Applying the Kelly Criterion to an event with existing bets

The Kelly Criterion is a well-known concept that instructs bettors how to optimally size their bet for a given event based on the odds offered and the probability that the bet will win. It provides a simple formula to find the fraction of a bettor's bankroll to wager in order to maximize the expected value (EV) of the logarithm of wealth, or in other words, to optimize the expected rate of growth of your bankroll. Your rate of growth is an exponential function that answers questions like "how can I double my bankroll" instead of questions like "how can I win $100?" By using this criterion correctly, bettors can grow their bankroll quickest without risking going broke. This should be the real goal of sharp bettors, but it's different than maximizing your EV in dollars.

A common form of the classic Kelly Criterion equation for independent bets is:

$$f = p - \frac{q}{b}$$

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 Kelly Criterion formula was derived for the case where bets are made sequentially (and therefore are independent) and the bettor's bankroll is adjusted for the outcome of the previous bet before making the next one. These are natural assumptions to make when flipping a coin over and over or betting on one sporting event at a time. But how do you apply it to the case where you have an existing bet and you're considering placing a new bet on either the same or the opposite side? The risk to your bankroll will depend not only on the result of the new bet but also on the potential payout of the existing bet. Here we'll see how to apply the Kelly Criterion to these situations.

Betting the opposite side

Suppose you have an existing bet and there's an option to place a new bet on the opposite side. For example, you're offered a promotion where you get a free bet if you bet a 3-team parlay card. You decide to play, and you've hit two winners in college games so far with only Tampa Bay left to cover -8.5 on Sunday against Detroit. You initially bet 1 unit (or 1%) of your bankroll getting 6-1 odds, so the potential payout is 7% of your roll. Should you bet on the Lions plus the points in case Tom Brady has a bad game? And if so, how much should you bet? Enough to "lock in" the same profit whether the Buccaneers cover or not, or maybe more or less?

This kind of bet is commonly known as a "hedge" and is generally frowned upon, especially when the EV of the new bet is negative. The conventional wisdom is that you should never take a -EV bet, but could reducing your variance by hedging help grow your bankroll the quickest even if you do? What does the math of the Kelly Criterion say about the issue? Since the outcomes of the existing bet and a hedge bet aren't independent (in fact, if one wins the other must lose), the simple Kelly formula for independent events doesn't apply. We could just simply crunch the numbers, but how can we derive a useful Kelly formula for this case?

The equation for calculating the expected growth of your bankroll when making an independent bet is:

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

Where:

E = expected growth (EG) of your bankroll

If we factor in the potential payout from an existing bet, the equation for calculating the EG of your opposite-side bet becomes:

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

Where:

w = payout returned if existing bet wins (as a fraction of your bankroll)

When your new bet wins (with probability p), your new bankroll is your current bankroll plus the payout from your bet (i.e. 1 + fb) in either equation. However, the opposite-side equation is subtly different because when your new bet loses (with probability 1 - p) your new bankroll is your current bankroll minus your bet plus the payout from your existing bet. Notice that "p" in this equation represents the probability that the new bet will win. Your EG does not depend on how likely your existing bet was to win when you first made it, and the chances that the existing bet will still win is the same as the probability that the new bet will lose.

The goal of the Kelly Criterion is to find the optimal value of f to maximize E. To calculate that value, we need to differentiate E with respect to f and set the derivative equal to 0. This is because when a function first increases and then decreases as its variable keeps increasing, its maximum value is when the slope of the tangent line is zero. The result of this calculation gives us the optimal fraction f*, i.e. the value for which you can do no better in the long run by either betting more or less, as illustrated here:

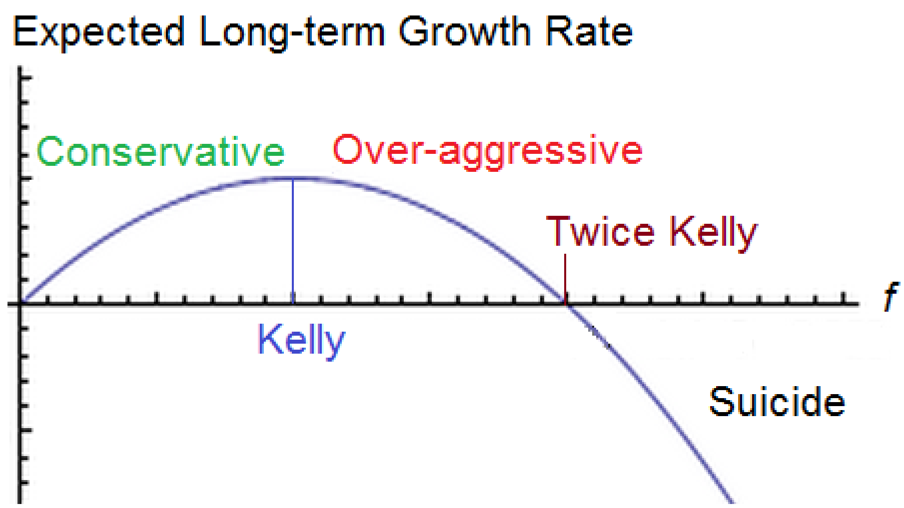

The derivation follows:

$$\frac{dE}{df}|f^{\circ} = \frac{pb}{1+f^{\circ}b} - \frac{(1-p)}{1-f^{\circ}+w} = 0$$

$$pb(1 - f^{\circ} + w) - (1 - p)(1 + f^{\circ}b) = 0$$

$$pb - pbf^{\circ} + pbw - 1 - bf^{\circ} + p + pbf^{\circ} = 0$$

$$-bf^{\circ} + pb + p + pbw - 1 = 0$$

$$bf^{\circ} = pb + p + pbw - 1$$

$$f^{\circ} = p + \frac{p-1}{b} + pw$$

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

Notice that this answer is similar to the independent solution to the Kelly Criterion shown above, but with the added factor pw. Let's define a hedge factor "h", such that h = pw:

h = probability that the new bet wins * potential payout of the existing bet

Intuitively, we'd think that it's more important to hedge your bets when their EV is already large. But we would be wrong! The hedge factor represents the difference between the payout of the existing bet and the current EV of it because if the new bet wins, the existing bet must lose. Existing bets that have a small chance of winning, but a very large upside, will yield a much larger hedge factor and have a much larger effect on your optimal stake. Therefore, we can state a corollary to the Kelly Criterion for a hedge bet as:

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

How does this result affect your optimal bet sizing? There are 3 possible cases to consider: when a hedge bet is +EV by itself, when the hedge is EV neutral, and when it's -EV by itself.

+EV hedge bets are worth their weight in gold

Clearly these bets are worth betting whether or not you have an existing bet, but how much is optimal to wager? Based on the opposite-side Kelly formula, it's the usual Kelly fraction plus the hedge factor. For example, if the stake recommended by the independent Kelly formula is 2% of your bankroll but the hedge factor is 4% of your bankroll, then your optimal bet size is 6%. This is a pretty common situation that can come up frequently, like with your 3-team parlay. Say by game time you think the Lions have a 55% chance to cover 8.5 because a couple injured players were upgraded. The best line you can find is -118 (or 0.85:1) on Lions +8.5, so you calculate that you should bet 2 units on the Lions since:

$$f = p - \frac{q}{b}$$

$$f = 0.55 - \frac{0.45}{0.85}$$

$$f = 0.55 - 0.53$$

$$f = 0.02 \text{ or } 2\%$$

But because betting on the Lions would be a hedge against your parlay card, your optimal bet size is almost 6% given a hedge factor of about 4%:

$$h = pw = 0.55 * 0.07 = 0.0385 \; or \; 3.85\%$$

The conclusion then is that, because you still cash for 7 units when your new bet loses, you can be much more aggressive when you hedge. With a bet size of almost 6% you can gain almost 3 times the EV as you safely could with an independent bet, so the standard mantra for Kelly sizing of "edge over odds" becomes "edge over odds, plus hedge."

Neutral hedge bets aren't neutral to your bankroll

For the case where you can hedge your existing bet without paying a vig, this play won't increase your EV but it does increase your EG. In our parlay example, your hedge would be EV neutral if the line on Lions +8.5 was -122, making the independent Kelly fraction = 0. In this spot, the optimal play is to lock in a profit simply by making your bet equal to the hedge factor (i.e. 55% of your potential payout).

-EV hedge bets can still be worth a look

In the case where the odds aren't high enough to make a +EV bet, hedging does in fact give away EV but it can still increase your EG – which is the whole point of the Kelly Criterion: to trade off some EV in order to reduce your risk of ruin. Since the optimal stake is the independent Kelly fraction (p – q/b) plus the hedge factor, a sufficiently large hedge factor can turn a -EV hedge into the correct play. It doesn't even need to be life-changing money on the line. For example, say that since the Lions have a better chance to cover +8.5, the line on that spread has moved to -125. How can you see if it's worth hedging? The independent Kelly bet size would be -1.3% of your bankroll (which would normally mean don't bet anything):

$$f = p - \frac{q}{b}$$

$$f = 0.55 - \frac{0.45}{0.80}$$

$$f = 0.55 - 0.563$$

$$f = -0.013 \; or \; -1.3\%$$

In your spot, however, the hedge factor is 3.85% of your bankroll so your optimal bet size is 3.85% - 1.3% = +2.55%. The conclusion then is that because the downside to your bankroll is limited by the payout of your existing bet, it effectively boosts the odds on your hedge to the point where it's better to bet than to stand pat.

Betting the same side

Now suppose you have an existing bet and there's an option to place a new bet on the same side after the odds and/or the probability of the outcome has changed. What does the Kelly Criterion have to say about the issue of increasing your current position or "doubling down?" Since the outcomes of the existing bet and the same-side bet are not independent, the simple Kelly formula for independent events still doesn't apply. So how can we derive a Kelly formula for the case of doubling down on your bet?

You still want to maximize your EG, but in this case we add the potential payout from the existing bet to your bankroll when the new bet hits. Again, we then differentiate E with respect to f and set the derivative equal to 0 to give us f*. The derivation follows:

$$\frac{dE}{df}|f^* = \frac{pb}{1+f^*b+w} - \frac{(1-p)}{1-f^*} = 0$$

$$pb(1 - f^{\circ}) - (1 - p)(1 + f^{\circ}b + w) = 0$$

$$pb - pbf^{\circ} - 1 - bf^{\circ} - w + p + pbf^{\circ} + pw = 0$$

$$-bf^{\circ} + pb + p + pw - w - 1 = 0$$

$$bf^{\circ} = pb + p + pw - w - 1$$

$$bf^{\circ} = pb + p - 1 + pw - w$$

$$bf^{\circ} = pb + (p - 1) + (p - 1)w$$

$$f^{\circ} = p + \frac{p - 1}{b} + \frac{(p - 1)w}{b}$$

$$f^{\circ} = p - (1 + w)\frac{q}{b}$$

Notice that this answer is similar to the independent Kelly formula, but with the added factor of (1 + w) multiplying q. How does this result affect your optimal bet sizing? Again, there are 3 possible cases to consider: when the double down bet is +EV by itself, EV neutral, and -EV.

Neutral and -EV double down bets still aren't a play

Because w must be positive (unless you're expecting a negative payout!), the Kelly fraction for double down bets will always be smaller than for independent bets. So, the fraction calculated for neutral and -EV bets will be negative and therefore mean "no bet."

+EV double down bets may be too risky

If you had no existing bet, all +EV bets would be worth betting with some fraction of your bankroll, but how much is optimal to wager on a double down? Based on the same-side Kelly formula, the independent Kelly fraction is reduced because "q/b" is multiplied by the factor (1 + w) before subtracting. For example, if you're offered a new bet with 11:1 odds and your existing bet now has a 10% chance of winning, then the independent Kelly fraction is about 2% as seen below:

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

$$f^{\circ} = 0.1 - \frac{0.9}{11}$$

$$f^{\circ} = 0.0182 \; or \; 1.82\%$$

However, if the potential payout of your existing bet is 10% of your bankroll then the optimal double down bet size is only 1%. Here it's still worth a play but for a smaller stake since your risk is all on one side of the game:

$$f^{\circ} = p - (1 + w)\frac{q}{b}$$

$$f^{\circ} = 0.1 - \frac{1.1 * 0.9}{11}$$

$$f^{\circ} = 0.01 \; or \; 1\%$$

Similarly, if you're offered a new bet with 1:1 odds and your side now has a 51% chance of winning then the independent Kelly bet size is also about 2% of your bankroll:

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

$$f^{*} = 0.51 - \frac{0.49}{1}$$

$$f^{*} = 0.02 \; or \; 2\%$$

However, for w = 10% of your bankroll then the optimal bet size is about -3%. As with the results from the independent Kelly formula, a negative percentage means no bet.

$$f^{*} = p - (1 + w)\frac{q}{b}$$

$$f^{*} = 0.51 - \frac{1.1 * 0.49}{1}$$

$$f^{*} = -0.029 \; or \; -2.9\%$$

In our example, say that after your parlay hit, you bet 1 unit on Tampa Bay to win the Super Bowl at 19-1. Even though you have a tough NFL season with your other bets, and you end up exactly break-even for the year, the Buccaneers make it to the big game and a win would pay out 20% of your roll. You figure the Bucs have a 53% chance to win and the best line you can find on them is +100 (or 1:1 odds). While this bet would definitely give you the best of it in terms of EV, it would also increase your current position. So, to decide whether or not a new bet would grow your bankroll, we can determine your optimal bet size like this:

$$f^{*} = p - (1 + w)\frac{q}{b}$$

$$f^{*} = 0.53 - \frac{1.2 * 0.47}{1}$$

$$f^{*} = -0.034 \; or \; -3.4\%$$

With a Kelly fraction of -3.4%, any bet you make on the Bucs would be a horrible play. In fact, they would need to have almost a 55% chance to win to make a double down correct. The conclusion then is that, because the upside to your bankroll isn't increased that much on a percentage basis when you increase your existing position, you should risk less than you would without an existing bet. Often, it isn't worth increasing your bet at all and you should be looking to hedge. I'll leave it as an exercise for the reader to figure out how large an optimal hedge against the Bucs would be with odds of -105 (or 0.95:1).

Seeing the big picture

It's hard to visualize just how different your betting strategy should be when making plays that are tied to existing bets from a couple examples and a bunch of formulas. To help readers see the bigger picture, I plotted the percent of your bankroll to bet versus your existing bet's payout to show how your betting strategy should change as your potential payout grows. If you can get fair odds of 1:1 on either side, and your hedge against the Bucs has a 47% chance to win (meaning a double down has a 53% chance to win), then you can follow the red and green lines in the graph below. If you have only a small existing bet (payout around 1-2% of your roll), you should double down with about 5% of your bankroll to snap up some +EV but never hedge. With a potential payout of around 12% your strategy should flip from double down to hedge, and with a payout of 40% you maximize your EG by hedging more than 10 units!

For a more realistic case of paying about a 5% vig at your chosen sportsbook (like with -110 odds on both sides of the spread), your staking should be guided by the yellow and blue lines. Between a 2% - 24% payout, no bet helps your EG more than standing pat. But with a potential payout in your portfolio or 30% or more, hedging your position is like Miracle Gro for your privet bushes.

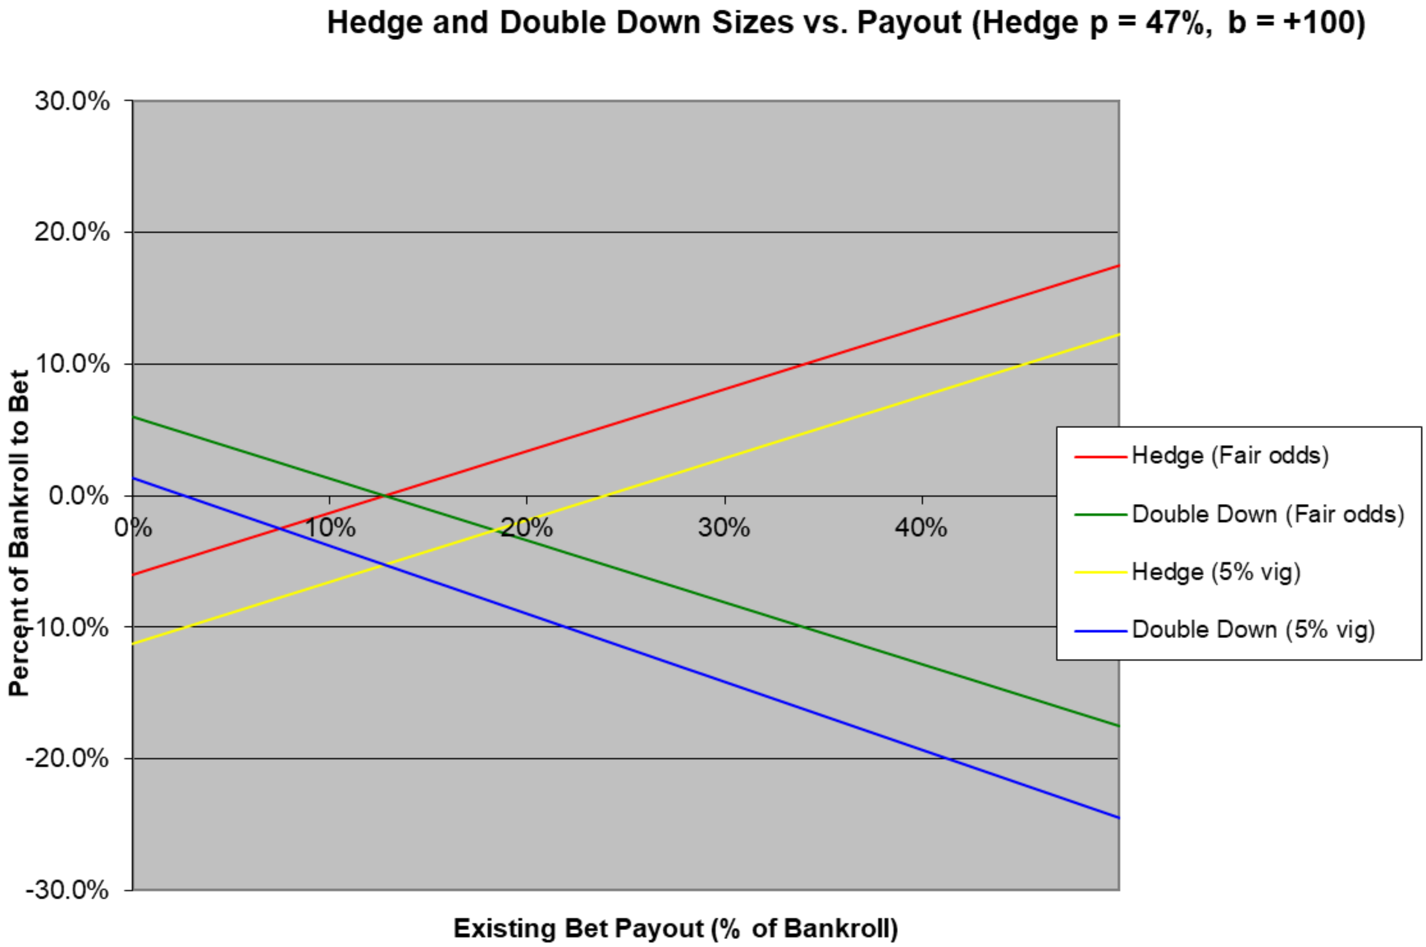

I also plotted your Kelly fraction versus the EV of the new bet, given a potential payout from your existing bet of 0%, 10%, and 20%. How should your betting strategy depend on the edge of the new bet? With a hedge that's 47% to win, then you can follow the red and yellow curves in the second graph. If you're lucky enough to have a spot with +EV then you should obviously be betting, but for w = 10% you can bet even with a standard -3% edge and for w = 20% you can optimize your bankroll growth by betting almost 10% of your roll even with no edge at all. For double down opportunities, even with a 53% chance to win, it's rare to find a bettable spot. With w = 20%, your double down bet would have to show an edge of over 10% to even be considered.

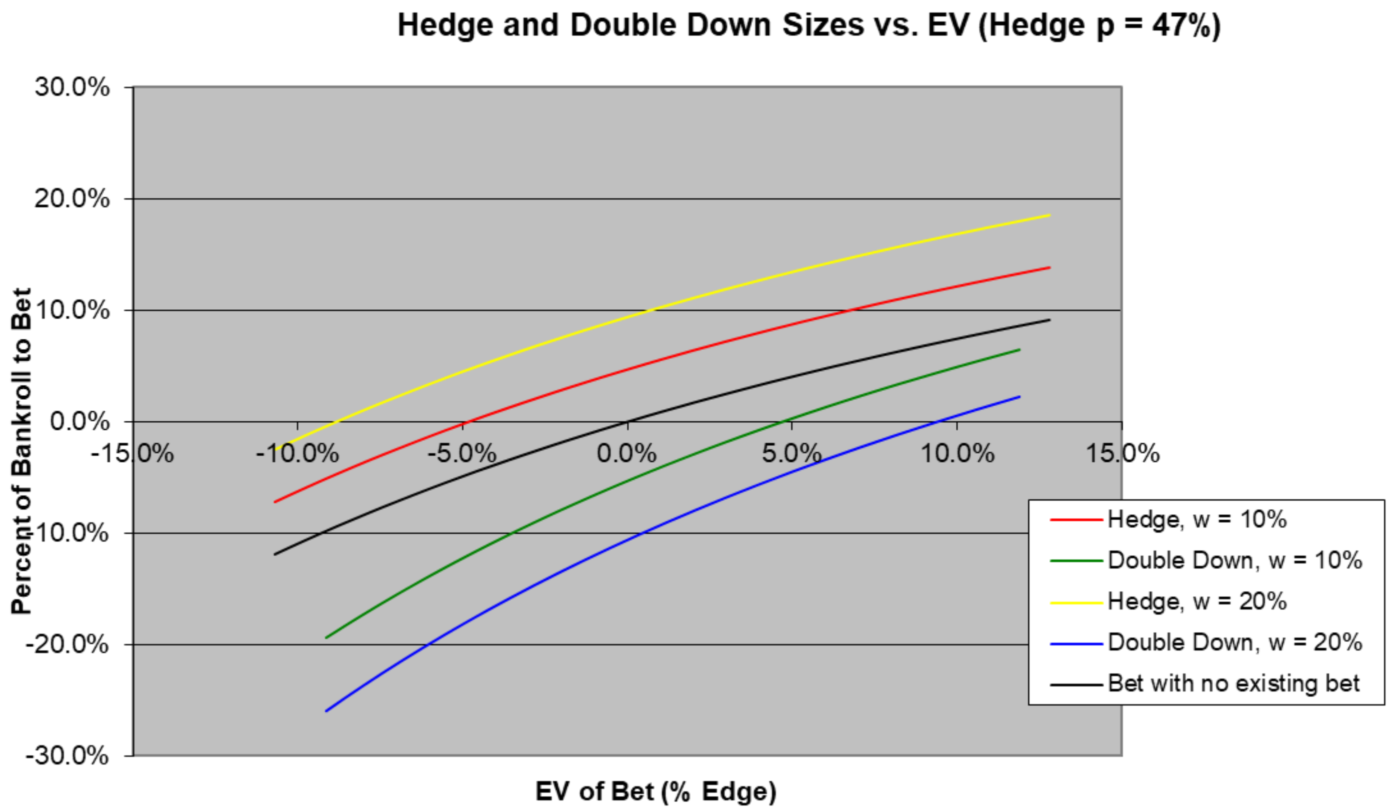

The conclusion I've come to is that while the Kelly Criterion as you know it is very important and powerful, it doesn't apply in all situations without some modification. In extreme spots, that are also extremely common, your optimal bankroll management may look very different than intuition would have you believe. That's ok. It took time to understand EV and to master how to use it to grow your bankroll by choosing profitable independent bets. And you can master these spots too.