

CHAPTER 1**Option Pricing**

In order to profitably trade options we need a model for valuing them. This is a framework we can use to compare options of different maturities, underlyings, and strikes. We do not insist that it is in any sense true or even a particularly accurate reflection of the real world. As options are highly leveraged, nonlinear, time-dependent bets on the underlying, their prices change very quickly. The major goal of a pricing model is to translate these prices into a more slowly moving system.

A model that perfectly captures all aspects of a financial market is probably unobtainable. Further, even if it existed it would be too complex to calibrate and use. So we need to somewhat simplify the world in order to model it. Still, with any model we must be aware of the simplifying assumptions that are being used and the range of applicability. The specific choice of model isn't as important as developing this level of understanding.

THE BLACK-SCHOLES-MERTON MODEL

We present here an analysis of the Black-Scholes-Merton (BSM) equation. The BSM formalism becomes the conceptual framework for an options trader: In the same way that we hear our own thoughts in English, an experienced derivatives trader thinks in the BSM language.

The standard derivation of the BSM equation can be found in any number of places (for example, Hull 2005). While good derivations carefully lead us through the mathematics and financial assumptions, they don't

8

VOLATILITY TRADING

generally make it obvious what to do as a trader. We must always remember that our goal is to identify and profit from mispriced options. How does the BSM formalism help us do this?

Here we approach the problem backwards. We start from the assumption that a trader holds a delta hedged portfolio consisting of a call option and Δ units of short stock.¹ We then apply our knowledge of option dynamics to derive the BSM equation.

Even before we make any assumptions about the distribution of underlying returns, we can state a number of the properties that an option must possess. These should be financially obvious.

- A call (put) becomes more valuable as the underlying rises (falls), as it has more chance of becoming intrinsically valuable.
- An option loses value as time passes, as it has less time to become intrinsically valuable.
- An option loses value as rates increase. Since we have to borrow money to pay for options, as rates increase our financing costs increase, ignoring for now any rate effects on the underlying.
- The value of a call (put) can never be more than the value of the underlying (strike).

As we have said, even before the invention of the BSM formalism, option traders were aware that directional risk could be mitigated by combining their options with a position in the underlying. So let's assume we hold the delta hedged option position,

$$C - \Delta S_t \quad (1.1)$$

where

- C is the value of the option
- S_t is the underlying price at time t
- Δ is the number of shares we are short

¹That this is a standard delta hedged portfolio should be common knowledge for option traders or indeed anyone who has previously seen any sort of derivation of BSM. Indeed, traders knew about delta hedging long before BSM. For an interesting history refer to Haug 2007a. But even if this is the first derivation of BSM the reader has seen, this shouldn't be seen as a remarkable fact. A call (put) option gains (declines) in value as the underlying rises. So in principle we can offset this directional risk with a position in the underlying. This should be obvious. The details of exactly how much of the underlying to hold are certainly not meant to be obvious.

Over the next time step the underlying changes to S_{t+1} . The change in the value of the portfolio is given by the change in the option and stock positions together with any financing charges we incur by borrowing money to pay for the position.

$$C(S_{t+1}) - C(S_t) - \Delta(S_{t+1} - S_t) + r(C - \Delta S_t) \quad (1.2)$$

The last term is written as positive because we know that the value of the long call/short stock portfolio will be negative (or at most zero when Δ is 1) and hence will have us lending money and thus receiving interest income. Note also that we assume the time step is small enough that we can take delta to be unchanged.

The change in the option value due to the underlying price change can be approximated by a second-order Taylor expansion. Also, we know that when other things are held constant, the option will decrease due to the passing of time by an amount denoted by θ .

So we get

$$\Delta(S_{t+1} - S_t) + \frac{1}{2}(S_{t+1} - S_t)^2 \frac{\partial^2 C}{\partial S^2} + \theta - \Delta(S_{t+1} - S_t) + r(C - \Delta S_t) \quad (1.3)$$

Or

$$\frac{1}{2}(S_{t+1} - S_t)^2 \Gamma + \theta + r(C - \Delta S_t) \quad (1.4)$$

where Γ is the second derivative of the option price with respect to the underlying.

Expression (1.4) gives the change in value of the portfolio, or the profit the trader makes when the stock price changes by a small amount. It has three separate components:

1. The first term gives the effect of gamma. Since gamma is positive, the option holder makes money. The return is proportional to half the square of the underlying price change.
2. The second term gives the effect of theta. The option holder loses money due to the passing of time.
3. The third term gives the effect of financing. Holding a hedged long option portfolio is equivalent to lending money.

Further, we see in Chapter 2 that on average

$$(S_{t+1} - S_t)^2 \cong \sigma^2 S^2$$

where σ is the standard deviation of the underlying's returns, generally known as *volatility*.

So we can rewrite expression (1.4) as

$$\frac{1}{2}\sigma^2 S^2 \Gamma + \theta + r(C - \Delta S_t) \quad (1.5)$$

If we accept that this position should not earn any abnormal profits because it is riskless and financed with borrowed money, the expression can be set equal to zero. Therefore the equation for the fair value of the option is

$$\frac{1}{2}\sigma^2 S^2 \Gamma + \theta + r(C - \Delta S_t) = 0 \quad (1.6)$$

Before continuing, we need to make explicit some of the assumptions that this informal derivation has hidden.

- In order to write down expression (1.1) we needed to assume the existence of a tradable underlying asset. In fact, we assume that it can be shorted and the underlying can be traded in any size necessary without incurring transaction costs.
- Expression (1.2) has assumed that the proceeds from the short sale can be reinvested at the same interest rate at which we have borrowed to finance the purchase of the call. We have also taken this rate to be constant.
- Expression (1.3) has assumed that the underlying changes are continuous and smooth. Further, we have considered second-order derivatives with respect to price but only first-order with respect to time.

But something about which we haven't made any assumptions at all is whether the underlying has any drift. This is remarkable. We may naively assume that an instrument whose value increases as the underlying asset rises would be dependent on its drift. However, the effect of drift can be negated by combining the option with the share in the correct proportion. As the drift can be hedged away, the holder of the option is not compensated for it. When we consider hedging later in Chapter 4, we find that in the real world, where the assumptions about continuity fail, directional dependence will reemerge.

However, note that while the price change does not appear in equation (1.6), the square of the price change does through the volatility term. So the magnitude of the price changes is central to whether the trader makes a profit with a delta hedged position. This is true whether or not returns are

normally distributed. As long as the variance of returns is finite, this result holds. In fact, if we had included higher-order price terms in the Taylor expansion, we would see that the option's price change also depended on higher-order price differences.

With appropriate final conditions, equation (1.6) holds for a variety of instruments: European and American options, calls and puts, and many exotics. It can be solved with any of the usual methods for solving partial differential equations.

In this exercise we have derived a form of the BSM equation by working backwards from our trader's knowledge of how options react to changes in underlying and time. In doing so, it has given us what we need to know to trade options from the point of volatility.

We have shown how the fair price for an option is related to the standard deviation of the underlying returns. Since at any time there is an option market and the underlying market, there are two ways we can proceed:

1. Using the quoted price of the option, calculate the implied standard deviation or volatility.
2. Using an estimate of the volatility over the life of the option, calculate a theoretical option price.

If our estimate of volatility differs significantly from that implied by the option market, then we can trade the option accordingly. If we forecast volatility to be higher than that implied by the option, we would buy the option and hedge in the underlying market. Our expected profit would depend on the difference between implied volatility and realized volatility. Equation (1.6) says that instantaneously this profit would be proportional to

$$\frac{1}{2}S^2\Gamma(\sigma^2 - \sigma_{implied}^2) \quad (1.7)$$

A complementary way to think of the expected profit of a hedged option is by considering vega. Vega is defined as the change in value of an option if implied volatility changes by one point (e.g., from 19 to 18 percent). This means that if we buy an option at $\sigma_{implied}$ and volatility immediately increases to σ we would make a profit of

$$vega(\sigma - \sigma_{implied}) \quad (1.8)$$

If we have to hold the option to expiration and realized volatility averages σ we will also make this amount, but only on average. The *vega profit*

is realized as the sum of the hedges as we rebalance our delta. This can be formalized by noting the relationship between vega and gamma,

$$vega = \sigma TS^2 \Gamma \quad (1.9)$$

So expression (1.7) can also be written as

$$\frac{vega}{\sigma T} (\sigma^2 - \sigma_{implied}^2) \quad (1.10)$$

The problem this presents is that the gamma is highly dependent on the moneyness of the option, which obviously changes as the underlying moves around. So the profit is highly volatile and path dependent. We examine this further in Chapter 5.

It is perfectly acceptable to make simplifying assumptions when developing a model. It is totally unacceptable to make assumptions that are so egregiously incorrect that the model is useless, even as a basic guide. So before we go any further we look at how limiting our assumptions really are.

- We assumed that the underlying was a tradable asset. While the BSM formalism has been extended to cases where this is not true, notably in the pricing of real options, we are primarily concerned with options on equities and futures, so this assumption is not restrictive. However, on many optionable underlyings liquidity is an issue, so *tradable* is not always a clearly defined quality.
- We assumed that the underlying pays no dividends or any other income. This changes the equations slightly, but the same principles still apply.
- We also assumed this asset was able to be shorted. This is not a problem where the underlying is a future but when it is a stock, shorting is often more difficult. For example, in the United States, stocks can only be shorted on an uptick. Further, even when shorting is achievable, the short seller rarely receives the full proceeds of the sale for investment, as fees must be paid to borrow the stock. This can be accounted for synthetically by assuming an extra dividend yield on the underlying, equal to the penalty cost associated with shorting the stock.
- Interest rates have a bid/ask spread. We cannot invest the proceeds of a sale at the same rate at which we borrow. Further, rates are not constant. However, the BSM is often used to price options on bonds and money market rates which would have no volatility if this assumption was valid. We can get away with this because the risk due to interest charges (*rho*) is insubstantial in comparison to other risks, at least for short-dated options.

- We have assumed that volatility is a constant, neither a function of time nor of the underlying price. Not only is this untrue but we will be actively trying to trade these changes. There are models that explicitly take into account volatility changes. However, we choose to recognize this limitation and learn to use the BSM model anyway. This is consistent with our philosophy of the model as a framework for organizing our thoughts rather than as an accurate depiction of reality.
- We assumed that volatility is the only parameter needed to specify the distribution of the underlying returns. The mean can be hedged away and we have ignored higher-order moments. This is the same as assuming a normal return distribution or a lognormal price distribution. The fact that this is incorrect leads to the well-known phenomenon of the *volatility smile*, where implied volatility is a function of strike. In essence, implied volatility is the wrong number we put into the wrong formula to get the correct option price. This can be rectified in several ways. In Chapter 3 we present methods of quantifying the implied skewness and kurtosis.
- We have assumed that the underlying's changes are continuous so we can continually adjust our hedge. This is not true. Sometimes the underlying has vast jumps. For example, it isn't uncommon for the shares of a biotech company to jump by 70 to 80 percent in one day. Modifications have been made to the BSM formalism to price options in these circumstances (Merton 1976), but this isn't really the point. These jumps cannot be hedged and the replication strategy fails utterly. We have to learn to hedge this risk with other options. This is the concept of semistatic hedging that traders need to use in practice.

The BSM model is remarkably robust. Most of the assumptions that we need to derive the equation can be loosened without destroying the model's utility, but note that we will only use the BSM paradigm as a pricing method, not a risk control method. It is useful to translate the fast-moving option prices into a slow-moving parameter, implied volatility, which can be compared to the estimated realized volatility. But risk control must be handled separately. Traders should never think about extreme risk in terms of the moments of the Gaussian distribution. Asking, "What happens if IBM moves five standard deviations?" is useful only in normal situations (where *normal* is defined as those times when moves are well described by the Gaussian distribution). We must also always be aware of what happens if IBM drops by 50 percent, despite the fact that this has never happened. Merton argued that these extreme jumps could be diversified away. The bad news is that traders have to hope he was correct. Tail risk can often be capped by trading far out-of-the-money options, and keeping individual positions to a small proportion of the total portfolio can

also help. But generally we get paid for taking risks. Just try to be aware of the risks you have an edge in and those you don't. And never estimate the magnitude of risks from within the same model that you priced them with.

SUMMARY

We always need to remember that a pricing model is not meant to be an accurate reflection of reality. It is a framework to think within, and a consistent intellectual backdrop that helps us build intuition. We eventually learn to express our trading thoughts in terms of the model. The most important things to remember about the BSM model are:

- The drift of the underlying can be hedged away.
- The magnitude of the underlying price moves cannot.
- The BSM model is a model for finding trades, not a model for controlling risk.
- The assumptions used in the derivation of the model need to be remembered at all times.

<http://www.pbookshop.com>