Two-State Options

John Norstad

j.norstad@mac.com
http://homepage.mac.com/j.norstad

January 12, 1999
Updated: April 27, 2005

Abstract

How options are priced when the underlying asset has only two possible future states. Studying these trivial options helps develop insight into how real options and the Black-Scholes equation work. We learn about arbitrage, the law of one price, hedging, risk-aversion, and risk-neutral valuation.

Contents

List of Figures

1 Two-State Options

This note explores some very simple examples of "two-state options," where the underlying asset has only two possible future prices. Our analysis goes into considerable detail and illustrates several important aspects and properties of option pricing.

A "call option" is a contract which grants the purchaser of the contract the right, but not the obligation, to purchase a specific item called the "underlying asset" from the seller of the contract at a specific date in the future called the "expiration date" at a specific price called the "strike price."

If the future price of the asset on the open market on the expiration date is greater than the strike price, we say that the call option expires "in the money." In this case, the buyer "exercises" the option, buys the asset from the seller at the strike price, sells the asset on the open market at the future price, and makes a profit on the difference.

If the future price of the asset on the open market on the expiration date is less than the strike price, we say that the option expires "out of the money." In this case, the option is worthless, and the buyer loses his entire initial investment in the option.

Our definitions above are for "European options," which can only be exercised on the expiration date. "American options" can be exercised on or at any time before the expiration date. The math for European options is simpler, so we use them here.

As an example which we'll use throughout this note, suppose an asset has a current price of $s = \$100$. There are two possible future prices: $s1 = \$120$ with probability $p1 = 3/4$ and $s2 = \$80$ with probability $p2 = 1/4$. Consider a call option on this asset with a strike price of $x = \$100$.

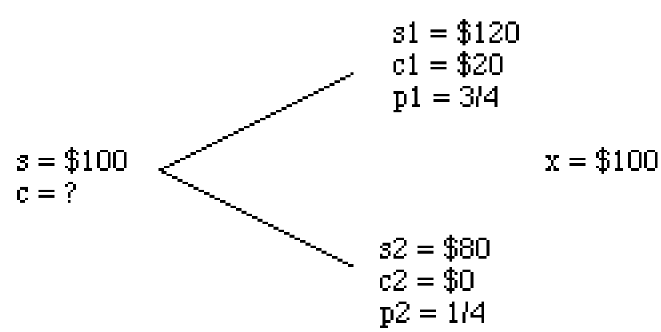

Figure 1: Example

If the future price is $120, the option expires in the money, and it is worth $c1 = s1 - x = \$20$. If the future price is $80, the option expires out of the money, and it is worthless ($c2 = 0$).

What is the proper price c of this call option?

The expected value of the option is $(3/4) * 20 + (1/4) * 0 = \$15$. This is not, however, the proper price of the option. The proper price of the option turns out to be $10. We give the standard proof of this from the literature.

The proof is by counter example, so we begin by assuming that you are willing to buy this call option from me for $15. I'm in fact eager to sell you an option at this price, because I have a sure-fire scheme to make a guaranteed profit of $5 from the deal. My scheme involves a trick called "hedging."

Here's how the scheme works:

In addition to selling you the option for $15, I also borrow $40 from the bank or somewhere else. I now have $55. I pocket $5 and use the remaining $50 to buy 1/2 share of the asset. This 1/2 share in the underlying asset is my hedge and is the key to making the scheme work. The fraction 1/2 is called the "hedge ratio," and we'll see how to compute it in section 2 where we work out the simple algebra.

In our very simple two-state universe there are only two possible outcomes. Either the asset increases in value to $120, or it decreases in value to $80.

First suppose that the ending price is $120. I sell my 1/2 share for $60. I use the proceeds of this sale to repay my $40 loan. (For simplicity, we assume we don't have to pay any interest on the loan. See section 2 for the details if we have to pay interest.) I'm left with $20. Because the ending price is greater than the strike price, you exercise your option to buy one share from me for $100. You pay me $100. I put that together with my $20, buy one share on the market for the current price of $120, and give you that share. This total transaction comes out even for me, and I still have the original $5 I pocketed at the beginning, so my net profit is $5. Note that in this case you make $20 on your option, which gives you a profit of $5 on your initial $15 investment in the option.

Now suppose that the ending price is $80. I sell my 1/2 share for $40 and use that to repay my $40 loan. Your option has expired out of the money and so you don't exercise it. Again, I still have the $5 I pocketed at the beginning, so my net profit is $5. In this case you lose your entire $15 investment in the option.

In either case, I still have the $5 I pocketed at the beginning. So I come out ahead $5 in either case. I have taken advantage of the mispricing of the call option to make a risk-free profit of $5.

Note that I do not have to come up with any of my own money to work this scheme. Given a large enough number of other investors who are willing to buy call options from me at a price of $15, and given a bank or someone else willing to loan me large amounts of money, I can make an arbitrarily large amount of money immediately with no risk at all, with a rate of return on my investment of infinity. This is the kind of opportunity investors dream about.

This kind of scheme to exploit mispriced assets to make a risk-free profit is called "arbitrage."

We can use the same arbitrage scheme to guarantee a risk-free profit for any option price above $10. The difference between the option price and $10 is the guaranteed profit.

If the option is priced below $10, we can employ a similar arbitrage scheme to guarantee a profit. In this case, we reverse everything and buy the option instead of selling it, we lend money instead of borrowing it (e.g., we buy a Treasury bill), and we sell a fraction of the asset short instead of buying it. The details can be found in any good textbook on investing.

Thus, if the option is priced at any value other than $10, we can use arbitrage to make a risk-free profit. So can anyone else who understands the trick and is able to find some sucker to sell the option to or buy the option from at the bad price. Rest assured that every professional in the financial world understands the trick and is constantly on the lookout for suckers to exploit (it's their job). In the professional financial world, if the option were priced at any value greater than $10, everyone would want to sell the option and nobody would want to buy it. Similarly, if the option were priced at any value less than $10, everyone would want to buy it and nobody would want to sell it.

Thus, the option's price must be exactly $10 in order for there to be any significant market in the option at all.

This completes our proof that the proper price for the option is $10.

At the beginning our proof we computed the expected value of the option to be $15. The proper price turned out to be $10. Why is the expected value of the option greater than its price? The explanation involves risk aversion. The option is risky, and the extra $5 is a risk premium. If the expected value were equal to the price, investing in the option would be a fair game, and risk-averse investors would not be willing to purchase it. The $5 risk premium compensates the investor for undertaking the risk of the investment.

As a brief digression, we must mention at this point that our example and our analysis is ideal and ignores several important real-world issues. For example, buying and selling fractional shares is often impossible or expensive. Also, buying and selling assets and options on assets is not free as we have assumed. Brokers, exchanges, advisors, dealers, and others all get fees to process these kinds of transactions. We have ignored these fees. Also, our scheme involves making and losing money in the financial markets. These capital gains and losses are taxable, and we've ignored the impact of these taxes. We have also made the important assumption that the market for the underling asset is "liquid," which means that we can easily buy and sell shares of the asset at any time at a well-defined "current market price." In the real world different kinds of assets have widely different liquidity. These kinds of pesky impediments to investing are often called "friction." Our ideal world is frictionless. The real world has friction. To adjust our example for friction, we should say that the option's price must be exactly $10 or "close to" $10 in order for there to be any significant market in the option at all. Making this notion of "close to" more precise by dealing formally with all the pesky frictions is a very difficult problem which we are ignoring completely.

Now let's return to our arbitrage argument and look at it from another point of view. Consider once again the strategy we used of first borrowing $40 and buying 1/2 share of the asset, then later selling the 1/2 share and repaying the loan. This combination of borrowing and buying a fraction of a share is called a "synthetic call option."

In our arbitrage argument, we constructed our synthetic call option so that it would have exactly the same payoffs in both possible outcomes as does the real call option ($20 in one case, $0 in the other).

The "law of one price" states that in any situation in which two different investments have exactly the same payoffs in all possible future states of the world, the investments must have the same price. The synthetic call option has a price of $(1/2) * \$100 - \$40 = \$10$. Thus the corresponding real option must also have a price of $10.

Thus using the law of one price gives us another shorter proof that the proper price for our option is $10.

Let's examine our example further to explore in a bit more detail these notions of "hedging," "arbitrage" and "the law of one price."

In our scheme, we borrow money and buy a fraction of a share of the underlying asset to offset the loss we might experience as a seller of the option. As a seller of the option, we lose if the price rises, and we gain if the price falls. As a buyer of the underlying asset, we gain if the price rises, and we lose if the price falls. The two investments perfectly offset or "hedge" each other.

Borrowing money to help buy the fraction of a share is called "leverage." In our example, we buy 1/2 share for $50 and we borrow $40. Thus 4/5 of the purchase price is borrowed. We say that our purchase is "80% leveraged."

What we've shown with our synthetic call is that in this particular example an 80% leveraged purchase of 1/2 share of the underlying asset is equivalent to one call option on the asset with a strike price of $100. When we say "equivalent" we mean that the two investments have exactly the same payoffs in all possible future states of the world.

In our example, the price of the synthetic call is $10, and the corresponding real option is priced at $15. Our arbitrage scheme boils down to simply selling the overpriced investment (the real option) and buying the underpriced investment (the synthetic call). Because the two investments hedge each other perfectly, we run no risk, and we pocket the difference in the prices, which is our $5 profit.

This is the essence of arbitrage. Whenever we find a situation where the law of one price is violated and two equivalent investments are selling at different prices, we exploit the situation by selling the higher priced investment and buying the lower priced one. The difference in the prices is our profit. Because the investments are equivalent, and because we bought one of them and sold the other one, they hedge each other perfectly, and we make our profit with no risk.

Significant arbitrage opportunities are difficult to find in the real financial world. The natural forces of the marketplace tend to eliminate them quickly when they do appear. For example, suppose investments A and B are equivalent and A is currently available on the market at a lower price than B. Arbitrageurs very quickly notice the opportunity and rush to buy A and sell B. This drives up the price of A and drives down the price of B until equilibrium is restored and the arbitrage opportunity disappears. In today's very efficient financial markets, this process happens very quickly, often in only minutes.

There is no regulatory body that sets option prices to their "proper price." This is done by the natural forces of buyers and sellers in the marketplace. The law of one price governs the marketplace. Arbitrageurs enforce the law. Punishment for breaking the law is losing tons of money.

We now change the topic and return to our example to explore the role of the probabilities of the possible outcomes. This is an important issue which we haven't covered yet.

Our arbitrage argument using our synthetic call has the interesting and important property that it is independent of the probabilities of the possible outcomes. The synthetic call constructed via a leveraged purchase of the underlying asset perfectly replicates the call option no matter what the probabilities of the outcomes are.

Thus the price of the option in this universe with only two possible future states is independent of the probabilities of those two states. Stated another way, in this simple universe, the price of the option is independent of the expected rate of return of the underlying asset. This seems strange at first glance, but it's true.

The price does, however, depend on the volatility of the underlying asset. For example, if our two possible ending prices are $140 and $80 instead of $120 and $80, the price of the call option is $13.33 instead of $10. In this case, the replicating synthetic call is constructed by borrowing $53.33 and buying 2/3 shares of the asset for $66.67. See section 2 below for details on how to calculate these values.

We see the same phenomenon in the Black-Scholes equation for much more complicated real-life options for which the underlying assets have an infinite number of possible future prices instead of just two. Once again, the option price given by the Black-Scholes equation is independent of the expected return on the underlying asset, but depends significantly on the volatility of the underlying asset. As in the simple two-state universe, this is counter-intuitive but nonetheless true. We'll return to discuss this problem a bit more later.

For yet another point of view on the probabilities of the outcomes, let's return to our original example, with the option mispriced at \$15. We can think of my selling the option to you at this price as a game we play. I use the arbitrage trick, so the game boils down to the following essentials:

With probability 3/4, I win \$5 and you win \$5.
With probability 1/4, I win \$5 and you lose \$15.

Which side of this game would sane people rather play? The answer is obvious.

Now let's look at this same game with the proper option price of \$10.

If we repeat the arbitrage scheme with the proper price of \$10, my sure-fire profit disappears. If I do the synthetic call hedging trick, I win \$0 in both cases. You win \$10 with probability 3/4 and lose \$10 with probability 1/4. This game is certainly more attractive to you as the option buyer, but it's completely unattractive to me as the option seller. There's no point in my playing this game with hedging, so I don't.

At the proper price of \$10 my arbitrage scheme no longer generates any profit, so I don't use it. Instead, however, I might simply sell you the option for \$10 and take my chances without hedging my bet by purchasing a fraction of a share of the asset. Let's examine this possibility. This is called an "unprotected" or "naked" call sale. For me, the outcome in this case is a loss of \$10 with probability 3/4 and a win of \$10 with probability 1/4.

To summarize, at the proper price of \$10 without the arbitrage scheme, the game is the following:

With probability 3/4, I lose \$10 and you win \$10.
With probability 1/4, I win \$10 and you lose \$10.

In this game, the advantage is clearly with you, the buyer, not with me, the seller. I would never play this game.

I might, however, play this game if I have a different opinion about the probabilities of the outcomes. For example, if I believe that the \$120 price has probability 1/3 and the \$80 price probability 2/3, I might be willing to sell you an option at the price of \$10.

To summarize, the probabilities of the outcomes do not enter into the determination of the option price. Investor's subjective beliefs about the probabilities do, however, determine who buys the options and who sells them.[1]

What if the probabilities are so overwhelmingly in favor of the higher future price possible outcome that everyone in their right mind considers the "proper price" for the option to be very attractive? On the face of it, one might think that this would tend to drive up the option price to balance buyers and sellers, and that our argument must be wrong, because the "proper price" as set in the marketplace for the option really does depend on the probabilities. This is why option pricing in general and the Black-Scholes equation in particular are so counter-intuitive.

Once again, we use our example to explore this further. Suppose, for example, that nearly all knowledgeable analysts agree that the probability of the $120 ending price is 99/100, and the probability of the $80 ending price is 1/100. Our option as priced at $10 is extremely attractive in this case, because we have a 99% probability of ending up in the money and making a profit of $10. We have only a 1% chance of losing our $10 investment in the option. Nearly everyone should want to buy this option because it's such a good deal. Shouldn't this drive the price up?

This is indeed a problem. To get to the bottom of it, let's ignore the option for a moment and instead focus on the underlying asset itself, which is currently priced at $100. In our scenario, the consensus opinion is that there's a 99% chance that the asset's price will increase to $120, and only a 1% chance that the asset's price will fall to $80. This makes the asset itself a very attractive investment, for exactly the same reason that the option on the asset is attractive. Many more investors will want to buy the asset than will want to sell it at its current price of $100. This will drive the asset's price up. Note that as the asset price rises, so does the price of the call option on the asset (see the equations in section 2). Eventually equilibrium is reached, where both the asset price and the option price reflect the consensus opinion of investors concerning the future prospects of the asset.

This resolves the conundrum. Yes, consensus probability beliefs do indeed affect option prices, but only indirectly through the price of the underlying asset. Stated another way, option pricing formulas for markets in equilibrium are independent of the expected return on the underlying asset because this information is already contained in the price of the underlying asset.

[1] There are other important reasons for buying and selling options. For example, I might sell you the option in our example because there is some other idiosyncratic characteristic of my financial life that exposes me to the risks of the underlying asset. My selling the option to you in this case is a way for me to hedge that other risk I face, but that you do not face. By exchanging risks in this way, we both benefit from the transaction. This kind of risk exchange is in fact the primary use of options in the financial world, primarily by large financial institutions like banks.

2 The Algebra of Two-State Call Options

Note: In our example we made the simplifying assumption that the interest rate for borrowing money is 0%. We did this to make the arithmetic easier for the sake of exposition. In this section we include interest as a factor. Also, in our example, the strike price happened to be the same as the current asset price. This is also a non-critical assumption which we do not make in this section.

Let:

$$
\begin{aligned}
s &= \text{current asset price} \\
c &= \text{call option price} \\
s_1 &= \text{larger of the two possible future asset prices} \\
s_2 &= \text{smaller of the two possible future asset prices} \\
c_1 &= \text{call option value if future asset price is } s_1 \\
c_2 &= \text{call option value if future asset price is } s_2 \\
r &= \text{current yearly continuously compounded risk-free interest rate} \\
t &= \text{time in years between the option sale and the option expiration} \\
a &= \text{hedge ratio for synthetic call} \\
b &= \text{amount borrowed for synthetic call}
\end{aligned}
$$

In the synthetic call, we buy a shares of the asset and borrow b dollars. The price of this combination is $as - b$. If the future asset price is s_1, the payoff is $as_1 - be^{rt}$. If the future asset price is s_2, the payoff is $as_2 - be^{rt}$.

In the real call option, if the future asset price is s_1, the payoff is c_1. If the future asset price is s_2, the payoff is c_2.

We want the synthetic call to replicate the real call. To accomplish this, we set the payoffs in each state equal:

$$
\begin{aligned}
as_1 - be^{rt} &= c_1 \\
as_2 - be^{rt} &= c_2
\end{aligned}
$$

Solving these equations for a and b gives:

$$
\begin{aligned}
a &= \frac{c_1 - c_2}{s_1 - s_2} \\
b &= (as_2 - c_2)e^{-rt}
\end{aligned}
$$

By the law of one price, the price of the real call option c must be the same as the price of the synthetic call:

$$c = as - b$$

In our example, we have $s = 100$, $s_1 = 120$, $s_2 = 80$, $c_1 = 20$, $c_2 = 0$, $r = 0$. Plugging in these values gives:

$$\begin{aligned} a &= \frac{20-0}{120-80} = \frac{1}{2} \quad \text{(hedge ratio)} \\ b &= \left(\frac{1}{2} \times 80 - 0\right) e^{-0 \times t} = 40 \quad \text{(amount borrowed)} \\ c &= \frac{1}{2} \times 100 - 40 = 10 \quad \text{(option price)} \end{aligned}$$

3 Risk-Neutral Valuation

In section 1 we briefly discussed how investors are risk-averse, and this is why the value of an asset is less than its expected value.

We can, however, imagine an alternate universe in which investors are "risk-neutral" instead of risk-averse. In this universe, investors don't care about risk, they don't demand risk premia, and the value of an asset is simply its expected value.

More properly speaking, because of the time value of money, in this universe an asset's value would be its future expected value discounted to present value using the risk-free interest rate. To make things simpler, in this section we're going to pretend once again that our universe doesn't charge interest. As we did in section 2, the reader is invited to rework the equations in this section to accommodate interest. If you do this, you will find that the main result of this section still holds.

How would our trivial two-state options be priced in this universe? The arbitrage argument we used in our universe still works in the alternate risk-neutral universe, so the equations in section 2 apply. Let's see what happens.

In the risk-neutral universe, the value of the underlying asset is its expected value:

$$s = p_1 s_1 + p_2 s_2 \quad \text{where } p_2 = 1 - p_1$$

Apply the equations we derived in section 2 for the value of the call option:

$$\begin{aligned}
c &= as - b \\
&= \frac{c_1 - c_2}{s_1 - s_2} s - \left(\frac{c_1 - c_2}{s_1 - s_2} s_2 - c_2 \right) \\
&= \frac{c_1 - c_2}{s_1 - s_2} [p_1 s_1 + (1 - p_1) s_2] - \frac{c_1 - c_2}{s_1 - s_2} s_2 + c_2 \\
&= \frac{c_1 - c_2}{s_1 - s_2} p_1 s_1 + \frac{c_1 - c_2}{s_1 - s_2} s_2 - \frac{c_1 - c_2}{s_1 - s_2} p_1 s_2 - \frac{c_1 - c_2}{s_1 - s_2} s_2 + c_2 \\
&= \frac{c_1 - c_2}{s_1 - s_2} p1 (s_1 - s_2) + c_2 \\
&= (c_1 - c_2) p_1 + c_2 \\
&= p_1 c_1 + (1 - p_1) c_2 \\
&= p_1 c_1 + p_2 c_2
\end{aligned}$$

Thus, in the risk-neutral universe, the value of the call option is also its expected value.

We can work out these same equations backwards. That is, if we start by assuming that the value of the call option is its expected value, we can derive our equations in section 2.

This is an interesting result. In a risk-neutral universe, assets are priced quite differently than they are in our own risk-averse universe. In particular, in the situation we are studying, both the underlying two-state asset and the two-state call option on the asset are priced quite differently in the two universes. What we've seen, however, is that the equation in section 2 which tells us how to compute the value of the option as a function of the value of the asset is exactly the same in the two universes!

This is called "the principle of risk-neutral valuation." It turns out that this principle is also valid for real-life options where the underlying asset can have an infinite number of possible ending values. We examine this in detail in reference [1].

References

[1] John Norstad. Black-scholes the easy way. http://homepage.mac.com/j.norstad/finance, Feb 1999.