

UCD CENTRE FOR ECONOMIC RESEARCH

WORKING PAPER SERIES

2023

**Fortune's Formula or the Road to Ruin?
The Generalized Kelly Criterion With Multiple Outcomes**

Karl Whelan
University College Dublin

WP23/09

April 2023

**UCD SCHOOL OF ECONOMICS
UNIVERSITY COLLEGE DUBLIN
BELFIELD
DUBLIN 4**

Fortune's Formula or the Road to Ruin? The Generalized Kelly Criterion With Multiple Outcomes

Karl Whelan*

University College Dublin

April 2023

Abstract

You can bet on an event where there are multiple possible winners but only one will actually win. At the odds offered, you think there may be multiple bets worth taking. How much do you place on each bet to maximize your expected utility? We describe how this problem can be solved for concave utility functions and illustrate the properties of the solution. The optimal betting strategy is more aggressive than strategies derived from considering each outcome separately such as the Kelly criterion. This strategy also recommends sometimes placing bets with negative expected returns because they act as hedges against losses on other bets. While this strategy maximizes the bettor's subjective expected utility, if betting odds incorporate a profit margin and reflect underlying probabilities correctly, then this more aggressive strategy loses more money and results in lower realized utility.

Keywords: Decision-making under uncertainty; optimal betting; Kelly criterion

JEL Classification: D81, G11

*karl.whelan@ucd.ie. Thanks to Tadgh Hegarty for comments on a previous draft.

1. Introduction

You can bet on an event where there are multiple possible winners but only one will actually win. At the odds offered, 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? This problem applies to gambling on sports or investing in prediction markets but it can also be generalized to business or investment opportunities where the market is perceived as having a winner takes all property. This paper shows how to solve the problem for concave utility functions, characterizes the properties of the solution and describes how those who apply this solution would do if the probabilities implied by the odds are correct.

One answer to our question is provided by the so-called Kelly criterion, derived in a famous 1956 paper by John L. Kelly, a Bell Labs scientist. The Kelly criterion provides a simple formula, based on the odds and your subjective probability of the bet's success, for how much to bet to maximize your expected geometric average of wealth. While some have argued for the merits of maximizing the expected geometric average of wealth as an optimal long-run investment strategy, this approach can only be interpreted as a utility-maximizing strategy if agents have log utility.¹ Since most people probably do not have log utility, Paul Samuelson (1971, 1979) argued vociferously against recommendations to use the Kelly criterion, including his 1979 paper containing only words with one syllable.

Despite Samuelson's efforts, the Kelly criterion is more influential than ever. Its popularity was boosted by William Poundstone's entertaining 2006 book *Fortune's Formula* with its stories of attempts by Claude Shannon, Edward Thorp and others to make money from Kelly's rule at Vegas roulette tables and investing in stock warrants. And in the modern world of online sports betting, now legal in the United States since the abolition of its federal prohibition in 2018, the internet provides many "Kelly calculators" illustrating how much people should bet.

An important limitation of the Kelly criterion—and its utility-maximizing equivalents that can be derived under the assumption of non-log CRRA utility—is that it assumes there are only two outcomes: Either an event happens or it doesn't happen. However, in many cases—such as betting on horse racing or a bet in December on which NFL team will win the Super Bowl—there are multiple possible outcomes but only one can occur. Someone who places two bets in December, one on the Kansas City Chiefs to win the Super Bowl and the other on the Buffalo Bills, needs to factor in that only one of these bets will win. To correctly solve this problem, the optimal amount to place on each bet needs to be a function of the probabilities of all N outcomes and of all N betting odds. We show how to derive this solution and present its implications

To my knowledge, only two pieces of academic research have addressed this question in a practical manner. Sung and Johnson (2010) describe solving this problem for optimizing log utility in the context of assessing whether the recommended strategy could potentially earn profits betting on UK

¹See, for example MacLean, Thorp, Zhao and Ziemba (2011)

horse racing but they do not describe their solution method or characterize the general properties of the optimal strategy. Smoczynski and Tomkins (2010) presented a solution method for the problem for maximizing the expected log of wealth. In the absence of an explicit formula for the optimal number of bets on each option, they recommended an iterative numerical method for calculating the optimal set of bets and presented two specific numerical examples. Their method starts with betting the Kelly criterion's recommendation for the bet with the highest expected profit and then checks whether adding other bets violates the problem's optimality conditions.²

This paper adds the following new elements. First, we characterize optimal solutions for the general case of concave utility, and present calculations for other CRRA utility functions as well as log utility. We do this both for when there are two possible outcomes and for when there are N possible outcomes. For the two outcome case, we show that while it has often been hypothesized that optimal strategies for non-log CRRA utilities are well approximated by fractionally scaled versions of the Kelly criterion with the coefficient of relative risk aversion as the scaling parameter, this is generally a poor description of optimal betting behavior.

Second, we explain the circumstances under which the optimal betting problem with multiple mutually exclusive outcomes does and does not have a unique solution. We show that when bookmakers' odds contain no profit margin and bettors can take both positive or negative positions—in other words “lay bets” are available as well as “back bets”—there is no unique solution to the optimal betting problem with N outcomes. We then explain how the problem has a solution when a profit margin for bookmakers is built into odds and when only bets backing events occurring can be taken. Our solution method works for all concave utility functions and can also be obtained from a simple Matlab program.

Third, we characterize how betting behavior differs when solving the true optimization problem relative to using the two-outcome rules implied by the Kelly criterion or its non-log CRRA alternatives. At first glance, it may seem that factoring in that there is only one winner would make betting less attractive than when using a two-outcome rule: The bets on both Chiefs and Bills may seem a bit less profitable relative to calculations that assumed both could win. However, we show that the optimal strategy is to bet more than recommended by two-outcome rules. This is because the optimal strategy takes into account the value of bets as hedges against losses on other bets. Losses on one bet boost the marginal utility in the state of the world where another bet wins and this raises the optimal amounts placed relative to when only the two outcomes of win/don't win are considered. In particular, the true optimal strategy can involve placing bets that have negative expected returns.

We use realistically calibrated simulation exercises to illustrate how much more aggressive the optimal strategy is than two-outcome rules and the likely prevalence of taking negative expected

²A separate but also interesting case, where there are N independent events that can be betted on, was analyzed using numerical methods by Whitrow (2007).

return bets. Sung and Johnson (2010) noted that negative expected return bets can be recommended by the optimal rule but that they were a small fraction of the recommended bets in their study and that their strategy did not suggest placing multiple bets on the same event. Our calculations suggests betting is much more aggressive under the optimal strategy than under the two-outcome rules, that multiple bets on one event would be common if this rule was followed and that bets with negative expected returns would also be commonly taken.

Finally, we illustrate the performance of the derived betting strategy when the odds posted by bookmakers are a correct reflection of the underlying probabilities. When applied by bettors who mistakenly believe they have an edge over bookmakers, the more aggressive “optimal strategy” results in larger losses and lower utility in judging probabilities than if they had adopted two-outcome optimal rules.

2. Optimal Betting With Two Outcomes

We begin with a model with two possible outcomes. Bettors are offered a bet at a price p that will return \$1 if an event occurs and zero if it does not. They can also place a “lay bet” backing that the event does not occur, which will cost $1 - p$ and pay out \$1 if the event does not happen and zero if it does. An agent with starting wealth w , utility function $U(w)$ (with $U' > 0$ and $U'' < 0$) and subjective probability π that the event will occur, solves the problem of how many bets x to place by maximizing their subjective expected utility³

$$\text{Max}_x \left[\pi U(w + x(1 - p)) + (1 - \pi) U(w - xp) \right] \quad (1)$$

where a positive x represents a bet that the event will happen and a negative x represents a bet that the event will not happen. The first-order condition is

$$\pi(1 - p)U'(w + x(1 - p)) = (1 - \pi)pU'(w - xp) \quad (2)$$

The marginal utility in the case of winning is monotonically falling in x and the marginal utility in the case of losing is monotonically rising in x , so any x that solves this equation will be unique. The ratio of marginal utilities is

$$\frac{U'(w + x(1 - p))}{U'(w - xp)} = \frac{1 - \pi}{\pi} \frac{p}{1 - p} \quad (3)$$

³One could object to the formulation of utility as a function of wealth, since utility is usually modeled as a function of consumption rather than wealth. However, for the CRRA utility functions considered in this paper, it is well known that when solving optimal consumption problems with dynamic programming, the value functions with respect to wealth for the lifetime present discount value of utility take the same functional form as the original consumption-based utility function. So this utility is better seen as a measure of the present discounted value of lifetime utility obtained from wealth rather than an instantaneous flow of utility from wealth.

which means that x depends positively on π and negatively on p . If $p = \pi$, then the marginal utilities are equal, which implies that $x = 0$, so no bet is placed. For CRRA utility,

$$U(w) = \begin{cases} \frac{w^{1-\sigma}-1}{1-\sigma} & \text{if } \sigma \neq 1, \\ \log w & \text{if } \sigma = 1 \end{cases} \quad (4)$$

the optimal betting strategy is

$$x = \left(\frac{1 - \left(\frac{1-\pi}{\pi} \frac{p}{1-p} \right)^{\frac{1}{\sigma}}}{(1-p) \left(\frac{1-\pi}{\pi} \frac{p}{1-p} \right)^{\frac{1}{\sigma}} + p} \right) w \quad (5)$$

For log utility, the number of bets placed is

$$x = \left(\frac{\pi - p}{p(1-p)} \right) w \quad (6)$$

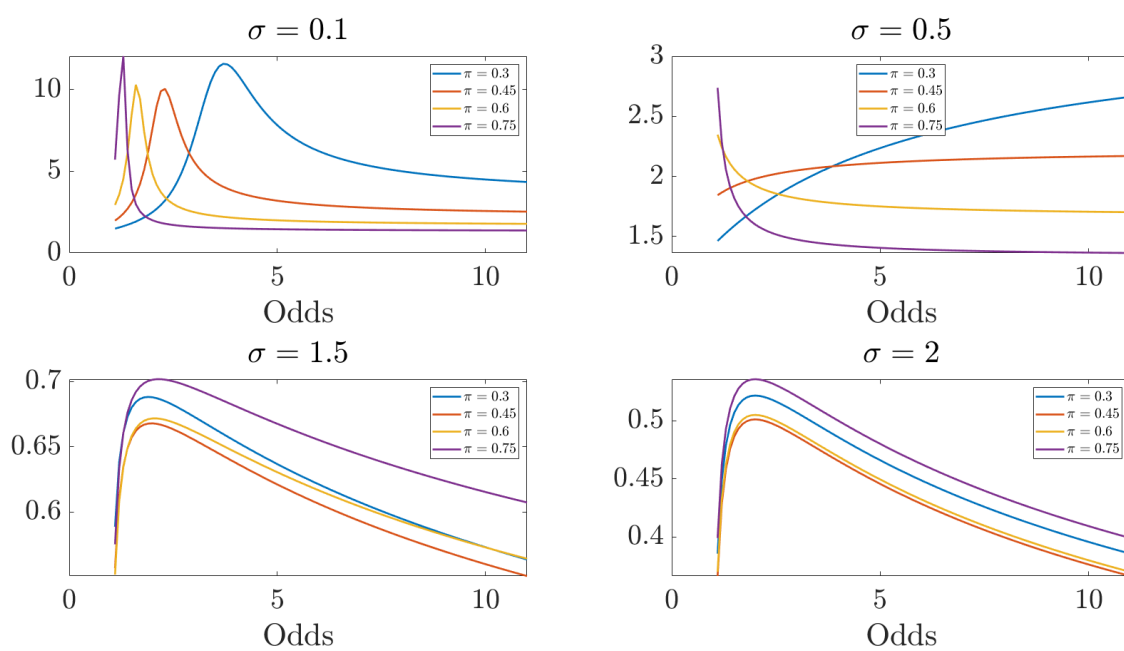
This numerator here is the expected profit $\pi - p$ on the bet and the denominator is the expected variance of these profits, $p(1-p)$. If we rephrase this in terms of fractional betting odds, so that you win O from a \$1 bet if outcome one occurs, meaning $O = \frac{1}{p} - 1$, the total fraction of wealth allocated on the bet is $px = \frac{\pi O - (1-\pi)}{O}$ which is the common ‘‘edge over odds’’ way of expressing the Kelly criterion, as used for example by Poundstone (2006).

As expected, the gross amounts of bets placed declines as the coefficient of relative risk aversion σ increases so for bettors with $\sigma < 1$, their optimal strategy is to be more aggressive than the Kelly criterion while those with $\sigma > 1$ should be less aggressive. Internet discussions about the Kelly criterion commonly feature the conjecture that the rule can be generalized so that those who have a different tolerance for risk than log utility should use a ‘‘fractional Kelly’’ strategy, setting their investment equal to the Kelly criterion amount divided by σ . However, while there are certain conditions under which this result holds, it is not an accurate description of the optimal betting problem.⁴

Figure 1 shows ratios of the numbers of bets implied by equation 5 to the corresponding number of bets implied by the Kelly criterion in equation 6 for four different values of σ , for four different values of π and for a set of decimal bettings odds ($O = \frac{1}{p}$) ranging from 1 (at which point the bettors take a negative position for each of the values of π and σ) up to 11 (at which point everyone takes a positive position). If a fractional Kelly betting strategy with scaling factor σ^{-1} was the optimal strategy then each of these lines would be flat and equal to σ^{-1} . Instead the figure shows large and systematic deviations from these values.

⁴See Maclean, Ziemba and Li (2005) for a discussion of fractional Kelly strategies when investing in assets that follow a Brownian motion.

Figure 1: Ratio of number of bets in equation 5 to number implied the Kelly criterion (equation 6) for various odds, beliefs (π) and preferences σ



3. Betting with N Outcomes: A Non-Uniqueness Result

Now consider a market where an event has N possible mutually exclusive outcomes with the bookmaker offering prices p_i for bets that win \$1 if outcome i occurs and the bettor can take as many or as few of the bets as they want. Those who rarely bet may imagine bettors don't often take up the opportunity to bet on multiple outcomes in the same event but I have been assured by a real-world bookmaker that it is quite common. In fact, there are many websites with "dutching" calculators that tell you how much to place on each of a number of options to result in the same profit should any of the options win.⁵

For now, we will assume for now that the bookmaker does not take a profit so she sets each price equal to her perceived probability of the outcome which means they expect both bettors and bookmakers to break even on each bet. This means

$$\sum_{i=1}^N p_i = 1 \quad (7)$$

For a bettor with a set of subjective beliefs π_i about the probability of each outcome i , faced with prices p_i , the expected utility maximizing strategy is obtained by solving

$$\text{Max}_{x_1, x_2, \dots, x_N} \sum_{i=1}^N \pi_i U \left(w + x_i - \sum_{k=1}^N x_k p_k \right) \quad (8)$$

where again a positive x_i denotes a bet backing outcome i and a negative x_i is a bet that outcome i will not occur. The first-order condition for the i th bet is

$$\pi_i (1 - p_i) U' \left(w + x_i - \sum_{k=1}^N x_k p_k \right) = p_i \sum_{n \neq i} \pi_n U' \left(w + x_n - \sum_{k=1}^N x_k p_k \right) \quad (9)$$

which can also be written as

$$\pi_i U' \left(w + x_i - \sum_{k=1}^N x_k p_k \right) = p_i \sum_{n=1}^N \pi_n U' \left(w + x_n - \sum_{k=1}^N x_k p_k \right) \quad (10)$$

Writing the marginal utilities obtained when outcome i occurs as

$$U'_i = U' \left(w + x_i - \sum_{k=1}^N x_k p_k \right) \quad (11)$$

⁵See for example <https://www.oddschecker.com/betting-tools/dutching-calculator>

The term comes from the supposed invention of this technique by Dutch Schultz, who was, among other things, apparently a believer in smoothing marginal utility across Arrow-Debreu states of the world.

an optimal set of bets must satisfy

$$U'_i = \frac{p_i}{\pi_i} \sum_{k=1}^N \pi_k U'_k \quad (12)$$

Marginal utility when outcome i occurs rises with the price of the bet p_i and falls with the belief π_i . This means any optimal solution will have the number of bets x_i varying positively with π_i and negatively with p_i , which is to be expected. What is perhaps less expected is that there is no unique solution for the optimal marginal utilities and thus for the number of bets x_i . Any set of marginal utilities such that

$$\frac{U'_i}{U'_j} = \frac{p_i \pi_j}{p_j \pi_i} \quad (13)$$

will satisfy the optimality conditions. This means that widely varying sets of betting positions can all satisfy the optimality conditions and deliver the same level of expected utility. In general, once the agent has chosen to bet a lot of money on outcome i occurring, the fact that this large bet has been lost in the case where outcome j occurs increases the marginal utility of wealth relevant for deciding how many bets to place on outcome j . Thus, the increased betting on one outcome can raise betting on the others and this leads to there being multiple optimal betting strategies.

A more formal way to express the absence of a single optimum for this problem is to write the first-order conditions in matrix form as

$$AU' = 0 \quad (14)$$

where

$$A = \begin{pmatrix} (1-p_1)\pi_1 & -p_1\pi_2 & -p_1\pi_3 & \dots & -p_1\pi_N \\ -p_2\pi_1 & (1-p_2)\pi_2 & -p_2\pi_3 & \dots & -p_2\pi_N \\ -p_3\pi_1 & -p_3\pi_2 & (1-p_3)\pi_3 & \dots & -p_3\pi_N \\ \dots & & & & \\ \dots & & & & \\ -p_N\pi_1 & -p_N\pi_2 & -p_N\pi_3 & \dots & (1-p_N)\pi_N \end{pmatrix} \quad U' = \begin{pmatrix} U'_1 \\ U'_2 \\ U'_3 \\ \dots \\ \dots \\ U'_N \end{pmatrix} \quad (15)$$


Each of the columns of A sum to zero because $\sum_{i=1}^N p_i = 1$, so this matrix has rank $N - 1$. This means there is an infinite number of combinations of marginal utilities that satisfy the optimality conditions. Technically, the set of marginal utilities that solves this equation is defined by the null space of A , which is all possible multiples of $\left(\frac{p_1}{\pi_1}, \frac{p_2}{\pi_2}, \dots, \frac{p_N}{\pi_N}\right)$.

For the case of two outcomes, the non-uniqueness of optimal demand is somewhat trivial. One can show that there is a unique optimal value for $x_1 - x_2$, so the conditions pin down the net position

of the bettor, if not their gross positions. This is why the case where there are two possible outcomes, but bets can only be taken for or against one them, has a unique solution. However, for three or more outcomes, the multiplicity of equilibrium is not trivial. For example, consider the case where subjective beliefs are $\pi_1 = \pi_2 = \pi_3 = \frac{1}{3}$ and prices are $p_1 = \frac{1}{3} + 0.1$, $p_2 = \frac{1}{3}$ and $p_3 = \frac{1}{3} - 0.1$. In this case, the Kelly criterion selects $(x_1, x_2, x_3) = (-0.56, 0, 0.41)$ and shares of wealth placed are $(p_1x_1, p_2x_2, p_3x_3) = (-0.13, 0, 0.18)$

By contrast, Figure 2 shows a set of betting positions on the three outcomes that all satisfy the optimality conditions for these beliefs and prices. These were each calculated by assuming a value between -1 and 1 for one of the x_i and then solving for the now-unique values of the other two bets. The conditions are satisfied by betting a large amount on each of the outcomes and by betting small amounts on each and also by betting a large amount against each of the outcomes. The line slopes up in each direction as the betting positions go from negative to positive, meaning there is a positive reinforcement effect so that taking large positions on one outcome encourages similar large positions on the other outcomes.

Figure 2: Set of optimal bets when $\pi_1 = \pi_2 = \pi_3 = \frac{1}{3}$ and $p_1 = \frac{1}{3} + 0.1$, $p_2 = \frac{1}{3}$ and $p_3 = \frac{1}{3} - 0.1$



4. Obtaining A Unique Solution

Two elements of real-world betting turn out to provide a unique optimal solution for our problem. First, bookmakers have costs to cover and need to earn a profit to stay in business. This means the prices of the bets will sum to more than one

$$\sum_{i=1}^N p_i > 1 \quad (16)$$

The percentage to which these prices sum to greater than one is often termed the “overround” in gambling circles and typical values in competitive betting markets tend to be around 4%. Second, bookmakers do not typically take “lay” bets. Just because they offer a price of p_i on outcome i , that does not mean they are offering a price of $1 - p_i$ on outcome i not occurring. So we will restrict our search for an optimal betting strategy to positive values of x_i . The bettor’s problem can now be expressed as a Karush-Kuhn-Tucker problem with Lagrangian

$$L = \sum_{i=1}^N \pi_i U \left(w + x_i - \sum_{k=1}^N x_k p_k \right) + \lambda_i x_i \quad (17)$$

The first-order conditions are now given by $AU' = \lambda$ where $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_N)'$, the constraints $x_i \geq 0$ and $\lambda_i \geq 0$ and the N complementary slackness conditions $\lambda_i x_i = 0$.

Both new elements of the problem are required for there to be a unique optimal solution and they interact with each other. The fact that $\sum_{i=1}^N p_i > 1$ means there will be at least one binding non-negativity constraint for betting quantities. To see this, note that the strategy of taking one of each of N bets will return a dollar for sure but will also be a loss-making strategy since the cost of these bets will be greater than a dollar. So consider a betting strategy for an event with three possible outcomes, (x_1, x_2, x_3) where all of the betting amounts are positive and x_3 is the smallest of the number of bets placed. This can be considered a combination of the betting strategy $(x_1 - x_3, x_2 - x_3, 0)$ and (x_3, x_3, x_3) . Since the latter is for sure a losing strategy, if the original strategy was expected to make a profit, then the option of reducing each bet size by x_3 must make a greater profit in all cases. This means that if $N - 1$ of the x_i are positive, then the N th must be zero so the non-negativity constraint will be binding for at least one of the bets.

The sum of the prices no longer equaling one also means the previous argument for A having a reduced rank no longer holds, so A is now an invertible matrix. But unless there are some binding constraints, the first-order conditions would just reduce to $AU' = 0$ which only has solution $U' = 0$, which does not have a defined solution for any set of x_i s for concave utility functions. However, because at least one of the λ_i are positive, $AU' = \lambda$ has a unique solution for the N state-contingent marginal utilities.

Once we have a solution for the optimal marginal utilities $\theta = (\theta_1, \theta_2, \dots, \theta_N)'$, we can then map these into x_i values for any concave utility function. For CRRA utility functions, this is done as follows. The solutions to the first-order conditions can be written as

$$\pi_i (1 - p_i) \left(w + x_i - \sum_{k=1}^N p_k x_k \right)^{-\sigma} = \theta_i \quad (18)$$

which can be re-written as $Ax = b$ where $b = \left(\left(\frac{\pi_1(1-p_1)}{\theta_1} \right)^{\frac{1}{\sigma}} - w, \left(\frac{\pi_2(1-p_2)}{\theta_2} \right)^{\frac{1}{\sigma}} - w, \dots, \left(\frac{\pi_N(1-p_N)}{\theta_N} \right)^{\frac{1}{\sigma}} - w \right)'$. This problem can be solved for unique non-trivial values of x because A is invertible and $b \neq 0$.

This problem has been previously discussed for the case of maximizing the expected log of wealth by Smoczynski and Tomkins (2010). They provide an formula for the amount of bets that should be placed among the sub-group of bets for which $x_i > 0$ is optimal, contingent on the other bets not being taken. Calling the group of bets that are taken S and using our terminology, the formula is

$$x_i = \frac{\pi_i}{p_i} - \frac{\left(\sum_{k \notin S} \pi_k \right)}{1 - \sum_{k \in S} p_k} \quad (19)$$

which collapses to the Kelly criterion when only one bet is placed. This formula, however, does not determine which bets belong in S . Smoczynski and Tomkins show that for log utility, for x_i to be positive, you need

$$\frac{\pi_i}{p_i} > 1 - \sum_{k \in S} p_k x_k \quad (20)$$

They recommend an iterative algorithm that starts with the bet with the highest value of $\frac{\pi_i}{p_i}$ and then looks to add other bets, provided equation 20 is then satisfied for these bets, stopping with the first bet that does not satisfy this condition.

In contrast, we use the Matlab function `fmincon`, which employs an interior point algorithm, to solve the Lagrangian optimization problem in equation 17 for any concave utility function.⁶ Matlab programs available on my website show that for log utility, this code provides the same answers as the two numerical examples in Smoczynski and Tomkins (2010) and, more generally, that the solutions for betting amounts with log utility satisfy equation 19. The code solves the model quickly for all possible CRRA utility functions as well as other concave functions such as exponential utility functions. It also allows for easy computation of the optimal betting strategy across a wide range of environments.

⁶See Wright (1992) for a discussion of interior point algorithms.

5. Illustration with a Simple Example

We first illustrate the properties of the optimal solutions and compare them with the decisions derived from a two-outcome approach with a simple example. We continue to assume that the bettor's subjective probabilities are $\pi_1 = \pi_2 = \pi_3 = \frac{1}{3}$. We assume prices for bets are set as a markup over the market's assessment of the underlying probabilities. These probabilities could be either the subjective assessment of a bookmaker in the case of fixed-odds betting or the probabilities implied by the public's beliefs in pari-mutuel betting odds. For simplicity, we will refer to these as bookmaker's probabilities. We will assume a constant markup across bets so that prices sum to $\frac{1}{0.96}$, meaning the bookmaker believes each bet will have an average loss rate for bettors of 4%.

We assume $p_2 = \frac{1}{0.96} \frac{1}{3}$, so the bettor views it as having an expected loss rate of 4%. The other prices are based on a range of bookmaker probabilities for the first bet that go from $\pi_1 = 1/3 - 0.2$ to $\pi_1 = 1/3 + 0.2$ while the range for the third bet goes from $\pi_3 = 1/3 + 0.2$ to $\pi_3 = 1/3 - 0.2$. This means the first bet goes from having an expected profit of 40% to an expected loss of 40% with the third bet having the opposite pattern as p_1 increases. Figure 3 shows the fraction of wealth put at risk by the optimal betting strategy for this range of prices for log utility and compares it with the recommended optimal strategy when considering only two outcomes, in this case the Kelly criterion, this time with the restriction $x_i \geq 0$ imposed.

Relative to the Kelly criterion with a non-negativity constraint, there are two key differences. First, the optimal strategy recommends larger amounts be placed on the bets perceived as profitable. For the two cases at the extremes of the graphs, the optimal strategy recommends betting about 8 percent more than recommended by the Kelly criterion. Second, and more strikingly, the optimal strategy recommends betting on the second outcome—which has a fixed expected loss rate of 4%—when p_1 is not close to its true value and for the amount placed on this bet increases as betting positions on the other options rise. For the most extreme gaps between prices and beliefs considered here, the optimal strategy recommends betting about 50 percent more money than the Kelly criterion.

Figure 4 shows how the optimal strategy generates higher expected utility than the Kelly criterion. While the two strategies give almost equal utilities when the prices are not too different from beliefs, at the extremes the differences are larger: For the largest differences between beliefs and prices shown here and for starting wealth of $w = 1$, the optimal strategy gives an expected log of post-betting wealth of 0.134 while the Kelly criterion bets give a value of 0.121.

It may seem counter-intuitive for any optimal strategy to involve placing a bet that you believe will, on average, lose money but the earlier logic that drove our non-uniqueness result still influences the now-unique optimal solution. Bettors decide how much to bet based on the marginal utility they will obtain if the bet wins. In the example in Figure 3, as the amounts staked on the first and third bets increase, the bettor knows that if outcome 2 occurs, they will have lost their large bet on either 1 or 3.

Diminishing marginal utility means that the marginal utility from taking the second bet is higher and so, in many cases, it is deemed worth taking despite its price being higher than the bettor's subjective belief about its probability of success.

Another way to understand the result is to note that expected arithmetic mean returns play no role in the decisions of someone with concave utility. For log utility, it is the geometric mean of wealth that matters. Consider, for example, the case where $p_1 = 1.04(1/3 - 0.2)$ and $w = 1$. In this case, the optimal set of bets is $(x_1, x_2, x_3) = (1.76, 0.32, 0)$. This generates an expected geometric mean of wealth for our bettor of

$$\mu^g = (2.40)^{\frac{1}{3}} (0.96)^{\frac{1}{3}} (0.65)^{\frac{1}{3}} = 1.143 \quad (21)$$

where the three numbers in the calculation are wealth when outcome 1 happens, when outcome 2 happens and when outcome 3 happens. You might think the bettor will be better off cutting out their negative expected value bets on the second option and just betting $(x_1, x_2, x_3) = (1.76, 0, 0)$. However, the geometric mean of wealth in this case is

$$\mu^g = (2.51)^{\frac{1}{3}} (0.76)^{\frac{1}{3}} (0.76)^{\frac{1}{3}} = 1.128. \quad (22)$$

which is lower than for the optimal betting strategy. Not taking the second bet boosted the return in the case of the first and third outcomes but reduced it when the second outcome occurs and this latter effect dominates.

Figures 5 and 6 repeat the same calculations for $\sigma = 2$ (more risk averse than log utility) and $\sigma = 0.5$ (less risk averse). In both figures, we compare the optimal betting strategy to the two-outcome optimal strategy given by equation 5 rather than the Kelly criterion, which is not optimal for these bettors under any circumstances. While the scale of betting is higher for $\sigma = 0.5$ and lower for $\sigma = 2$, the same patterns are evident. Relative to the recommendation of equation 5, the optimal strategy calls for more betting and all agents are willing to take the negative expected value bet on the second outcome once p_1 has departed sufficiently from one-third.

Figure 3: Optimal fraction of wealth placed on bets for $\sigma = 1$ compared with two-outcome recommendations as π_1 goes from 0.139 to 0.555, while π_3 goes from 0.555 to 0.139 and $\pi_2 = 0.347$.

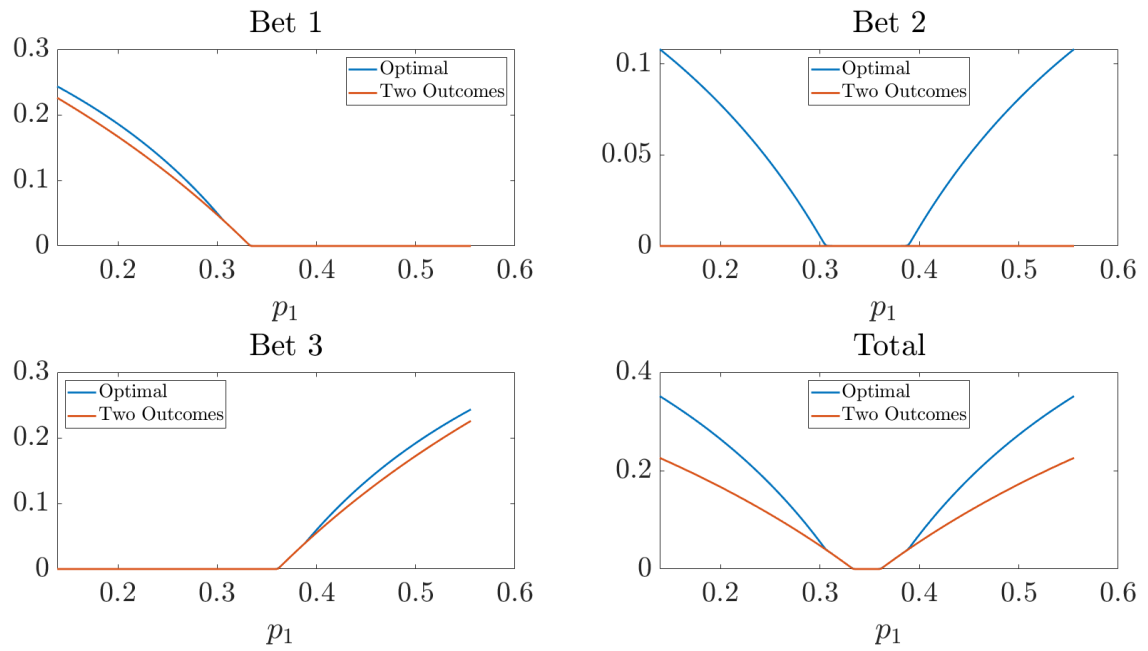


Figure 4: Log of wealth for optimal bets for $\sigma = 1$ compared with two-outcome recommendations as π_1 goes from 0.139 to 0.555, while π_3 goes from 0.555 to 0.139 and $\pi_2 = 0.347$.




Figure 5: Optimal fraction of wealth placed on bets for $\sigma = 2$ compared with two-outcome recommendations as π_1 goes from 0.139 to 0.555, while π_3 goes from 0.555 to 0.139 and $\pi_2 = 0.347$.

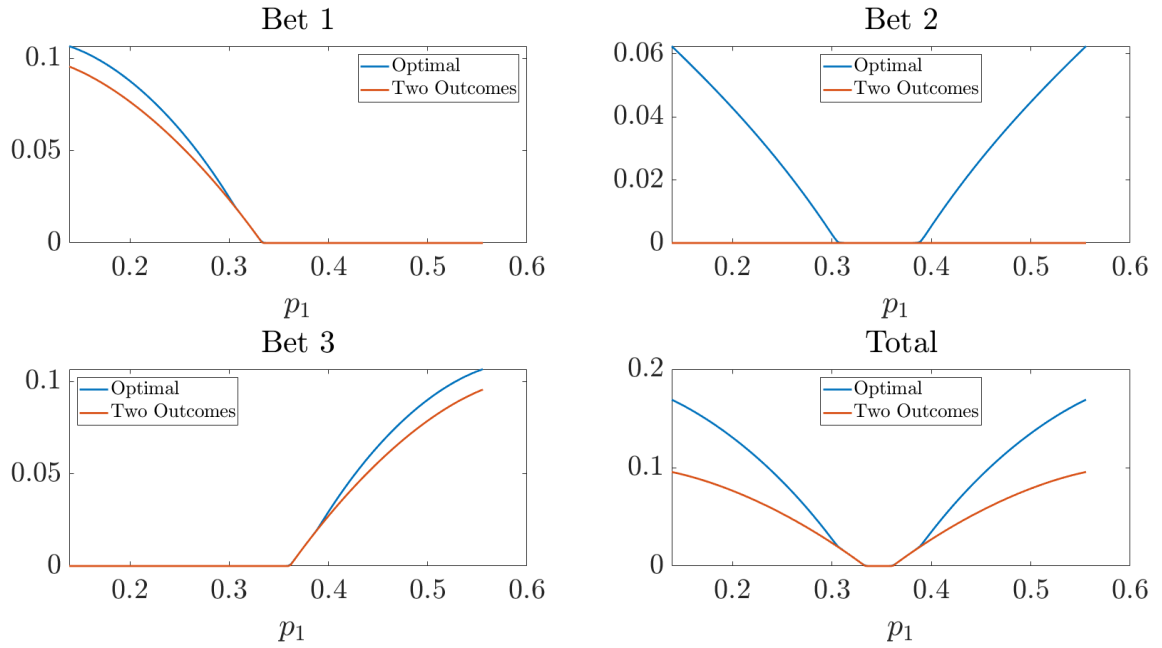
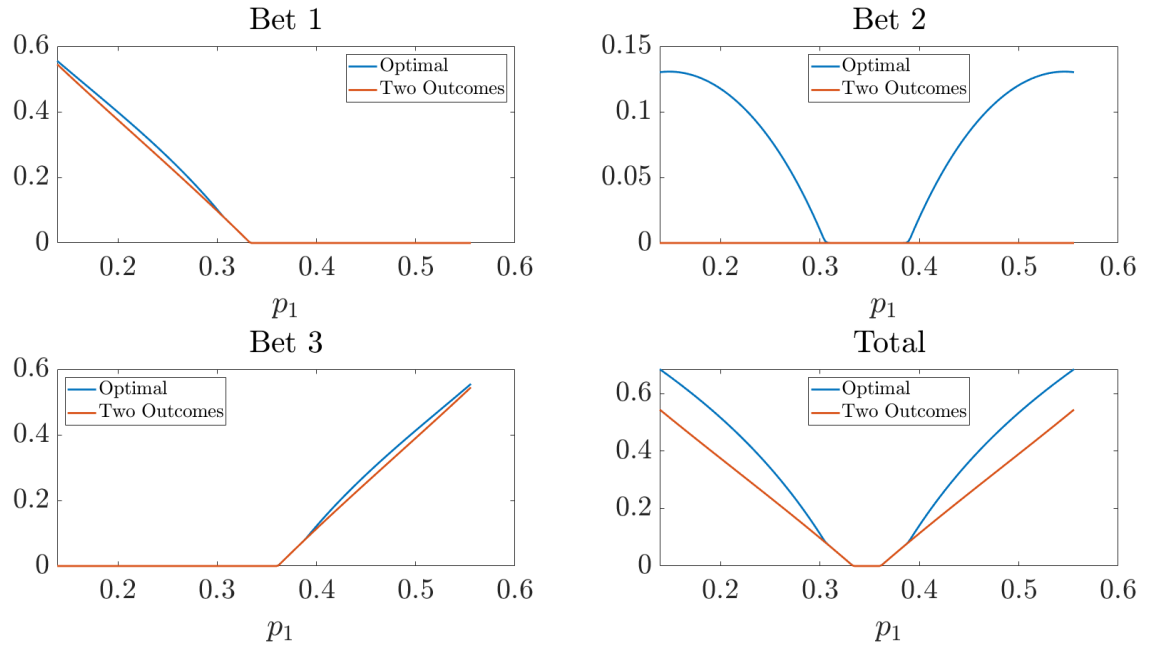


Figure 6: Optimal fraction of wealth placed on bets $\sigma = 0.5$ compared with two-outcome recommendations as π_1 goes from 0.139 to 0.555, while π_3 goes from 0.555 to 0.139 and $\pi_2 = 0.347$.



6. A More Realistic Simulation Exercise

The previous stylized example gives some intuition for the optimal betting strategy but here we will provide a more realistic simulation model to give a sense of how much more aggressive the optimal strategy is than the two-outcome rule, the extent to which bets with negative expected returns will be taken and how these will vary with the magnitudes of the gaps between bettor's beliefs and those implied by the odds.

The only previous study that examines a closely related question is Sung and Johnson's paper on using the generalized optimal rule for betting on UK horse racing. In light of the well-known pattern of favorite-longshot bias in horse racing such that bets on favorites tend to lose less than bets on longshots, they use a logit model to map the odds into better measures of the underlying probabilities. They then use these probability estimates to maximize equation 8 for log utility subject to bets having to be positive. Across the 1,910 races on which they assess this strategy, they report the optimal strategy recommended placing 565 bets on 492 races, so their method mainly recommended not betting and rarely proposed taking multiple bets on the same race. Also, while Sung and Johnson noted that bets with a negative expected return were sometimes recommended, few occurred in practice and they conclude "the Kelly wagering strategy ensures that, in general, only those horses for which the predicted probability of winning exceeds the odds implied probability are bet".

These findings, however, are likely to be specific to Sung and Johnson's method which generally produced estimated probabilities that were close to those implied by the bookmaker's odds. But the fact that sporting events generally attract bettors on all the available options suggests it is common for bettors to have beliefs that differ sufficiently from the probabilities implied by the odds for them to be willing to take a bet. Here, we show that for even modest levels of disagreement between a bettor's subjective probabilities and those implied by the odds, the generalized optimality conditions recommend a much larger scale of betting than when only two outcomes are considered. They also commonly recommend taking on multiple bets and regularly prescribe taking on bets with a negative expected return.

We simulate an event with 5 possible outcomes. Prices are set as a markup on a set of bookmaker's estimated probabilities of the i th outcome occurring

$$p_i = \mu\pi_i \quad i = 1, \dots, 5 \quad (23)$$

where again we set $\mu = \frac{1}{0.96}$. For each simulation, four of the five events were picked at random and the bookmaker's probabilities set as

$$\pi_i = 0.2 + \epsilon_i \quad (24)$$

where $\epsilon_i \sim N(0, \gamma)$. We have used $\gamma = 0.02$ in the simulations reported here. The fifth probability was then set to make the bookmaker's probabilities sum to one. If the sum of the first four randomly

chosen probabilities was over 1, the final probability was set to a small number just above zero and the remaining probabilities were re-scaled so all probabilities summed to one. This calibration means the average event is perceived by the bookmakers as having five equally likely winners but there are variations such that the smallest win probability in large simulations tends to be about 0.05 and the largest to be about 0.35. However, the value of γ does not matter for the results that we present here. This has been introduced for realism rather than its use in illustrating the properties of the optimal strategy.

The beliefs of bettors are then modeled as follows. Beliefs about four of the outcomes are selected at random and set as

$$\tilde{\pi}_i = \pi_i + \eta_i \quad (25)$$

where $\eta_i \sim N(0, \theta)$ and the same technique as before is used to rule out negative subjective probabilities and ensure these probabilities sum to one. This specification means that the bettors on average have the same beliefs as the bookmaker but differ according to the randomly drawn η_i series (at this point, no judgment is being made on which of the parties are more likely to be correct). Figure 7 illustrates a typical pattern of bookmaker and bettor beliefs generated by the values $\theta = \gamma = 0.02$ with a scatter plot of 50,000 beliefs and probabilities from a simulation with 10,000 runs.

Table 1 summarizes the optimal betting strategy for different preferences for risk and for differing values of θ , illustrating the size of the bettor's typical disagreement with the bookmaker. The reported figures are each averaged over 10,000 repetitions. The upper panel shows how much more betting the optimal strategy recommends. For $\theta = 0.01$, the optimal strategy recommends about one-third more betting than the two-outcome rules but this rises rapidly as θ increases. For $\theta = 0.03$ and $\sigma = 2$, the optimal rule recommends almost two and a half times as much betting as the two-outcome rule. The middle panel shows that some of this larger betting comes from placing bets that have negative expected value. For $\theta = 0.01$, these constitute about 10 percent of all bets, rising to about 17 percent for $\theta = 0.02$ and about 21 percent for $\theta = 0.03$. Finally, the bottom panel shows that placing multiple bets is common, with two bets more common than one when $\theta = 0.01$ and $\theta = 0.03$ implying an average of three bets from the five options.

Figure 8 illustrates the optimal betting strategy from 100,000 simulations for log utility by plotting a scatter plot of the fraction of wealth placed at risk for each individual bet and against the bettor's subjective expected return from placing \$1 on these bets, which is $\tilde{\pi}_i - p_i$. It does this for both for the optimal strategy (left panel) and for the Kelly criterion (right panel). The figure illustrates how the optimal strategy is more aggressive both in betting more when bets have positive expected value but also in placing often quite large fractions of wealth on bets with negative expected returns. In fact, the largest absolute betting positions are for bets with expected returns that are either just above or below zero. Inspection of these cases shows they occur when large amounts have been placed on multiple other outcomes. The "twin peaks" pattern in the graph stems from large positions on low

or negative profitability bets being optimal when other big bets have been placed.

Figure 7: Simulated probability beliefs of bookmakers and bettors with $\gamma = \theta = 0.02$ and 10,000 repetitions




Table 1: Characteristics of the optimal betting strategy with $w = 1$ and $\gamma = 0.02$ for three values of θ (standard deviation of difference between bettors' beliefs and true probabilities) and three values of σ . 10,000 repetitions.

Ratio of total amounts at risk relative to the two-outcome rule

	$\theta = 0.01$	$\theta = 0.02$	$\theta = 0.03$
$\sigma = 0.5$	1.29	1.7	1.86
$\sigma = 1$	1.33	1.86	2.19
$\sigma = 2$	1.33	1.95	2.44


Fraction of bets with negative expected returns

	$\theta = 0.01$	$\theta = 0.02$	$\theta = 0.03$
$\sigma = 0.5$	0.102	0.169	0.209
$\sigma = 1$	0.119	0.175	0.212
$\sigma = 2$	0.139	0.185	0.218

Average number of bets placed

	$\theta = 0.01$	$\theta = 0.02$	$\theta = 0.03$
$\sigma = 0.5$	1.65	2.6	3.04
$\sigma = 1$	1.73	2.65	3.03
$\sigma = 2$	1.83	2.68	3.11

Figure 8: Expected returns and fractions of wealth at risk for log utility for both optimal strategy and the Kelly criterion. $\gamma = \theta = 0.02$ and 100,000 repetitions.



7. Implications When Odds Are Accurate

The optimal strategy presented here may be of use to a certain type of person. That person would have to be smart enough to be better than bookmakers at figuring out the probabilities for certain events but not so smart that they already know how to use this information to optimize their utility. One might question whether such people exist but the considerable demand for online Kelly criterion calculators suggest many people consider themselves to be this kind of person.

There is another explanation for this demand which is that there are many people who believe they are better than bookmakers at calculating probabilities but are actually not. Consider betting markets on sports, which have been widely examined by academic research. As noted above, these markets tend not to be perfectly efficient and tend to exhibit a favorite-longshot bias. This has been extensively documented for pari-mutuel betting markets, as summarized by Snowberg and Wolfers (2008), but this pattern has also been found for fixed-odds betting markets where bookmakers set the odds.⁷ However, these inefficiencies relate to some people systematically losing more than others, not to some strategies making profits.

Of course, there is no body of evidence that any reliable system available to the general public exists to systematically make profits betting on sports. I say “available to the general public” because there are some well-known professional gamblers who appear to consistently make profits. The most famous, Tony “the Lizard” Bloom, runs the secretive firm StarLizard which takes money from millionaires to place bets on sports based on complex statistical algorithms. Starlizard places its bets in Asian markets which do not deduct tax from winnings and have low bookmaker’s margins.⁸ Reports suggest that StarLizard look to make average profit margins of 1% to 3% on their bets.⁹ Thus, it seems likely that Bloom’s edge over bookmakers is pretty small. And most people searching the internet for Kelly calculators are not like Tony Bloom, so their chances of having an edge are likely to be very small.

To illustrate what happens when bettors do not actually have an edge over bookmakers, we describe how the different strategies in the simulations just reported will perform if the bookmaker’s probabilities are correct. We report outcomes for the subjective optimal strategy, the two-outcome rule and a benchmark rule of randomly betting 3% of wealth. Table 2 reports these results for log utility. It shows that as the beliefs of bettors become less accurate (meaning θ rises), losses increase for the two-outcome rule but rise even more for our aggressive “optimal” strategy. With $\theta = 0.01$, so beliefs are relatively accurate, losses as a fraction of wealth are worse for the optimal strategy but still relatively small. However, as a simple metric for how badly this strategy performs when $\theta = 0.03$, a

⁷Two recent studies documenting this are Moscovitz and Vasudevan (2022) for betting on US basketball and football games and Hegarty and Whelan (2023) for betting on European soccer.

⁸Whelan (2023) describes how federal income taxation of gambling winnings makes it essentially impossible for anyone to consistently make a post-tax profit betting on sports in the US.

⁹See this story by Business Insider <https://www.thejournal.ie/tony-bloom-starlizard-2597458-Feb2016/>

weekly set of bets placed by someone adopting this strategy will see them lose 26% of their wealth after a year.

These higher losses are not because the more aggressive bettors pick particularly bad bets. None of these bettors have either particularly good or bad insights and the average loss rates across their bets all equal 4% due to the bookmaker's margin, so total losses are driven by how much money is put at risk. With $\theta = 0.03$, the log-utility bettor is placing 14.7% of their wealth at risk, compared with the 6.7% recommended by the Kelly criterion. For reference, the strategy of using 3% of wealth to place bets randomly will lose about 6% over the course of a year with this value of the bookmaker's margin.

Tables 3 and 4 repeat these calculations for CRRA coefficients $\sigma = 2$ and $\sigma = 0.5$, again comparing the optimal strategy with the strategy recommended by equation 5 rather than the Kelly criterion. As expected, losses from all strategies are lower for the more risk-averse bettor with $\sigma = 2$ and higher for the less risk-averse bettor with $\sigma = 0.5$. In the worst case scenario of a bettor with $\sigma = 0.5$ and with inaccurate beliefs such than $\theta = 0.03$, expected losses over the course of a year of weekly bets would be 40% of wealth. We don't report realized utilities in the tables but these utilities are always lower for optimal strategy for maximizing subjective expected utility than for the two-outcome rules.

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		

Table 3: Loss implications for betting strategies with CRRA coefficient $\sigma = 2$, $w = 1$ 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	0.030%	0.147%	0.319%
Average percent of wealth lost for two outcome strategy	0.022%	0.074%	0.129%
Average percent of wealth at risk for optimal strategy	0.75%	3.67%	7.98%
Average percent at risk for two outcome strategy	0.56%	1.86%	3.2%
Expected wealth after 52 bets with the optimal strategy	0.984	0.926	0.847
Expected wealth after 52 bets with the two-outcome strategy	0.978	0.934	0.935

Table 4: Loss implications for betting strategies with CRRA coefficient $\sigma = 0.5$, $w = 1$ 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	0.120%	0.539%	0.976%
Average percent of wealth lost for two outcome strategy	0.093%	0.320%	0.571%
Average percent of wealth at risk for optimal strategy	3.03%	13.47%	24.42%
Average percent at risk for two outcome strategy	2.34%	8.01%	14.26%
Expected wealth after 52 bets with the optimal strategy	0.939	0.755	0.600
Expected wealth after 52 bets with the two-outcome strategy	0.952	0.846	0.743

8. Conclusions

Opportunities to bet on contests where only one of multiple bets can win are common. Beyond the obvious examples in sports betting, it is also possible to bet on events such as which politician will win an election or which movie will win the Oscar for best picture. We have shown how to solve the general problem of how much to bet for people with concave utility when market prices include a bookmaker's margin and it is not possible to "short" any of the options.

We have shown that the optimal betting strategy is more aggressive than suggested by separately considering the two outcomes for each bet of win/don't win. It suggests placing more on bets and also placing often substantial amount of money on bets with negative expected returns due to their values as hedges. In the hypothetical case where bettors have a superior understanding of the probabilities to those that set the odds, this optimal strategy can be used to generate higher utility.

It is important to stress, however, that under the more likely scenario in which a bettor does not have superior insight into the underlying probabilities, the strategies derived here will result in greater losses than from considering just two outcomes. In praising the Kelly criterion, MacLean et al (2011) concluded that "In addition to maximizing the asymptotic rate of long-term growth of capital, it avoids bankruptcy and overwhelms any essentially different investment strategy in the long run." However, they also noted (page 97) in relation to calculating the means and variances that go into the Kelly criterion calculation "What's clear is that getting means correctly estimated is crucial for portfolio success." This is true for the Kelly criterion and it is even more true for the strategies derived here. Caveat emptor.

References

- [1] Hegarty, Tadgh and Karl Whelan (2023). Forecasting Soccer Matches With Betting Odds: A Tale of Two Markets, CEPR Discussion Paper No. 17949.
- [2] Kelly, John L. (1956). "A New Interpretation of Information Rate." *Bell System Technical Journal*, Volume 35, pages 917–926.
- [3] Maclean, Leonard, William T. Ziemba and Yuming Li (2005). "Time to Wealth Goals in Capital Accumulation," *Quantitative Finance*, Volume 5, pages 343–355.
- [4] Maclean, Leonard, Edward Thorp, Yonggan Zhao and William Ziemba (2011). "How Does the Fortune's Formula Kelly Capital Growth Model Perform?" *Journal of Portfolio Management*, Volume 37, pages 96-111.
- [5] Moscovitz, Tobias and Kaushik Vasudevan (2022). Betting without Beta. Yale University Working paper.
- [6] Poundstone, William (2006). *Fortune's Formula: The Untold Story of the Scientific Betting System That Beat the Casinos and Wall Street*, Hill & Wang.
- [7] Samuelson, Paul (1971). "The Fallacy of Maximizing the Geometric Mean in Long Sequences of Investing or Gambling," *Proceedings of the National Academy of Sciences*, Volume 68, pages 2493-2496.
- [8] Samuelson, Paul (1979). "Why We Should Not Make Mean Log of Wealth Big Though Years to Act Are Long," *Journal of Banking and Finance*, Volume 3, pages 305-307.
- [9] Smoczynski, Peter and Dave Tomkins (2010). "An Explicit Solution to the Problem of Optimizing the Allocations of a Bettor's Wealth When Wagering on Horse Races", *Mathematical Scientist*, Volume 35, pages 10-17.
- [10] Snowberg, Erik and Justin Wolfers, (2008). Examining Explanations of a Market Anomaly: Preferences or Perceptions? in *Handbook of Sports and Lottery Markets*, pages 103–136. Elsevier.
- [11] Sung, Ming-Chien and Johnnie Johnson (2010). "Revealing Weak-Form Inefficiency in a Market for State Contingent Claims: The Importance of Market Ecology, Modelling Procedures and Investment Strategies," *Economica*, Volume 305, pages 128-147.
- [12] Whelan, Karl (2023). "US Taxation of Gambling Winnings and Incentives to Bet," *Journal of Gambling Studies* <https://doi.org/10.1007/s10899-023-10189-z>.

- [13] Whitrow, Chris (2007). "Algorithms for Optimal Allocation of Bets on Many Simultaneous Events," *Journal of the Royal Statistical Society. Series C. (Applied Statistics)*, Volume 56, pages 607-623.
- [14] Wright, Margaret (1992). "Interior Methods for Constrained Optimization," *Acta Numerica*, Volume 1, pages 341–407

UCD CENTRE FOR ECONOMIC RESEARCH – RECENT WORKING PAPERS

- [WP22/12](#) Dimitrios Bermpereoglou, Yota Deli, Sarantis Kalyvitis: 'Investment Tax Incentives and Their Big Time-to-Build Fiscal Multiplier' April 2022
- [WP22/13](#) Tadgh Hegarty, Karl Whelan: 'The Wisdom of No Crowds: The Reaction of Betting Markets to Lockdown Soccer Games' April 2022
- [WP22/14](#) Manthos D. Delis, Yota D. Deli, José-Luis Peydró, Adele Whelan: 'Education and Credit: A Matthew Effect' April 2022
- [WP22/15](#) Diane Pelly: 'Worker well-being and quit intentions: is measuring job satisfaction enough?' April 2022
- [WP22/16](#) Kevin Devereux, Margaret Samahita: 'Gender, Productivity, and Promotion in the Irish Economics Profession' June 2022
- [WP22/17](#) Manuel E. Lago, Santiago Lago-Peñas, Jorge Martinez-Vazquez: 'On the Effects of Intergovernmental Grants: A Survey' June 2022
- [WP22/18](#) Judith M. Delaney, Paul J. Devereux: 'Gender Differences in STEM Persistence after Graduation' June 2022
- [WP22/19](#) Karl Whelan: 'Where Do We Stand With "Whatever It Takes"?' July 2022
- [WP22/20](#) Neil Cummins, Cormac Ó Gráda: 'The Irish in England' July 2022
- [WP22/21](#) Yota Deli, Manthos D. Delis, Iftekhar Hasan, Panagiotis N. Politsidis, Anthony Saunders: 'Corporate tax changes and credit costs' July 2022
- [WP22/23](#) Karl Whelan: 'US Taxation of Gambling Winnings and Incentives to Bet' August 2022
- [WP22/23](#) Karl Whelan: 'Risk Aversion and Favorite-Longshot Bias in a Competitive Fixed-Odds Betting Market' August 2022
- [WP22/24](#) Diane Pelly, Orla Doyle: 'Nudging in the workplace: increasing participation in employee EDI wellness events' October 2022
- [WP22/25](#) Margaret Samahita, Kevin Devereux: 'Are economics conferences gender-neutral? Evidence from Ireland' October 2022
- [WP22/26](#) Michele Gubello, Nora M. Strecker: 'Weakly Progressive: Disproportionate Higher Education Attendance and the Structure of Income Taxes' November 2022
- [WP22/27](#) David Madden: 'Mental Health in Ireland During the Covid Pandemic: Evidence from Two Longitudinal Surveys' December 2022
- [WP22/28](#) Ronald B. Davies, Zuzanna Studnicka: 'Tariff Evasion, the Trade Gap, and Structural Trade' December 2022
- [WP23/01](#) Xidong Guo: 'An Analysis of a Rural Hospital's Investment Decision under Different Payment Systems' January 2023
- [WP23/02](#) Emanuele Albarosa, Benjamin Elsner: 'Forced Migration and Social Cohesion: Evidence from the 2015/16 Mass Inflow in Germany' January 2023
- [WP23/03](#) Xidong Guo: 'Hospital Efficiency and Consultants' Private Practices: Analysing the Impact of a Voluntary Reform' February 2023
- [WP23/04](#) Tadgh Hegarty, Karl Whelan: 'Calculating The Bookmaker's Margin: Why Bets Lose More On Average Than You Are Warned' February 2023
- [WP23/05](#) Tadgh Hegarty, Karl Whelan: 'Forecasting Soccer Matches With Betting Odds: A Tale of Two Markets' February 2023
- [WP23/06](#) Karl Whelan: 'How Do Prediction Market Fees Affect Prices and Participants?' March 2023
- [WP23/07](#) Kevin Devereux, Zuzanna Studnicka: 'Local Labour Market Concentration and Wages in Ireland' March 2023
- [WP23/08](#) Kevin Devereux, Zuzanna Studnicka: 'When and Where do Minimum Wage Hikes Increase Hours? Evidence from Ireland' April 2023

UCD Centre for Economic Research
Email economics@ucd.ie