33 allow for jumps and stochastic volatility, and their system determines the optimal hedge ratio that minimizes the P&L variance under all of these scenarios. The fair price of an option is the cost of the hedging strategy that works best, however imperfectly. Ito33.com can replicate vanilla as well as exotic options and can calibrate their underlying model of stock movements to vanilla options as well as to any available exotic prices. In the FX options world where one-touch option are fairly liquid, they calibrate also to these; in equity markets where the price of forward-starting options are available, they calibrate also to these. It's a computationally intensive and complex picture of the world, and has the potential for moving a step or two closer to reality, but it's a complex reality whose successful implementation depends upon how well these underlying discontinuities reflect those of actual markets.

Conclusion: Smiling in the Dark?

So, where do we stand, 15 years after the appearance of the smile?

I came here to Amsterdam from a derivatives conference in Barcelona, which I also attended one year ago. At that time, I chaired a discussion group on the volatility smile. Years ago, when I first became aware of the smile, we hoped to find the “right” model, and when I met people from other trading firms I used to ask them which model they thought was correct. But now there is such a profusion of models that I have begun to ask more practical questions – not “What do you believe?” but “What do you do?” It makes a difference. Local volatility models tend to produce hedge ratios for vanilla options that, when calibrated to the index options smile, are smaller than Black-Scholes hedge ratios. Stochastic volatility models tend to produce hedge ratios that are greater than Black-Scholes. And for exotic options, they vary from each other even more dramatically.

The reason I became interested in this more direct question was that when I worked in Firmwide Risk at Goldman Sachs, I became aware that each trading desk in the firm developed and used their own smile models, and the models differed from each other. FX traders used different models from fixed-income options traders. Even Nikkei equity index options traders in Tokyo used heuristic hedging strategies that differed from those used by S&P options traders in New York. Some groups used Black-Scholes, perhaps lazily, and if they were clever added their own guess as to how to modify the hedge ratio. Others used stochastic volatility or local volatility models, or a mix of them. Our electronic market makers used empirical regression models for the hedge ratio.

So, the question I asked people at the round table discussion was: “If you had to hedge a vanilla option on the S&P 500, would you use a hedge ratio greater than Black-Scholes, equal to Black-Scholes, or smaller than Black-Scholes?”

When I went around the table in Barcelona asking this question, more than ten years after the smile became a modeling issue, there was still no consensus on the right hedging strategy, let alone on the right model. Almost everyone had a different opinion, and some had no opinion at all. And the question isn’t even completely sensible, because it assumes that you can hedge an index option with the index alone, and if volatility is stochastic you may need other instruments to complete your hedge.

I think our dream of a perfect replacement for Black-Scholes is only a dream. There isn’t a uniformly good model. Since Black-Scholes is the market’s language for quoting options prices, local volatility is a natural way to quote forward volatility in terms of the values of portfolios of options spreads, just as forward rates are a natural way to think about the future interest rates. Which model is right depends on your market.

In markets where there are obvious scale- or level-dependent effects, when underlier volatility is highly correlated with underlier level, local volatility may be a reasonable model for dynamics. This is likely the case for modeling low-credit stocks where low equity levels correspond to greater firm leverage, higher volatility and the increased likelihood of default. Similarly, in the fixed-income world when certain interest rates or exchange-rate levels correspond to some cen-

tral-bank or economically significant support levels, local volatility may be a good picture.

In FX markets more generally, volatility is indeed highly stochastic and this may be the dominant mode of volatility change.

In equity index markets jumps and the fear of jumps are certainly important in the short term. It all, unfortunately, depends on what you think. No model is going to blindly free you of responsibility for deciding what you think the future regime of the market will be.

Therefore, it's good to avoid as many assumptions as possible, and practitioners on trading desks do this increasingly. One good strategy in attempting to value exotic options that are sensitive to the smile is to try to avoid modeling the dynamics of volatility as much as possible. In this approach, one tries to decompose an exotic option into a portfolio of liquid vanilla options chosen so that their joint payoff closely matches the payoff of the exotic. Sometimes you can achieve this so-called static replication with great accuracy and few modeling assumptions, and at other times you can achieve it only by making some unavoidable dynamical assumptions. But, if you can do it, then all quantities can be calculated in terms of vanilla options and their observed smile.

This is an approach I like, an honest recognition of our broad ignorance. Academics in finance departments often think that options theory is a solved problem. In fact, 15 years after the appearance of the smile, we are still in many ways, like the smile, laughing in the dark. There are too many models.

When you do research in options, you have to use advanced mathematics. If you are a practitioner you must never forget that you are moving through lawless roads where the local inhabitants don't respect your customs. All financial models are wrong, or at best hold only for a little while until people change their behavior. Aristotle, in his *Nicomachean Ethics*, says that in all endeavors one should adopt a degree of precision appropriate to the subject under consideration. In options theory, there has been a tendency to violate this maxim, to use methods that assume a precision much greater than the intrinsic degree of reliability in the field. As in most social sciences, the big and interesting battle in options theory and the smile is to avoid being Utopian, and instead to try to pick methods and models whose results depend as little as possible on unverified, indeed unverifiable, assumptions.