ECON30580 Economics of Betting Markets
17. The Kelly Criterion

Karl Whelan

School of Economics, UCD

Spring 2026

Part I

Betting with an Edge

What To Do When You Have An Edge?

- There is a bet with decimal odds D that you think has probability p of winning where $pD > 1$. In other words, you think you have an edge over the bookmaker. What do you do?
- Up to now, our models have limited bet size to equal one unit, so the answer so far has been "place your one unit bet."
- But what if you get to choose how much to bet?
- Suppose you decide to invest a fraction x of your wealth W.
  - If you win, your wealth increases to $(1 + x(D - 1))W$.
  - If you lose, your wealth falls to $(1 - x)W$.
- The expected value of your wealth will be

$$[p(1 + x(D - 1)) + (1 - p)(1 - x)]W = [1 + (pD - 1)x]W$$

- Because $pD - 1 > 0$, this is maximised by setting x as high as possible.
- You maximise expected wealth by betting as much as you can.

Problems with Betting it All

- Ok, suppose you decide to bet all your wealth.
- Your expected wealth from this will be pDW which is greater than W.
- If you could bet all your wealth for N periods in a row on bets like this, your expected wealth after period N would be $(pD)^N W$.
- As N gets larger, your expected wealth heads towards infinity, mwah-ha-ha!
- Sounds great. What's the catch?
- Well your probability of losing all your money is $1 - p^N$. As N gets big, this tends towards one.
- On average, you get infinite wealth and yet most likely you end up with nothing.
- This can't really be an optimal strategy, particularly if you are risk averse.

John Kelly

- A famous 1956 paper by mathematician John L. Kelly suggested the best strategy was not to maximise wealth itself but instead to maximise the log of wealth.
- He noted that if you started with wealth of W, placed N bets risking a fraction x of your wealth and won M of these, your wealth would be

$$\left[(1 + x(D-1))^{M}(1-x)^{N-M}\right] W$$

- Given this, he suggested maximising

$$(1 + x(D-1))^{p}(1-x)^{1-p}$$

- Because $\log(y)$ is increasing in y, this is the same as maximising

$$p \log(1 + x(D-1)) + (1-p)\log(1-x)$$

- And actually, maximising this is just the same thing as maximising the expected log of wealth each period.
- The rule suggested by Kelly is known today as "the Kelly criterion" and it is very widely discussed in gambling circles and (somewhat less) in financial markets.

Kelly's Motivation

- Kelly worked for the American Telephone and Telegraphs company (AT&T). The connection is not as random as you might think.
- In the early years of telegraphs, there was a big demand for reporting results of horse races.
- Having a faster "wire" sometimes allowed gamblers to bet on horses that they knew had won.
- Kelly's motivation in his paper was the idea that the signal obtained by the gambler could be "noisy" so it only had a probability p of being correct.
- Alternatively, bookmakers were very keen on being able to take bets on horses that they knew had lost.
- There is a highly entertaining book about the origin of the rule and its history called "Fortune's Formula" by William Poundstone.
- It has lots of entertaining stories, including about how the controversies over the mob's wire business that sent results to illegal bookmakers caused lawmakers to work harder to restrict gambling in the US and how various people have tried to use Kelly's rule to make money at blackjack, roulette and in investments.

Deriving the Kelly Criterion

- Remembering that

$$\frac{d \log x}{dx} = \frac{1}{x}$$

we can calculate the x that maximizes

$$p \log \left(1 + x \left(D - 1\right)\right) + \left(1 - p\right) \log \left(1 - x\right)$$

from the first-order condition

$$\frac{p \left(D - 1\right)}{1 + x \left(D - 1\right)} - \frac{1 - p}{1 - x} = 0$$

- Re-arranging it, we get the solution

$$x = \frac{pD - 1}{D - 1}$$

- The formula is sometimes described as "**edge over odds**" because $pD - 1$ is the expected profit on a one dollar bet (your edge) and $D - 1$ is the fractional odds.

Examples

- Suppose the bet is "even money", meaning $D = 2$.
- The Kelly criterion predicts you should bet the following fraction of your wealth

$$x = 2p - 1$$

- So, for example, if you think $p = 0.6$ then $D = 2$ implies betting one-fifth of your wealth because $x = 2(0.6) - 1 = 0.2$.
- Alternatively, suppose the odds are $D = 3$ but you think the chance of winning is $p = 0.5$. In this case, the optimal fraction to bet is

$$x = \frac{(0.5)(3) - 1}{3 - 1} = \frac{0.5}{2} = 0.25$$

- The graph on the next page shows the expected log of wealth for different values of x for this case. Invest too little and you don't take enough advantage of your edge. Invest too much and you run too big a risk of losing all your money.

Expected Log Wealth Starting from $W = 1$ with $p = 0.5$ and $D = 3$

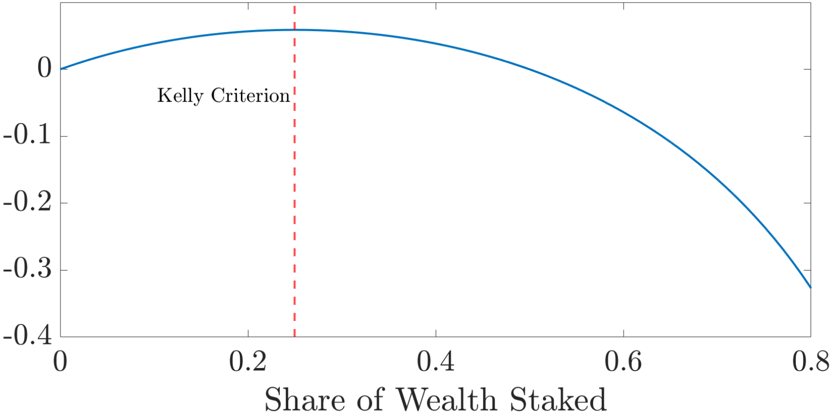

Positives and Negatives of the Kelly Criterion

- The Kelly criterion is known to have various positive properties.
  - Let W_n^a be the wealth after n periods of any other non-terminating strategy, meaning a strategy that does not allow your wealth to go to zero. Then as n gets large, the ratio of wealth from adopting the Kelly criterion to W_n^a gets ever higher.
  - The expected time for your wealth to achieve a specified goal is minimised by adopting the Kelly criterion.
- But it also has its critics.
  - Legendary Nobel-prize winning MIT economist Paul Samuelson criticised recommendations to use the Kelly criterion because, by definition, maximising the expected log of wealth is only the same thing as maximising expected utility if the person has a log utility function $U(w) = \log w$. And this is a very specific utility function which probably doesn't describe most people's preferences.
  - Others feel that the Kelly criterion's recommendations are too aggressive (i.e. they recommend betting too much). These people are probably more risk averse than implied by log utility.
  - Some argue for using a scaled "fractional" version of the Kelly criterion.

Maximising a More General Utility Function with an Edge

- Going beyond log utility, it is relatively easy to derive the criteria for the fraction of wealth to bet that maximises expected utility when you have an edge.
- You want to pick x to maximise

$$E(U) = pU(1 + x(D-1)) + (1-p)U(1-x)$$

which has first-order condition

$$p(D-1)U'(1 + x(D-1)) - (1-p)U'(1-x) = 0$$

- This reduces to a condition on the ratio of marginal utilities for the two possible outcomes, i.e. winning and losing

$$\frac{U'(1 + x(D-1))}{U'(1-x)} = \frac{1-p}{p(D-1)} = \frac{1}{p(D-1)} - \frac{1}{D-1}$$

- If the utility function is concave, then marginal utility is declining in wealth and the fraction on the left side depends negatively on x. This means that for all concave utility functions x depends positively on both p and D.

Solution with CRRA Utility

- A reminder that the CRRA utility function is

$$U(W) = \begin{cases} \frac{W^{1-\gamma}-1}{1-\gamma} & \text{if } \gamma \neq 1, \\ \log W & \text{if } \gamma = 1 \end{cases}$$

- For this utility function, after a bunch of algebra, you can show that the optimal fraction of wealth to bet is

$$x = \frac{\left(\frac{p(D-1)}{1-p}\right)^{\frac{1}{\gamma}} - 1}{\left(\frac{p(D-1)}{1-p}\right)^{\frac{1}{\gamma}} + D - 1}$$

- When $\gamma = 1$, this reduces to the Kelly criterion formula.
- I won't be asking anyone to derive this formula on the exam.

About Fractional Kelly Rules

- There is an "urban legend" in online discussions of the Kelly criterion that non-log CRRA utility perhaps justifies "fractional Kelly" rules.
- For example, if $\gamma = 2$, so the coefficient of relative risk aversion is twice that of someone with log utility, then perhaps the optimal rule is to bet half the amount predicted by the Kelly criterion.
- In my paper "On Optimal Betting Strategies With Multiple Mutually Exclusive Outcomes" I show that this is not the case.
- In general, it is true that the more risk averse someone is (the higher their value of γ) then the less they will bet.
- But fractional Kelly rules are not optimal.
- The chart on the next page shows the ratio of the optimal betting amount to the Kelly criterion level for different values of γ, for different decimal odds and for different values of the subjective belief about the probability of the bet's success (here labelled π).
- If the optimal rules were fractional Kelly, then all these lines would be flat but none of them are.

Ratio of optimal fraction of wealth staked to the fraction implied by the Kelly criterion for various odds, beliefs (π) and preferences γ

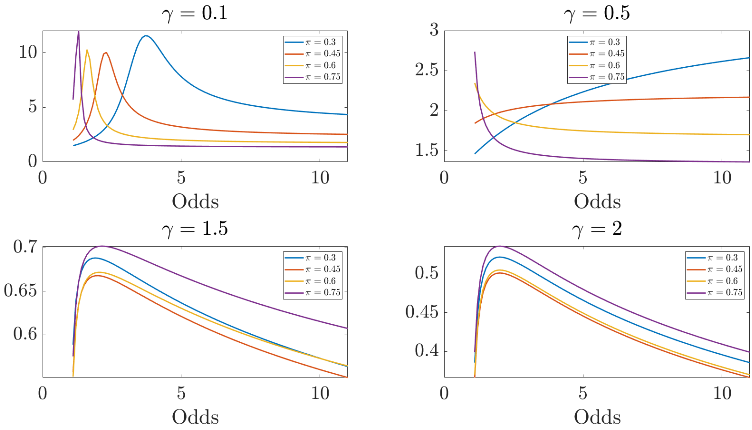

Multiple Possible (But Mutually Exclusive) Outcomes

- The Kelly criterion describes how to approach a single bet.
- But many sporting events have multiple possible winners but only one will actually win.
- If may be that for the odds offered on such an event, your assessment of the probabilities suggests there may be multiple bets worth taking.
- How much do you place on each bet to maximize your expected utility? The traditional Kelly rule no longer describes the optimal strategy. You have to factor in that if one bet wins, then the others won't.
- My paper on this topic (On Optimal Betting Strategies With Multiple Mutually Exclusive Outcomes) uses numerical methods to solve for the optimal strategy in this case and to illustrate its properties.
- Two interesting patterns emerge about the optimal betting rule in this case
  - It recommends higher bet volume than the Kelly criterion.
  - It sometimes recommends taking bets that have a negative expected value. This is because bets on one outcome act as a "hedge" against the outcome of another bet losing. In a situation where I have lost a lot of money betting on other options, the negative expected value bet has a utility value as a form of "insurance."

A Simulation Exercise

- To illustrate the optimal betting rules, I simulated odds for events with 5 possible outcomes, each with an average chance of 0.2 of occurring.
- I randomly assigned probabilities to each outcome in an event (adjusting to make sure they still summed to one).
- I assumed bookmakers know the true probabilities and set odds based on a 4% profit margin.
- For each probability, we then modelled people's beliefs about the probabilities as equalling the true probability plus some random "noise" distributed as $N(0, \theta)$.
- The chart on the next page shows the probabilities for all the simulated bets and also the corresponding beliefs.
- The page after shows expected returns and fractions of wealth staked for the Kelly criterion and for the optimal strategy with log utility for all available bets.
- The optimal strategy is more aggressive and makes a surprisingly large amount of negative expected value bets. In the example shown here, 17% of bets placed have negative expected values.

Simulated probability beliefs of bookmakers and bettors ($\theta = 0.02$)

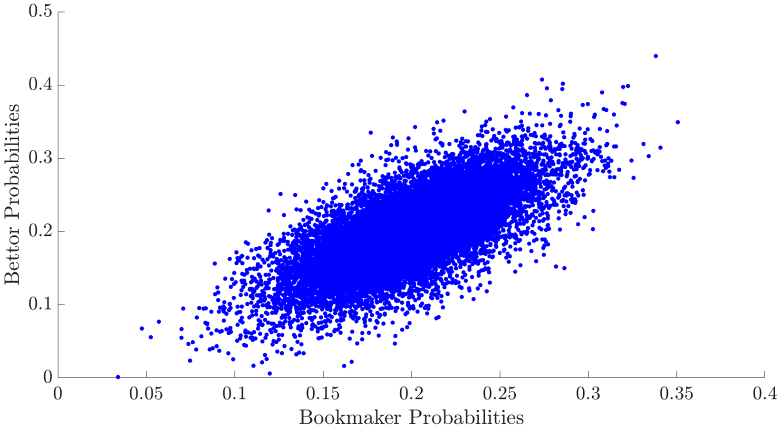

Expected returns and fractions of wealth at risk for log utility for both optimal strategy and the Kelly criterion.

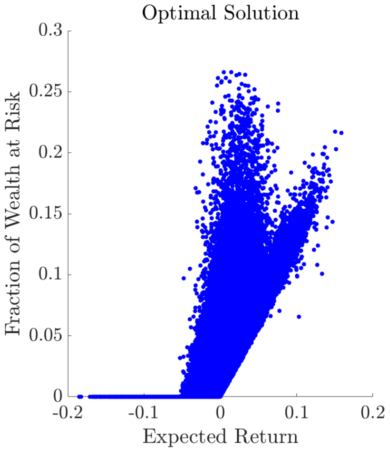

The Real Problem with the Kelly Criterion

- I do not recommend you bet on sports.
- So obviously I don't recommend that you use the Kelly criterion or any of the CRRA variants described here.
- The reason: **You almost certainly do not have an edge**.
- Bookmakers base their odds on complex statistical modelling. What are the chances that you've found a bet where they are systematically wrong and you're right?
- The table on the next page shows the losses incurred by the bettors in the simulations just described.
- The θ parameter represents the typical size of the random incorrect element of the beliefs of bettors. Beliefs equal the true probabilities plus a term that is distributed $N(0, \theta)$.
- Using 52 bets as a benchmark (a bet a week), the case we have shown sees the bettors following the "optimal strategy" losing 14% of their wealth compared with losing 8% of it following the traditional Kelly criterion.
- The real optimal strategy is not betting at all.

Losses from the "optimal strategy"

Table 2: Loss implications for betting strategies with log utility, $w = 1$ and $\gamma = 0.02$ for three values of θ (standard deviation of difference between bettors' beliefs and true probabilities) for the five-outcome case. 10,000 repetitions.

		$\theta = 0.01$	$\theta = 0.02$	$\theta = 0.03$
Average percent of wealth lost for optimal strategy		0.060%	0.280%	0.590%
Average percent of wealth lost for two outcome rule		0.046%	0.150%	0.270%
Average percent of wealth at risk for optimal strategy		1.51%	7.11%	14.66%
Average percent at risk for two outcome rule		1.14%	3.80%	6.73%
Expected wealth after 52 bets with the optimal strategy		0.969	0.862	0.737
Expected wealth after 52 bets with the two-outcome strategy		0.976	0.924	0.869
Average loss from randomly betting 3% of wealth	0.12%			
Expected wealth after 52 bets randomly betting 3% of wealth	0.939			

Supplementary Material in the Draft Book

- Chapter 28: Risk, Ruin and the Kelly Criterion