Option
Trading
Session Five: Hedging and Trade Evaluation

This is an adapted rendition of Dr. Euan Sinclair's lecture notes

About me

- Varun Kumar Pothula
- Options trader, global macro analyst, algo trading strategist
- Current Role: Quantitative Analyst at Quantra

Let's start learning!

Session Five Overview

- Hedging in practice.
- Expiration trading.
- Early exercise.
- Trade planning and evaluation.
- Risk measures.
- An example that summarizes the volatility trading process.

Hedging in Practice

- In order to convert an option into a volatility trade, we need to maintain delta neutrality.
- In theory, we hedge continuously.
- In practice that leads to infinite transaction costs.

Volatility Trading

- Theoretical profit:

$$PL = Vega(\sigma_I - \sigma_R)$$

- But only on average…

v10.1.1

Volatility Trading: Replication

- We own the one year 100 call on a $100 stock with volatility of 30%.
- It is worth $11.92 and has a delta of 0.56 so to hedge we sell short 0.56 shares.
- Now the stock jumps to $110. The call is $18.14, and the delta increases to 0.68.
- So, we need to sell 0.12 shares to stay hedged.
- At expiration this process captures the difference between implied and realized volatilities.

Hedging: First Idea

- Whenever the option is in the money, hedge it as a 100- delta option.
- Example: Stock is \$100, and we sell a 101 call.
- Don't hedge because option is OTM.
- When stock goes above \$101, buy a share.
- If it drops again, sell the share.
- If S<X at expiration: keep entire premium, C.
- If S>X at expiration: Option P/L=C-(S-X)

: Hedge P/L = S-X

: Total P/L = C

v10.1.1

Hedging: First Idea

- Sadly, this also leads to a very volatile P/L.
- (also, infinite transaction costs in the limit)
- Simulation: long a one-month 101 call with implied and realized volatility both 0.3.
- Initial call value is 3.00.
- Theoretical PL=0 – costs and slippage

v10.1.1

Hedging: First Idea

- Assume zero transaction costs
- Average PL=$4
- Median PL= -$27

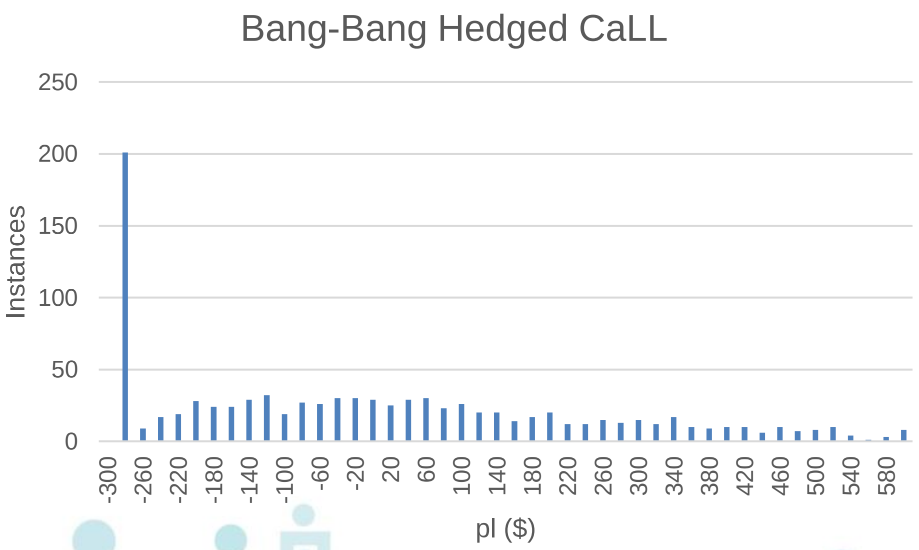

v10.1.1

Hedging: First Idea

- Generally, this is cheating. It can *look* good, particularly when short options, but in the long-run it will hurt.
- In spite of this, this is my recommended way of hedging daily straddles. It is a risk but not a huge one because of the very short time frame.
- Need something more subtle and closer to continuous.

v10.1.1

Hedging Heuristics

- Continuous hedging leads to infinite transaction costs.
- Hedge at a given time.
- Hedge at a given price move, either in terms of $ or standard deviations.
- Hedge to a set delta band.
- All are sub-optimal in terms of cost relative to risk reduction.

v10.1.1

Utility Based Methods

- Utility is the concept of balancing risk and reward.
- How much are we prepared to pay to reduce risk?
- Since BSM really prices replication, we can do same thing with costs.
- At some point, you will be indifferent to risk and costs of removing it.
- But now personal preferences matter.

v10.1.1

Utility Based Methods

- Many functional choices for utility but all are increasing and convex.

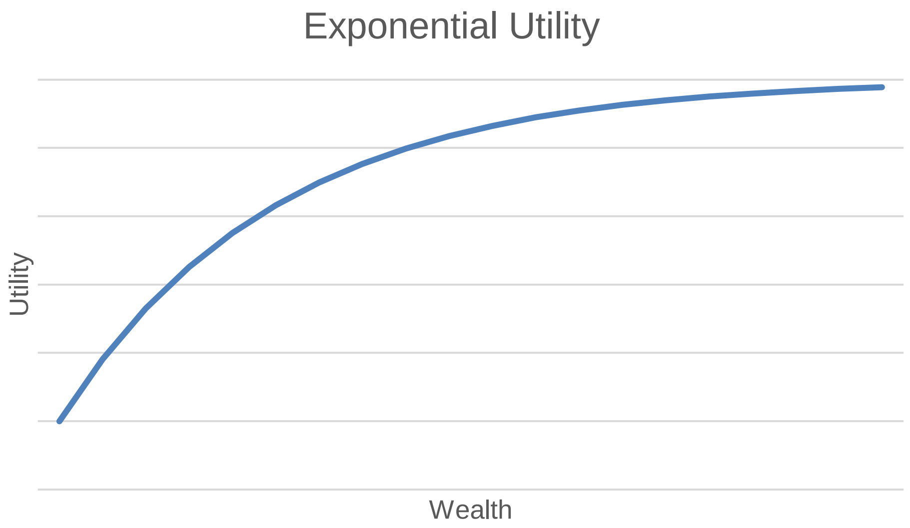

v10.1.1

Utility Based Methods

- Assume wealth is either $100 with probability 0.6, or $0 with probability 0.4.
- What guaranteed amount would you take instead of this bet (the "certainty equivalent")?
- Here my number is $55.

v10.1.1

Utility Based Methods

- We need expected utility of the certain value to be the same as the expected utility of the gamble.

$$U(W) = -exp(-\gamma W)$$

$$U(bet) = -0.6exp(-\gamma 100) - 0.4exp(-\gamma 0)$$

$$U(certain) = -exp(-\gamma 55)$$

- Equate this to get gamma (0.0041).

v10.1.1

Utility Based Methods

- Redo the BSM model but allow for a spread in the underlying.
- Hedge to the edge of a band that is gamma dependent.
- Hedge short gamma more aggressively.
- (play aggressive defense and let profits from long gamma run).
- High costs=> wider band.
- High risk aversion=> narrow band.

v10.1.1

Utility Based Methods

- This theory results in having to solve a nasty PDE.
- Can’t be done in real time (or even close).
- But, using an asymptotic expansion, we can come up with an easily used approximation.

Utility Based Methods

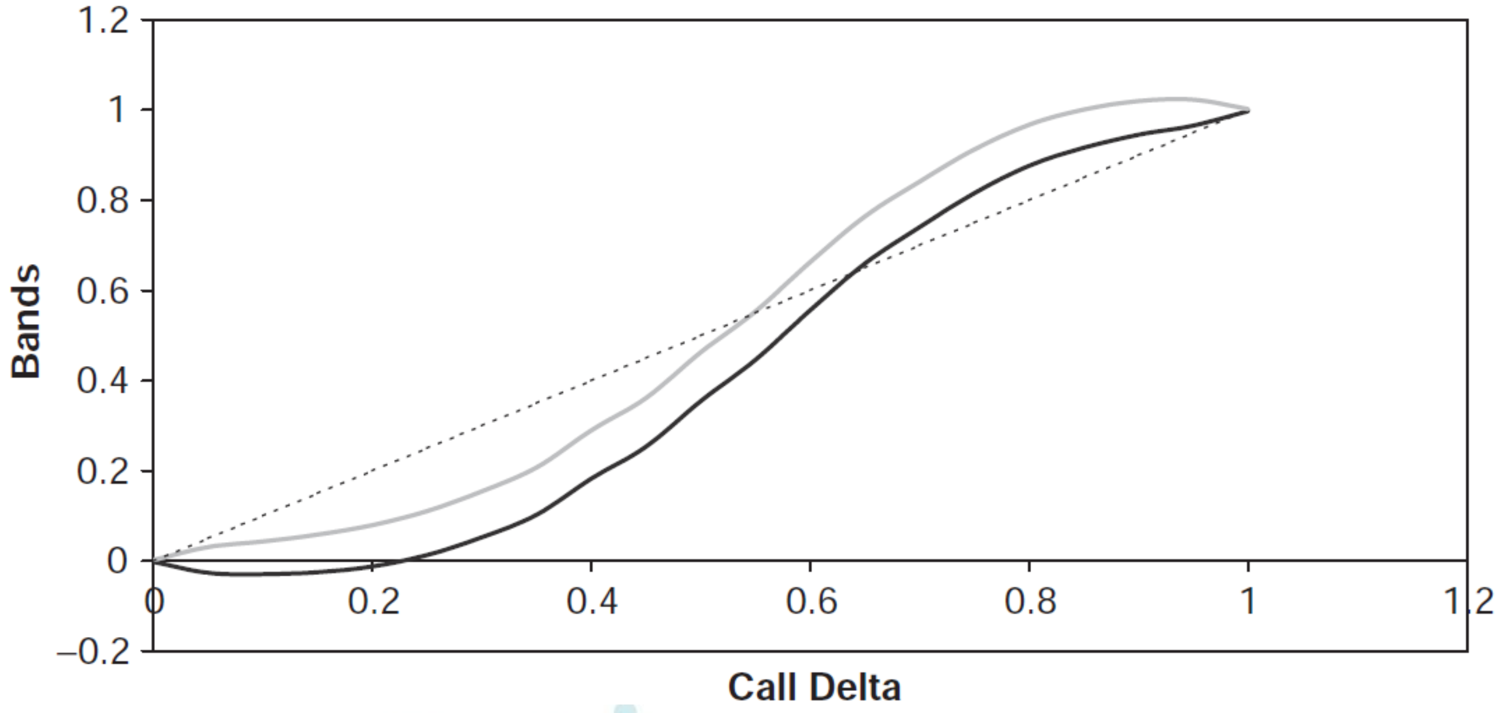

Utility Based Methods

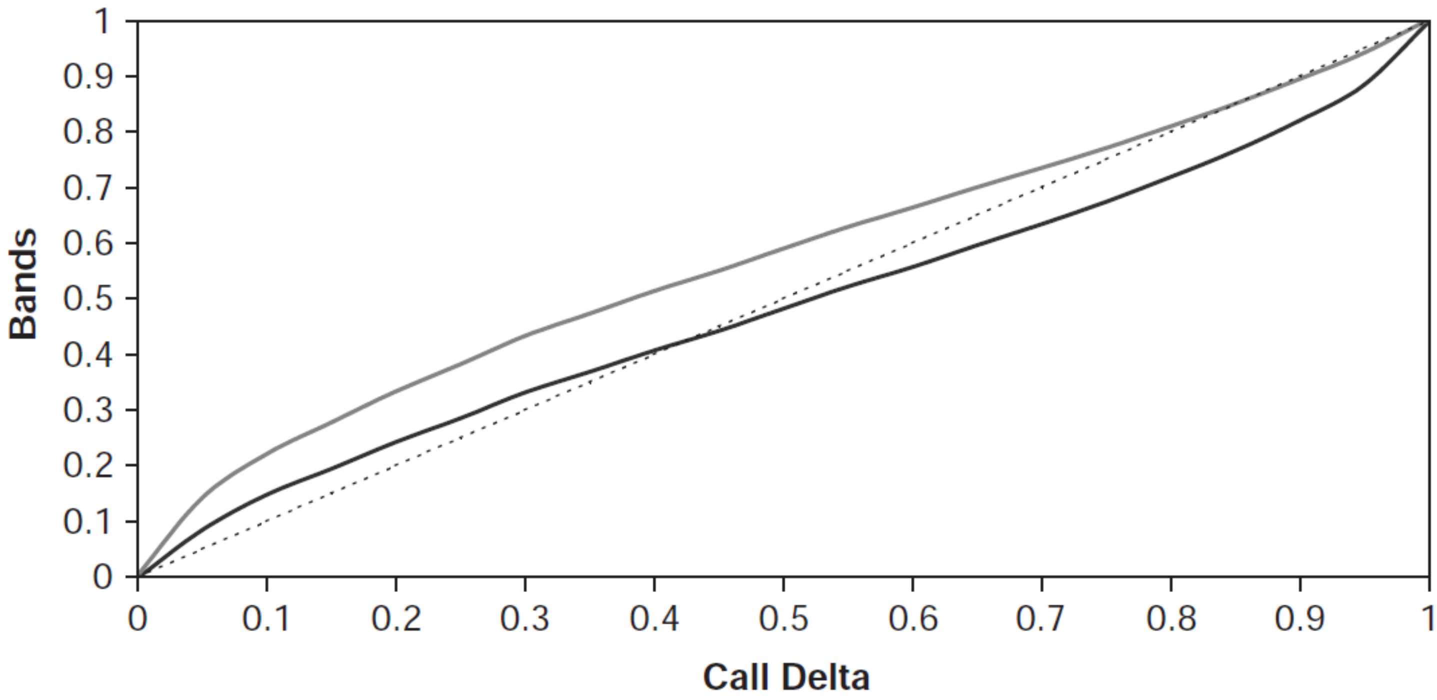

v10.1.1

Utility Based Methods

- Short option positions have a narrower band.
- This is because of utility: the trade off between certainty and risk.
- For shorts, we have a “head start” and are willing to pay costs to defend this lead.

Utility Based Methods

- For small delta options, the bands may not span the BSM delta.
- E.g. for a long call the BSM delta might be 10, but the band may be between 6 and 8.
- The long option “sees” a lower volatility and $d\Delta/d\sigma > 0$.

v10.1.1

Wilmott-Whalley Approximation

$$\Delta = \frac{\partial V}{\partial S} \pm \left(\frac{3}{2} \frac{exp(-r(T-t))\lambda S \Gamma^2}{\kappa} \right)^{\frac{1}{3}}$$

- Lambda is proportional transaction cost.
- Kappa is risk aversion.
- So, choose a risk aversion parameter for a position you understand, then use this for all positions.
- This can save about 10% of costs over ad hoc methods.

v10.1.1

Wilmott-Whalley Approximation

- Costs are insidious because in the short term they are too small to notice.
- Say we are trading one-month options on 100 stocks. If volatility is 30%(a typical number) and stock price is $30 (US average), a typical daily move is about $1.
- If we have, on average, 1,000 straddles each day our delta changes by about 30,000 share per stock or 3,000,000 shares.
- Bid/ask of 0.5 cents a share gives a cost of $15,000 a day.
- Saving $1,500 (10% of costs) a day adds up…

Wilmott-Whalley Approximation: Example

- Use a reference position to calibrate.
- Assume I'm very comfortable trading 30-day SPY options. I'm going to use these to calibrate my risk.
- Currently I will let a zero-delta straddle become 5-delta.
- With volatility at 30%, index at 200, gamma is 0.07, and transaction cost is 0.03/200.
- This implies a risk aversion of 0.24.

Wilmott-Whalley Approximation: Example

- Now assume I'm trading one-year options on an unfamiliar index.
- With volatility at 50%, index at 100, gamma is 0.015 and transaction costs are 0.05/100.
- Using my implied risk aversion of 0.24 I get a hedging band of 0.04.

v10.1.1

Hedged Position Results

- Volatility edge means you win on average.
- Any particular trade is hugely variable, even if your vol. forecast is right.
- For example, if implied is 20% and realized is 15% and you sell options, you can (easily) still lose.

v10.1.1

Path Dependency

- Discrete hedging creates path dependence.
- Simple example: a big move the day of expiration (when an option has a lot of gamma) will give a different PL than the same move a year out.
- In each case the volatility is the same, but the result isn't.
- Dispersion is roughly inversely proportional to square root of number of hedges.
- Note: unhedged European options are not path-dependent.

v10.1.1

Hedging Frequency

- Dispersion is roughly inversely proportional to square root of number of hedges.
- For example, $1000 vega of one-year options traded with no volatility edge has expected P/L of zero.

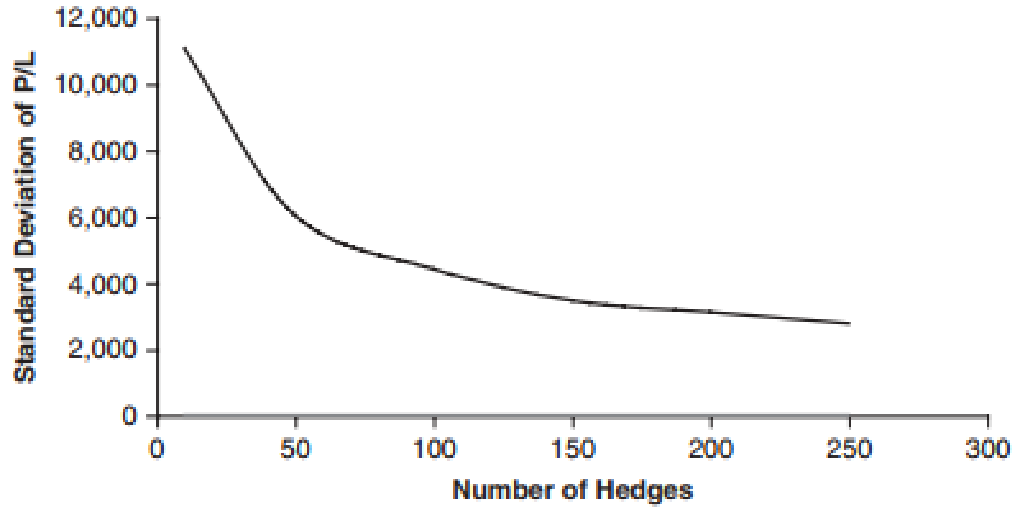

v10.1.1

- Hedging more gives lower variance but more costs.

.

Choosing a Hedging Volatility

- You need to put a volatility into your pricing model.
- The volatility you use will determine delta and gamma and hence when and how you hedge.
- You are free to choose whatever you want…
- Implied vol. => low MTM variance but an uncertain final number.
- (True) Realized vol. => higher MTM variance but a guaranteed profit (if your forecast is correct).

v10.1.1

Hedging at Implied Volatility

- One-year options ($1000 vega) sold at 40% vol, hedged at 40% vol, realized vol 30%.
- Daily standard deviation: $140.

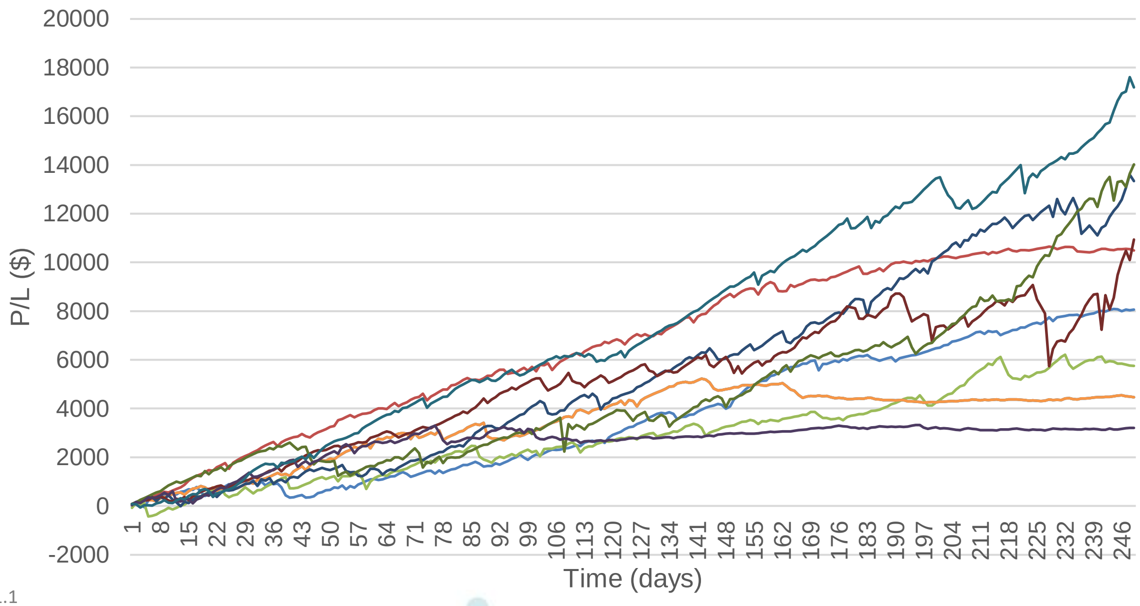

v10.1.1

Hedging at Realized Volatility

- One-year options ($1000 vega) sold at 40% vol, hedged at 30% vol, realized vol 30%.
- Daily standard deviation: $231.

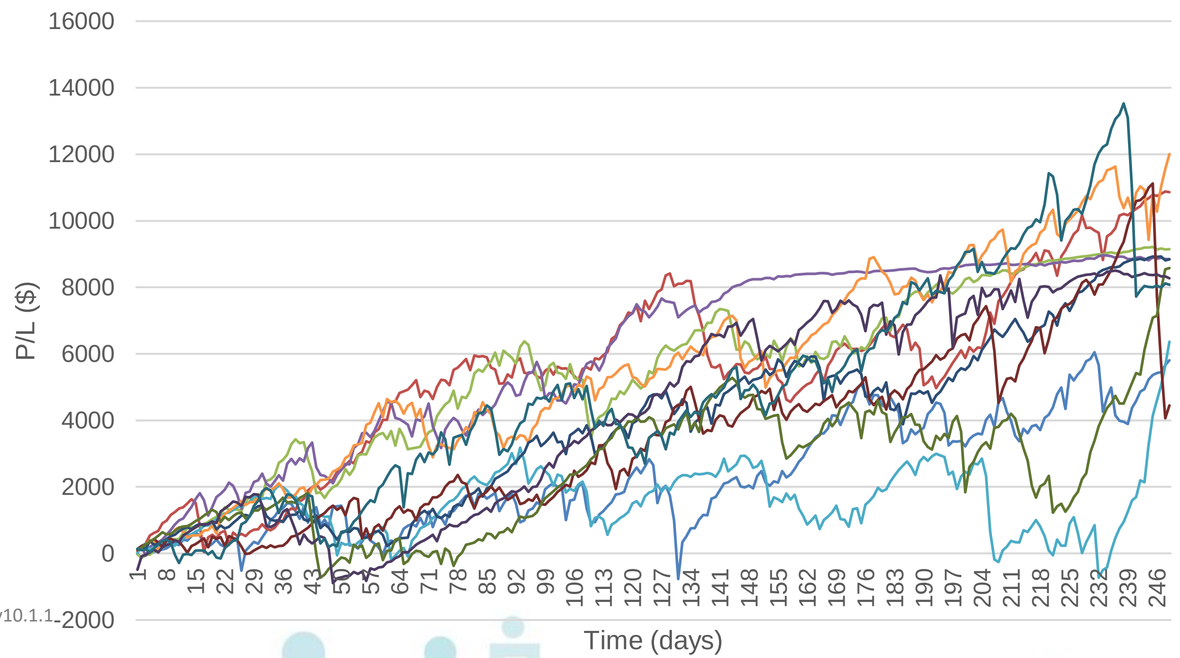

v10.1.1

Hedging at Implied Volatility

- One-year options ($1000 vega) sold at 40% vol, hedged at 40% vol, realized vol 30%.
- Standard deviation of final P/L: $4,460.

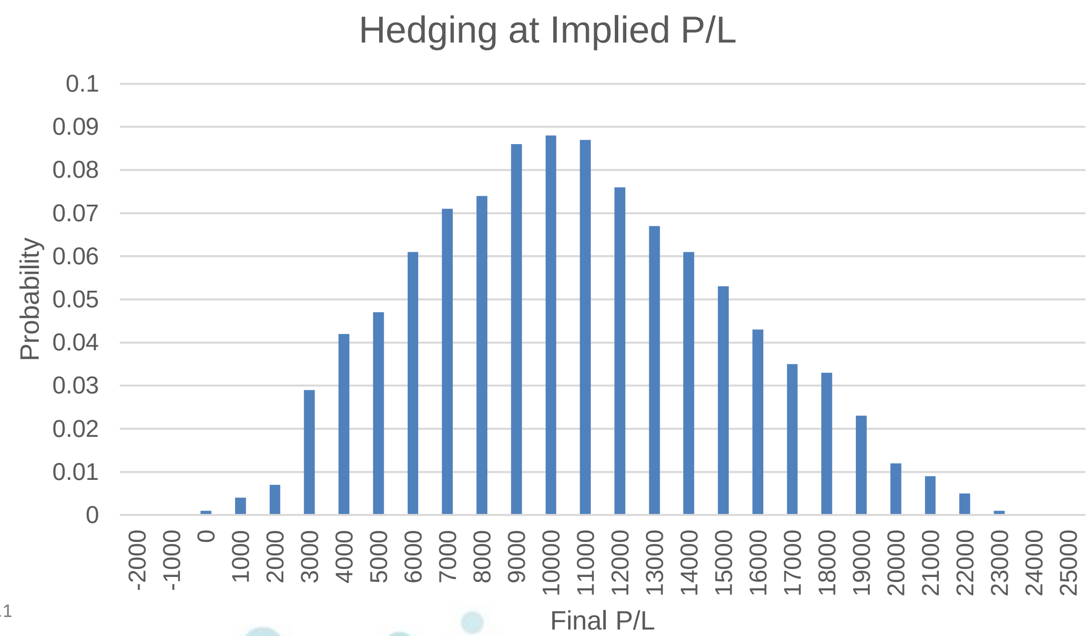

v10.1.1

Hedging at Realized Volatility

- One-year options ($1000 vega) sold at 40% vol, hedged at 30% vol, realized vol 30%.
- Standard deviation of final P/L: $2,840.

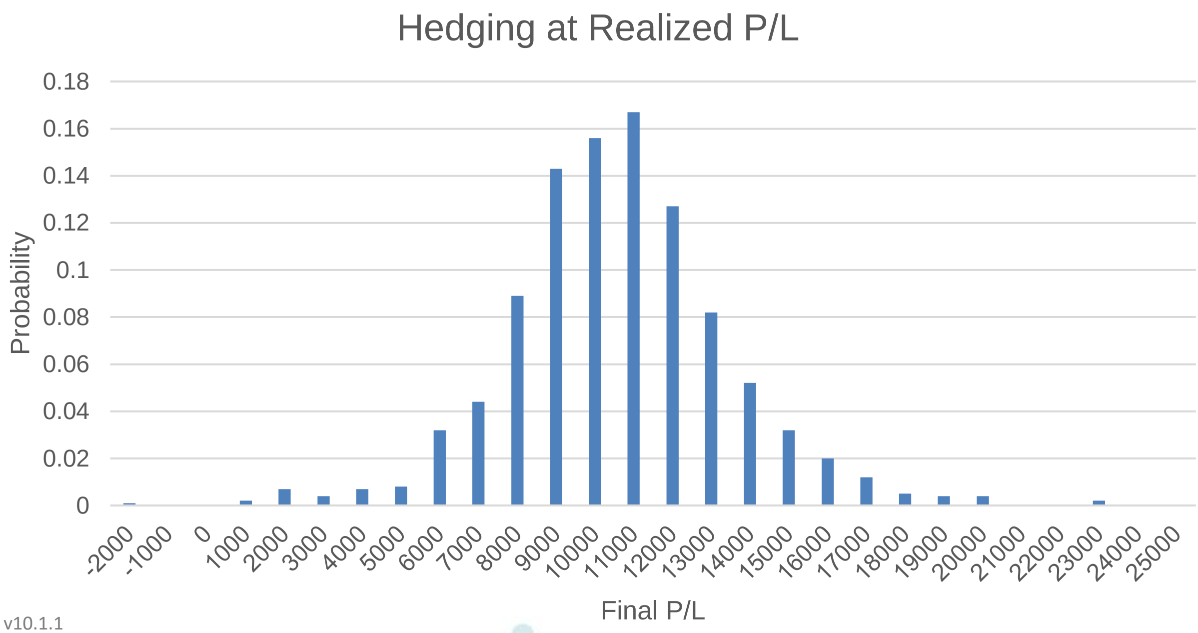

v10.1.1

Drift VS No Drift

- When the underlying is drifting and we are long gamma, our hedges are going to be losers. We will be selling into a rising market. So, we want to hedge less often.
- If we use a higher volatility, we see a lower gamma, so we hedge less.
- Traders refer to this hedging trick as *letting their deltas run* if they are long in a trending market,
- Or *hedging defensively* if they are short gamma in a trending market.

Drift VS No Drift

Position	Market	Volatility Bias for Hedging
Short Gamma	Trending	Low
Short Gamma	Range Bound	High
Long Gamma	Trending	High
Long Gamma	Range Bound	Low

v10.1.1

Early Exercise

- Failing to exercise correctly is just a way to leave money on the table.
- You exercise early if what you will receive is worth more than what you forego.
- For example, when exercising a call you get the stock, but you miss out on any remaining optionality.
- Compare the alternatives and choose the better one.
- Note: Sometimes we need to consider the OTM option of the same strike price as well. We may need to buy it to cover risk.

Early Exercise

- Expected duration of American option is
- $\frac{\rho_A}{\rho_E}$
- Example: S=X=100, T=1, r=0.1 and volatility=50%, for the puts
- $\frac{\rho_A}{\rho_E}=\frac{0.293}{0.489}=0.6$
- 40% chance of being exercised early.

Exercising a Put on a Stock

- Exercised to avoid the interest costs of holding the shares until expiration.
- Underlying is $100,we own the 120 put. T=30 days r= 5%.
- If we don’t want to change our short delta position in the market, we can either:
- Hold the option
- Sell the option and sell the underlying
- Exercise the option

Exercising a Put on a Stock

- Exercise and invest the cash. This earns interest=$120 x 0.05 x 30/365 = $0.49
- By not exercising we forgo this money so we should exercise if the option is trading less than 49c above intrinsic.

Exercising a Call on a Stock

- Turning a call into a share is a way to get dividend.
- Exercise as close as possible to the record date.
- Exercise-> get intrinsic and dividend but lose time value.
- Hold-> value the option against ex div price and a day later.

Exercising a Call on a Stock

- An Example:
- 80 strike call on a $100 stock is worth 20.10.
- Stock is going to pay a $1 dividend.
- Exercise and get $19 intrinsic and the $1 dividend.
- Don't exercise and the next day the call is worth $19.10
- So exercising is 90c better.
- (dividend is bigger than extrinsic value).

Dividend Capture

- Buy a deep ITM call spread.
- Exercise all of your long calls correctly, collecting 100% of the dividend.
- Some of your shorts probably won’t be exercised so you won’t need to pay the dividend out on all of them.
- Ideally, your short strike will have large open interest to give more people an opportunity to make a mistake.

Exercising a Call on a Future

- Turning a call into a long future is a way to avoid paying interest.
- Exercise-> can buy futures and not have to pay interest on the option premium.
- So, exercise if interest income is more than time value.
- Need to check every day.

Expiration Trading

- Very close (hours) to expiration, options lose most optionality.
- If ITM behave as underlying.
- If OTM behave as worthless.
- The tricky part is what happens at the strike.
- Note that pricing models don't "break" at expiration. They are still a helpful guide.

Expiration Trading: Pinning

- The underlying will settle near a strike more than randomness would imply.
- This is due to market makers hedging long gamma, as a result shorts don’t need to hedge.
- Sell ATM options close to expiration, ideally those with high open interest.

v10.1.1

Expiration Trading: Pin Risk

- If the market settles at a strike, it can be very difficult to predict which options will be exercised.
- Be aware that slightly OTM options can be exercised.
- Sometimes done to squeeze the underlying or to avoid slippage in liquidating a large underlying position.

Expiration Trading: Cash Settlement

- Cash settled options will be hedged with futures.
- At expiration, the futures position will still be there.
- Need to allow for unwinding.
- Example: S=100, long the 90 call and short one share.
- If this is stock settled, at expiration I exercise my call and get a share. This nets out with my short and I have no position.
- If this is cash settled, I exercise my call and get $10 but I still have a short stock position.

Expiration Trading: Greeks

- Greeks are still correct but may be misleading.
- For example, a one-hour straddle will often have theta 20 times higher than its value.
- Also, gamma is actually infinite right at the strike.
- Sometimes useful to set implied volatility to zero to avoid possible confusion.

Trade Evaluation– What is a Good Trade?

- A good trade is one that we would repeat no matter the result.
- Requires positive expectation and an acceptable level of risk.
- EV is risk-neutral but "acceptable risk" varies between traders and firms.
- Multi-dimensionality and path dependency of options mean results don't tell whole story in any single case.
- Repeated trades of the same strategy eventually give us enough data to analyze.

A “Good Trade Strategy” Example

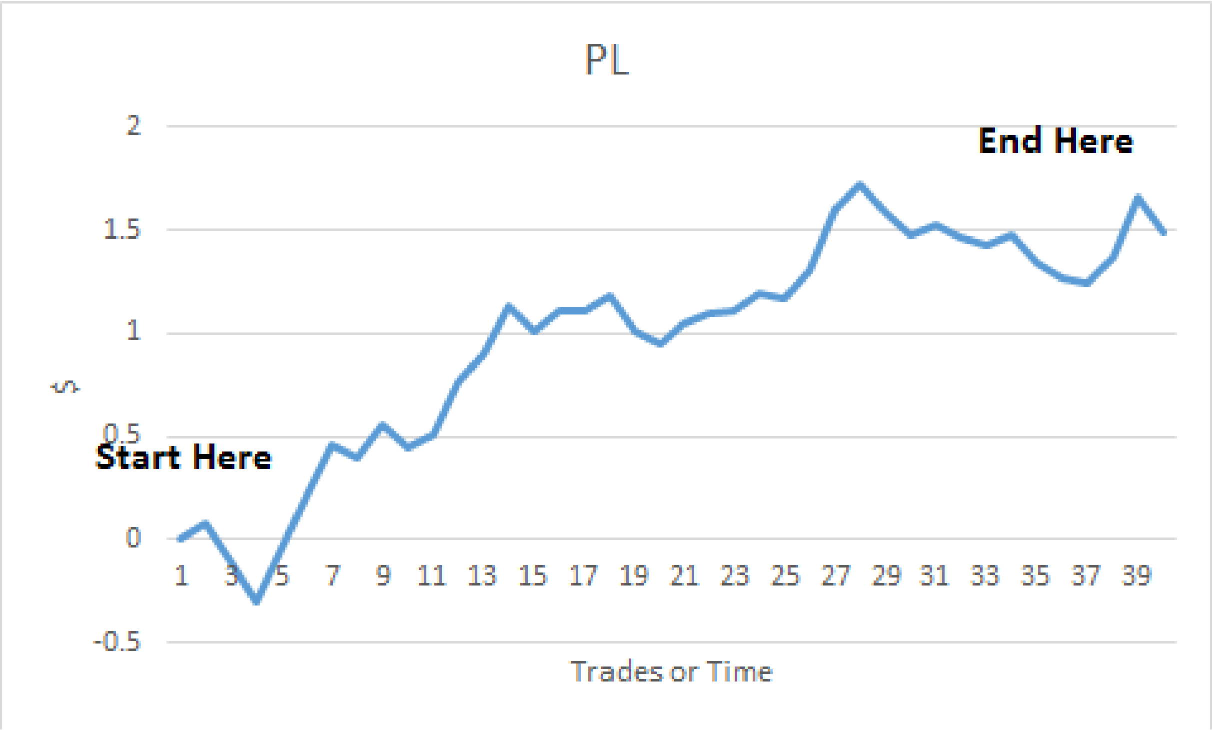

v10.1.1

Pre-Trade Planning

- You need a baseline. If you don't know what to expect, you can't tell if you have succeeded.
- If you don't know this baseline, you are just guessing.
- Back-testing.
- Academic results.
- Seeing another trader's results.
- This should give you an idea of returns and variability.

v10.1.1

Evaluating a Single Trade

- Can’t really be done.
- Luck dominates any one idea or trade.
- Need to keep records of all trades of a certain type.
- “Keep a trading journal” is most said and least followed trading advice.
- If you don’t keep records, you can’t improve. You won’t know what needs improving or if you are improving.

v10.1.1

Back-tested and Future Results

- Your real results will probably not be as good as back-tests or published results.
- "Publication decay" is a real effect. Assume a half-life of edge of a year for a daily strategy.
- Published or tested results will have data-mining biases.
- Two related issues:
- 1. Is the *back-tested* strategy good enough?
- 2. Are the strategy's *implemented* results good enough.

v10.1.1

Evaluating a Strategy

- No one thing is enough. It is better to have 10 different metrics than attempt to find one all encompassing one.
- Minimum: trade frequency, average PL, win %, worst loss in a given period, worst draw down, Sharpe.
- If sample is big enough the entire distribution might be informative.

A Strategy Comparison

Statistic	Strategy A	Strategy B
Profit/Year	$2,000,000	$600,000
Average Profit/Trade	$10,000	$500
Win %	30%	90%
W/L	4/1	1/8
Sharpe Ratio	1.2	2.0
Max Draw Down	40%	6%

It is legitimate to prefer a strategy based on a number of characteristics.

v10.1.1

Performance Metrics

- If there is only one measure of success, we can evaluate by a composite measure.
- In sports, player or team performance is all about score differential. So, calculate a player's contribution to this.
- Soccer: XG
- Golf: shots gained
- Baseball: WaR
- Trading has too many different *risks* to do this.

v10.1.1

Risk Ratios

- All risk ratios are similar. They measure success per unit of risk. They differ in what they consider success or risk.
- Numerator is something good like return or excess return.
- Denominator is something bad like volatility, downside volatility, mean deviation, draw down, averaged draw down.

v10.1.1

The Sharpe Ratio

- Sharpe is excess return divided by volatility.
- William Sharpe: Nobel prize winner in 1990.
- The Sharpe ratio has weaknesses.

1. Penalizes big wins.
2. Has large sampling errors.

v10.1.1

Sharpe Ratio Weaknesses

- Volatility penalizes large up moves as heavily as down moves.
- Ignores higher order moments. E.g positive skew should be rewarded but isn't.

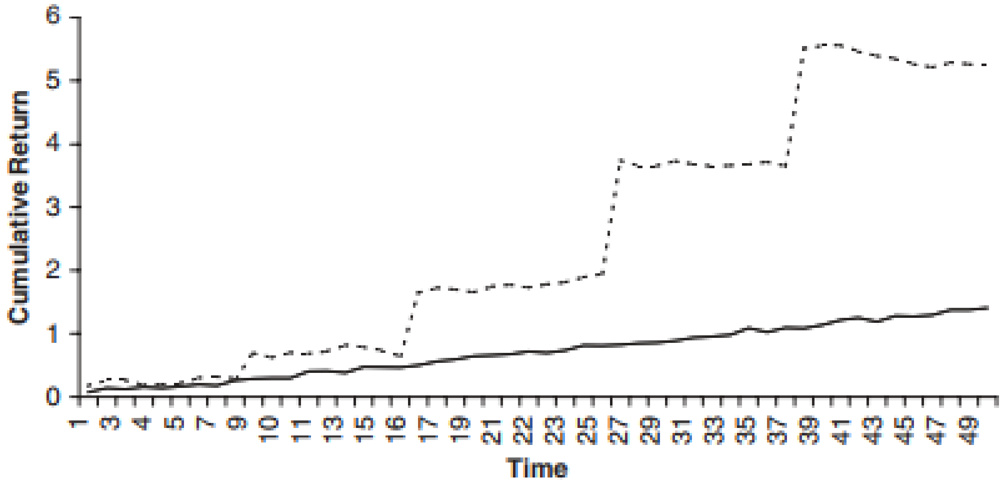

v10.1.1

Sharpe Ratio Weaknesses

- Fairly large sampling error, because it is a composite statistic and the components each have sampling errors.
- This simulated strategy has a true Sharpe of 0.5.

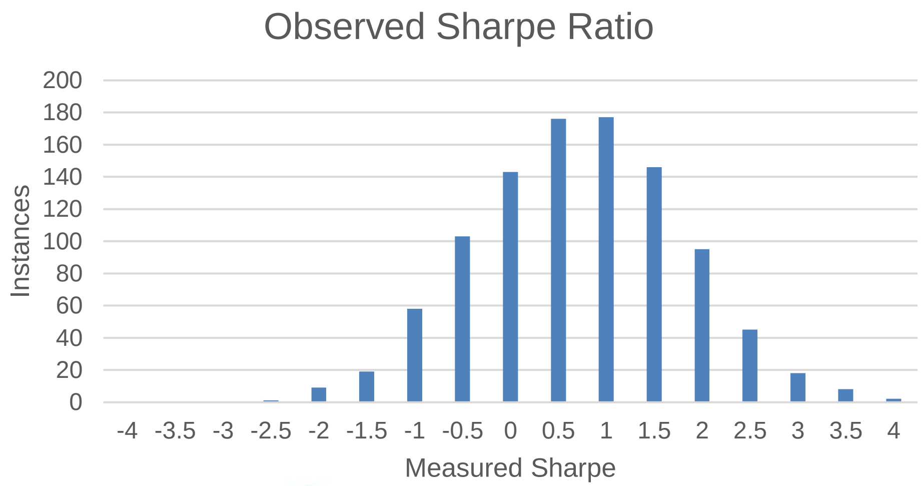

v10.1.1

Alternatives

- Sortino: Excess return divided by the standard deviation of losses.
- Calmar: Excess return divided by the maximum drawdown.
- Sterling: Excess return divided by the average of three largest drawdowns.
- Generalized Sharpe Ratio: Incorporates skewness and kurtosis into the Sharpe ratio.

Performance Persistence

- Are results changing over time?
- A good way to measure skill is repeatability of performance and separating skill from luck is the central problem in evaluating strategies and traders.
- Can just use a t-test to compare means of different periods (simple but weak).
- Kolmogorov-Smirnov test (no distributional assumptions).

Performance Persistence

- Are results changing over time?
- Trying to see if an idea has "stopped working" from data is not ideal.
- You need a lot of data, often as much as you had when it "was working".
- Vastly better to know the reason for the edge. If the reason disappears, stop doing the trade.

v10.1.1

Meta Analysis

- Work on strengths and avoid weaknesses. You don’t get paid to be an all rounder. Trading is basketball, not golf.
- “Play” with new ideas in small size. You will go off script anyway and this way might lead to something.
- When things look like they have deteriorated, look for concrete reasons.

Putting it All Together

- We are going to sell SPY options to collect the variance premium.
- We need to select an expiration.
- We need to select a structure and strikes.
- We need to choose a hedging methodology.
- We need to establish an expected profit and the expected variance around that amount.
- After the trade, we do a post-trade analysis.

Choosing an Expiration

- On Monday, June 3rd, 2019 the SPY implied volatility term-structure was essentially flat. The one-month, three-month and six-month options had the same ATM implied volatility.
- If shorter dated options had higher volatilities, I would have anticipated short term turbulence, but this wasn't the case.
- I sold the options expiring on the 28th of June. This was a compromise between high VP for shorter dated options and a longer time period to allow some luck to be averaged away.
- The ATM implied volatility was 18.8%.

v10.1.1

Choosing a Structure

- I sold 100 of the 262/286 strangles (20/19 delta) for 3.09. The put strike had an implied volatility of 23.5% and the call strike had an implied volatility of 16%.
- SPY was at 275.
- This gave me this risk profile.

Spy Change	-30%	-25%	-20%	-15%	-10%	-5%	0%	5%	10%
Vega	-$0	-$13	-$130	-$710	-$1925	-$3170	-$4220	-$3950	-$1520
P/L ($)	-$564,100	-$457,600	-$345,600	-$228,600	-$123,400	-$38,200	0	-$35,200	-$143,500

Choosing a Structure

- To cap my crash risk, I bought 100 of the 232 puts (2 delta). This changes my risk profile to:

Spy Change	-30%	-25%	-20%	-15%	-10%	-5%	0%	5%	10%
Vega	$670	$1,480	$2,150	$1,780	$15	-$2,100	-$3820	-$3,830	-$1,490
P/L ($)	-$262,200	-$252,900	-$222,000	-$167,600	-$99,000	-$31,500	0	$37,200	-$145,200

Choosing a Structure

- Most trading systems can display this visually.

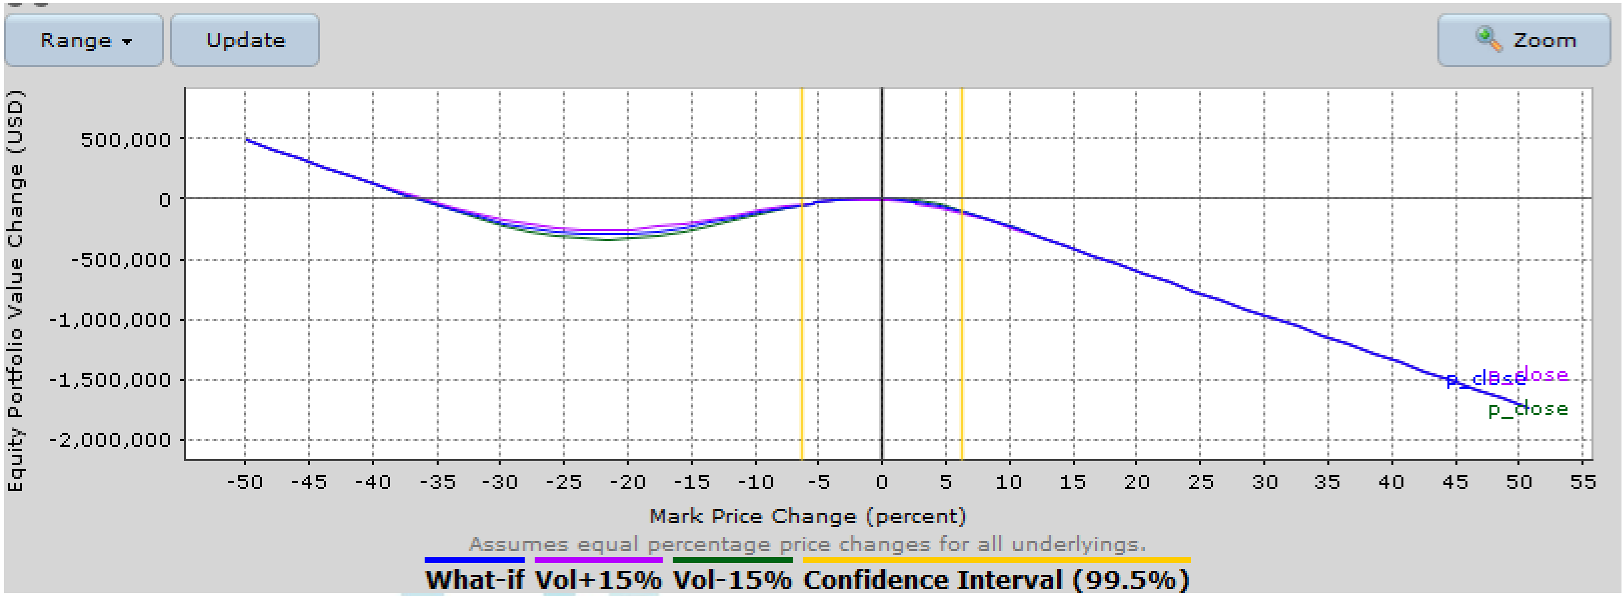

Hedging Methodology

$$\Delta = \frac{\partial V}{\partial S} \pm \left(\frac{3}{2} \frac{exp(-r(T-t))\lambda S \Gamma^2}{\kappa} \right)^{\frac{1}{3}}$$

- I use the previously estimated risk aversion parameter of 0.24 to calculate my current hedging band.
- The band will change mainly due to my gamma.

v10.1.1

Hedging Methodology

- I'm not primarily concerned about MTM variance of P/L so I'm going to calculate delta and gamma using my forecast volatility.
- At this VIX level, I expect a variance premium of about 25%. This gives a forecast of 14.1%.
- My GARCH model gave a forecast of 13.7%.
- My hedging volatility is the average of these: 13.9%.

v10.1.1

Expectations

- Theoretical profit:

$$PL = Vega(\sigma_I - \sigma_R)$$

- For short put: $2040x(23.5-13.9)
- For short call: $1875x(16-13.9)
- For long put: $330x(13.9-34)
- So total expected P/L is $16,900 (assuming our forecast volatility is correct).
- To get standard deviation of results, run a monte-carlo simulation. This gives $8,200.

v10.1.1

Results

- Actual profit: $24,218.
- We did better than our forecast because realized volatility was 10.6% (much lower than our forecast).
- The average profit with a 10.6% realized volatility would be $28,700.

v10.1.1

Could we have Done Better?

- We got lucky with realized volatility. We can't consistently expect a better result than forecast.
- But each day, we updated our volatility forecasts and never saw any reason to either exit or to trade bigger.
- SPY rallied consistently and smoothly from $275 to $293. Had we known we were short-gamma in a trending market we would have biased our volatility forecast to hedge more often i.e. to reduce our gamma.
- So we should have used a lower than expected volatility to calculate hedging numbers.
- So the low realized volatility helped our vega but hurt our hedging strategy.
- BUT none of this was known in advance.

Psychology: How do you feel?

- Probably bad.
- A winner makes you wish you had traded bigger.
- A loser makes you wish you hadn't traded.
- Regret is the dominant emotion of trading, not greed or fear.
- Deal with it or get a therapist (Seriously. This is actually an area where sensible psychology can help.)

Psychology: Accepting Imperfection

- To be a good trader:
  - Find an edge.
  - Find a strategy and business model that monetizes that edge.
  - Manage risk, so you can get to the long run.
  - But knowing how to be good won't mean you *always* are good.
  - Sometimes you will be tired or sick.
  - You need to accept that your real results will always be "worse than they should be".

v10.1.1

Conclusion

v10.1.1