TN04-01:
THE VOLATILITY SMILE

Version date: August 7, 2008 C:\Classes\Teaching Notes\TN04-01.doc

Whenever I see your smiling face
I have to smile myself
James Taylor
"Your Smiling Face"
©1977 Country Road Music, Inc.

If James Taylor had been an options trader, he might well have not written that second line. Options "smile" at traders, but that smile is not a source of happiness. It is in fact a source of considerable confusion and misunderstanding. In this note, we explore how this confusion manifests in the pattern of different implied volatilities for different options on the same stock, a pattern commonly known as the volatility smile.

The following table shows the closing ask prices on January 9, 2004 of options on the exchange-traded fund called the QQQ, which attempts to replicate the NASDAQ 100 index, for options traded on the American Stock Exchange (with small volume on the CBOE). These options are among the most active of all listed options. All of the options are European-style. The January options expire on January 16, and the February options expire on February 20.[1]

	Calls		Puts	
Exercise Price	January	February	January	February
34	3.90	4.10	0.05	0.25
35	2.80	3.20	0.05	0.35
36	1.85	2.35	0.10	0.55
37	1.00	1.65	0.25	0.85

[1]Correction: When I originally created this example, it was my intention to use the option on the NASDAQ 100, known as the NDX, which is a European-style option that trades at the Chicago Board Options Exchange and (with light volume) at the American Stock Exchange. The QQQ contract is the option on the NASDAQ-100 Index tracking stock, which is a portfolio that attempts to replicate the NASDAQ-100 and is priced at approximately 1/40 the value of the index. Volume on the QQQ option is much greater than on the NASDAQ-100, so I switched to it. I did not, however, notice that the option on the QQQ is American-style. In the example here, I continued to treat it as European-style. All of the principles illustrated in this teaching note are valid, but in practice, the option would be American-style.

38	0.35	1.05	0.60	1.25
39	0.10	0.60	1.40	1.90
40	0.05	0.35	2.35	2.60

Suppose we wished to know which of these options is the most expensive. First, let us consider the calls. We know that February calls have to be more expensive than January calls. For a given expiration, however, which of the options is the most expensive? In that case, we know that the calls with the lowest exercise price are the most expensive. Likewise, we know that February puts should be more expensive than January puts.[2] We also know that the put with the highest exercise price should be, and indeed is, the most expensive.[3]

But suppose we would like to know *which options are the most expensive after taking into account any differences in time to expiration and exercise price*. To answer that question, we need to know something about how option prices are determined. For that, of course, we turn to the Black-Scholes model. We shall not go into the details of the Black-Scholes model here, as they are covered quite well in numerous other references. We can, however, use some fairly basic information from the model to tell us which of the above options is the most expensive.

The Black-Scholes model says that the price of an option is a function of the stock price, exercise price, risk-free rate, time to expiration, volatility, and any dividends on the stock over the life of the option. The options are all on the same stock so we cannot look to the stock price to tell us which option is more expensive. Dividends affect option prices but do so by effectively lowering the stock price by the present value of the dividends. Since none of the options differ by the dividends on the underlying, dividends cannot explain which option is more expensive. We already know the effect of the

[2]Because the puts are European, it is technically possible for the longer-term puts to be worth more. This situation occurs only for deep in-the-money puts and is part of the motivation for why holders of these options would prefer to exercise them early if they could. None of these options is sufficiently deep in-the-money for this effect to take hold, and we shall just treat them as if the longer-term options should logically be worth more.

[3]We do observe one anomaly in the put prices. The January 34 and 35 puts have the same ask price. These prices were the closing quotes, but the January 34 put had only one trade that day and the January 35 put had none. Thus, the liquidity of these two options is very low, and at least one of these quotes must be incorrect or they do not represent the prices at which market makers would trade.

exercise price and time to expiration. We want to account for any effect on the option price over and above the exercise price and time to expiration. Each option would have a risk-free rate derived from a risk-free investment maturing at the time of the option expiration. Hence, the January and February options could have different prices because of different risk-free rates. But it is well-known that the risk-free rate has only a minor effect on the price of an option, so we cannot look to the risk-free rate for an explanation. Because all of the options are on the same stock, the same volatility is used to price all of the options. Hence, we cannot use volatility to determine which option is more expensive.

Or can we? It turns out that we can.

Of these six variables, the stock price, exercise price, and time to expiration are easily observable. Hence, we could not possibly measure these with any appreciable error. The risk-free rate is largely observable, and its impact is small. The dividends are not observable, but they are not too difficult to measure accurately. The volatility, however, is almost completely unobservable. But even so, all of these options should be priced using the same volatility. It is, after all, the volatility of the underlying stock. The stock cannot have more than one volatility.

The Concept of the Implied Volatility

Suppose we use the Black-Scholes model to infer the volatility used by option traders to price the option. We search for the volatility that makes the model give us an option price that corresponds to the market price. The model will then be said to *imply* a volatility, giving rise to the concept of the *implied volatility*.[4] But again, would each option not imply the same volatility? Let us see.

We wish to estimate the implied volatilities of these options. We need the stock price, which is the value of the NASDAQ 100 tracking index at the close, which is 37.73. The exercise prices are obviously known. The time to expiration is the number of days to

[4]The mechanics of calculating the implied volatility are not straightforward. The formula is sufficiently complex that it is not possible to simply rearrange it to isolate the volatility on one side of the equation. The volatility can typically be derived by iterating with different volatilities until the market and model prices are sufficiently close. Convergence tricks, such as Newton-Raphson, tend to work well. Alternatively, there are a number of other methods that can obtain good results quickly.

expiration dividend by 365. Let us focus only on the February options.[5] They expire on February 20, which is 42 days later. Thus, the time to expiration is 42/365 = 0.1151. The risk-free rate is the continuously compounded yield of the U. S. Treasury bill maturing as close as possible to February 20. This rate is 0.83%.[6] The dividend yield is obtained as 0.18%.[7] We then find the volatility as the standard deviation that equates the price of the option to the Black-Scholes price. The table below shows our results:

Exercise Price	Calls	Puts
34	0.3230	0.2900
35	0.2904	0.2634
36	0.2592	0.2493
37	0.2455	0.2369
38	0.2279	0.2198
39	0.2156	0.2279
40	0.2181	0.2206

The results are graphed below.

[5]We could look at the January options, but we should not compare the January and February options. The implied volatility over the life of a January option is likely to differ from the implied volatility over the life of a February option, because the model specifies that the volatility should be the volatility over the life of the option. Differences in volatility for different expirations is referred to as the *term structure of volatility*.
[6]This information is obtained by first taking the average of the bid and ask discount yields of 0.80 and 0.81. We then compute the price of a 42-day Treasury bill based on a discount of 0.815. We find the continuously compounded rate at which the Treasury bill price grows to par value over 42 days. That rate is 0.83%.
[7]This information was obtained form the web site www.indexarb.com.

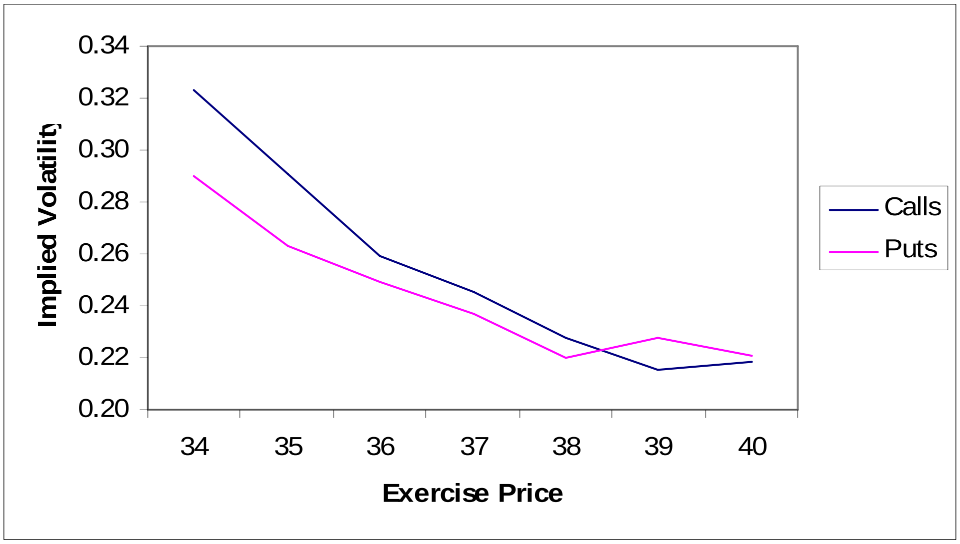

Obviously we do not obtain the same volatility for each option. Moreover, we observe a pattern in which the implied volatilities almost monotonically decrease with higher exercise price.[8] The relationship between the option exercise price and the implied volatility has been documented since at the least the time of the crash of 1987. When first observed, the implied volatilities were u-shaped, giving the appearance of a smile. Hence this relationship was named the *volatility smile*. In more recent years, the smile has mostly disappeared, and the relationship has sometimes been referred to as a skew or even a smirk, which might well describe the picture we see here for puts.[9]

Note also that we obtain different implied volatilities depending on whether we are looking at calls or puts. There is no a priori reason why puts and calls should have different implied volatilities but obviously they do.

The existence of multiple implied volatilities, regardless of whether they arrange themselves in a smile-like pattern, should be somewhat disconcerting. How can the market be telling us that there is more than one volatility for the stock? Clearly there is something wrong with the Black-Scholes model, which is that it fails to consider all of the factors that enter into the pricing of an option. It accounts for the stock price, the

[8]We say *almost* because the monotonicity breaks with the highest exercise prices.

[9]In some cases, the smile is illustrated in a two-dimensional graph in which the third dimension is the time to expiration. As noted in footnote 4, the volatilities could differ by time to expiration, giving rise to the notion of a term structure of volatilities. When the implied volatilities are illustrated with exercise price along one axis and time to expiration along the other, the relationship is known as the *volatility surface*.

exercise price, the time to expiration, the dividends, and the risk-free rate. The implied volatility is more or less a catch-all term, capturing whatever variables are missing, as well as the possibility that the model is improperly specified or blatantly wrong.

What we learn, however, is what we wanted to know: which options are the most expensive. We see that the calls and puts with the lowest exercise price are the most expensive options. But what is so puzzling is that in the Black-Scholes world, no option should be more expensive than any other option. Any option is a perfect substitute for any other option.

The Perfect Substitutability of Options

To see this, consider an option priced at w_1 and a second option on the same stock priced at w_2. These options could have different exercise prices and times to expiration. They could differ in that one is a call and one is a put. It does not matter. They are simply two options on the same stock, which itself is priced at S.

Suppose we wished to buy the first option. Under the assumptions of the Black-Scholes model, we could create the same result by combining the second option with risk-free bonds. The math to demonstrate this result requires essentially a derivation of the Black-Scholes model.[10] We can avoid these technical details by using some simple results from the Black-Scholes formula itself.

We can write a general expression for the value of these options as follows:

$$w_1 = h_{s1}S + h_{b1}X_1e^{-rT_1}$$
$$w_2 = h_{s2}S + h_{b2}X_2e^{-rT_2}.$$

The quantities h_{s1} and h_{s2} represent the holdings of stock necessary to replicate options 1 and 2, and the quantities h_{b1} and h_{b2} represent the holdings of bonds necessary to replicate options 1 and 2.

If option 1 is a call, the Black-Scholes formula is, of course,

$$w_1 = SN(d_1^{(1)}) - X_1e^{-rT_1}N(d_2^{(1)}),$$

[10]The standard derivation of the Black-Scholes model involves constructing a risk-free hedge consisting of a long position in shares and a short position in call options. By dynamically adjusting this combination, a riskless position is maintained. Hence, the long stock-short call combination replicates a risk-free bond. Likewise, a long-stock, short risk-free bond can be shown to replicate a long option. Hence, the stock and bond are perfect substitutes for the option.

where $N(d_1^{(1)})$ and $N(d_2^{(1)})$ are the well-known standard normal probabilities, typically identified as $N(d_1)$ and $N(d_2)$ when working with a single option. These values are computed using the appropriate exercise price X_1 and time to expiration T_1 in the standard formula for d_1 and d_2.[11] Then h_{s1} is $N(d_1^{(1)})$ and h_{b1} is $-N(d_2^{(1)})$.

If option 1 is a put, the formula is

$$w_1 = X_1 e^{-rT_1} N(-d_2^{(1)}) - SN(-d_1^{(1)}).$$

Then h_{s1} would be $-N(-d_1^{(1)})$, and h_{b1} would $N(-d_2^{(1)})$.

If option 2 is a call, the Black-Scholes formula is

$$w_2 = SN(d_1^{(2)}) - X_2 e^{-rT_2} N(d_2^{(2)}),$$

with the $N(d_1^{(2)})$ and $N(d_2^{(2)})$ terms defined in the manner previously described. Then h_{s2} would $N(d_1^{(2)})$, and h_{b2} would be $-N(d_2^{(2)})$.

If option 2 is a put, the formula is

$$w_2 = X_2 e^{-rT_1} N(-d_2^{(2)}) - SN(-d_1^{(2)}),$$

and h_{s2} would be $-N(-d_1^{(2)})$ and h_{b2} would be $N(-d_2^{(2)})$.

From the general expressions, we can see how to replicate any option with any other option on the same stock. Take the general expression for w_2 and solve for S to obtain:

$$S = \frac{w_2 - h_{b2} X_2 e^{-rT_2}}{h_{s2}}.$$

Then substitute this result for S into the expression for w_1:

$$w_1 = \left(\frac{h_{s1}}{h_{s2}}\right) w_2 - \left(\frac{h_{s1}}{h_{s2}}\right) h_{b2} X_2 e^{-rT_2} + h_{b1} X_1 e^{-rT_1}.$$

We see that the first option can be replicated by holding (h_{s1}/h_{s2}) units of option 2, $-(h_{s1}/h_{s2})h_{b2}$ units of a risk-free bond paying X_2 at the expiration of option 2, and h_{b1} units of a risk-free bond paying X_1 at the expiration of option 1.[12] Of course, this position must be dynamically adjusted through time. Nonetheless, any option can be used to

[11]Given an exercise price of X, time to expiration of T, and volatility of σ, the formulas for d_1 and d_2 are

$$d_1 = \frac{\ln(S/X) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = d_1 - \sigma\sqrt{T}.$$

[12]If the two options had either the same time to expiration or the same exercise price, the formula for the holdings would simplify a little.

replicate any other option on the same stock, regardless of exercise price, time to expiration, or whether it is a call or a put. Hence, any option is a perfect substitute for any other option on the same stock.

That being the case, there is no rationale for why any one option should be more expensive than any other option, after accounting for time to expiration, exercise price, and whether it is a put or a call. But clearly the existence of multiple implied volatilities tells us that some options are more expensive than others, and the smile tells us that there is a pattern to the relative costs of these options. Typically the most expensive options are the deep-in-the-money calls and deep out-of-the-money puts.

Unfortunately, no one really knows why certain options are more expensive than others. Deep out-of-the-money puts have oftentimes been viewed as a form of insurance against large drops in the market. Fear of a crash, it is said, leads some investors to pay relatively more for the protection these options afford than for other options. Yet, our Black-Scholes theory says that any option should substitute perfectly for any other option. Yet, theory and reality diverge. Perhaps the dynamic replicating strategy is too complex or too costly to implement. For whatever reason, perfect substitution does not hold in reality.

The Black-Scholes model is a wonderful theory, but it tells us nothing about why anyone would hold an option. It cannot motivate the holding of options, because any option serves as well as any other option. In reality some options are more desirable than others. Whatever factors motivate the holding of options are simply not captured in the Black-Scholes model. Hence, these factors show up hidden within the implied volatility.

One must then wonder why we do not simply throw out the Black-Scholes model. There are two reasons why we do not. First, the model is attractively simple. Although the mathematical details that support the derivation of the model are complex, there are a number of simple approaches to understanding the model.[13] Practitioners have shown that they not only can understand the model, but that they can also accept it. The computational demands of the model are also quite modest. Perhaps its greatest virtue is

[13]For example, almost anyone can understand the binomial model. The Black-Scholes model can then be seen as what happens when the fixed life of the option is divided into an increasingly larger number of time steps. The Black-Scholes model can also be seen as giving the price of an option as the discounted expectation of the option payoff when the holder is risk neutral.

that it requires so little information. But therein lies its vice: by requiring so little information, the model almost surely misses factors that explain why options are held and why some investors would pay more for certain options.

But the model is admired so much for its simplicity that practitioners continue to use it. The volatility smile is the price paid for oversimplification.

Other Attempts to Explain the Implied Volatility Smile

When the volatility smile was first observed, some researchers believed that the explanation was liquidity. The true "smile" appearance meant that out-of-the-money options had the highest implied volatilities. These options were also the least liquid; hence, it was argued that the prices observed for these options of low liquidity reflected the thinness of their markets. But this explanation would suggest that highly liquid options – typically those trading nearly at-the-money – would have the same implied volatilities. In fact, they do not and never did. Moreover, when the smile turned into a skew, the moneyness argument fell by the wayside.

Other researchers believe that the smile reflects stochastic volatility. Volatility is surely not constant as assumed in the Black-Scholes model. If volatility is stochastic, researchers argue that the smile reflects the failure of the Black-Scholes model to capture the random nature of volatility.

Others argue the Black-Scholes model, which assumes that stock prices fluctuate in a smooth and continuous manner, fails to capture the true nature of stock price movements, which are observed to have discrete jumps.

The arguments of stochastic volatility and jumps have a great deal of appeal, because they in some sense preserve many of the essential features of the Black-Scholes model. These arguments do not require that the model motivate the holding of options and the preference for some investors for certain options over others. The essentially argue that if the Black-Scholes model were re-derived under looser assumptions, the smile would go away. Unfortunately, once these looser assumptions are introduced, the mathematical tractability of the model is lost, and the process of pricing an option becomes one of making other heroic assumptions or imposing severe computational demands.[14] It is fair to say that mathematicians have devoted excessive hours of human and machine time to researching the smile with little if any regard to the reasons why the smile exists.[15]

How Practitioners Use the Implied Volatility Smile

Practitioners seem capable of operating in a world of volatility smiles. They even use the smile to simplify how they trade. For example, they oftentimes quote option prices not in terms of the actual price but in terms of the implied volatility. For example, a dealer might indicate an intention to sell the January 36 call by quoting a volatility of 25.92. This statement is interpreted to mean that the actual price is derived from the Black-Scholes model using a volatility of 25.92. Assuming agreement on the dividends and risk-free rate, such a quote for this option would lead to a price of $1.85. By quoting prices this way, traders can immediately see which options are truly more expensive, that is, after accounting for moneyness, time to expiration, and whether the option is a call or a put.

Traders also will often express the smile in terms of the implied volatility associated with the option's delta. For example, an at-the-money call has a delta of about 0.50. Traders would often then refer to this as a 50 delta call. The volatilities are then graphed with the deltas on the horizontal axis. The resulting smile would be the same image as we have seen here.

But in general, traders use the smile as a way of determining the values of options relative to each other, after taking into account obvious differences related to moneyness, time to expiration, and the type of option (call or put).

Concluding Comments

The whole notion of implied volatility and the existence of the volatility smile are results of using a model that does not capture everything that affects the prices of options. Practitioners and academics largely accept the limitations of the model and consider the smile a means of forcing the model to reveal information it is not designed to reveal.

[14]For example, one heroic assumption is that the risk arising from stochastic volatility is non-priced risk. That is, the risk associated with uncertain volatility is a risk that does not concern investors. They are either neutral toward that risk or the risk is uncorrelated with their other holdings, meaning that the risk is diversifiable. Once these assumptions are invoked, the financial economics are lost and what remains is simply an exercise in computational finance.

[15]Perhaps if these reasons were found, the mathematicians would be out of work.

However flawed the model may be, the advantages of the Black-Scholes model and its attendant defects may outweigh the disadvantages of other more complex models.

References

The seminal articles on the smile are:

Derman, E. and I. Kani. "Riding on a Smile." *Risk* 7 (February, 1994), 32-39.

Dupire, B. "Pricing with a Smile." Pricing with a Smile." *Risk* 7 (January, 1994), 18-20.

Rubinstein, M. "Implied Binomial Trees." *The Journal of Finance* 69 (1994), 771-818.

An excellent treatment of implied volatility and the smile is in:

N. Chriss. *Black-Scholes and Beyond: Option Pricing Models* Chicago: Irwin (1997), Chs. 8, 9, 10.