

## Chapter 6

# Puzzles and Bugs

There are a lot of common mistakes, or bugs, made when developing or testing advanced derivative valuation models, and many of those same bugs crop up even when the model is plain old Black-Scholes. Many highly experienced derivatives traders and researchers will miss at least one of the puzzles and bugs in this chapter. If you understand them here, then you will immunize yourself to them in more complicated situations, and be a step ahead of the competition.

### 6.1 Why Did We Lose All That Money?

What do we do when things go wrong? Many traders facing mounting losses are often relegated to the pit of despair known as “P&L decomp.”

P&L decomp is short for profit and loss decomposition, the process of explaining why certain positions made or lost money.

It asks questions like: “How much should we have made from hedging? How much did we actually make? And how much should we have made from changes in parameters? How much did we actually make?”

“Lucy, you got some ‘splainin’ to do!”

*Ricky Ricardo, “I Love Lucy” (1951)*

At the end, after asking all possible reasonable questions, whatever is left over, if anything, is called the *residual* P&L. This number ought to be near zero, otherwise you haven’t decomposed enough.

Another way of thinking about it is that you start with your residual P&L equal to your true P&L, the amount of money you actually made or

lost, and then you split it up into buckets for better understanding, until hopefully your residual bucket is quite small.

But we can't even begin to properly do P&L decomp until we understand the "Greeks," the various mathematical derivatives of financial derivatives.

"I fear the Greeks, even when they bring gifts."

*Virgil*

Suppose you own a call option with a delta of 0.4, but you made a mistake hedging. Instead of selling 0.4 shares, you actually bought 0.4 shares. So your net delta, which should be zero if you are hedging correctly, is instead long 0.8 shares of the underlying. This is a big blunder.

How much money will you lose tomorrow?

- (a) You will lose 0.8 times the percentage change in the underlying.
- (b) You will lose 0.8 times the dollar change in the underlying.
- (c) You will lose 0.8 in the notional units of the underlying, e.g. \$0.80 if the underlying is denominated in dollars.

The answer is none of the above. If the underlying price goes up, you are net long delta, so you will actually make a profit, not a loss.

"Okay," you sigh, "then let's suppose the underlying falls tomorrow. Then what is the correct answer?"

The correct answer is still none of the above. However, of the choices given, (b) is the closest because delta by definition is the change in the derivative for a unit change in the underlying price:

$$\Delta = \partial V / \partial S \tag{6.1}$$

But (b) is not the whole story. This is not your total P&L; it is only the delta P&L, or the P&L that is due to the initial net delta having been different from zero:

$$\text{Delta P\&L} = \Delta(S_1 - S_0) \quad (6.2)$$

Your total profit or loss will *not* exactly equal your net delta times the change in the underlying price.

Why not?

For one reason: because your net delta will change right along with changes in the underlying price.

Remember gamma? Gamma is the change in delta for a unit change in the underlying price:

$$\Gamma = \partial\Delta/\partial S = \partial^2 V/\partial S^2 \quad (6.3)$$

At the end of the day tomorrow, if the price has dropped from  $S_0$  to  $S_1$ , then your delta will have dropped from  $\Delta$  down to  $\Delta + \Gamma(S_1 - S_0)$ .

So, what would your profit or loss tomorrow be?

- (a) Your initial net delta times the change in the underlying.
- (b) Your final net delta times the change in the underlying.
- (c) The average of (a) and (b).

Again, the correct answer is none of the above, because your gamma would change too. In principle, we could keep going to higher and higher orders of partial derivatives, and for big moves, those higher order terms may be important. But in most cases, they are not.

So ignoring the higher order terms, the closest correct answer would be (c). In other words, the change in the underlying price times your initial net delta increased by half the change in the delta due to gamma. But bear in mind that the change in delta depends on the change in the underlying itself. That's why the gamma P&L, or the portion of your total P&L that is due to the delta changing because of gamma, depends on the square of the change in the underlying price:

$$\text{Gamma P\&L} = \frac{1}{2}\Gamma(S_1 - S_0)^2 \quad (6.4)$$

Another way to see that is to write the sum of delta P&L and gamma P&L together:

$$\text{Delta P\&L} + \text{Gamma P\&L} = \left( \Delta + \frac{1}{2} \Gamma (S_1 - S_0) \right) (S_1 - S_0)$$

But it would still be the wrong answer to the question of *total* P&L, even if there are no higher order terms than gamma.

Why?

Because time passed.

The valuation of every option is subject to time. As time passes, options usually lose money because they have a shorter expiry. This sensitivity to time is called theta:

$$\Theta = \partial V / \partial t \tag{6.5}$$

And the portion of your total P&L that is explained by theta is your theta P&L:

$$\text{Theta P\&L} = \Theta (t_1 - t_0) \tag{6.6}$$

Assuming that nothing else changes, i.e., the implied volatility, interest rate, and dividends remain the same, then we can decompose your total P&L into the parts due to delta, gamma, and theta:

$$\text{Total P\&L} = \text{Delta P\&L} + \text{Gamma P\&L} + \text{Theta P\&L} + \text{Residual P\&L}$$

Notice the all-important residual P&L term at the end. That is where all the higher-order terms go to die. It includes the P&L due to such sensitivities as  $\frac{\partial^2 V}{\partial S \partial t}$ ,  $\frac{\partial^3 V}{\partial S^3}$ , and so on.

The residual P&L term is the most critical one. That is the term you hope is close to zero, not too positive and not too negative. A lot of unexplained profits is almost as bad as a lot of unexplained losses because it suggests you are running an uncertain risk that you are not aware of.

“I have known uncertainty: a state unknown to the Greeks.”

*Jorge Luis Borges*

This is a crucial point. The one overarching thing you want to do in finance is make money. But only slightly less important is to make it for the reasons you expected. Otherwise you may as well be gambling.

We have thus decomposed our total P&L into a delta P&L, a gamma P&L, and a residual P&L. We know that part of the residual P&L will be the higher order terms that we have ignored, but we think they will be usually quite small.

But, to answer all our earlier questions, what *is* our total P&L?

The total P&L is an accounting number. You just calculate how much money you would have today if you were able to liquidate all of your positions in an orderly manner, and subtract the similarly calculated number from yesterday. That's your total mark-to-market profit.

In other words, your total P&L is your market P&L, and the definitions and derivations above have served to do nothing more and nothing less than define the residual P&L as quite literally the remaining, unexplained P&L.

“Your life is the sum of a remainder of an unbalanced equation inherent to the programming of the matrix. You are the eventuality of an anomaly, which despite my sincerest efforts I have been unable to eliminate from what is otherwise a harmony of mathematical precision. While it remains a burden assiduously avoided, it is not unexpected, and thus not beyond a measure of control.”

*The Architect*, “The Matrix Reloaded” (2003)

## 6.2 Should We Mark to Market or Model? To Bid or Bid?

In the introductory P&L decomposition above, we assumed that the total P&L was the market P&L, computed as the realized difference in the orderly liquidation value of your entire portfolio. But what is the orderly liquidation value of a derivative that is never sold?

It is an unobserved, and unobservable, number. It can be tested somewhat by liquidating a portion of your position. You could decide to sell and then buy back one percent of all of your positions every day.

Suppose it costs you  $\$X$  to sell one percent of your positions and  $\$Y$  to buy them back. You could then argue that the true value of your entire

position is  $\$100 \cdot (X + Y)/2$ . This is typically called marking-to-mid, because you are marking the value of your portfolio to the mid-point of the market.

If you instead mark your position each day to  $\$X$ , then you are marking-to-bid.

Is marking-to-bid a more conservative measure of profit and loss than marking-to-mid?

- (a) Yes, because you are using a more realistic measure of your liquidation value.
- (b) No, because it is much easier to manipulate the bid price than the mid price.

The answer, as usual, is none of the above. Marking-to-mid is certainly more conservative on the first day, because your portfolio will be marked to a lower number.

But on each subsequent day, the difference may be higher, lower, or zero, depending on how the bid-offer spread changes with respect to changes in the underlying price.

As an example, consider a liquidity story. In this case, when the underlying price drops, liquidity dries up. When liquidity dries up, the bid-offer spread widens, because there are fewer people willing to provide tight markets.

In this case, marking-to-bid results in an even larger loss than marking-to-mid because the bid falls even more than the mid.

Conversely, in up markets, the bid-offer spread presumably tightens, so marking-to-bid results in an even larger profit than marking-to-mid because the bid rises even more than the mid.

Thus, it is not clear whether marking-to-bid is necessarily a better strategy than marking-to-mid.

“When Greeks joined Greeks, then was the tug of war.”

*Nathaniel Lee*

Further, from the perspective of fairness, if occasionally new investors come in and old investors drop out, then marking-to-bid is a net

benefit to new investors at the expense of old investors, because the new ones hold positions with relatively more expected return than if the positions had been marked-to-mid.

However, when they take their money out, they will also take it out at bid, so later investors may decide that it is approximately equivalent.

But the very first investors will get out at bid even though they entered, effectively, at mid, and were forced to bear the full cost of half of the bid-offer spread. There are lots of issues about fairness when it comes to marking to mid or to bid.

In doing P&L decomp, one approach might be to mark everything to liquidation value, and segregate out in one of the decomposition buckets the difference between the liquidation value and the mid-market value, and in another bucket the difference between the mid-market value and the model value.

Changes in the first would reflect changes in the underlying liquidity or the overall position held, while changes in the second would reflect changes in model cheapness or richness, and the remaining terms, being all model terms, would decompose the model changes into its own constituent parts.

### 6.3 Why Does the Hedged Profit Differ from Zero?

We saw in the previous chapter how a five-day Black-Scholes simulation of an at-the-money five-day call on a 20-vol non-dividend-paying stock in a world with zero interest rates could generate returns ranging from  $-100\%$  to  $+100\%$  when Black-Scholes asserts it should be zero always, along every path. How can that be?

Your first instinct might be that five days is just too coarse a time period. Black-Scholes assumes continuous movements and hedging, and daily is just not good enough, particularly when there are only five days until expiry, so that each interval carries enormous weight.

How can we test that instinct? Let's change the time intervals from days to hours. Instead of assuming we have five business days left with intervals of  $dt = 1/252$ , let's look at an option with just five hours left with intervals of  $dt = 1/252/10 = 1/2520$ . This implicitly assumes

there are ten trading hours per day, which is acceptable as an approximation for now.

In that case, we can update our spreadsheet from the last chapter with the following formulas. Define some special cell to be  $dt$ . To do this, type  $= 1/2520$  into your chosen cell, for example A1, and then click on the Name Box — the part of Excel to the left of the formula bar where it says “A1” — and change it to say “dt” instead. Then you can refer to it in formulas as “dt” directly, making them easier to read.

Here’s what the formulas for the second row look like, with the remaining ones being essentially the same:

Cell	Title	Formula
A5	Expiry	=A4-dt
B5	Stock	=B4*LOGNORM.INV(RAND(), 0-1/2*(0.2*SQRT(dt))^2, 0.2*SQRT(dt))
C5	Delta	=BSCallDelta(100,A5,B5,.,2,0,0)
D5	Hedge P&L	=-C4*(B5-B4)

The total P&L is as before the sum of the total hedging P&L across all time periods and the change in the option price from the initial premium on hour one to its intrinsic value on hour five. The initial premium is about \$0.32. What is the total P&L?

Sometimes it is near zero. But sometimes it is as low as -\$0.30 and as high as +\$0.30. In other words, it can still range from -100% return to +100% return. So your instinct does not seem to be validated. So let’s go back to looking at daily instead of hourly changes.

But do not fret. Proving an instinct wrong is almost as valuable as proving it right. In some cases, it can be even more valuable, especially if it helps you to avoid mistakes in the future that you otherwise would have made.

“Results! Why, man, I have gotten a lot of results. I know several thousand things that won’t work.”

*Thomas Edison*



What is your next instinct? If the answer is not the size of the time difference in each interval, could it be the number of intervals? Maybe five of anything — be it hours or days — is just not enough to reduce the noise to the levels predicted by Black-Scholes.

A good financial hacker not only generates ideas but tests them too. Let's rerun our spreadsheet with 253 rows instead of 5. That would represent a year's worth of hedging on a daily basis, plus the initial day's values.

If Black-Scholes continues to fail in this scenario, and particularly if it continues to fail because the intervals are neither continuous enough nor plentiful enough, then Black-Scholes is not a useful model and we would have to explore other alternatives, because daily hedging is quite reasonable in the real world, particularly for a year-long option.

As a reminder, this exercise has ramifications far beyond Black-Scholes. Any valuation model on any more complicated derivative should still provide useful values even when hedging is not continuous.

To evaluate this possibility, we need to merely add more rows to our spreadsheet, change the initial maturity to be one year, and (this is a common Excel oversight) make sure we change our summation and other summary cells at the top to include the new, longer range of cells. Specifically, we need to make sure our intrinsic value refers to the last row, and the total hedging P&L sums across all of the rows, not just the first five.

Having done that, we can then try, as before, to recalculate numerous times to get a sense of how the total P&L changes. However, if we want to get a list of several hundred such total P&L across different paths, we need to copy the last three columns several hundred times and store the results somewhere.

This is where Excel can get cumbersome, and where it is useful to be able to rely on other tools. It is certainly possible to write Visual Basic code to run the simulation numerous times, and each time print the resulting total P&L to some special worksheet.

But using a higher level language such as Mathematica lets us automate and generalize a lot of the work.

For example, the very general code for simulating lognormal paths is:

```
Simulate[payoff_, n_: 10 000,  $\mu$ _: 0,  $\sigma$ _: .2, dt_: 1/252, S0_: 100, maturity_: 1] :=
  payoff[FoldList[Times, S0, #]] &/@
  RandomReal[LogNormalDistribution[( $\mu$  -  $\sigma^2$  / 2) dt,  $\sigma$  Sqrt@dt],
    {n, Ceiling[maturity / dt]}]
```

The first line of the code above lists the input parameters and their default values. For those defaults, the third line generates 10,000 paths of 252 random daily returns. The second line computes the implied price paths from each of those 10,000 sequences of returns, and also applies an arbitrary function as the payoff to that sequence of prices. It uses the special “/@" shorthand notation for the Map function, to indicate a function being applied to each element of a list.

For example, if the payoff function were to simply return the last price in each path, then plotting the resulting histogram should show the terminal stock price distribution.

And indeed it does:

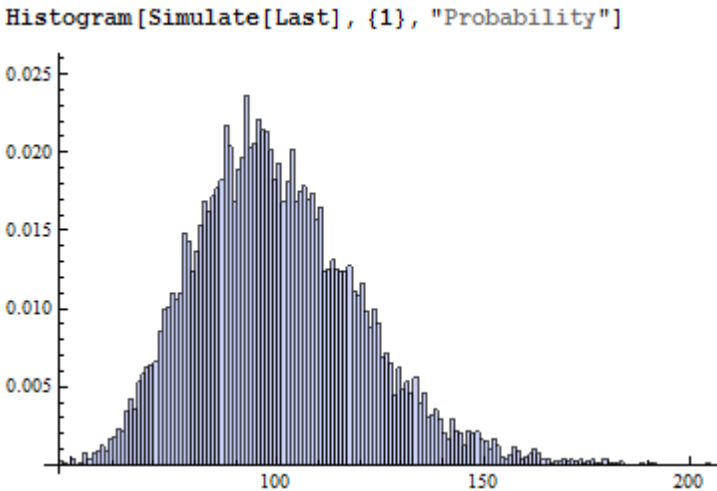


Figure 6.1. Histogram of simulated last price

We shall return to the general Simulate function later when we discuss exotic options. For now, let us explore what happens if we delta hedge each path.

All we need to do is define a payoff function that returns the total P&L from delta-hedging the option along that path according to the Black-Scholes delta at each point. We also want to express the total P&L as a portion of the initial option premium.

Here's the payoff function:

```
BSReturn[prices_,  $\sigma$  : .2, dt_ : 1/252, S0_ : 100, maturity_ : 1,
  strike_ : 100] :=
(Max[Last[prices] - strike, 0] -
  MapIndexed[BSCallDelta[strike, maturity - (First[#2] - 1) dt,
    #,  $\sigma$ , 0, 0] &, Most@prices].Differences[prices]) /
  BSCall[strike, maturity, S0,  $\sigma$ , 0, 0] - 1
```

For the default parameters, and given a path of prices, the function calculates the intrinsic value and the initial option premium in the second and fourth lines above, and the total hedging P&L in the third line (and the first half of the fourth). The total hedging P&L is calculated as the dot product, meaning the total of the pairwise products, of the Black-Scholes call delta at each point in time and the subsequent change in underlying price.

It is not immediately obvious that we are multiplying by the subsequent difference in price rather than the preceding difference in price. This is a drawback of languages such as Mathematica and R. It is much easier to tell in Excel if you are using information from the future. That is why it is a good idea to first build some small test cases in Excel before automating it in a more formal language.

Here, we can eventually tell that we are using the subsequent price because the differences start with the second price minus the first price, and the deltas are calculated for each price starting with the first price. How do we know that it starts with the first price? Because we are looking at “Most@prices” and the Most function drops the last element, just as in Excel we never calculate the P&L from the last delta.

Another function we used in the definition of BSReturn was MapIndexed. This is similar to the Map function whose shorthand “/@” we used in the definition of Simulate, but which also passes the number of the element so that we can compute the remaining expiry.

The remaining expiry is calculated as the maturity minus the product of  $dt$  and what is effectively the current row number minus one. The current row number is accessed through `MapIndexed` as `First@#2`, meaning it is the first element of the second argument to the function that we are applying to the list. The slot “#” indicates the first argument, which is filled by the particular price at that point in the list.

You do not need to know or even understand Mathematica code. All of these examples could have been coded in R or Visual Basic or C or any other language. Indeed, it is a good exercise to try to code it in other languages. The presentation here is done with Mathematica for four reasons: because of its flexibility and generality, because of its numerous built-in functions that help leverage a financial hacker’s time, because of its recent innovations in allowing interactive computations, and because its code is both relatively legible to someone unfamiliar with the language, and concise.

If we simulate 10,000 paths, what will the distribution of the total P&L be as a percentage of the initial option premium?

We hope that it is a percentage near zero, and indeed it seems to be an acceptable range:

```
Histogram[Simulate[BSReturn, 10 000], {.01}, "Probability"]
```




Figure 6.2. Histogram of returns from dynamically hedged option portfolio

The return from buying a one-year at-the-money call and hedging it daily using the Black-Scholes formula for the delta is zero percent return

on average, with a range of about plus or minus 10 percent. The standard deviation of the return is 5 percent.

Why is it not exactly zero?

### 6.3.1 How does the error depend on the hedging frequency?

We can combine your two instincts and test what happens if we have  $252 \times 10$  intervals of length  $dt = 1/2520$ ; in other words, what happens as we hedge closer to continuously.

That histogram for 10,000 paths is shown in Figure 6.3.

```
Histogram[Simulate[BSReturn[#, .2, 1/2520] &,
  10000, 0, .2, 1/2520], {.0025}, "Probability"]
```




Figure 6.3. Histogram of returns from more frequently hedged option portfolio

The standard deviation of the hourly-hedging return is 1.8 percent, less than half of what it was for the daily-hedging return. Clearly the more frequently we hedge, the closer we get to the Black-Scholes prediction of zero P&L along every path. The noise around the zero percent comes from non-continuous hedging.

Though we are now hedging ten times as frequently, the standard deviation of our return due to the fact that it is still not continuous did not decline ten times but only by about two-thirds.

What is the relationship there? Is there a simple rule of thumb we can follow?

Figure 6.4 shows the standard deviation of returns for a few hedging frequencies ranging from once (daily) to one hundred (about every few minutes).

It looks like it declines proportionally to some inverse power of the hedging frequency. The best-fit curve shown that matches the data is the inverse square root.




Figure 6.4. Standard deviation of returns to hedged portfolio vs. hedging frequency

Of course! That makes sense, doesn't it? We know volatility grows by the square root of time. So if we hedge ten times more frequently, each interval is one-tenth the size, and so the overall volatility should be the square root of one-tenth, which is about one-third.

We could have gotten that insight directly. If you did, that is great. But insight can be slippery, and failed insight can leave you right where you started. Hacking on the other hand is reliable, and even a failed attempt at financial hacking will have given you some progress to work from. For example, if you had guessed it declined like the reciprocal, or an exponential function, you would have quickly seen the fit is not very good, and you could have easily tweaked your code to test other variants.

In short, the inability to hedge perfectly continuously impacts your trading by introducing a random risk. That risk decreases if you hedge

more frequently, but only as fast as a square root, so if you want to halve your risk, you have to hedge four times as often.

### 6.3.2 *Is the error too large?*

Is it too much? Is the five percent error from hedging daily too much of a risk? Do we need to throw away Black-Scholes because it is insufficiently useful as an approximation when hedging occurs daily rather than continuously? Is continuous time just a fun mathematical trick that has no bearing on reality?

How can we address this question? Five percent can sound like a lot or a little, depending on the context. What is the appropriate context?

In the land of options, the appropriate context is always volatility. What is the vega of a one-year at-the-money call option, or the sensitivity of the option price to a one point change in volatility?

We can calculate it as 0.397, or we can recall the approximation from Section 5.2.1 about the value of an at-the-money option, and take its derivative with respect to the volatility  $\sigma$  to note that for  $T = 1$ , the vega should be about 0.4.

The vega is usually expressed in dollar terms. The price of the option, either calculated or using the approximation, is about \$8. So a one-point increase in volatility from 20% to 21% would raise the option price by  $0.4/8 = 5\%$ .

In other words, the noise from hedging a one-year option on a daily basis instead of continuously is about the same as one volatility point. So as a first test, we can ask if one vol point of noise renders Black-Scholes unusable. On first blush, it does not seem like too much noise; after all, when we estimate future volatility to be 20 percent, and it ultimately delivers 19 percent or 21 percent, we would not likely conclude that we were wrong. There is noise in realized volatility too.

As a first step, that's helpful, but it's not enough. We can also ask: is one volatility point too much extra risk?

Here's one way to think about it. Suppose this was the only risk in your options position. Would you still put on the trade?

To answer this question, you need to know how much money you expect to make on the option. Usually if you are making volatility bets

and intend to hold the option to maturity, you'd like to make at least one vol point, maybe two or more. An exception might be if you are specifically in the high-volume, high-turnover business of making markets in options, and would try to get out earlier when possible.

If you are making one volatility point in expected profit, and the standard deviation of your profit is one volatility point, then your Sharpe ratio is about one. (The Sharpe ratio is the expected excess profit over the risk free rate divided by the standard deviation, but we are assuming the risk free rate is zero anyway.)

A Sharpe ratio of one is a pretty good trade. Not the best, not the worst. Pretty good. If the only risk from doing the option is the noise that comes from hedging daily rather than continuously, then this is likely a trade you will still want to do.

So in this second context, the noise is not too much to hinder a good trade. Therefore the Black-Scholes model is still useful.

A third way of looking at it is to compare it to transactions costs. The bid-offer spread of a one-year at-the-money call option is on the order of one volatility point in many markets; sometimes less. So the noise introduced into the model by hedging daily rather than continuously is on par with the market impact cost of trading the position. That doesn't seem to be too much either. So again, Black-Scholes is still useful.

Finally, the fourth way of looking at it is through the lens of diversification. The risk you are running here is market risk, but with the direction determined by the particular path followed by the underlying. If you have multiple options in different directions on the same underlying, your risk can be lessened, because the correlations among all those risks are less than one, and so diversification can make the portfolio of many options trades look better than each individually. Even more valuable would be to diversify your options book among different underlyings, particularly ones that are less correlated. Then your risk again diminishes.

So for these four reasons, we can conclude that yes, there is a risk to not hedging continuously, but even on a daily basis, it is not too large, it is comparable to the bid-offer spread or the expected profit, and it can be largely diversified away.



## 6.4 Should You Hedge to Model or Market?

You are working on an equity derivatives proprietary trading desk at a top international bank. One morning, a few minutes before a big research meeting with all the traders, the head of the desk tells you he wants you to talk for a few minutes at the meeting about whether options traders should hedge to model or to market. You nod okay but you're not really sure what he's talking about. Before you can ask any follow-up questions, the head of the desk has stepped out to get coffee. You know that when he comes back, it will be time for the meeting, and he will ask you to present to everybody. What do you do?

### 6.4.1 *The financial mathematician*

The financial mathematician would start deriving partial differential equations and imposing different boundary conditions.

He would attempt to first set everything up just right, make some simplifying assumptions, and combine complicated equations into a closed form solution, and then evaluate that closed form solution for the answers.

There is almost no way this can be done in a few minutes, and even if it somehow is done, it adds very little intuition relative to the long, detailed chain of mathematical work, where a single mistake in any of the links in the chain mean the entire result is wrong.

### 6.4.2 *The financial hacker*

But our approach as financial hackers is much more straightforward. First, we need to figure out what we are really talking about. What does it mean to hedge to model or market? How can you figure out what it means?

Always start with the basics. Let's say we bought a one-year at-the-money call option. What do the words "model" and "market" mean?

In general, with any option trade you put on and hedge to maturity, you must implicitly think the market price is incorrect. In this case, you bought the call because you thought it was cheap.

So your model price is \_\_\_\_\_ than the market price:

- (a) higher
- (b) lower

It must be higher. If your model price was lower than the market price, then that would mean you think the market is paying too much for the option, and you would sell it.

Okay, so the model price is what you think it ought to be worth and the market price is what the market says it is worth. Fine. But what does it mean to “hedge to” model versus market?

Recall what we did in the previous section in testing Black-Scholes. We bought an option, simulated the path of the underlying, and continually rebalanced our delta. That rebalancing of our delta was, of course, the hedging. In that example, we simulated stock returns having a 20 percent volatility, the implied volatility of our option was 20 percent, and the input to our delta calculation was also 20 percent. Which of those “20 percent” numbers is the hedging volatility?

The hedging volatility is:

- (a) the implied volatility of the market price of the option on the day you buy it.
- (b) the implied volatility of the model price of the option on the day you buy it.
- (c) the recently realized historic volatility of the underlying.
- (d) the volatility of the future simulated returns of the underlying.

The answer is none of the above. If you hedge assuming Black-Scholes, then the hedging volatility is the input to your call option delta formula. Recall that the delta of a call is a function like:

$$\Delta = \text{BSCallDelta}(K, T - t, S, \sigma, q, r) \quad (6.7)$$

where the fixed parameters  $K$  and  $T$  are the strike and maturity of the option,  $S$  is the current underlying price,  $t$  is the current time,  $r$  is the interest rate,  $q$  is the dividend yield, and  $\sigma$  is the volatility.

That input  $\sigma$  to the delta formula is your hedging volatility. You change that, and you change the volatility that you are hedging to.

So now the question is starting to take shape and make sense. If you hedge to model, then your  $\sigma$  in the delta formula would be the model volatility that you predicted for the stock on day one. If you hedge to market, then your  $\sigma$  in the delta formula would be the market implied volatility from day one.

For example, if we buy a call option at 20 vol, thinking it will realize 30 vol, should we hedge to 20 vol or 30 vol? And does the answer depend on whether the true vol is 20 or 30?

That's it! We've articulated an exact question. That is at least half the battle.

"If I had an hour to solve a problem and my life depended on it, I would use the first 55 minutes to formulate the right question, because as soon as I have identified the right question, I can solve the problem in less than five minutes."

*Albert Einstein*

Now, when you go to the morning meeting, you can say, "So here's my thinking on the question of hedging to model versus market. If we just look at a Black-Scholes world to start, what happens if we buy at 20 thinking it will be 30, and hedge to 20 versus 30, and see how our P&L looks, both when we are right and when we are wrong. And here are my results..."

Let's quickly get those results. Conceptually, we want the following columns in a spreadsheet for each point along a single path of the underlying (and then we will simulate thousands of paths):

- (1) The time
- (2) The stock price
- (3) The delta of the option if we had hedged to 20 vol (model delta)
- (4) The delta of the option if we had hedged to 30 vol (market delta)
- (5) The hedging P&L on the model delta (yesterday's model delta times the change in stock price from yesterday to today)
- (6) The hedging P&L on the market delta (calculated similarly)

Then, as before, the total P&L of the position would be the final intrinsic value of the option, minus the initial option price, plus the sum of the hedging P&L. We will assume zero interest rates and dividend yields to make things easier without losing any of the important intuition.

But how should we generate the stock prices? Should those simulations be done at 20 vol or at 30 vol?

The financial hacker doesn't think too long. There's no time. The financial hacker says, "Let's try both."

In fact, the financial hacker notices that there is no longer any concept of "model vol" or "market vol" in the simulations he's about to do. The key variables are rather the simulation vol which generates the stock prices and the hedging vol which determines the delta.

So really all we need to do is run all four combinations of simulation vol and hedging vol each being either 20% or 30% and see what the effect on the P&L is.

How can we easily measure the effect?

First, we can look at a histogram of the P&L, as we did in Section 6.3.

Second, we can look at how the total P&L is affected by the last price in the simulated stock price path.

And third, we can look at how the total P&L is affected by the realized volatility along the simulated stock price path.

The next page shows these three graphs for all four possibilities.

The relevant code is only a few simple lines:

```
modelVmarket[ohedge_, osimul_, n_:1000, dt_:1. / (252 * 4),
  S0_:100, K_:100, r_:0, q_:0, T_:1] :=
Module[{initpx = BSCall[S0, K, r, q, osimul, T],
  Ts = ReplacePart[Range[T, 0, -dt], -1 -> €]},
  Simulate[
    {Last@#, StandardDeviation@Differences@Log@# / Sqrt@dt,
      (BSCall[Last@#, K, r, q, osimul, €] - initpx -
        Dot[Most[BSCallDelta[#, K, r, q, ohedge, Ts]],
          Differences[#]]) / initpx} &, n, r - q, osimul,
    dt, S0, T]];
```

### 6.4.3 The histogram of the P&L

There's a lot to learn from these graphs. Let's start with the P&L histogram when both the simulation and hedging volatilities are the same, either both 20% or both 30%. Then as before we see that the P&L is centered around zero and has a standard deviation of about 5 percent. It's nice to test new results by making sure they continue to agree with what we already know: it both verifies your code and provides context for the graphs to make it easier to spot the new differences.

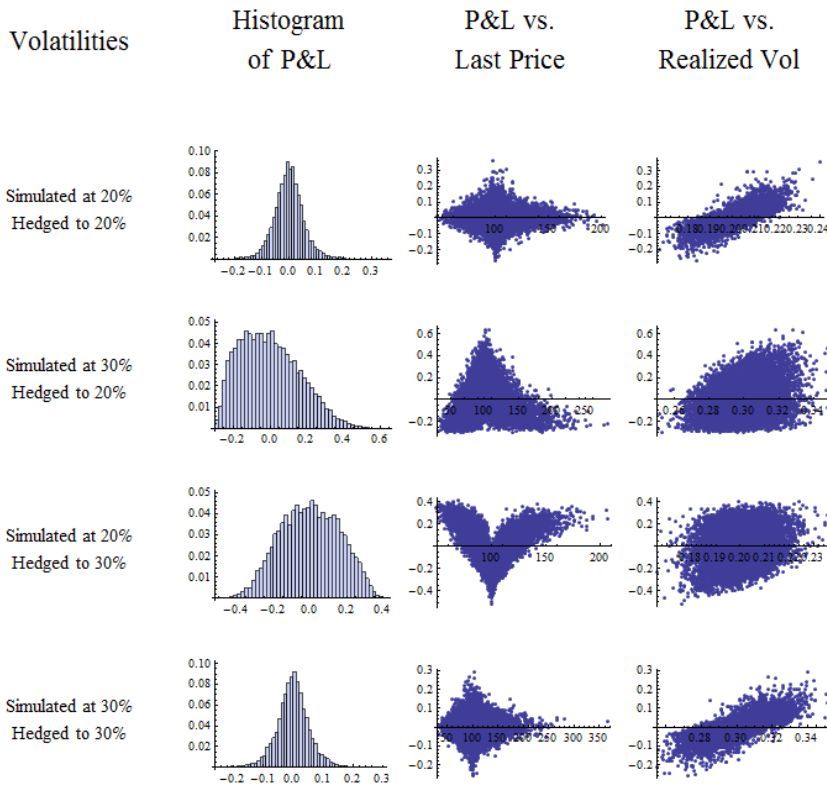


Figure 6.5. Market vs. model

Compare those P&L histograms when we hedge to the truth to those when we hedge away from the truth. When the underlying is drawn from a 30 vol distribution but we hedge to only a lower 20 vol, the standard

deviation of the return is three times as high, at more than 15 percent. Even curiously, the average is no longer centered around zero: the median return is negative. What is going on?

(The case when we hedge to the higher 30 vol on a 20 vol underlying is similar: the standard deviation is again 15 percent but the mean and the median are both positive. If we understand one, we would understand the other.)

#### **6.4.4 *The scatterplot of the P&L versus the last price***

To make progress, we need to look at the second column of charts, plotting the P&L versus the last price in the underlying path. When we hedge to the truth at either 20 or 30 vol, the P&L is symmetric around the x-axis. This means that no matter where the stock price ends up at maturity, we are equally likely to have made money as to have lost money.

Why do the points taper off at the sides and bunch up and down in the middle? Imagine the extremes. If the stock ends up at a very low price, the delta of the call probably went to zero very quickly. There was therefore very little risk left. Once the delta became zero, there was no more hedging, and so it didn't matter whether you didn't hedge on a daily basis or you didn't hedge on a continuous basis.

Similarly, if the stock price ends up very high, then the delta of the call probably went to one very quickly. Again, there was therefore very little risk left. Once the delta became one, there was no more rebalancing of the hedge, and so it didn't matter whether you didn't rebalance on a daily basis or you didn't rebalance on a continuous basis.

But the plots in the cases where we didn't hedge to the truth is different. Let's focus now on the third row where we hedge to a higher vol than the truth. This is the prototypical case of hedging to the model. Presumably we believe the future realized volatility should be higher than the implied volatility of the price we bought it at. So we might consider hedging to our best forecast.

In that case when the hedge vol is 30 and the simulation vol is 20, the P&L as a function of the last price looks like a heart. It is not symmetric around the x-axis and it doesn't taper off at the sides: for terminal stock

prices far from the initial stock price, we are guaranteed to have made a profit, and for terminal stock prices near the initial stock price, we are much more likely to have lost.

Why does it have this odd heart-shaped pattern?

- (a) It is a mistake. We made an error somewhere.
- (b) It is random. Run the simulations again and it will go away.

The answer is neither. The heart is real.

To figure out what is happening, again consider the extremes.

What does it mean if the stock ends up around 200? It probably means that the path was by and large straight up, because there is not enough time or volatility for the stock to have fallen first. The paths that end at the highest attainable levels of the underlying price must have attained those levels by going there as fast as possible, just like the winner in a sprint probably didn't dilly-dally along the way.

So if the underlying grew for much if not most of the history of the path, what kind of risk were you running by being mishedged?

Remember, you are hedging to a higher volatility than is being realized.

So is your delta along such a path higher or lower than it ought to have been?

- (a) The model delta of the call, using 30 vol, was higher than the market delta of the call, using 20 vol, along the paths that ended up at high terminal values of the underlying price.
- (b) No, the model delta was usually lower than the market delta.

To mix things up a little, the answer this time is not none of the above. It is just plain old (b).

To understand why on an intuitive level, recall that the delta can be thought of as approximately the probability of ending up in-the-money. Along a path that rises from 100 to 200, the option struck at 100 was almost always in-the-money.

So think of an in-the-money option, in general. If the volatility of the underlying suddenly increases, what do you expect to happen to its probability of ending up in-the-money?

- (a) It will go down.
- (b) It will remain about the same.
- (c) It will go up.

The probability should go down, because there is a greater chance that the stock will drop so far that it will actually go beneath the strike price.

If that doesn't make sense to you, think of the opposite: if the volatility falls drastically, say all the way to zero, what would happen to the probability of ending up in-the-money, if it is already in-the-money? The stock would no longer move, so naturally the probability would increase all the way to one.

If the intuition is not immediately obvious, you can hack it. Simply plot the sensitivity of the delta of a call option to the implied volatility as a function of the underlying price.

In the crazy language of option Greeks, this sensitivity is called “vanna,” which, if you are a careful student of history and linguistics, you may recognize as having never in the history of the world been an actual Greek letter.

You may then rightly wonder why option “Greeks” are called that. You will then learn that there is no good reason.

In the world of Black-Scholes, vanna can be calculated equivalently as either the derivative of delta with respect to the volatility or the derivative of vega with respect to the spot price.

In general, though, all we really need to plot is the difference between the model delta and the market delta, literally the delta at a higher vol minus the delta at a lower vol. Figure 6.6 shows what it looks like as a function of the spot price.



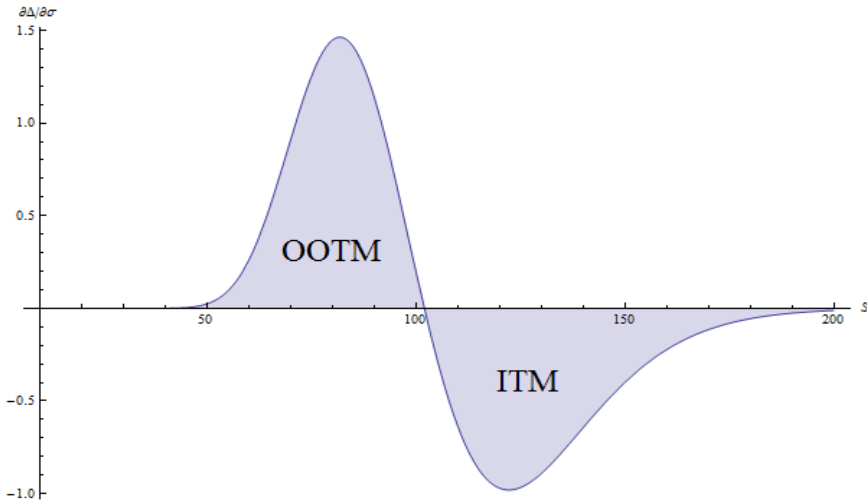


Figure 6.6. Vanna vs. underlying

The sensitivity of delta to vega is positive when the option is out-of-the-money and it is negative when the option is in-the-money. In other words, when you are hedging to a higher vol than the true vol, and your call is in-the-money, the delta you calculate will be smaller than it would have been if you had used the true vol.

So here's what's happening on paths that end up on high levels. You hold an in-the-money option. You hedge to a higher volatility than the truth. So your model delta is too low. You hedge by selling the model delta, but because it was too low, you did not sell enough of it to truly hedge yourself. Thus you are net long delta. And the market continues to climb. Along the whole path then, you are long the market, and the market is going up. Of course you end up making money.

That explains the right part of the heart.

The left part is similar. The paths that end up at the lowest final stock prices must have been falling the whole way. So your call option was out of the money the whole time. And because you were hedging to a higher than true vol, your delta was too high. You sold that delta; hence you were net short the market, and the market kept falling, so you again made money.

“As a child, when I was having that alphabet soup, I never thought that it would pay off.”

*Vanna White*

So we've explained the left and right parts of the heart. What about the middle? There are two ways to think of it. The easy way is to note that so far we have not had any negative P&L at all, and so the middle should compensate for that, by being skewed towards losses. This is a reasonable way of looking at it.

The more complicated way gives a little bit more intuition. How can a volatile path end up where it started? Only if it is mean reverting, meaning that whenever it goes up, it tends to go back down the next day, and vice versa. Think of a path that constantly jumps between 102 and 98. At 102, your call is in the money, and you are net long delta because you are hedging to a higher-than-true vol. But then the market declines the next day, and so you lose money. At 98, your call is out of the money, and you are net short delta because you are hedging to a higher-than-true vol. But then the market soars the next day, and so you lose money. Each day, you lose money. That's why the middle is so pointy and lossy.

A final, somewhat oversimplified way of remembering the intuition is that as a trader you are betting that volatility will be high, that is to say, the stock will move far away from where it started. When it does that, you end up making money; when it doesn't, you lose money.

The upside-down heart when you hedge to a lower-than-true vol that appears in the second row of Figure 6.5 is not as sharply defined. Here, the true vol is 30 and your hedging vol is only 20. The reason the heart is less well-defined is because your error is more bounded. Let's walk through this logic.

When you are in-the-money, your 20-vol delta is typically above 0.5. Let's say on average it is about 0.75. Your 30-vol delta is going to be lower, and there is a lot of room for how much lower it can be; it can go all the way down to zero. Similarly, when you are out-of-the-money, your 20-vol delta is typically below 0.5. Let's say on average it is about

0.25. Your 30-vol delta is going to be higher, and there is a lot of room for how much higher it can be; it can go all the way up to one.

But now think if you are hedging to a lower vol, like in the second row of Figure 6.5 where the true vol is 30 and you are hedging to 20. When you are in-the-money, your 30-vol delta will be a little lower than in the above example, but let's say it averages about 0.7. The 20-vol delta will be higher, but there is not much room for how much higher it can be; the highest it can be is one. Similarly, when you are out-of-the-money, your 30-vol true delta will average about 0.3, and your 20-vol model delta will be lower, but there is not much room for how much lower it can be; the lowest it can be is zero. Thus, the difference between the model and the market delta is muted. It is still there, and you can still make out an upside-down heart, but it is less well-defined because the differences are smaller.

#### 6.4.5 *The scatterplot of the P&L versus the realized vol*

Let's go back to the world where we hedge to the truth.

“The truth is not for all men, but only for those who seek it.”

*Ayn Rand*

Look at the third column of plots in either the first or the fourth rows, i.e. where the hedge vol and simulation vol are either both 20 or both 30. It looks like the P&L from dynamically hedging an options position depends linearly on the realized volatility. The realized volatility, of course, is just the annualized standard deviation of the simulated log-returns.

If we believe Black-Scholes, shouldn't the P&L be randomly near zero? In other words, we can accept noise around the theoretical zero profit that we should see, noise coming from the failure to hedge perfectly continuously. But that noise should be random.

In the histogram, and in the scatterplot relative to the final stock price, it is indeed random: both overall and for any given final price, you are just as likely to make money as to lose money.

But when it comes to the realized volatility, it seems that the profit depends very heavily and very linearly. When the realized volatility is more than 22, you are guaranteed to make a profit, and when it is less than 18, you are guaranteed to lose.

Why does the P&L of a hedged options position depend so strongly and linearly on the realized volatility of the simulated path?

- (a) Because the assumptions behind Black-Scholes are wrong.
- (b) Because the derivation of Black-Scholes was wrong.
- (c) Because we made an error in our simulation.
- (d) It is just random. If we resimulate, it will go away.

The answer is none of the above. The anomaly is real and replicable.

Monty Brewster: What are you gonna vote?

Crowd: [in unison] None of the above!

*Brewster's Millions* (1985)

At this point, you should object.

“Hold on,” you should say. “We don’t know ahead of time what the realized volatility will be. It is just as likely to end up at 22 and generate a profit as it is to end up at 18 and generate a loss. So it’ll still be random.”

Your statement is true. But remember that Black-Scholes claims that you should have zero profit along every single possible path. We can accommodate some noise due to discontinuous hedging, but that noise better be random and independent.

“Maybe it’s the same error,” you opine. “Maybe the discontinuous hedging is what is causing the realized volatility to differ from the true volatility.”

Aha! You are right. If we look at thinner and thinner slices of returns, going from days to hours to minutes to seconds, the range of the possible realized volatility compresses. In the continuous limit, the realized volatility is always exactly 20.

So it is true that you cannot make money with the observation that the P&L is related to the realized volatility, because even if it has realized 22

after a few months, you don't know if the future realized volatility, being just another noisy draw from a population with a true vol of 20, will be above or below 20.

“It ain't what you don't know that gets you into trouble. It's what you know for sure that just ain't so.”

*Mark Twain*

Here is why the P&L is positive for higher realized volatilities: it is because those paths are indistinguishable from paths generated from a higher true vol. Consider one specific path that was simulated from a 20 vol underlying but that randomly exhibited a realized vol of 22. That path could just as easily (more easily, really) come from a simulation of a 22 vol underlying.

Now, if you had a 22 vol underlying, and you hedged it to 20 vol, that would introduce some noise to your profit. But, on the other hand, you must have gotten a good deal on your initial options purchase, because you bought it at 20 implied vol, rather than at 22.

Thus the profit you have made is due to the 2 implied volatility point discount on the initial option, coupled with the noise from mishedging to the wrong vol. The explanation for the losses when the realized volatility is low is similar.

The short-form intuition is this: you bought a call and hedged it. So you are betting on higher volatility. When volatility ends up higher, even if only for random reasons, you benefit, and when it ends up lower, you lose.

That intuition continues to hold even if you hedge at the wrong vol. If, for example, the true vol is 30 but you hedge to 20, you are just introducing noise. The slope between your P&L and the realized vol is still positive, but not as sharply defined.

#### **6.4.6 *The verdict and the missing piece of the puzzle***

So what's the answer? Should we hedge to model or to market?

There are two possibilities. If the model is correct, then we have seen above that hedging to the truth reduces the overall risk as well as the

asymmetry depending on the final stock price. So if the model is also the truth, it seems we should hedge to the truth too. And if the market is the truth, then we should hedge to the market. Seems a little simplistic, doesn't it? Especially after all this work.

“Ye shall know the truth, and the truth shall make you mad.”

*Aldous Huxley*

Three responses to that.

First, we know how just how much more risk hedging to the market is relative to the model, given that the model is right: it triples the hedging risks.

Second, we know what our risk looks like when we hedge to market instead of the true model: it looks like we are long a straddle, making more money when the underlying moves far away. This is already very useful.

But there is a third response, and that is that there is a wrinkle, something we have missed in all of the preceding analysis. What is this missing piece of the puzzle?

Is it the initial option price? Note that nowhere did we ask about what price we bought the option for. Could that be what is missing? If in one scenario we get a better deal, shouldn't that influence our decision?

No, because that has been assumed away. The question has been about how best to hedge an existing position, not whether to put it on or not. The position is on the books; how should you hedge it?

The question of what price you bought it at is irrelevant. The position has already been marked to market, so your initial execution brilliance or idiocy has already been accounted for.

That is why all of our graphs have been essentially centered around zero; the effect of buying the option cheaper would merely be to shift the center to a positive P&L without affecting any of the discussions about the risk.

The missing piece of the puzzle is the implicit assumption that all we care about is the terminal risk, the risk of our position at the end of its life.

Imagine you are a trader and you are holding a position in ten-year calls. The chances that you will still be at the same firm ten years from now is relatively small. Further, your bonus and your employment will be reviewed within one year, less than ten percent of the way into your option trade. The trade is not likely to converge by then.

What is the risk you are worried about? Is it really that your total P&L after ten years might have some noise on it? Or is it that your mark-to-market P&L may fluctuate?

“Markets can remain irrational a lot longer than you and I can remain solvent.”

*A. Gary Shilling*

Of course, you are likely more worried about the mark-to-market P&L. From that perspective, does it make a difference whether you hedge to model or to market?

Yes, it does.

If the underlying goes up a dollar, the option price will go up by the market delta. That’s nearly the definition of the market delta, and it follows exactly if the market implied volatility remains unchanged. If your hedge is not equal to the model delta, then you will experience some volatility in your mark-to-market P&L.

Thus, if you want to minimize your mark-to-market P&L, you may choose to hedge to the market even if you think the market volatility is wrong.

How do you trade-off these two risks, the mark-to-market risk versus the at-maturity risk? Ultimately, you probably will decide based on the maturity of the option you are hedging.

If the option will expire in a month or two, you will almost surely be able to weather any intermittent mark-to-market volatility, so you will lean towards hedging to model.

If the option will expire in many years, you will likely lean towards hedging to market, at least until the expiry gets closer.

And what do people do in practice? They hedge their bets on how to hedge.

One common rule of thumb is to hedge halfway between the model and the market delta. Then you're never exactly hedged, but you're never too far away either.

"It is better to be vaguely right than exactly wrong."

*Carveth Read*