

Option Pricing

It is possible to trade options without any valuation model. For example, traders might buy a call option because they think the underlying will rally further past the strike than the price they have paid. This is the simplest, most direct use of options. At a level of complexity only slightly greater than this we can trade volatility without a model. Traders might sell a straddle because they think the underlying will expire closer to the strike than the value of the straddle. There are an enormous number of option positions like this where traders can attempt to profit from their opinion of the future distribution of the underlying. However, if we want to express an opinion based on the behavior of the underlying before expiration, we will need a model.

A model 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 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 Black-Scholes-Merton Model

We will present 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 thoughts in English, an experienced derivatives trader thinks in the BSM language. This is an important difference between the models used by traders and the models used in a hard science such as physics. Models in physics aim to make statements about the world that are at least in some sense true, and then use the model to make predictions. The degree of truth needn't be consistent between all models. There are some successful theories that are based on highly simplified phenomenological models. An example would be Rutherford's model of the atom, which assumes that electrons orbit the nucleus like planets orbit the sun. This contains some truth: The atom consists of electrons and nuclear particles, but the planetary model isn't an accurate depiction of atomic structure.

Trading models are fundamentally different. The BSM model isn't good because it is an accurate representation of reality. It is actually fairly poor in this regard, with most of the model's assumptions being gross oversimplifications. It is a good model because the weaknesses are well understood and the model gives results that are intuitively sensible. The model fits its purpose. It is useful. It makes as little sense to say it is correct or incorrect as to say that German is incorrect and French is correct.

The standard derivation of the BSM equation can be found in any number of places (for example, Hull 2005). Although good derivations carefully lead us through the mathematics and financial assumptions they don't 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 backward. We start from the assumption that a trader holds a delta-hedged portfolio consisting of a call option and Δ units of short stock. We then apply our knowledge of option dynamics to derive the BSM equation.

That this portfolio is delta hedged should be obvious to option traders. Actually, 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 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 obvious.

Even before we make any assumptions about the distribution of the underlying's 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.
- The value of a call (put) can never be more than the value of the underlying (strike).
- An option loses value as time passes, as it has less time to become intrinsically valuable.
- An option must have positive dependence on uncertainty. If the underlying had no risk there would be no need to pay for a product that only has value in certain states of the world. Options only have value because we are uncertain about the future, so it follows that the more uncertain we are the more valuable the options will be.
- An option loses value as rates increase. Because 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.
- Dividends (and storage or borrowing costs) have different effects on calls and puts. The holder of an option does not receive the dividend. This means that a dividend lowers the effective value of the underlying stock for the purposes of option valuation. So a dividend increases the value of a put and lowers the value of a call.

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 , and Δ is the number of shares we are short. 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.

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

To see why the last term is positive we need to consider our cash flows. We bought the option, so we need to finance that cost, but we shorted stock so we receive money for this. Over a single time step we gain $r\Delta S_t$ from this.

Note also that we assume that 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 .

At this point in our argument we have assumed that we need to consider second derivatives with respect to price but only first derivatives with respect to time. Why is either of these choices valid? Ignoring higher derivatives with respect to price really cannot be justified at this point. We have only done it because we are trying to recover the BSM equation. In a more formal derivation this would be related to the assumption of a normal distribution of underlying returns. This is a major simplification that I am not ignoring. I’m postponing the discussion until later. The assumption that we need fewer derivatives with respect to time is easier to justify. Underlying price changes are stochastic and so they are a source of *risk*. Time change is predictable and the effect of time on options is merely a *cost*.

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. Equation 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 the next chapter 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 Equation 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 equation 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.

- To write down Equation 1.1, we needed to assume the existence of a tradable underlying asset. In fact we assumed that it could be shorted and the underlying could be traded in any size necessary without incurring transaction costs.
- Equation 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.
- Equation 1.3 has assumed that the underlying changes are continuous and smooth. And as we mentioned earlier, we have considered second order derivatives with respect to price but only first order with respect to time. This is a very limiting assumption and will be returned to in some depth.

But something that we haven't made any assumptions about 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. Later in the chapter on hedging we see that in the real world, where the assumptions about continuity fail, directional dependence reemerges.

However, note that although the price change does not appear in Equation 1.6, the square of the price change does. 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 returns are normally distributed or not. This result holds as long as the variance of returns is finite. 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. The closed forms for these solutions (when they





exist) can be found in a number of texts (e.g., Hull 2005; Sinclair 2010). A trader needs to understand how the solutions depend on changes in the pricing variables and the volatility parameter. I assume deep familiarity with this behavior.

In this exercise we have derived a form of the BSM equation by working backward from our trader's knowledge of how options react to changes in underlying and time. In doing so, it has given us much of 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's returns. Because we have assumed that at any time there is an option market and the underlying market, there are two ways we can use what we have learned:

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

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 *partial derivative of the option price with respect to implied volatility*. It is generally expressed 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

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

The relationship between the instantaneous profit of Equation 1.7 and the total profit of Equation 1.8 could be proved by integrating 1.7 over time and using the relationship between gamma and vega,

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

but this provides little insight.





Instead imagine we have a call, C , originally priced with volatility, $\sigma_{implied}$, and this changes to σ . Define $\delta = \sigma^2 - \sigma_{implied}^2$

To first order in variance

$$C(\sigma_{implied}^2 + \delta) = C(\sigma_{implied}^2) + \delta \frac{\partial C}{\partial(\sigma^2)} \quad (1.10)$$

and

$$\frac{\partial C}{\partial(\sigma^2)} = \frac{\partial C}{\partial \sigma} \frac{\partial \sigma}{\partial(\sigma^2)} = vega \times \frac{1}{2\sigma} \quad (1.11)$$

So the second term of the Equation 1.10, the profit and loss (P/L or P&L) term, is

$$\begin{aligned} \delta \times vega \times \frac{1}{2\sigma} &= \frac{vega}{2\sigma} (\sigma^2 - \sigma_{implied}^2) \\ &= \frac{vega}{2\sigma} (\sigma - \sigma_{implied})(\sigma + \sigma_{implied}) \\ &\approx vega(\sigma - \sigma_{implied}) \end{aligned} \quad (1.12)$$

where the last line follows from the fact that the initial and final volatilities are comparable in size. This derivation is not rigorous but the result holds in general.

This form of the P/L equation is the more useful to traders, who are generally more interested in total profit than instantaneous profit. It is also easier to think about as it is linear in volatility. 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.

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 7.

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.

■ Modeling Assumptions

Existence of a Tradable Underlying

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. If we encounter situations where we are unable to trade the underlying in the size we need, we will be in trouble.

Absence of Dividends or Storage Costs

We assumed that the underlying pays no dividends or any other income. Note that in Equation 1.2 we associated the risk-free rate, r , with both the financing of the call premium and the hedge portfolio, ΔS . This need not be the case.

- If the underlying pays a dividend with a yield of q , the second term would need to be associated with $r - q$ instead.
- A continuous dividend yield is often an appropriate approximation for indices but stocks pay discrete dividends. Here we would need to modify the approach by assuming that the true underlying is the stock minus the discounted value of the dividends. This complicates the equations but does not modify the spirit of the argument.
- A short seller rarely receives the full proceeds of a sale for investment. Shorting a stock is a privilege a broker extends to customers and this generally needs to be paid for. This can be accounted for synthetically by assuming a fake dividend yield to reflect these costs.
- If the underlying is a physical commodity that incurred a storage cost at a rate of q^* , the rate associated with the hedge would need to be $r + q^*$.
- If the underlying was a future the hedge would be costless to finance. In this case the rate associated with the hedge would be zero.

Ability to Short the Underlying

This is not a problem where the underlying is a future but when it is a stock, shorting is often more difficult. 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, up to the penalty cost associated with shorting the stock.





The Existence of a Single Constant Interest Rate

Interest rates have a bid/ask spread. We cannot invest the proceeds of a sale at the same rate at which we borrow. The BSM formalism can be modified to take this into account (Bergman 1995) but the equations become intractable.

Further, rates are not constant. Even though this is an assumption of BSM the theory is still often used to price options on bonds and money market rates, which would have no volatility if this assumption were valid. We can get away with this because at least for short-dated options, the risk due to interest charges (ρ) is insubstantial in comparison to other risks.

Absence of Taxes

We have assumed that there are no taxes. In reality the fact that different market participants may have different tax liabilities can create trading opportunities and pitfalls. This occurs most frequently with dividends where foreign investors are often taxed at significantly different rates to domestic investors. Traders should always remember that they must value the option based on what it is worth to them, not the marginal investor, which is where the market will be pricing it.

The Underlying Can Be Traded in Any Size

We have already stated that problems will occur if we need to trade larger than the market can handle. But the derivation also assumes that we can trade as small as we need, including fractions of shares. Clearly this is impossible and if our brokers are charging minimum ticket charges it may be uneconomical to trade smaller than blocks of 100 shares. This practical limitation will be addressed in Chapter 6 when we examine methods for hedging at discrete intervals.

It Is Costless to Trade the Underlying

This is closely related to the preceding point. Trading the underlying always incurs costs: brokerage, clearing, or bid-ask spreads. These costs will dampen our desire to hedge continuously (even if this was possible) as





our risk reduction from hedging now needs to be balanced against the costs of doing so. We study this extensively in Chapter 6.

Volatility Is Constant

In our derivation of BSM, we have assumed that volatility is a constant, neither a function of time nor the underlying price. The fact that we have started to discuss vega, the effect of a change in implied volatility should highlight the inconsistency of this approach. Not only is the basic assumption 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.

Assumptions about the Distribution of Returns

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 log-normal price distribution. In Chapter 3, we look at the statistics of real markets and see that this is not the case. 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 5, we present methods of quantifying the implied skewness and kurtosis.

We have also 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 semi-static hedging that traders need to use in practice.





■ Conclusion

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 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. It is also well suited to comparing different options to each other. Even if most of our assumptions are incorrect it is likely that they will impact the pricing of a 50-delta call and a 40-delta call in similar ways. This makes our estimate of the spread somewhat more robust than our estimate of the individual options. With a little less confidence we can extend the argument to compare options on different underlyings.

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 where 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 (Merton 1976). Unfortunately, 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

Models are not magic. In particular, option pricing models don't really "price" options. They instead map option prices to a slower moving parameter, implied volatility. This simplification allows us to compare options with different strikes, maturities, and underlyings.

The BSM model is one of the oldest most tested models. With enough ad hoc modifications, it can be used to price most exchange traded options. It isn't essential that this model is chosen, but I recommend it for its



robustness, simplicity, and the fact that it is almost the universal language of the option markets. Its most important properties are:

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