Options

A Summary

Option Pricing Models

> *Since all models are wrong the scientist must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad.*
>
> —Box (1976)

Some models are wrong in a trivial way. They clearly don't agree with real financial markets. For example, an option valuation model that included the return of the underlying as a pricing input is trivially wrong. This can be deduced from put-call parity. Imagine a stock that has a positive return. Naively this will raise the value of calls and lower the value of puts. But put-call parity means that if calls increase, so do the values of the puts. Including drift leads to a contradiction. That idea is trivially wrong.

Every scientific model contains simplifying assumptions. There actually isn't anything intrinsically wrong with this. There are many reasons why this is the case, because there are many types of scientific models. Scientists use simplified models that they know are wrong for several reasons.

A reason for using a wrong theory would be because the simple (but wrong) theory is all that is needed. Classical mechanics is still widely used in science even though we now know it is wrong (quantum mechanics is needed for small things and relativity is needed for large or fast things). An example from finance is assuming normally distributed returns. It is doubtful anyone ever thought returns were normal. Traders have long known about extreme price moves and Osborne (1959), Mandelbrot (1963), and others studied the non-normal distribution of returns from the 1950s. (Mandelbrot cites the work of Mitchell [1915], Olivier [1926], and Mills [1927], although this research was not well known.) The main reason early finance theorists assumed normality was because it made the equations tractable.

Sometimes scientists might reason through a stretched analogy. For example, Einstein started his theory of the heat capacity of a crystal by first assuming the crystal was an ideal gas. He knew that this was obviously not the case. But he thought that the idea might lead to something useful. He had to start somewhere, even if he knew it was the wrong place. This model was metaphorical. A metaphorical model does not attempt to describe reality and need not rely on plausible assumptions. Instead, it aims to illustrate a non-trivial mechanism, which lies outside the model.

Other models aim to mathematically describe the main features of an observation without necessarily understanding its deeper origin. The GARCH family of volatility models are phenomenological, and don't tell us why the GARCH effects exist. Because these models are designed to describe particular features, there will be many other things they totally ignore. For example, a GARCH process has nothing to say about the formation of the bid-ask spread. The GARCH model is *limited,* but not *wrong*.

The most ambitious models attempt to describe reality as it truly is. For example, the physicists who invented the idea that an atom was a nucleus around which electrons orbited thought this was actually what atoms were like. But they still had to make simplifying assumptions. For example, when formulating the theory, they had to assume that atoms were not subject to gravity. And, in only trivial situations could the equations be analytically solved. The Black-Scholes-Merton (BSM) model was meant to be of this type.

But it isn't used that way at all.

The inventors of the model envisaged that the model would be used to find a fair value for options. Traders would input the underlying price, strike, interest rate, expiration date, and volatility and the model would tell them what the option was worth. The problem was that the volatility input needed to be the volatility over the life of the option, an unknown parameter. Although it was possible for a trader to make a forecast of future volatility, the rest of the market could and did make its own forecast. The market's option price was based on this aggregated estimate. This is the implied volatility, which became the fundamental parameter. Traders largely didn't think of the model as a predictive valuation tool but just as an arbitrage-free way to convert the quickly changing option prices into a slowly changing parameter: implied volatility. For most traders, BSM is not a predictive model; it is just a simplifying tool.

This isn't to say that BSM can't be used as a pricing model to get a fair value. It absolutely can. But even traders who do this will think in volatility terms. They will compare the implied volatility to their forecast volatility, rather than use the forecast volatility to price the option and compare it to the market value. By using the model backwards, these traders still benefit from the way BSM converts the option prices into a slowly varying parameter.

We need to examine the effects of the model assumptions in light of how the model is used. Although the assumptions make the model less realistic, this isn't important. The model wasn't used because it was realistic; it was used because it was useful.

Obviously, it is possible to trade options without any valuation model. This is what most directional option traders do. We can also 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. However, to move beyond directional trading or speculating on the value of the underlying at expiration we will need a model.

The BSM model is still the benchmark for option pricing models. It has been used since 1973 and has direct ancestors dating to the work of Bachelier (1900) and Bronzin (1906). In terms of scientific theory, this age makes it a dinosaur. But just as dinosaurs were the dominant life form for about 190 million years for a reason, BSM has persisted because it is good.

We want an option pricing model for two reasons.

The first is so we can reduce the many, fast-moving option prices to a small number of slow-moving parameters. Option pricing models don't really *price* options. The market prices options though the normal market forces of supply and demand. Pricing models convert the market's prices into the parameters. In particular, BSM converts option prices to an implied volatility parameter. Now we can do all analysis and forecasting in terms of implied volatility, and if BSM was a perfect model, we would have a single, constant parameter.

The second reason to use a pricing model is to calculate a delta for hedging. Model-free volatility trading exists. Buying or selling a straddle (or strangle, butterfly, or condor, etc.) gives a position that is primarily exposed to realized volatility. But it will also be exposed to the drift. The most compelling reason to trade volatility is that it is more predictable than returns (drift) and the only way to remove this exposure is to hedge. To hedge we need a delta and for this we need a model. This is the most important criterion for an option trader to consider when deciding if a model is good enough. Any vaguely sensible model will reduce the many option prices to a few parameters, but a good model will let us delta hedge in a way that captures the volatility premium.

In this chapter we will examine the BSM model and see if it can meet this standard. By *BSM model* I mean the partial differential equation rather than the specific solution for European vanilla options. The particular boundary conditions and solution methods aren't a real concern here.

Derivations of BSM can be found in many places (see Sinclair, 2013, for an informal derivation). Here we will look at how the model is used.

Option Trading Theory

Here we will very briefly summarize the theory of option pricing and hedging. For more details refer to Sinclair (2010; 2013).

An option pricing model must include the following variables and parameters:

- Underlying price and strike; this determines the moneyness of the option.
- Time until expiration.
- Any factors related to carry of either the option or the underlying; this includes dividends, borrow rates, storage costs, and interest rates.
- Volatility or some other way to quantify future uncertainty.

A variable that is not necessary is the expected return of the underlying. Clearly, this is important to the return of an option, but it is irrelevant to the instantaneous value of the option. If we include this drift term, we will arrive at a contradiction. Imagine that we expect the underlying to rally. Naively, this means we would pay more for a call. But put-call parity means that an increase in call price leads to an increase in the price of the put with the same strike. This now seems consistent with us being bearish. Put-call parity is enough to make the return irrelevant to the current option price, but (less obviously perhaps) it is also enforced by dynamic replication.

This isn't an option-specific anomaly. There are many situations in which people agree on future price change, but this doesn't affect current price. For example, Ferrari would be justified in thinking that the long-term value of their cars is higher than their MSRP. But they can build the car and sell it at a profit right now. Their replication value as a manufacturer guarantees a profit without taking future price changes into account. Similarly, market-makers can replicate options without worrying about the underlying return. And if they do include the return, they can be arbed by someone else.

The canonical option-pricing model is BSM. Ignoring interest rates for simplicity, The BSM PDE for the price of a call, C, is

$$\frac{\sigma^2 S^2}{2}\frac{\partial^2 C}{\partial S^2} + \frac{\partial C}{\partial t} = 0 \tag{1.1}$$

where S is the underlying price, σ is the volatility of the underlying, and t is the time until expiration of the option.

Or using the standard definitions where Γ is the second partial derivative of the option price with respect to the underlying and θ is the derivative of the option price with respect to time,

$$\frac{\sigma^2 S^2}{2}\Gamma + \theta = 0 \tag{1.2}$$

This is then solved using standard numerical or analytical techniques and with the final condition being the payoff of the particular option.

This relationship between Γ and θ is crucial for understanding how to make money with options. Imagine we are long a call and the underlying stock moves from S_t to S_{t+1}. The delta P/L of this option will be the average of the initial delta, Δ, and the final delta, all multiplied by the size of the move. Or

$$P/L = \left(\frac{\Delta}{2} + \frac{(\Delta + \Gamma(S_{t+1} - S_t))}{2}\right)(S_{t+1} - S_t) \tag{1.3}$$

If this option was initially delta-hedged, the P/L over this price move would be

$$P/L = \frac{\Gamma(S_{t+1} - S_t)^2}{2} \tag{1.4}$$

Next note that

$$(S_{t+1} - S_t)^2 \simeq \sigma^2 S_t^{\,2} \tag{1.5}$$

so that the profit in hedging over each time interval is

$$P/L = \frac{\Gamma \sigma^2 S_t^{\,2}}{2} \tag{1.6}$$

(Although equation 1.6 is only asymptotically true, if we worked with an infinitesimal price change, this derivation would be exact.

This is the first term of the BSM differential equation. Literally, BSM says that these profits from rebalancing due to gamma are exactly equal to the theta of the option. Expected movement cancels time decay. The only way a directionally neutral option position will make money is if the option's implied volatility (which governs theta) is not the same as the underlying's realized volatility (which determines the rebalancing profits). This is true no matter which structure is chosen and the particulars of the hedging scheme.

If we can identify situations where this volatility mismatch occurs, the expected profit from the position will be given by

$$P/L = \text{Vega}|\sigma_{implied} - \sigma_{realized}| \tag{1.7}$$

This is the fundamental equation of option trading. All the "theta decay" and "gamma scalping" profits and losses are tied up in this relationship.

Note also that this vega P/L will affect directional option trades. If we pay the wrong implied volatility level for an option, we might still make money but we would have been better off replicating the option in the underlying.

The BSM equation depends on a number of financial and mathematical assumptions.

- The underlying is a tradable asset.
- There is a single, risk-free interest rate.
- The underlying can be shorted.
- Proceeds from short sales can be invested at the risk-free rate.
- All cash flows are taxed at the same rate.
- The underlying's returns are continuous and normally distributed with a constant volatility.

Traders have devised various workarounds to address these limiting assumptions (see Appendix One). The most important of these is the concept of the implied volatility surface. If the BSM were an accurate descriptive theory, all options on a given underlying would have one volatility. This is not true. For the BSM equation to reproduce market option prices, options with different strikes have different implied volatilities (the *smile*) and options with different maturities have different implied volatilities (the *term structure*). These implied volatilities make up the IV surface. An example is in shown in Figure 1.1.

The IV surface exists partially because the BSM is mathematically misspecified. The underlying does not have returns that are continuous and normally distributed with a constant volatility. However, even a model that perfectly captured the underlying dynamics would need a fudge factor like the implied volatility surface. Some of the reasons for its existence have nothing to do with the underlying. Different options have different supply and demand, and these distort option prices. Because of this, there is often an edge in selling options with high volatilities relative to others on the same underlying (see the section on the implied skewness premium in Chapter Four).

Equation 1.7 gives the average PL of any hedged option position, but there is a wide dispersion of results for this mean, and the spread of this distribution decreases with the number of hedges. Figure 1.2 shows the PL distribution of a short straddle that is never re-hedged, and Figure 1.3 shows the distribution when the straddle is re-hedged every day. The underlying

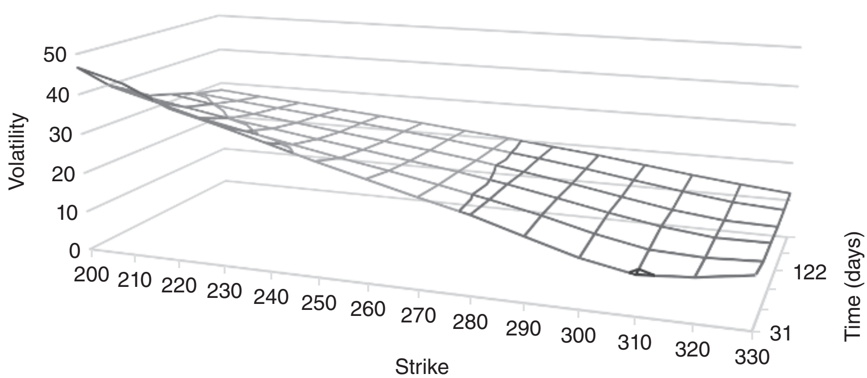

FIGURE 1.1 The implied volatility surface for SPY on September 10, 2019.

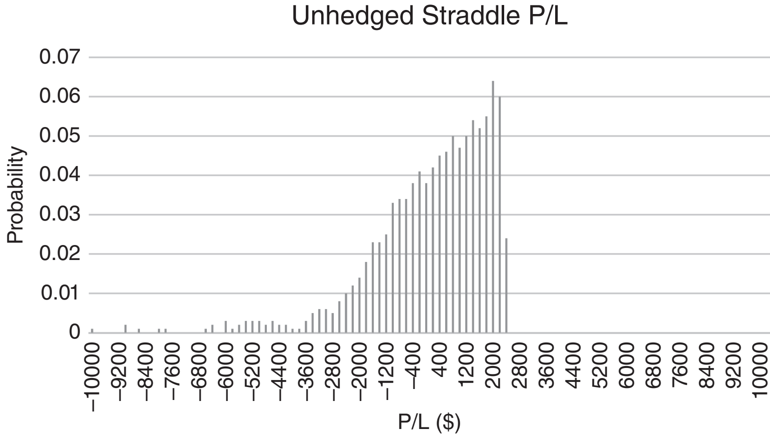

FIGURE 1.2 The terminal PL distribution of a single short one-year ATM straddle that is never re-hedged. Stock price is $100, rates are zero, and both realized and implied volatilities are 30%.

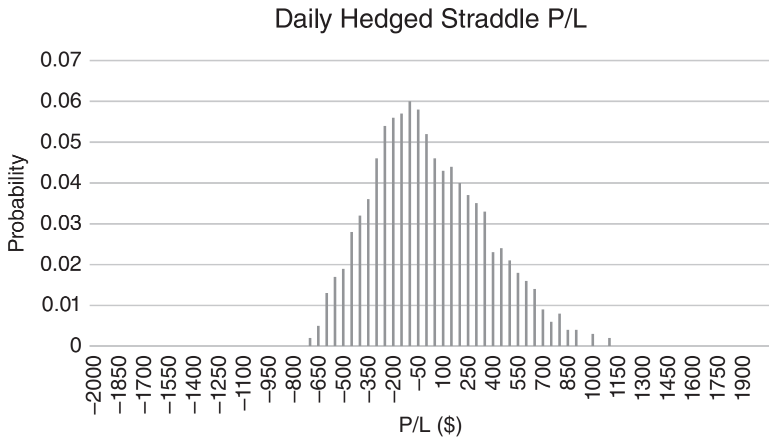

FIGURE 1.3 The terminal PL distribution of a single one-year ATM straddle that is hedged daily. Stock price is $100, rates are zero, and both realized and implied volatilities are 30%.

paths were generated from 10,000 realizations of a GBM. The implied and realized volatility were equal so we expect an average PL of zero.

The dependence of the standard deviation of the PL distribution on the number of hedges is shown in Figure 1.4.

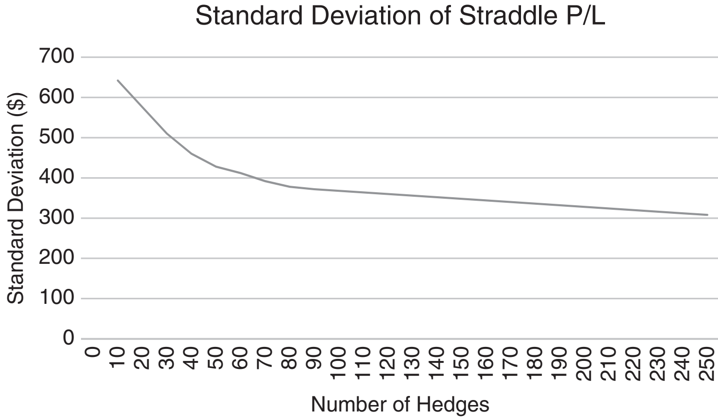

FIGURE 1.4 The standard deviation of the terminal PL distribution of a single one-year ATM straddle as a function of the number of hedges. Stock price is $100, rates are zero, and both realized and implied volatilities are 30%.

The reason to hedge less frequently and accept a wider standard deviation of results is that hedging costs money. All hedges incur transaction costs (brokerage, exchange fees, and infrastructure costs). Costs like this are an easily forgotten drain on a portfolio. Individually they are small, but they accumulate. To emphasize this point, Table 1.1 compares the summary statistics of results for the daily hedged short straddle when there is a transaction cost of $.10 per share and when hedges are costless.

The difference between these two cases is roughly equivalent to misestimating volatility by two points.

In practice, aggressive re-hedging is done by market-making firms and some volatility specialists. The vast majority of retail and buy-side users

TABLE 1.1 **Statistics for the Short One-Year ATM Daily Hedged Straddle With and Without Hedging Costs (stock price is $100, rates are zero, and both realized and implied volatilities are 30%.)**

Statistic	Costless Hedges	$.10/Share Hedges
Average	−$6.10	−$121.54
Median	−$49.85	−$111.68
Percent profitable	44%	30%

seldom or never hedge. The relevant theory for those hoping to approximate continuous hedging is discussed in Sinclair (2013). In this book we will generally assume that no re-hedging takes place. These results are also applicable to those who hedge infrequently. They can just assume that the original position has been closed and a new one opened. So, a one-year position that is hedged after a month would thereafter have the expected distribution of an 11-month option.

Conclusion

The BSM model gives the replication strategy for the option. The expected return of the underlying is irrelevant to this strategy. The only distributional property of the underlying that is used in the BSM model is the volatility. A hedged position will, on average, make a profit proportional to the difference between the volatility implied by the option market price (by inverting the BSM model) and the subsequent realized volatility. The choice of the option structure and hedging scheme can change the shape of the PL distribution, but not the average value. These choices are far from immaterial, but successful option trading depends foremost on finding situations in which the implied volatility is mispriced.

Summary

- Arbitrage-free option pricing models do not include the underlying return. BSM includes only volatility.
- Inverting the pricing model using the option's market price as an input gives the implied volatility.
- The average profit of a hedged option position is proportional to the difference between implied volatility and the subsequent realized volatility.
- Practical option hedging is designed to give an acceptable level of variance for a given amount of transaction costs.