Applying the Kelly Criterion to an event with existing bets, part 2

In part 1 of this series, I took a look at how to apply the Kelly Criterion correctly to spots where you already have existing bets. The point of that article was to show you how to optimize the expected growth (EG) of your bankroll given the bets you already have by reducing your risk. This is done by hedging against them based on the new probability that your hedge bet will win, and being much less likely to double down on your existing side even if the odds on your existing bets have improved. The point of this article will be to show you how to use advanced hedging techniques to generate far more EV than you can by merely finding edges on independent bets.

Before we take another deep dive into the math of the Kelly Criterion, it's important to introduce some new terminology for edges in sports betting. Recall that the simple Kelly Criterion equation for independent bets is usually written as either:

$$f = p - \frac{q}{b} \;\; or \;\; f = \frac{bp - q}{b}$$

Where:

f = fraction of bankroll to bet

b = net fractional odds of the bet (for American odds +200 → +200/100 = 2, -200 → -100/-200 = 0.5)

p = probability that the bet wins

q = probability that the bet loses, or 1 - *p*

This second form is usually referred to with the shorthand "edge over odds," because the numerator represents your mathematical edge in terms of EV. Unfortunately, this kind of shorthand can lead to the assumption than an edge is an edge is an edge. But that's not true! While the odds used in the equation are easily known (since they're listed right on the bet slip), the true probability of your side winning is actually unknown and unknowable. The best we can do is to estimate this factor, but depending on the kind of model we use our degree of certainty can vary widely.

These concepts are among the very important points made in a Pinnacle blog article called "*Toward a theory of everything*" by another practitioner of sports betting analytics who goes by "@PlusEVAnalytics." His conclusion is that when using a model to estimate the true probability of an event, any errors are likely to be with your odds and not the market odds. Therefore, he recommends staking with a fractional Kelly method in order to build confidence in the model and avoid over betting. Many successful bettors already do this intuitively, but in order to truly maximize your EG we need to examine our assumptions about edges to see when this technique is helpful and when it isn't. Once we do, you'll see that you can use full Kelly staking with much more confidence when hedging.

A Sharp Edge vs. A Square Edge

The type of edge that you calculate by modeling an event and coming up with a different true probability than the implied probability can be thought of as a "sharp edge." In other words, your advantage comes from being sharper than the sportsbook. This is the type of edge that's truly unknowable and much more likely to be overestimated, so bettors must be very careful when wagering the Kelly fraction calculated using a sharp edge. On the other hand, we can think of an edge that you generate by using the market implied probability to calculate your estimate of *p* and *q* as a "square edge." Even square bettors can find this kind of edge, because all it takes is using the "vig-free" line to come up with a win percentage and then line shopping to find the best odds offered. How should these differences between a sharp edge and a square edge affect your staking?

First off, if your model gives you an estimate of the true probability, it's easy to plug that into a Kelly Criterion calculator to determine the optimal fraction to bet. With a square edge, it takes several steps to get the same answer because we need to calculate the implied probability first. But since this probability is simply based on the "true odds," we can derive a handy shortcut to calculate your optimal stake. I won't bore you with the math this time, but the formula is:

$$f = \frac{b - b_0}{b(b_0 + 1)}$$

and a very close approximation is:

$$f = \frac{b - b_0}{b^2 + b}$$

Where:

f = fraction of bankroll to bet

b = net fractional odds of the bet

b_0 = true odds implied by the market

When you hedge your bets, you'll often use a square edge to determine most or all of your staking size. On top of this, it's almost impossible to lose capital when hedging a bet! Sure, you may win less money than if you let it ride, but from an EG perspective that's a very different animal. There's no risk of going broke or crippling your bankroll if you go on a losing streak, so there's less incentive to use fractional Kelly. By "playing with house money" and not having to reduce your stake size due to parameter uncertainty, you can bet far more on hedge bets than you can on independent bets and gain many times more EG.

Take for example two bets, both at 1:1 odds and with a 2% edge. Bet #1 is an independent bet on Manchester City (+100) to beat Liverpool, and your model says they have a 51% chance to win. Bet #2 is a hedge against your existing "Tampa Bay to win the Super Bowl" future with a potential payout of 20% of your bankroll, and while most books have the line at Tampa -105/Kansas City -115, one of your outs has KC to win the game at +100. The no-vig line on KC gives them an implied probability of 51% in this case and gives you a square edge of 2%. Of course, with such a large payout you already have a lot of EG if you stand pat with Tampa Bay, but by making an optimal hedge you can gain additional EG that's much greater than you can get with bet #1.

In the chart below, I plotted the EG of both bets after subtracting out the EG of standing pat. For bet #2, your optimal EG is about 0.62% at a bet size of about 12% of your bankroll. Notice that, since this stake is less than your existing potential payout, you can't end up with a smaller bankroll when making this bet. Plus, since your optimal stake is based on a square edge, there's no need to reduce your stake size much (if at all) from full Kelly to guard against over betting. On the other hand, bet #1's optimal EG is a mere 0.02% when using a full Kelly stake of 2% (see the math below if you dare). Since this bet uses a sharp edge, most bettors would wisely wager only half or less of full Kelly and reduce their EG even more. Even though bet #2 has the same edge and same odds, it lets you bet much more money with lower variance than Bet #1, so it's almost 50 times more valuable.

$$EG = e^{p*\ln(1+fb)+(1-p)*\ln(1-f)} - 1$$

$$EG = e^{.51*\ln(1.02)+.49*\ln(0.98)} - 1$$

$$EG = e^{.0002} - 1$$

$$EG = .0002 \text{ } or \text{ } 0.02\%$$

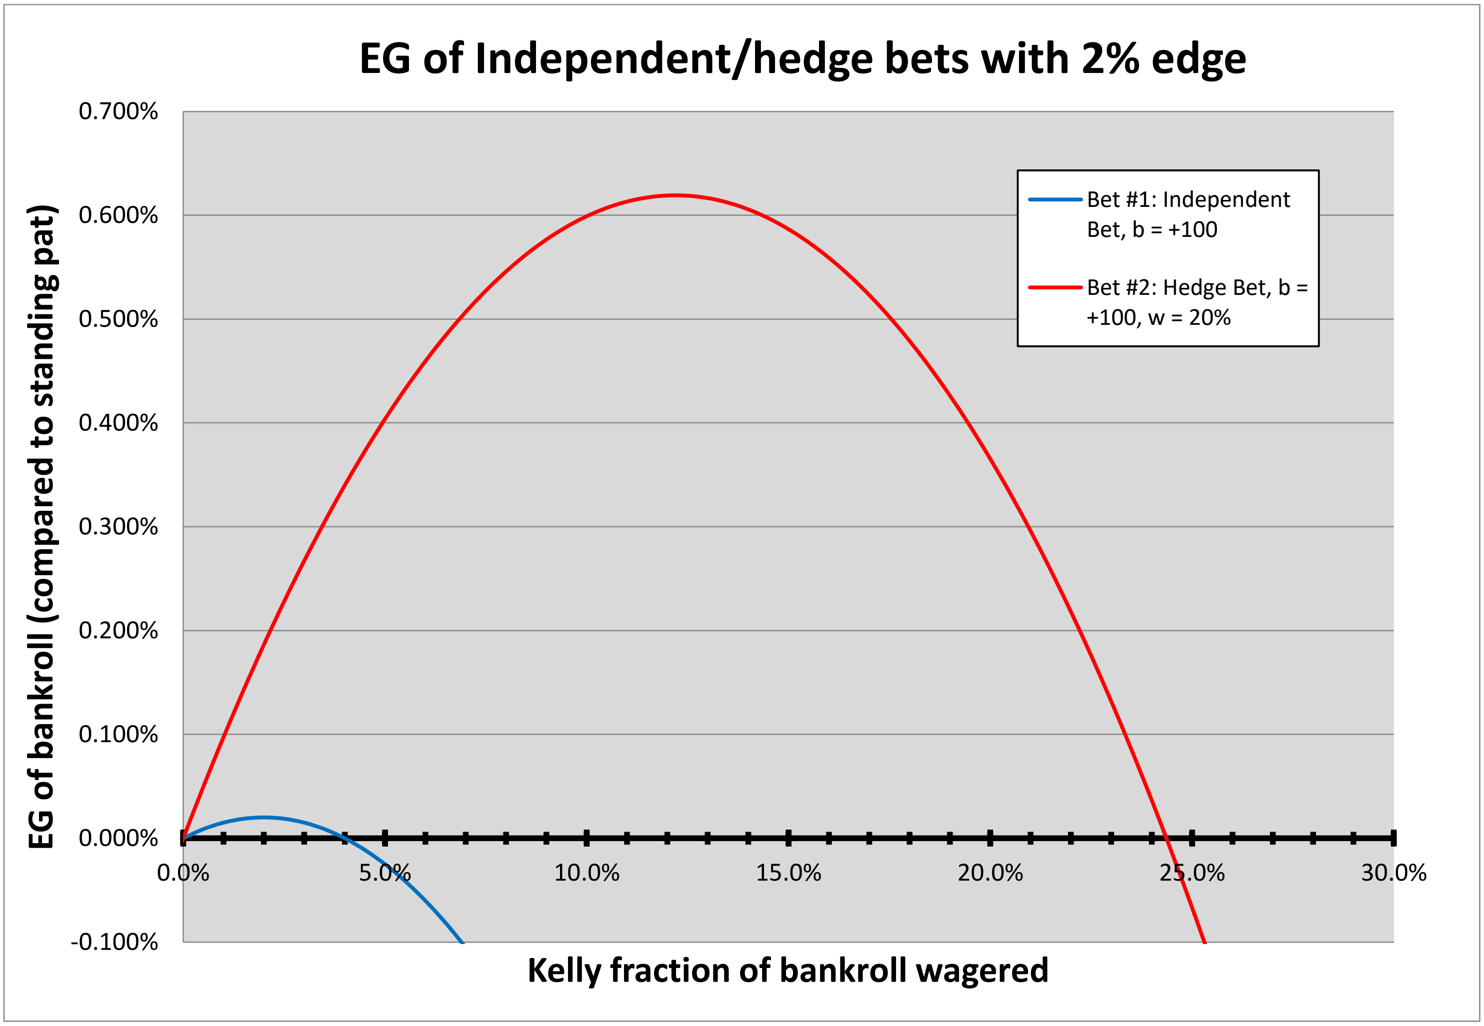

The Neutral Hedge Gambit

Now suppose you're looking at a futures bet on Kansas City to win the Super Bowl at +200. Your best estimate is that they have a 60% chance to win the AFC championship, and then will average a 60% chance to win the Super Bowl if they get there. Since the futures bet is equivalent to a parlay of those two events which have a 36% chance to both occur, you have an 8% edge (do you see why?). Definitely a good spot, but how much should you bet on it?

This exact example was brought up on the Two Plus Two forums in a discussion about part 1 of this series. Poster "PokerHero77" suggested that bettors can do better by increasing their initial bet, assuming they can find a hedge at neutral EV before the bet is graded. Is he right? In order to find out we have to compare your optimal EG when staking it as an independent bet with the EG you gain by betting big and taking the chance that you can efficiently hedge out later.

If we treat this example as an independent bet, the optimal stake is your edge over your odds or 4% of your bankroll. If the Chiefs make the Super Bowl, you can then optimize your EG by hedging your bet before the game if you can find good enough odds to do so. You can use the formula I showed you in Part 1 to determine your optimal hedge by using p = 40% (the probability of Kansas City's opponent winning) and w = 12%, since that would be your payout if you let it ride and KC wins. If you're fortunate enough to be able to find fair odds on the game (i.e., NFC team at +150), and thus make a neutral EV hedge bet, your optimal hedge is:

$$f^{\circ} = p - \frac{q}{b} + pw$$

$$f^{\circ} = 0.4 - \frac{0.6}{1.5} + (0.4) * (0.12)$$

$$f^{\circ} = 0.048 \; or \; 4.8\%$$

Note that in this special case, the independent Kelly fraction is 0% so your optimal bet exactly balances your payout. This is equivalent to cashing out your bet for the current EV, since your ending bankroll will be the same regardless of which team wins the game. However, can you do any better by trying a gambit to increase your initial bet and then make a neutral EV hedge to greatly reducing your risk? Let's see how the math of the Kelly Criterion works out if we're able to pull off that manoeuver.

As a reminder, the equation for calculating the expected value of the growth of your bankroll when making an independent bet is:

$$E = p * \log(1 + fb) + (1 - p) * \log(1 - f)$$

If we factor in the potential neutral EV cash out of your futures/parlay bet, the equation becomes:

$$E = p_1 * \log(1 - f + p_2 f(b + 1)) + (1 - p_1) * \log(1 - f)$$

Where:

p_1 = probability that the first leg wins

p_2 = probability that the second leg wins (so $p = p_1 * p_2$)

The argument of the first logarithm accounts for your bankroll if the first leg wins, i.e., your current bankroll minus the fraction you will bet plus the proceeds of cashing out your ticket. A neutral EV cash out is $p_2 * f * (b+1)$, since it must equal the probability that KC wins the game (i.e., the second leg wins) times your potential payout, which is the fraction you will bet times (odds+1). The argument of the second log is exactly the same as for independent bets, because you lose your entire bet if the first leg doesn't hit. Again, in order to maximize your EG we differentiate E with respect to f and set the derivative equal to 0. The derivation follows:

$$\frac{dE}{df}|f^{\circ} = \frac{p_1 p_2(b+1) - p_1}{1 - f^{\circ} + p_2 f^{\circ}(b+1)} - \frac{(1-p_1)}{1-f^{\circ}} = 0$$

$$p_1 p_2(b + 1)(1 - f^{\circ}) - p_1(1 - f^{\circ}) - (1 - p_1)(1 - f^{\circ} + p_2 f^{\circ}(b + 1)) = 0$$

$$p(b + 1) - p(b + 1)f^{\circ} - p_1 + p_1 f^{\circ} - 1 + f^{\circ} - p_2 f^{\circ}(b + 1) + p_1 - p_1 f^{\circ} + pf^{\circ}(b + 1) = 0$$

$$bp + p - bpf^{\circ} - pf^{\circ} - p_1 + p_1 f^{\circ} + f^{\circ} - bp_2 f^{\circ} - p_2 f^{\circ} + p_1 - p_1 f^{\circ} + bpf^{\circ} + pf^{\circ} - 1 = 0$$

$$bp + p + f^{\circ} - bp_2 f^{\circ} - p_2 f^{\circ} - 1 = 0$$

$$f^{\circ} - bp_2 f^{\circ} - p_2 f^{\circ} = -bp - p + 1$$

$$f^{\circ}(1 - bp_2 - p_2) = -bp - p + 1$$

$$f^{\circ} = \frac{-bp - p + 1}{1 - bp_2 - p_2}$$

$$f^{\circ} = \frac{bp + p - 1}{bp_2 + p_2 - 1}$$

$$f^{\circ} = \frac{bp - (1 - p)}{bp_2 - (1 - p_2)}$$

$$f^{\circ} = \frac{bp - q}{bp_2 - q_2}$$

Notice that this answer is in the form of your edge on the initial bet divided by your virtual edge if the parlay bet still had the same odds, but the probability of winning was that of the second leg only. Alternately, we can divide each edge by *b* to simplify it even further like this:

$$f^{\circ} = \frac{(bp - q)/b}{(bp_2 - q_2)/b}$$

$$f^{\circ} = \frac{f^{\circ}{}_i}{f^{\circ}{}_p}$$

Where:

f^*_i = the Kelly fraction for an independent bet

f^*_p = the "parlay Kelly fraction", i.e., the optimal fraction to bet if the parlay had the same probability of winning as the second leg only

Expressing the Kelly fraction this way makes it easier to calculate by simply running a Kelly calculator twice, and it also illustrates how this gambit concentrates your edge into the first leg by removing the risk of the second leg. If the second leg is a sure thing, then there's no risk associated with it and the parlay Kelly fraction would equal 1 (meaning don't adjust your Kelly fraction at all, because you can't gain anything by hedging). The riskier the second leg is, the smaller the parlay Kelly fraction will be and the more you gain by removing that risk from the equation.

Examining our example again, the parlay Kelly fraction can be calculated using the expected win equity for KC in the Super Bowl (i.e., p_2 = 60%) and the odds of the whole bet (b = 2) to get:

$$f^{\circ}{}_p = p_2 - \frac{q_2}{b}$$

$$f^{\circ}{}_p = 0.6 - \frac{0.4}{2} = 0.4$$

Given that result, the optimal fraction to bet initially is:

$$f^{\circ} = \frac{f^{\circ}{}_i}{f^{\circ}{}_p} = \frac{0.04}{0.4} = 0.1 \; or \; 10\%$$

And the optimal hedge bet against the Chiefs is:

$$f^{\circ} = p - \frac{q}{b} + pw = 0.4 - \frac{0.6}{1.5} + 0.4 * 0.3 = 0.12 \; or \; 12\%$$

So, it works out that dividing by the parlay fraction boosts your effective edge from 8% to 20% in order to make a 10% stake optimal, because losing money only 40% of the time is far better than losing 64% of the time, even if your maximum win is lower. If instead you took another lower risk approach of betting Kansas City to win the AFC, and found a 5% edge with odds of -137, the full Kelly stake in that spot would only be about 7% (do you see why again?) Even though you lose 40% of the time either way, the neutral hedge technique still yields much more EV (as well as EG) because it combines your edges on both legs into just the first one.

But what if you're wrong? Your estimate of the Chiefs' win equity in the Super Bowl could be off, and so could your expected odds on the hedge bet. How do errors in these variables affect the optimal bet size of the initial bet, and how do they affect your EG if you choose to bet 10% up front anyway?

A Slippery Slope

In order to calculate the effect of different hedge odds and second-leg win equities on your overall EG, I enlisted the help of excel and put together a table and a couple of charts. The table shows the parameters for the odds of both the initial and the hedge bets, KC's equity to win the Super Bowl game, and the EG associated with the optimal futures bet and hedge bet sizing. When the hedge odds vary from +155 to +135 (in the middle rows), your optimal bet size decreases from 11.5% to 4.5%. At fair odds of +150, the initial bet size is 10% as predicted by my formula, and with +145 odds like you might get from a sharp book, the optimal futures bet size is still 8.5% (which corresponds to an EG of 0.286% instead of 0.404% like we get with fair odds). If you bet 10% on the initial bet, but then can only get +145 odds when it's time to hedge, you pay a slight penalty in EG (0.276% vs. 0.286%) as shown in the 5th column. The scenario gets progressively worse for your bankroll as your hedge odds get worse, but even if you bet 10% initially and then hedge at +140, you do about the same as betting 4% on the initial bet and not hedging at all.

KC win Super Bowl @ +200							
Futures odds =	200						
Hedge odds for SB =	150						
KC win equity in SB =	60.0%	_			Normal bet		
futures bet =	10.0%			full Kelly =	4.0%		
hedge bet =	12.0%			hedge bet =	0.0%		
		EG				EG normal	
KC no SB	40.00%	-0.04214				0.1581%	
KC win title	36.00%	0.02771					
KC lose SB	24.00%	0.01847					
Sum	100.0%	0.00403					
		0.4041%					
Hedge odds	Future Amt	Hedge Amt	EG @ opt bet	EG @ 10% bet			
155	11.5%	15.1%	0.552%	0.543%			
150	10.0%	12.0%	0.404%	0.404%	**Fair odds**		
145	8.5%	8.8%	0.286%	0.276%	**Sharp line**		
140	6.5%	4.9%	0.202%	0.159%	**Square line**		
135	4.5%	1.0%	0.161%	0.056%			
No hedge	4.0%	0.0%	0.158%	0.158%			
KC win equity in SB	Future Amt	Hedge Amt	EG @ 10% bet	EG @ 4% bet			
55.0%	10.0%	12.0%	-0.435%	-0.093%			
57.5%	10.0%	12.0%	-0.015%	0.082%			
60.0%	10.0%	12.0%	0.404%	0.257%	**Model p**		
62.5%	10.0%	12.0%	0.822%	0.432%			
65.0%	10.0%	12.0%	1.238%	0.607%			
67.5%	10.0%	12.0%	1.654%	0.781%			

What about if you can hedge at fair odds but your Super Bowl win equity changes? This effect is shown in the bottom rows of the table, both for 10% and 4% initial bets. Notice that as the win equity varies, your EG varies widely and makes it seem like a terrible bet in some cases but a gold mine in others. As PokerHero77 points out, though, if the distribution of win equity averages out to 60%, then the optimal Kelly fraction will be the same as when the win equity is exactly 60%. So if you're wrong about the win equity only because its exact value is unknowable, then there's no effect on your average EG. Now, if you only bet 4% initially, then your average EG is lower but with less variance. Your EG is also negative for win equities below 57% if you don't hedge at all, and the upside when the win equity is above 60% is far less than with the neutral hedge gambit.

You can visualize how the different permutations of hedge odds and Super Bowl win equity work out, versus staking with the simple Kelly formula from the start (or betting simple Kelly initially and then hedging before the Super Bowl), by looking at the charts below.

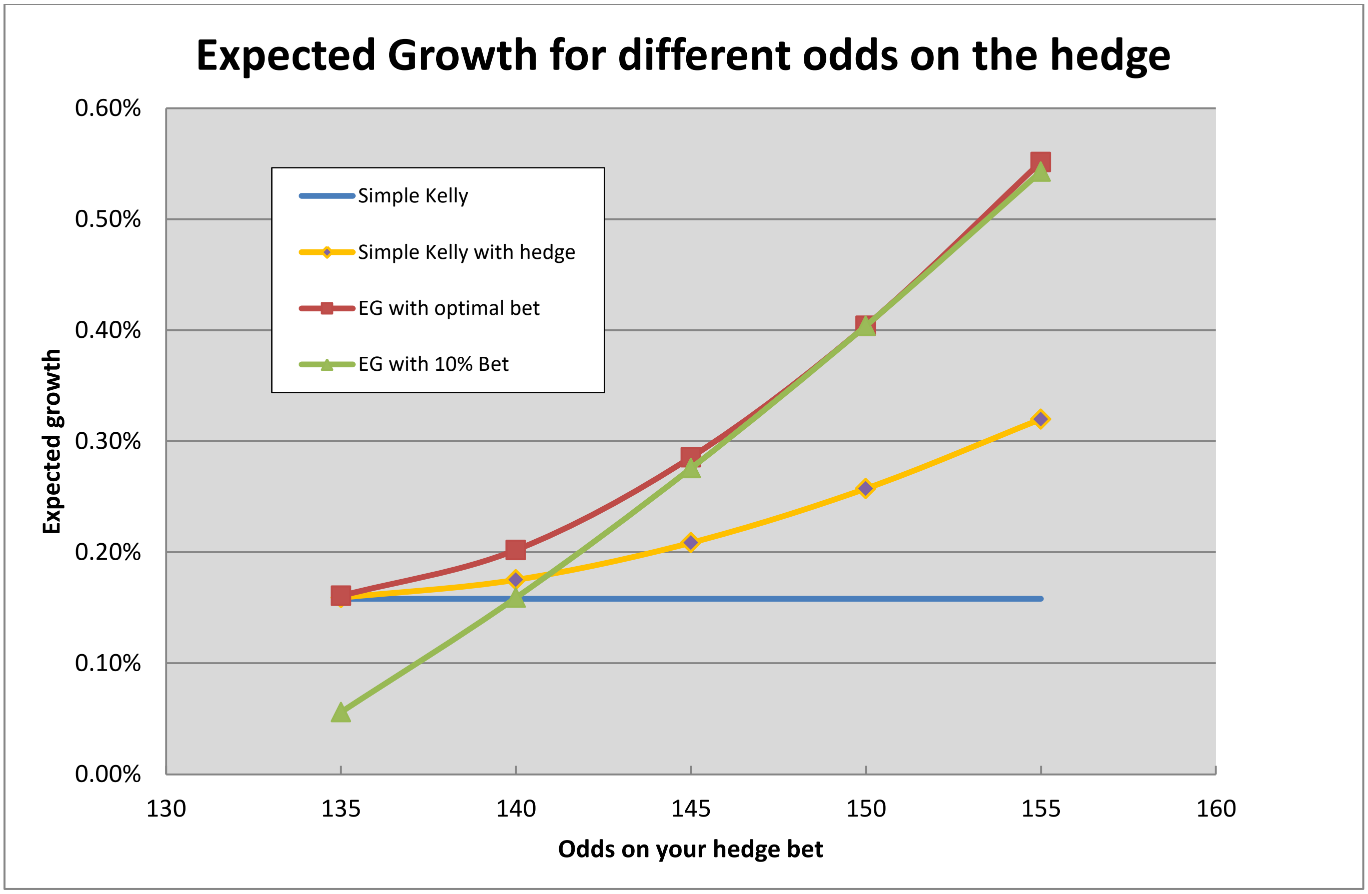
Expected Growth for different odds on the hedge
0.60%
0.50%
0.40%
0.30%
0.20%
0.10%
0.00%
Expected growth
Simple Kelly
Simple Kelly with hedge
EG with optimal bet
EG with 10% Bet
130
135
140
145
150
155
160
Odds on your hedge bet

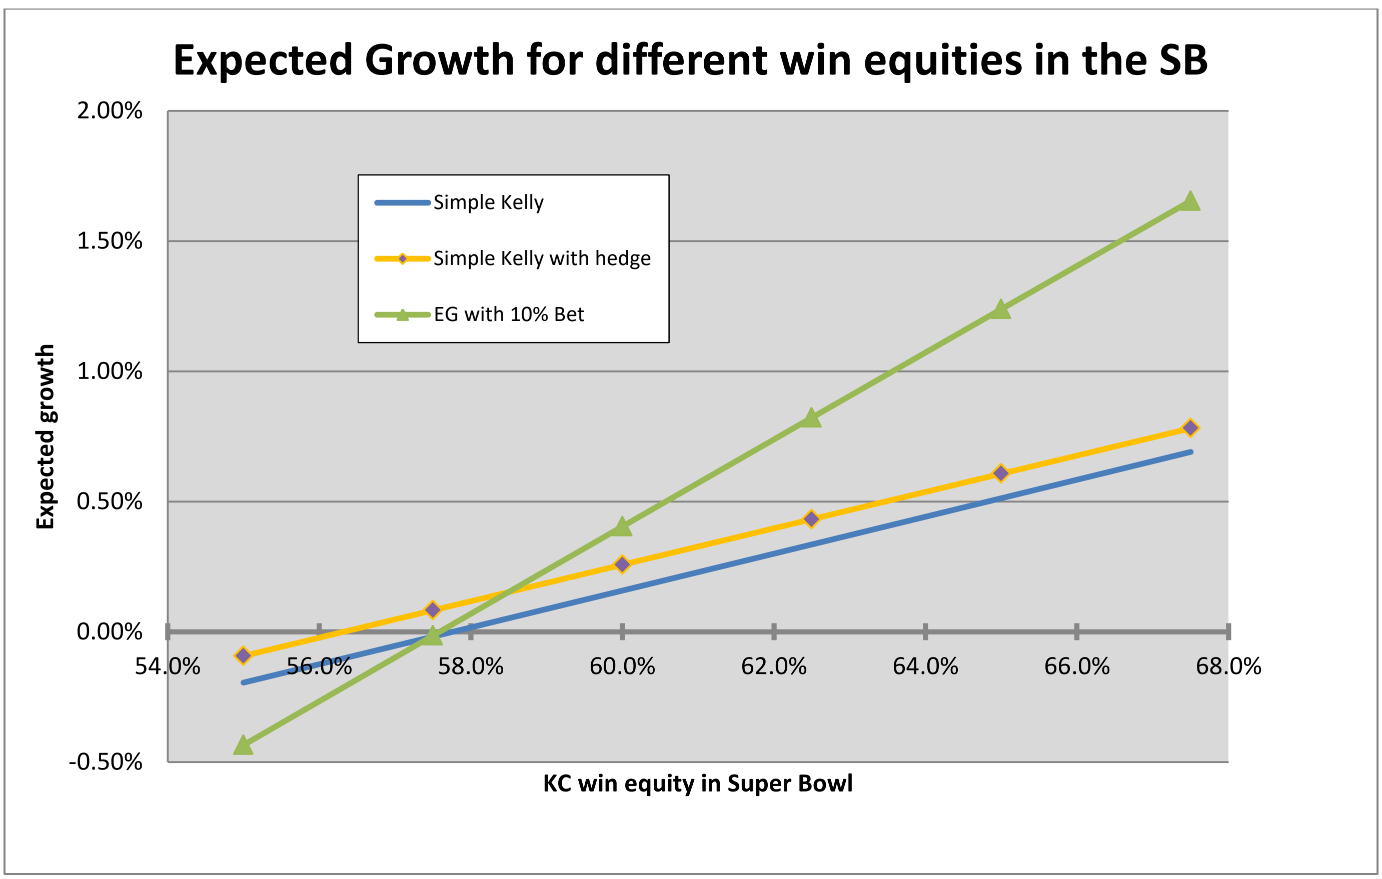
Expected Growth for different win equities in the SB
2.00%
1.50%
1.00%
0.50%
0.00%
-0.50%
Expected growth
Simple Kelly
Simple Kelly with hedge
EG with 10% Bet
54.0%
56.0%
58.0%
60.0%
62.0%
64.0%
66.0%
68.0%
KC win equity in Super Bowl

A Final Warning

While this gambit can allow you to bet far more than you would otherwise by boosting your effective edge, there is still a much bigger risk associated with bigger stake sizes. If you calculate your edge on the futures bet based on a sharp model, then consider it a sharp edge and reduce your staking size accordingly. Of course, the same warning applies to betting on that future using the simple Kelly formula, so your stake in that case should be cut too. Any way you slice it, planning for a hedge from the outset can almost double or triple your EG if you can find a way to get something you probably have passed up many, many times -- a neutral EV bet.