

DISCUSSION PAPER SERIES

DP17518

**Risk Aversion and Favorite-Longshot
Bias in a Competitive Fixed-Odds Betting
Market**

Karl Whelan

FINANCIAL ECONOMICS

CEPR

Risk Aversion and Favorite-Longshot Bias in a Competitive Fixed-Odds Betting Market

Karl Whelan

Discussion Paper DP17518

Published 30 July 2022

Submitted 29 July 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Financial Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Karl Whelan

Risk Aversion and Favorite-Longshot Bias in a Competitive Fixed-Odds Betting Market

Abstract

Research on sports betting markets has generally found a favorite-longshot bias, the empirical pattern for loss rates for bets on longshots to be higher than for favorites, which implies the odds do not reflect the underlying probabilities. The existing literature focuses largely on pari-mutuel betting but favorite-longshot bias is also evident in fixed-odds online betting markets of the type that are growing rapidly around the world. Explanations for this bias in previous work on pari-mutuel markets cannot explain why it would be a feature of competitive fixed-odds betting markets. We show how disagreement among gamblers and risk-aversion on the part of bookmakers in a competitive market can produce a pattern of favorite-longshot bias resembling the empirical evidence.

JEL Classification: D81, G14, L83

Keywords: Sports Betting, Gambling, Favorite-longshot bias, risk aversion

Karl Whelan - karl.whelan@ucd.ie
University College Dublin and CEPR

Risk Aversion and Favorite-Longshot Bias in a Competitive Fixed-Odds Betting Market

Karl Whelan*

University College Dublin

July 29, 2022

Abstract

Research on sports betting markets has generally found a favorite-longshot bias, the empirical pattern for loss rates for bets on longshots to be higher than for favorites, which implies the odds do not reflect the underlying probabilities. The existing literature focuses largely on pari-mutuel betting but favorite-longshot bias is also evident in fixed-odds online betting markets of the type that are growing rapidly around the world. Explanations for this bias in previous work on pari-mutuel markets cannot explain why it would be a feature of competitive fixed-odds betting markets. We show how disagreement among gamblers and risk-aversion on the part of bookmakers in a competitive market can produce a pattern of favorite-longshot bias resembling the empirical evidence.

Keywords: Sports Betting, Gambling, Favorite-Longshot Bias,

JEL Classification: D81, G14, L83

*karl.whelan@ucd.ie. Thanks to Ron Davies, Tadgh Hegarty and Chris Jepsen for helpful comments and Martina Lawless for discussion and inspiration.

1. Introduction

This paper presents a model of competitive fixed-odds betting markets in which bookmakers make offers such as “You win \$3 if your bet wins and lose your \$1 bet otherwise.” Online fixed-odds betting markets have grown rapidly across the world in recent years. This trend in sports betting is nowhere more evident than in the US. A 2018 Supreme Court decision declaring the federal prohibition on sports betting unconstitutional has opened a new market for legal sports betting. By June 2022, thirty states had legalized sports betting in various forms. Betting in these markets is already large, with \$136 billion placed in legal sport betting markets since the Supreme Court ruling, and the market seems likely to grow substantially over the next few years.¹

There is a large theoretical and empirical literature on how sports betting markets behave but most of the empirical research (and almost all of the theoretical work) has focused on pari-mutuel betting. Pari-mutuel betting pools the funds of gamblers and then pays them out (minus a fraction to cover costs and profits) to those who picked the winner in proportion to the size of their bet. This concentration on pari-mutuel markets was because federal restrictions on gambling had provided exceptions for pari-mutuel betting at racetracks so, for many, this was the only legal way to bet on sports in the US.² However, the online fixed-odds markets that are now growing rapidly do not use the price-setting mechanism employed in pari-mutuel markets.

Research on the efficiency of sports betting odds has focused on the favorite-longshot bias. Bets on longshots tend to lose more than bets on favorites, meaning market odds are not unbiased measures of the underlying probabilities. Importantly, this bias has been documented widely for both pari-mutuel and fixed-odds betting markets. We provide an explanation for how favorite-longshot bias emerges from competitive fixed-odds markets due to risk-aversion by bookmakers. Our explanation does not rely on the theoretical devices used to explain favorite-longshot biases in pari-mutuel betting such as some gamblers liking high variance bets or being particularly bad at estimating the probabilities of unlikely events or perhaps just betting randomly. It also does not rely on the presence of insiders with superior knowledge of the outcome of the game to either the bookmaker or the general public, as described in the fixed-odds betting models of Shin (1991, 1992).

We start with a model of risk-neutral bookmakers who set odds to meet a market-determined expected rate of return on bets after covering costs.³ We then consider the implications of bookmakers having some aversion to risk. Our model assumes that gamblers disagree about the probabilities of a team winning a game but the population’s distribution of beliefs centers around the true probability, in the manner first proposed by Ali (1977). Betting markets exist in this model because people have

¹Data from <https://www.legalsportsreport.com/sports-betting/revenue/>

²Of course, there was always an illegal sports betting trade and the possibly of making this market legal and thus a possible source of state tax revenues has been an important motivation for states deciding to legalize sports betting.

³We will use the term bookmaker because this term is generally used in the academic literature on sports betting. In the current US environment, legal betting firms that take bets on sports are generally called sportsbooks

differing assessments of the underlying probabilities.

We use the risk-neutral bookmaker version of the model to illustrate how various explanations for favorite-longshot bias in pari-mutuel betting would not generate such bias in competitive fixed-odds betting markets. This version of the model also generates an interesting prediction. Competitive betting markets in which the odds include a margin to cover costs and/or obtain a required profit rate will take more bets on longshots than favorites. To explain this result, consider a betting market with no administrative costs and a zero profit requirement for a game where the home team has a probability p of winning. gamblers with subjective probability belief \tilde{p} will take a bet on the home team with odds O_h (meaning it pays out $\$O_h$ on a $\$1$ bet) as long as $\tilde{p} \geq \frac{1}{O_h}$. Competitive markets will set odds of $O_h = \frac{1}{p}$ on the home team and $O_a = \frac{1}{1-p}$ on the away team. Potential gamblers will bet on the home team if they have $\tilde{p} \geq p$ and on the away team otherwise. All potential gamblers choose to place bets and with a symmetric distribution of beliefs around the true probability, an equal number of bets are placed on both teams.

Introducing a profit margin per bet θ and bookmaker cost per bet of μ means competitive odds change to $O_h = \frac{1-\mu-\theta}{p}$ and $O_a = \frac{1-\mu-\theta}{1-p}$ and the condition for betting on home becomes $\tilde{p} > \frac{p}{1-\mu-\theta}$ and the condition for betting on the away is $1 - \tilde{p} > \frac{1-p}{1-\mu-\theta}$. This means some potential gamblers do not bet. When $p = 0.5$ and $\mu + \theta = 0.05$, these belief thresholds are $[0.47, 0.53]$ and symmetry of beliefs means equal amounts are still bet on both teams. But when $p = 0.8$, so the home team is a strong favorite, belief thresholds change to $[0.79, 0.84]$. The introduction of the bookmaker's margin squeezes out more bets on the favorite than the longshot.

Intuitively, introducing margins multiplies the odds downwards and this multiplies upwards the threshold levels of beliefs needed to place bets. Without margins, those who believe there is at least an 80% chance of the favorite winning will go for that option and those who believe there is at least a 20% chance of the longshot winning will go with that option. Introducing margins multiplies these figures but since the threshold probability for those who bet on longshots is smaller, fewer gamblers are crowded out relative to potential longshot gamblers.

In the second part of the paper, we introduce risk aversion. Risk-neutral competitive bookmakers take more bets on longshots as the probability of the favorite winning rises but their profits on longshot bets have higher variance. Maintaining a free entry requirement but introducing a preference for lower variance profits, a competitive market's equilibrium will have worse odds on longshots and better odds on favorites than in the risk-neutral case. The ultimate risk-avoidance strategy is to quote odds inversely proportional to betting volumes, thus reproducing the outcome of pari-mutuel betting where the bookmaker's profit rate has zero variance, which Ali (1977) showed implies favorite-longshot bias. Adopting the more realistic assumption that bookmakers are willing to trade off risk and return, we illustrate how this still likely produces favorite-longshot bias. Incorporating bettors with realistic levels of risk-aversion for small-stake bets is also consistent with this result.

2. The Favorite-Longshot Bias

The evidence on favorite-longshot bias in pari-mutuel betting dates back to Griffith (1949) who noted that while actual win rates for horses with short odds corresponded reasonably well with the probabilities implied by betting odds, this was not the case for horses with longer odds, so gamblers on average lost more money if they picked longshots rather than favorites. This finding has since been corroborated in numerous studies of racetrack betting.⁴

Our focus, however, is on fixed-odds markets. There are fewer studies of the bias in these markets, both because their legality was limited by US federal laws and because, when offered legally they have often taken the form of “points spread” bets that tend to roughly equalize the chances of each side of the bet winning. Still, there are a number of studies that have found evidence of favorite-longshot bias in fixed-odds betting markets, with examples including UK horse races (Gabriel and Mardsen, 1990, Snowberg and Wolfers, 2010), US college football and basketball (Berkowitz, Depken and Gandarc, 2017), European soccer (Cain, Law and Peel, 2000 and Buhagiar, Cortis and Newell, 2018) and tennis (Forrest and McHale, 2007).

Tennis provides a useful example in the context of this paper both because it fits with our model’s structure in which there are two contestants and one winner (there are no ties in tennis) and because perceived mis-matches where there is a strong favorite are common, thus giving a wide range of market-implied probabilities. Figure 1 illustrates the pattern of favorite-longshot bias in professional tennis by illustrating the average payout on a \$1 bet for various market-based probabilities. If the odds are that a \$1 bet earns $\$O_a$ if A wins and $\$O_b$ if B wins, we calculate the market probabilities as


$$P_i = \frac{1/O_i}{1/O_a + 1/O_b} \quad (1)$$

Data are from www.tennis-data.co.uk and cover all 56,004 matches played in the men’s ATP and women’s WTA professional tournaments between 2011 and June 2022. Each point in the figure refers to one of twenty quantiles of the distribution of odds-implied win probabilities. If betting odds were set according to the actual probabilities of winning, loss rates should be equal across all bets. However, as found by Griffith many years ago, we can see that loss rates on tennis betting increase as the odds-based measure probability of the player winning falls, with particularly large loss rates for extreme longshots. We will show that this pattern can be reproduced by the introduction of risk aversion for bookmakers in a competitive betting market.

There is also a large theoretical literature on favorite-longshot bias, summarized comprehensively by Ottaviani and Sørensen (2008). We will describe some of the theories as we discuss our model, noting only for now that most of this theoretical work has focused on pari-mutuel markets.

⁴Snowberg and Wolfers (2010) provide an excellent summary of the findings of this literature

Figure 1: Average payout rates for a \$1 bet on professional tennis matches



3. A Competitive Fixed-Odds Betting Market with Unbiased Beliefs

Here we describe a competitive fixed odds betting market with risk-neutral bookmakers and risk-neutral gamblers who disagree about the underlying probabilities for the outcome of a game.

3.1. Competitive fixed odds

A game features two teams, home and away. The outcome is either home wins or away wins and the probability the home team wins is p . Bookmakers offer odds of O_h and O_a so that a \$1 bet on home returns O_h if the home team wins and a \$1 bet on away returns O_a if the away team wins. These odds are fixed in the sense that once a bet is placed, the amount earned for winning the bet is fixed and does not depend on the subsequent actions of other gamblers. We assume bookmakers know the true value of p (we will return to this assumption later).

There are a number of bookmakers who compete against each other for business and the competitive equilibrium requires each firm to make an expected profit rate of θ on each bet, reflecting the need to earn a “normal” rate of profit. If expected profit rates are higher than this, new bookmakers enter the market and offer lower odds to bring the profit rate back to θ . Costs associated with research, taxation and administration are a fraction μ of the total amount of bets. The competitive equilibrium profit rate of θ implies the following equality for revenue minus costs on a \$1 bet

$$1 - \mu - pO_h = \theta \quad (2)$$

so the equilibrium odds on the home team are

$$O_h = \frac{1 - \mu - \theta}{p} \quad (3)$$

and similarly odds for the away team are

$$O_a = \frac{1 - \mu - \theta}{1 - p} \quad (4)$$

These odds fully reflect the underlying probabilities with the ratio of odds given by

$$\frac{O_h}{O_a} = \frac{1 - p}{p} \quad (5)$$

Without describing yet how potential gamblers behave in our model, we can note already that the most popular theoretical devices used to explain favorite-longshot bias in pari-mutuel markets would not change the odds offered in this market. Gamblers may have locally risk-loving preferences and enjoy making high-variance low-probability bets, as suggested by Weitzman (1965). Alternatively, gamblers may systematically overestimate the probabilities of unlikely events, as suggested by Kahneman and Tversky's (1979) prospect theory. There could be "noise gamblers" who bet randomly on the two teams, as modeled by Hurley and McDonough (1995) or perhaps gamblers who choose to back the team they are a supporter of.

In a pari-mutuel setting, each of these features drive additional betting volume towards the longshot and since the pari-mutuel structure requires winning odds to be inversely related to betting volumes, this reduces the winning odds on the longshot team and increases them on the favorite. However, these factors would have no influence on odds in a competitive market because these odds are not determined by betting volumes. In each case, these mechanisms imply the existence of gamblers who would back the longshot at lower odds than the competitive prices. But getting those gamblers to accept these poorer odds would generate rates of return for bookmakers above θ and incentivize other bookmakers to undercut those making these excess profits. People who enjoy high variance bets or over-estimate the probabilities of unlikely events or like to back their home team are still capable of checking which bookmaker offers the best odds. Provided they do so, a competitive market with risk-neutral bookmakers will set the odds that reflect the underlying true probabilities.

3.2. Beliefs of gamblers and betting volumes

We now turn to how betting decisions are taken in our model. There is a continuum of risk-neutral potential gamblers of size 1 who can choose to bet or not. When they do bet, they place equal sized bets, normalized to equal one. Potential gamblers have a subjective belief \tilde{p} about the true probability p that the home team wins. As in Ali (1977), beliefs are characterized by a cumulative distribution

function $F(\tilde{p})$ with mean p , meaning the gamblers are, on average, correct about the probability of a home win and beliefs are symmetrically distributed around p . Potential gamblers are risk-neutral and only bet if the odds imply an expected return greater than or equal to zero and will place their bets with the bookmaker offering the best odds on a team.

We assume gamblers differ in their beliefs about the likely outcome of the game but these differences are not related to inside information that is correlated with the outcome of the game. This contrasts, for example, with the models of Ottaviani and Sørensen (2009, 2010), in which gamblers receive a signal drawn from a distribution that changes depending on the state of the world, where in one state of the world the home team wins and in the other it loses. Much of the early literature on sports betting focused on horse racing, where those who work in the industry may have inside information on how well horses can run. However, given the extensive media coverage of the major league sporting events that attract most betting, it seems unlikely that many people have useful inside information about these games. When a team wins the Super Bowl with a last second interception, the gamblers in our model see it as a random event of the type that makes people love sport rather than one they had been given a signal about before the game.

Betting volumes in the model are determined as follows. Potential gamblers with belief \tilde{p} will bet on home if

$$\tilde{p}O_h \geq 1 \quad (6)$$

and they will bet on away if

$$(1 - \tilde{p})O_a \geq 1 \quad (7)$$

At competitively-priced odds, those with a sufficiently high belief in the probability of a home win ($\tilde{p} > \frac{p}{1-\mu-\theta}$) will bet on the home team and those with a sufficiently low belief ($\tilde{p} < 1 - \frac{1-p}{1-\mu-\theta}$) will bet on away. Thus, demand for bets on the home and away teams will be

$$D(O_h) = 1 - F\left(\frac{1}{O_h}\right) = 1 - F\left(\frac{p}{1-\mu-\theta}\right) \quad (8)$$

$$D(O_a) = F\left(1 - \frac{1}{O_a}\right) = F\left(1 - \frac{1-p}{1-\mu-\theta}\right) \quad (9)$$

Some potential gamblers, those with

$$1 - \frac{1-p}{1-\mu-\theta} < \tilde{p} < \frac{p}{1-\mu-\theta} \quad (10)$$

will chose not to bet so the total volume of bets will be

$$V = 1 - F\left(\frac{p}{1-\mu-\theta}\right) + F\left(1 - \frac{1-p}{1-\mu-\theta}\right) \quad (11)$$

If $\mu + \theta = 0$, then everyone bets and the symmetry of the distribution of beliefs means an equal amount of bets are placed on each team. As $\mu + \theta$ increases above zero, the quantity of bets declines and the amount of bets placed on the favorite and the longshot becomes unequal. More bets are placed on the longshot than on the favorite and this asymmetry increases as both $\mu + \theta$ and p get larger.

As discussed in the introduction, the intuition for these results is relatively simple. When $p = 0.5$ and there are no bookmaker costs or profits, those with \tilde{p} beliefs above 0.5 bet on the home team and the identical number of people with beliefs below 0.5 bet on the away team. Introducing a positive value of $\mu + \theta$, then the belief thresholds for the $p = 0.5$ game become $\left[1 - \frac{0.5}{1-\mu-\theta}, \frac{0.5}{1-\mu-\theta}\right]$. The fractions of potential gamblers above and below these thresholds remain equal because the distribution is symmetric around the mean. However, suppose $p = 0.8$ so the home team is a strong favorite. Now the thresholds are $\left[1 - \frac{0.2}{1-\mu-\theta}, \frac{0.8}{1-\mu-\theta}\right]$ which are no longer equally distant from the mean. This tilt in volume towards longshots can be large. With $p = 0.8$ and $\mu + \theta = 0.05$, the belief thresholds are $[0.789, 0.842]$. For any reasonable probability distribution for beliefs with a mean of 0.8, there will be a lot more probability mass between 0.8 and 0.842 than there will be between 0.789 and 0.8. This result is not driven by the profit rate requirement, since this pattern of betting volumes would also occur in a zero profit betting market provided bookmakers incur administrative and taxation costs.

To provide a tractable example, consider the case where \tilde{p} is distributed uniformly on $[p - \sigma, p + \sigma]$. In this case, the parameter σ indicates the extent of disagreement among potential bettors. For this specification of the distribution of beliefs, the demand functions generating the total amount of home and away bets are

$$D(O_h) = 1 - \frac{1}{2\sigma} \left[\frac{1}{O_h} - p + \sigma \right] \quad (12)$$

$$D(O_a) = \frac{1}{2\sigma} \left[1 - \frac{1}{O_a} - p + \sigma \right] \quad (13)$$

This means the total volume of betting is given by

$$D(O_h) + D(O_a) = 1 - \frac{1}{2\sigma} \left(1 - \frac{1}{O_h} - \frac{1}{O_h} \right) \quad (14)$$

The total volume of bets does not depend on p but does depend negatively on the extent to which the sum of the inverses of the odds exceeds one. In sports betting circles, the amount by which the sum of inverses of the odds exceeds one is known as the “overround.” This is the most common way of calculating the bookmaker’s gross profit margin on a game. In this model, a higher overround leads to more gamblers concluding they will on average lose money and so fewer bets are placed.

We can solve for betting volumes on the two teams by inserting the equilibrium odds into these

demand functions

$$D(O_h) = 1 - \frac{1}{2\sigma} \left[\frac{p}{1 - \mu - \theta} - p + \sigma \right] = \frac{1}{2} - \left(\frac{\mu + \theta}{1 - \mu - \theta} \right) \frac{p}{2\sigma} \quad (15)$$

$$D(O_a) = \frac{1}{2\sigma} \left[1 - \frac{1 - p}{1 - \mu - \theta} - p + \sigma \right] = \frac{1}{2} + \left(\frac{\mu + \theta}{1 - \mu - \theta} \right) \frac{p - 1}{2\sigma} \quad (16)$$

Summing these, we get total betting volume as

$$V = D(O_h) + D(O_a) = 1 - \frac{\mu + \theta}{2\sigma(1 - \mu - \theta)} \quad (17)$$

As would be expected, volumes depend negatively on $\mu + \theta$ because these terms raise the bookmaker's overround. Volume depend positively on σ , which measures the extent of disagreement: More disagreement means more gamblers who have strong enough beliefs in either the home or away team winning that they believe they can cover the bookmaker's margin. The share of bets placed on the home team is

$$s_h = \frac{1}{2} + \frac{\mu + \theta}{2\sigma(1 - \mu - \theta) - \mu - \theta} \left(\frac{1}{2} - p \right) \quad (18)$$

As the level of disagreement falls, the share of bets on longshots increases. To understand this result, consider again the case where $p = 0.8$ and belief thresholds are $[0.789, 0.842]$. If $\sigma = 0.1$ so beliefs are uniform on $[0.7, 0.9]$, then 29% of gamblers go for the favorite and 44% back the longshot. If we reduce disagreement to $\sigma = 0.05$ so beliefs are uniform on $[0.75, 0.85]$, then only 8% of gamblers back the favorite while 39% back the longshot.

Figures 2 and 3 illustrate the behavior of betting volumes. Figure 2 describes the model with uniform beliefs and $\sigma = 0.1$ for different values of $\mu + \theta$. It shows that even relatively modest bookmaker margins can induce a significant imbalance in betting volumes when the probability of the favorite winning is high. Figure 3 describes the model with $\mu + \theta = 0.05$ and various values of σ . The share of bets on the favorite can drop dramatically as the level of disagreement falls. The patterns shown here are not driven the uniform distribution specification. Solving the model with normally distributed beliefs produces very similar results.

Figure 2: Betting volumes as bookmaker cost per bet plus required return ($\mu + \theta$) rises ($\sigma = 0.1$)

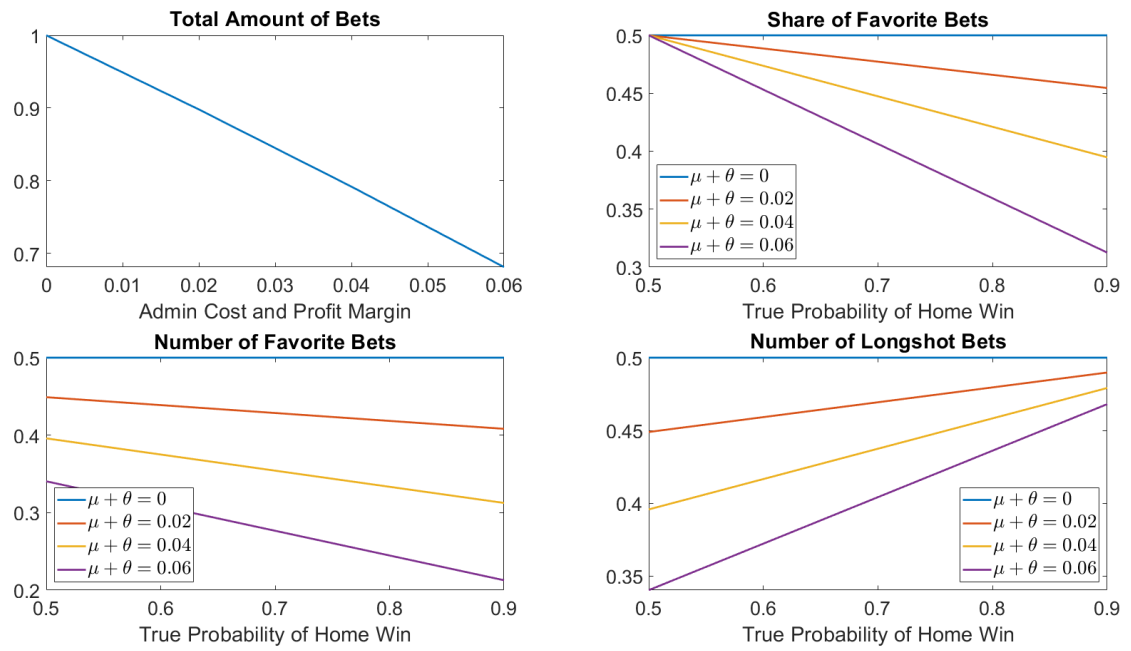
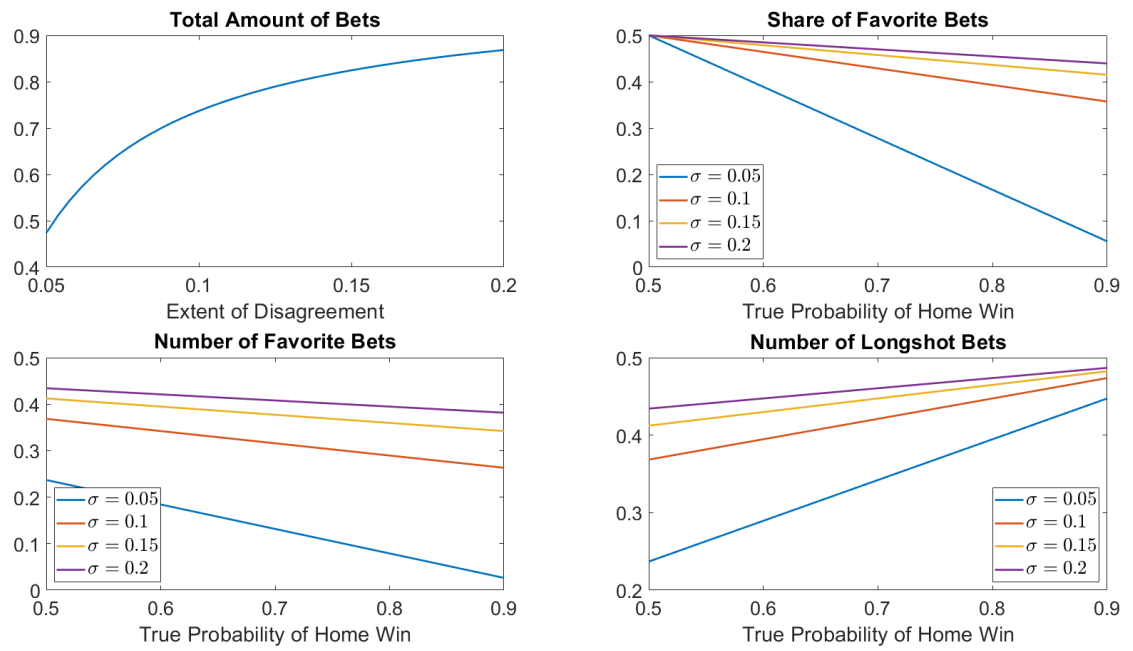


Figure 3: Betting volumes as disagreement among gamblers, σ , rises ($\mu + \theta = 0.05$)



3.3. Uncertainty about probabilities and out-of-equilibrium behavior

We have assumed that bookmakers know the correct value of p . It is useful, however, to consider what would happen to a bookmaker who did not know p and posted odds different from the ones above. If all bookmakers post the competitive equilibrium odds, then gamblers are indifferent between them and betting volumes for specific bookmakers are indeterminate. However, the model is clear about what would happen to volumes if a bookmaker did not quote the competitive equilibrium odds. Suppose a bookmaker offered odds with an implied home win probability p^b that differed from the true value p . In other words, they post odds of

$$O_h = \frac{1 - \mu - \theta}{p^b} \quad (19)$$

$$O_a = \frac{1 - \mu - \theta}{1 - p^b} \quad (20)$$

while other bookmakers continued quoting the competitive equilibrium odds. In this case, if $p^b < p$, then the bookmaker will attract $F\left(\frac{1-\mu-\theta}{p^b}\right)$ bets on the home team and no bets on the away team (because the competitive odds prices on away are better). Expected profit on these bets will equal

$$\left[(1 - \mu) - p \left(\frac{1 - \mu - \theta}{p^b} \right) \right] F \left(\frac{1 - \mu - \theta}{p^b} \right) \quad (21)$$

The profit rate equals the required rate of θ when $p_b = p$ and depends positively on p_b . This means that as p_b falls below p , the profit rate falls further below the required rate and the bookmaker will eventually incur losses. These losses increase nonlinearly as the bookmaker's implied probability falls further below the true level because the odds will attract even more gamblers and the payouts on those bets when the home team win will be more generous.

Alternatively, if $p^b > p$, then the bookmaker will attract no bets on the home team (because the competitive odds prices on home are better) but will get $1 - F\left(1 - \frac{1-\mu-\theta}{1-p^b}\right)$ bets on the away team. Expected profit will be

$$\left[1 - \mu - (1 - p) \left(\frac{1 - \mu - \theta}{1 - p^b} \right) \right] \left(1 - F \left(1 - \frac{1 - \mu - \theta}{1 - p^b} \right) \right)$$

The same argument shows that the profit rate will fall as p_b rises above p .

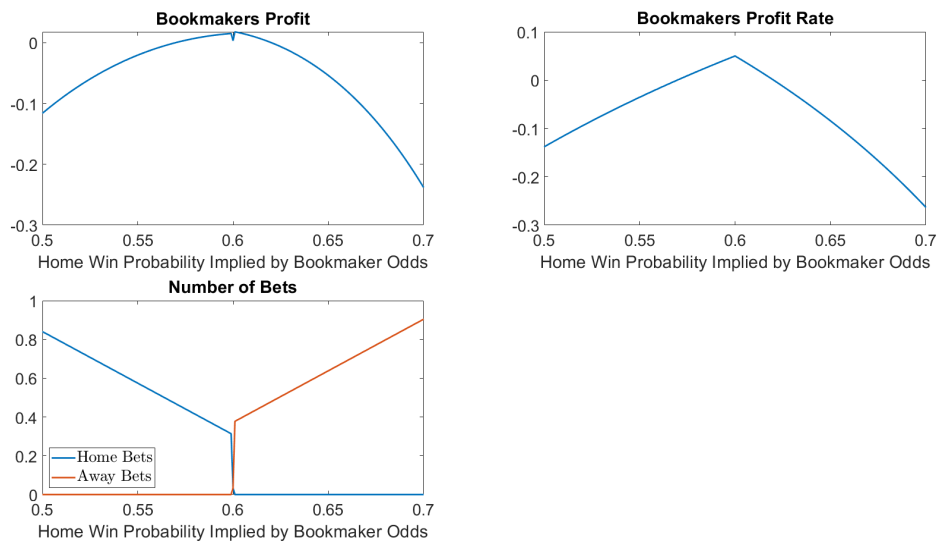
Figure 4 illustrates profits and betting volumes for a bookmaker posting odds that differ from competitive equilibrium odds. The x -axis shows various values of the implied probability from the bookmakers odds, with the true probability assumed to be $p = 0.6$. We assume beliefs are uniform on $[0.5, 0.7]$, set $\mu + \theta = 0.06$ and assume there are 10 bookmakers and that bets are split equally among them when they all set the competitive equilibrium odds. A useful result, in light of the extension of the model in the next section, is that losses incurred by a bookmaker quoting odds that differ from

competitive equilibrium are asymmetric. The bookmaker loses more on longshot bets when it underestimates the longshot's chance of winning than it does on favorite bets when it underestimates the favorite's chance of winning. If $\mu + \theta = 0$, the bookmaker that estimates $p = 0.7$ when the true $p = 0.6$ is willing to offer odds of $O_a = 3.33$ on the longshot when the equilibrium odds are $O_a = 2.5$. This involves a greater loss than the case where they estimate $p = 0.5$ leading them to offer odds of $O_a = 2$, when the equilibrium odds are $O_a = 1.67$. Payouts for the generously-priced bets are 33% too high in the first case and 20% too high in the second case, even though the extent of mis-estimation of the true probability is the same. We will return to the risk inherent in longshot bets in the next section.

These calculations illustrate how a competitive fixed-odds betting market can aggregate the public's information efficiently. If a bookmaker does not know the correct probability, the large volume of bets triggered by initially misaligned prices (and concern about the losses that taking high volumes of bets on one side can entail) will likely trigger them to adjust odds. Indeed, Ed Miller and Matthew Davidow's (2019) fascinating book, *The Logic of Sports Betting*, describes the odds formulation process of real-world online bookmakers in exactly these terms. They stress that bookmakers often place restrictions on betting volumes during the early "price discovery" stage of setting odds, thus restricting potential losses due to posting odds based on inaccurate probabilities and using the pattern of early betting to adjust odds. These forces tend to encourage competitively set odds to accurately reflect the underlying probabilities even in the absence of extensive research by bookmakers.

One can generalize the point made here. Even if no bookmakers actually know the true value of p , setting prices that differ from those implied by the average belief of the public will attract large volumes of bets placed on one team in a game. Once risk management is considered, it is unlikely that many bookmakers will want to set prices that differ significantly from their competitors.

Figure 4: Profit rates for a bookmaker setting odds based on an incorrect probability. True $p = 0.6$



Notes: Assumes a true home win probability of $p = 0.6$, uniform beliefs among gamblers over $[0.5, 0.7]$, $\mu + \theta = 0.06$ and one-tenth of bets placed with bookmaker when its implied probability equals the true probability.

3.4. On beliefs and disagreement

Two other issues are worth discussing before we introduce risk aversion for bookmakers.

First, in this risk-neutral-bookmakers version of our model, disagreement among gamblers influences the pattern of betting volumes but does not, in itself, generate favorite-longshot bias. This contrasts with Ali's (1977) demonstration that disagreement in beliefs of the form we are assuming was sufficient to produce favorite-longshot bias in a pari-mutuel betting market in which odds and volumes are simultaneously determined.⁵ The explanation for Ali's result is as follows. For pari-mutuel odds to reflect the true probability that a favorite will win, it is necessary for the ratio of betting on the favorite to betting on the longshot to be proportional to the teams' chances of winning. The competitive market odds described above will not deliver these volumes. Indeed, once $\mu + \theta > 0$, the competitive market sees more bets on the longshot than the favorite. To be consistent with more people betting on the favorite and fewer to on the longshot, the pari-mutuel odds offered on the favorite have to be better than the competitive odds, and the odds on the longshot have to be worse.⁶ While we use a similar formulation of beliefs to Ali (1977), our explanation for favorite-longshot bias in a fixed-odds market, based on risk aversion from bookmakers, is quite different.

Second, an obvious question to ask about the gamblers in our model is why they don't use the market price to infer the correct odds. Shouldn't those who have a strong belief that one team will win change their mind when the bookmaker's odds show the team is actually a longshot? One modeling device that could address this is to have gamblers accept that the market price summarizes the public's beliefs about the underlying probabilities but not accept that the public's distribution of beliefs is centered on the true probability. Instead, gamblers could believe the information they have used to formulate their initial beliefs still retains value as a signal even after seeing the market odds.

More generally, however, the question is perhaps best answered by the very existence of betting markets. Bookmakers make profits. Accepting market prices as fairly representing the true probabilities means accepting that all bets have a negative expected return. And most people do not bet, most likely for that reason. But some people do bet. This may partly be fueled by the utility they get from placing a bet, perhaps in a way that increases their enjoyment of watching sporting events. But even then, there is the question of why they choose to bet on one team rather than the other. When one person bets on the home team and another person bets on the away team, the explanation for their differing decisions is likely to be that they disagree with each other about the probability of their picks winning.

⁵In practice, gamblers do not know what the final odds will be in a pari-mutuel market. However, pari-mutuel betting operators display indicative odds based on betting volumes and update these figures frequently, often as quickly as every 30 or 45 seconds. While a flood of late bets on a particular competitor can change the odds from the indicative ones that were displayed, it is probably closer to true than false to say that those betting in pari-mutuel markets have a good idea of their odds.

⁶Quandt (1986) also derived a similar result though his derivation assumed that gamblers are risk lovers which is not a necessary assumption.

4. Risk Aversion for Bookmakers

We have assumed so far that bookmakers are risk-neutral but there is plenty of evidence that real-world bookmakers are concerned about risk. For example, gambling expert Joseph Buchdahl (2016) documents many different practices that bookmakers adopt to manage their risks.⁷ Adjusting odds on popular bets downwards to limit exposure is the most obvious technique—and this is what will emerge from our model—but there are many other aspects to risk management by bookmakers, including setting limits on how much can be bet at certain times and on certain teams as well as profiling customers and shutting accounts of “sharp” gamblers who are earning high profits.⁸

In this section, we discuss the risk to bookmakers of offering bets on longshots and then describe the joint risk inherent in offering bets on both teams in a game. We derive a formula for odds that generate a market-determined rate of return on each game while having zero variance in profit rates and illustrate how these odds exhibit favorite-longshot bias. Finally, we discuss how odds are set when competitive bookmakers are willing to trade off higher risk for higher returns and explain why the outcome of this process is still likely to produce favorite-longshot bias.

4.1. Variance of profits on bets and games for competitive risk-neutral bookmakers

Consider a \$1 bet with odds of O on a team that has probability p of winning and is expected to earn a profit of θ for the bookmaker, who incurs costs per dollar bet placed of μ . The variance of bookmaker profits on this bet is

$$\Sigma_p = \frac{1}{2} \left[p(1 - \mu - \theta - O)^2 + (1 - p)(1 - \mu - \theta)^2 \right] \quad (22)$$

Inserting the competitive equilibrium odds for this bet, this becomes

$$\Sigma_p = \frac{(1 - \mu - \theta)^2}{2} \left(\frac{1}{p} - 1 \right) \quad (23)$$

so the variance of profits rise as p falls. If betting markets operated with bookmakers competing against each other to offer standalone bets on individual teams, all priced precisely to provide a risk-adjusted required rate of return, then this result would be sufficient to produce favorite-longshot bias. At competitive equilibrium odds, longshot bets offer the same expected return as bets on favorites but generate higher variance. To induce risk-averse bookmakers to offer longshot bets, the odds for longshots should be lower than those for favorites.

However, the idea of completely separate markets with free entry conditions for each bet is not a good description of the actual betting market. Real world bookmakers offer bets on both teams in a

⁷Buchdahl is also the person behind the tennis data website used to do our calculations in Section 2.


⁸This is one justification for our approach of assuming gamblers are symmetric and there are no gamblers with superior information. If these sharper gamblers enter the market and start to make large profits, they tend to be shown the door by modern online bookmakers.

game and the variance due to making profits or losses on one team will be at least partially offset by the profits and losses earned on the other team. Rather than price bets on longshots as stand-alone products, bookmakers will likely consider the combined risk profile of the bets on both sides together to assess the overall amount of risk being taken on each game.

Figure 5 illustrates how the riskiness of a competitive risk-neutral bookmaker's profits on a game depends upon the extent to which one side is a favorite, again using the uniform distribution for beliefs. The role played by offsetting risks is illustrated in the top-right panel. The standard deviation of profits per game for a bookmaker taking \$1 bets on both teams is always less than the standard deviation of profits earned on taking only the longshot bet and if the game is a toss-up ($p = 0.5$), the variance of the profit rate is zero because the payout is the same no matter which team wins. However, as the probability of the favorite winning increases, the strategy of taking \$1 bets on both teams becomes increasingly risky. The rise in the variance of profits on bets on the longshots outweighs the smaller decline in variance of profits on the favorite.

We have also shown that competitive bookmakers will not take equal bets on both sides and the bottom-left panel shows how the increased amount of bets taken on longshots as p rises produces a large increase in total risk associated with longshot bets. This is only partly offset by the declining risk associated with offering bets on favorites. The bottom-right panel illustrates how the shift in the share of bets towards longshots leads to the variance of profits per bet increasing faster than they would if they were taking an equal volume of bets on both teams.

Figure 5: Betting volumes and standard deviation of profits ($\mu = 0.01, \theta = 0.04, \sigma = 0.05$)



Notes: Uniform distribution of beliefs over $[p-\sigma, p+\sigma]$ is assumed. “Both Bets Offered” illustrates the standard deviation to a bookmakers profits if they offer \$1 bets on both teams. “Equally-Weighted” shows the standard deviation of the profits of the market of competitive bookmakers when they take the same volume as occurs in equilibrium but instead have equal volumes of betting on each side.

4.2. Zero variance odds

Our calculations suggest that risk-averse bookmakers will not offer competitive equilibrium odds when one team is a strong favorite. Even allowing for the offsetting variance due to favorites winning, the possibility of losing long-odds bets raises variance and makes games with unbalanced probabilities unattractive. How could bookmakers adjust their strategy to mitigate this risk?

A common theme in gambling lore is that bookmakers should manage risk by taking equal volumes of bets on both sides of a game. The toss-up game example above seemed to conform to that claim. However, in fixed-odds bookmaking, a balanced amount of bets placed on each side only eliminates risk if the same odds are offered on both teams, which is not the case in most fixed-odds betting situations. In fact, the toss-up game example was a special case of a different result. The variance of a bookmaker's profit rate per game, given an expected profit rate of θ , is given by

$$\Sigma_p(\Pi_p) = \frac{1}{2} \left[p[(1 - \mu - \theta)V - s_h V O_h]^2 + (1 - p)[(1 - \mu - \theta)V - (1 - s_h) V O_a]^2 \right] \quad (24)$$

where V is the total volume of bets and s_h is the fraction placed on the home team. From this, we can see the variance of profits on the game will equal zero, while still achieving an expected rate of return of θ if

$$O_h = \frac{1 - \mu - \theta}{s_h} = \frac{(1 - \mu - \theta)(D(O_h) + D(O_a))}{D(O_h)} \quad (25)$$

$$O_a = \frac{1 - \mu - \theta}{1 - s_h} = \frac{(1 - \mu - \theta)(D(O_h) + D(O_a))}{D(O_a)} \quad (26)$$

In this situation, bookmakers use their knowledge of the demand functions for bets to mimic the outcome of pari-mutuel betting, in which odds are inversely proportional to the share of bets placed on the team. This inverse relationship between odds and betting volumes means payouts to winners are the same fraction of the total pool of bets placed every time, leaving the bookmaker with a certain profit rate of θ .

Without specifying a distributional form for beliefs, we know the odds that deliver this outcome will exhibit favorite-longshot bias for the same reasons discussed above in relation to Ali's result that disagreement among gamblers produces this bias in pari-mutuel markets. Risk-neutral competitive equilibrium odds do not produce volumes that are inversely proportional to probabilities, so the bookmakers re-balance volumes and reduce risk by setting odds that are more generous to the favorite and less generous to the longshot.

Again using a uniform distribution to provide a concrete illustration, we can derive a closed-form solution for zero-variance odds from inserting equations 25 and 26 into the demand functions

for bets

$$D(O_h) = 1 - \frac{1}{2\sigma} \left[\frac{D(O_h)}{(1 - \mu - \theta)(D(O_h) + D(O_a))} - p + \sigma \right] \quad (27)$$

$$D(O_a) = \frac{1}{2\sigma} \left[1 - \frac{D(O_a)}{(1 - \mu - \theta)(D(O_h) + D(O_a))} - p + \sigma \right] \quad (28)$$

These solve to give

$$D(O_h) = \frac{1}{1 + \frac{1}{2\sigma(1 - \mu - \theta - \frac{\mu + \theta}{2\sigma})}} \left(\frac{1}{2} + \frac{p}{2\sigma} \right) \quad (29)$$

$$D(O_a) = \frac{1}{1 + \frac{1}{2\sigma(1 - \mu - \theta - \frac{\mu + \theta}{2\sigma})}} \left(\frac{1}{2} + \frac{1 - p}{2\sigma} \right) \quad (30)$$

The odds can then be calculated as

$$O_h = \frac{(1 - \mu - \theta)(1 + 2\sigma)}{p + \sigma} \quad (31)$$

$$O_a = \frac{(1 - \mu - \theta)(1 + 2\sigma)}{1 - p + \sigma} \quad (32)$$

and the ratio of home odds to away odds is

$$\frac{O_h}{O_a} = \frac{\frac{1}{2} + \frac{1-p}{2\sigma}}{\frac{1}{2} + \frac{p}{2\sigma}} = \frac{1 - p + \sigma}{p + \sigma} \quad (33)$$

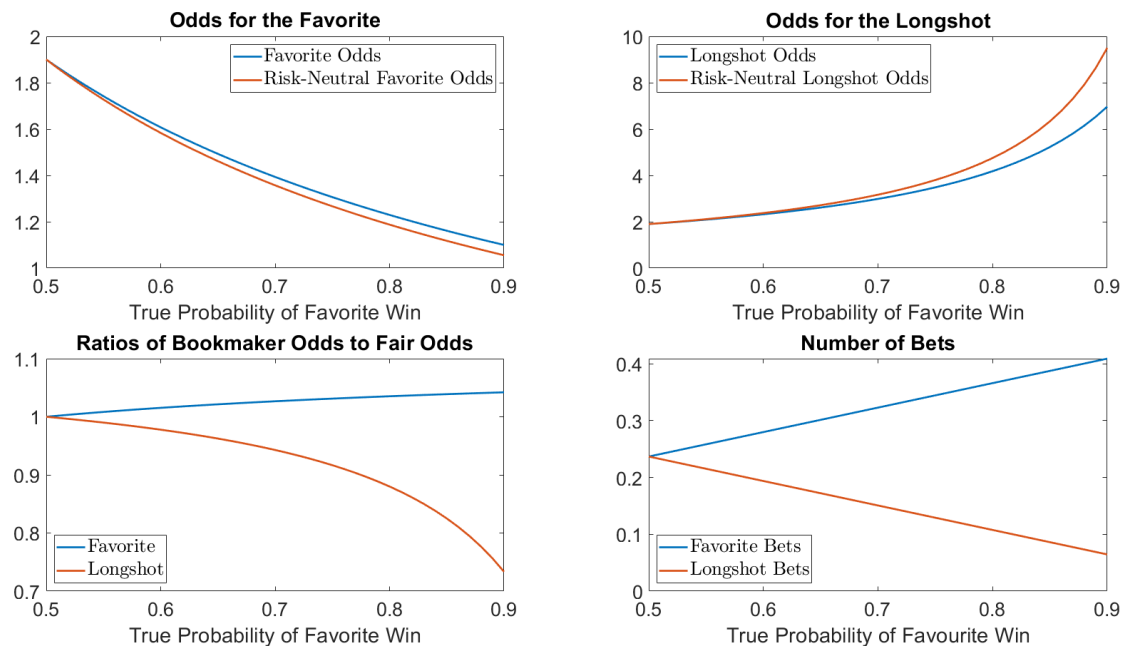
This formula collapses to the risk-neutral competitive odds when $\sigma = 0$, though in our model there is no market when there is no disagreement: Given the bookmaker's margins, nobody would choose to bet. However, once σ is positive, the betting market exists and the odds exhibit favorite-longshot bias that increases with the size of σ and also with the value of p . For example, under risk-neutral pricing with $p = 0.7$, the ratio of favorite odds to longshot odds would be $\frac{0.3}{0.7} = 0.43$. However, the zero-variance odds imply a ratio of 0.467 if $\sigma = 0.05$ and 0.5 if $\sigma = 0.1$. If the true probability of the favorite winning is $p = 0.8$, the risk-neutral odds ratio is 0.25 while the zero-variance odds imply a ratio of 0.294 if $\sigma = 0.05$ and 0.333 if $\sigma = 0.1$. Disagreement of $\sigma = 0.1$ raises the relative odds of the favorite by a third when $p = 0.8$ and by one-sixth when $p = 0.7$.

Figure 6 illustrates the zero-variance odds and the volumes they generate, using the same underlying parameters as for the risk-neutral case described in Figure 5. While the odds clearly move in the same direction as the risk-neutral odds as p changes, the favorite's odds are a bit more generous and the longshot's odds become considerably less generous as its probability of the winning falls. Bookmakers earn more profits from bets on longshots with these odds than they do under risk-neutral

pricing with these excess profits precisely offset by lower profits from bets on favorites.⁹ Because the excess profits on longshots are earned from a smaller number of gamblers, the profit rate per bet for longshots becomes much higher than for favorites as p rises: See Figure 7 for an illustration.

Of course, the flip side of profits for bookmakers is losses for gamblers. Figure 8 shows the pattern of payout rates for \$1 bets by true probability of bet win generated by these odds. The pattern is highly nonlinear—with payout rates falling rapidly as the probability of the winning bet falls below 0.3—and looks remarkably like the evidence presented here in Figure 1 for tennis and elsewhere for other sports. One thing to note is that this simulation sets $\sigma = 0.05$. This shows that it doesn't take much disagreement to generate patterns of favorite-longshot bias that match the data.

Figure 6: Betting odds and volumes for zero-risk odds ($\mu = 0.01, \theta = 0.04, \sigma = 0.05$)



⁹This raises the question of whether it may be a good idea for one of the bookmakers to seek to siphon off these longshot profits by offering slightly better odds on the longshots, thus making the equilibrium collapse. The answer to that will depend on how much risk bookmakers are willing to tolerate. More on this below.

Figure 7: Profit rates on favorite and longshot bets for zero-risk odds ($\mu = 0.01, \theta = 0.04, \sigma = 0.05$)





Figure 8: Average payout rates for gamblers for zero-risk odds ($\mu = 0.01, \theta = 0.04, \sigma = 0.05$)



4.3. Trade-offs between risk and return

The result that bookmakers may set odds to ensure risk-free returns is an extreme one. In a practical sense, even bookmakers who do wish to minimize variance have to deal with not actually knowing the true form of the demand function for bets and with bets arriving in gradually over time. These factors may make it difficult for the average odds on the final set of bets on a game to be inversely proportional to betting volumes. More importantly in the context of our model, setting risk to zero is clearly not the only option open to bookmakers. Is there a middle ground between setting odds without regard to risk and setting them to avoid risk altogether?

To illustrate the alternative options, we will assume that equilibrium odds for the case where bookmakers are somewhat (but not completely) risk averse will be a weighted average of those set by risk-neutral competitive bookmakers (which occur when risk is not factored into odds-setting) and the zero-risk odds (which will occur if bookmakers are unwilling to take any risk), both of which provide an expected profit rate of θ . In other words, we consider odds of the form

$$O_h = \omega \left[\frac{(1 - \mu - \theta)(1 + 2\sigma)}{p + \sigma} \right] + (1 - \omega) \left[\frac{(1 - \mu - \theta)}{p} \right] \quad (34)$$

$$O_a = \omega \left[\frac{(1 - \mu - \theta)(1 + 2\sigma)}{1 - p + \sigma} \right] + (1 - \omega) \left[\frac{(1 - \mu - \theta)}{1 - p} \right] \quad (35)$$

Figure 9 illustrates, for some different values of p , the expected profit rates on games and their standard deviations as ω goes from zero to one. For each value of p , the locus of mean-variance combinations is a near-inverse-U-shape (they are in fact slightly asymmetric). They show higher profit rates but also higher standard deviations as odds move away from the zero-variance option. Profit rates peak and then fall to the market-determined rate earned by the zero variance odds when odds reach the risk-neutral competitive level, while the standard deviation keeps increasing. The range of expected profit rates and standard deviations for games with a strong favorite (such as $p = 0.9$) are clearly far more widely dispersed than for games with a weak favorite (such as $p = 0.6$).

This chart illustrates why risk-averse bookmakers are unlikely to set odds that are close to the competitive risk-neutral odds. For each point on the right-hand side closer to the risk-neutral case there is a corresponding point closer to the zero-variance case that has the same expected profit rate and a lower variance. Risk-averse bookmakers will prefer the lower-variance option.

To illustrate this result in a more concrete way, consider the case where potential bookmakers have mean-variance preferences¹⁰

$$U(\pi) = E(\pi) - \frac{\gamma}{2} Var(\pi) \quad (36)$$

¹⁰Mean-variance preferences are closely related to constant relative risk aversion utility $U(r) = \frac{r^{1-\gamma}}{1-\gamma}$. For the case of log-normally distributed returns, the two are equivalent. See Ang (2014) or any other standard introductory graduate finance textbook.

We previously assumed that without concern about risk, free entry required an expected profit per bet of θ . Any reasonable free-entry condition incorporating risk-aversion will require that bookmakers only consider entering if the combination of risk and return at least matches the utility from the risk-free return of θ implied by the zero-variance odds. Figure 10 illustrates the level of utility for bookmakers implied by the combinations of means and variances shown in the curves in Figure 9 for a range of values for γ , sorting the odds combinations from left to right based on increasing the weight on the zero-variance odds. The extensive empirical literature on risk-aversion points to values of γ of about 1 as its typical estimate, so we use values of γ ranging from 0.5 to 2.¹¹

Figure 10 shows that no matter which level of risk-aversion is used, odds close to the risk-neutral competitive odds generate low levels of utility for bookmakers and would not trigger market entry. While a little hard to see in the chart, utility levels eventually turn positive as ω rises and peak above θ before falling to θ when $\omega = 1$. For each level of risk aversion, the exact values of ω that bring utility levels up to θ vary according to the probability of the favorite winning but the variations are small. For $\gamma = 0.5$, utility for each bet reaches θ when ω is above 0.74, for $\gamma = 1$, this occurs for values of ω above 0.84, for $\gamma = 1.5$, it occurs for values of ω above 0.89 and for $\gamma = 2$, it occurs for values of ω above 0.91.

This suggests something like $\omega = 0.84$ as a reasonable benchmark for the odds that will be set by a realistically risk-averse bookmaker. Figure 11 describes the outcomes from these odds. They very similar to those shown for the zero-variance case in Figures 6, 7 and 8. In particular, the pattern of payout rates by win probability still closely matches the empirical evidence provided in Figure 1.


We have derived these results as a symmetric equilibrium: Each bookmaker enters the market setting the same odds to achieve the same risk-return outcome per bet. But would this be a stable equilibrium? We showed previously in the risk-neutral case that bookmakers who deviated from competitive equilibrium odds received large volumes of bets skewed towards one or other of the teams and ended up making poor returns or outright losses. What would happen to a bookmaker that choose to deviate from the odds described in Figure 11?

At first look, there is a deviating option that looks like it might work: Offer poorer odds on favorites (which earn a lower return) and offer slightly more attractive odds on the more-profitable longshot bets. This would cause all longshot gamblers to move over to the deviating bookmaker. Figure 12 illustrates what would happen to a deviating bookmaker offering 1% higher odds on longshots and worse odds on favorites. They would generate huge betting volume from longshot gamblers and earn higher profits but would do so at the expense of much higher variance. The bottom right-hand panel shows the ratio of additional profits earned to additional variance is below 0.1 across the range of values of p . For realistic levels of risk aversion, this deviating strategy would reduce, rather than

¹¹See Chetty (2006), Layard, Mayraz and Nickell (2008) and Gandelman and Hernandez-Murillo (2013) for various estimates of γ using different methodologies.

raise, the bookmakers' utility. This suggests a competitive equilibrium of the type described in Figure 11 will be stable in the sense that no individual bookmaker will be incentivized to deviate from these odds.

Figure 9: Risk-return combinations for various values of p ($\mu = 0.01$, $\theta = 0.04$, $\sigma = 0.05$)



Notes: The chart shows expected profit rates and standard deviations of profit rates for a set of odds that go from competitive odds for risk-neutral bookmakers (the points on the extreme right) to the zero-variance case where odds are inversely proportional to betting volumes (the points on the extreme left).


Figure 10: Mean-variance utility for various values of γ ($\mu = 0.01, \theta = 0.04, \sigma = 0.05$)

Figure 11: Outcomes for odds priced with $\omega = 0.84$ ($\mu = 0.01, \theta = 0.04, \sigma = 0.05$)


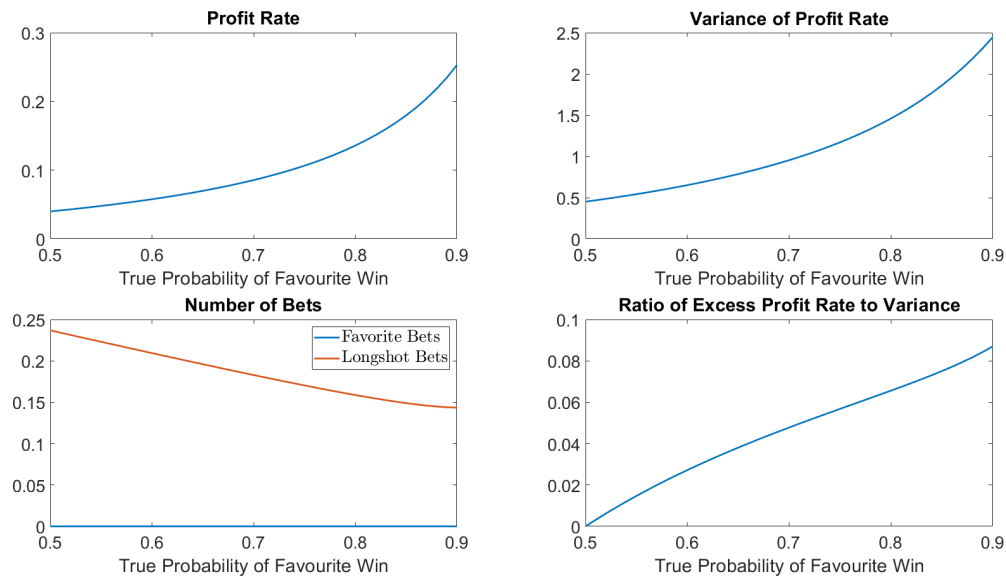


Figure 12: Profit rate and its variance for a bookmaker offering 1% better odds on the longshot and taking no bets on the favorite ($\gamma = 1, \mu = 0.01, \theta = 0.04, \sigma = 0.05$)



5. Risk Aversion for Gamblers?

We have described a model with risk-averse bookmakers but risk-neutral gamblers. What if gamblers were also risk averse? Some of the literature on sports betting has assumed that gamblers must be risk-lovers because they take on risky bets with negative expected value.¹² In our framework, however, it would not be necessary for gamblers to be risk-lovers to be willing to take bets. If a gambler's expected profit on a bet, based on their subjective belief about the probability of a team winning, was big enough then they could be willing to take a bet involving a lot of risk even if they disliked the risky element of the bet.

Consider a gambler making a \$1 bet on a team they believe has a probability \tilde{p} of winning with odds O . Their expected payout for the bet is $\tilde{p}O$ and their perceived variance for the payout is $\frac{\tilde{p}(1-\tilde{p})O^2}{2}$. A risk-averse bettor with mean-variance preferences of the same form as equation 36 will bet if

$$\tilde{p}O - \frac{\gamma\tilde{p}(1-\tilde{p})O^2}{4} > 1 \quad (37)$$

In other words, they will take the bet if the risk-adjusted return from it exceeds the certain return from not betting. This reduces to a quadratic inequality

$$\tilde{p}^2 + \left(\frac{4}{\gamma O} - 1\right)\tilde{p} - \frac{4}{\gamma O^2} > 0 \quad (38)$$

This quadratic function has one positive root and one negative root, with the latter not being relevant. This means for any given odds offered on a team, there is a threshold belief level \tilde{p} above which risk-averse bettors will choose to bet.

We can use these threshold belief levels to calculate how risk-averse gamblers would respond to the odds posted by risk-neutral competitive bookmakers. For a range of true values of a bet winning, the left panel of Figure 13 shows the threshold probabilities under risk-neutral competitive odds for three different values of the risk-aversion coefficient: $\gamma = 0.01$, meaning effectively risk-neutral, $\gamma = 0.5$, meaning relatively low risk-aversion compared with many empirical results and $\gamma = 1$, equivalent to log utility, often used as a benchmark for attitudes to risk.

For $\gamma = 0.01$, the threshold probabilities are close to a 45 degree line. As the gambler becomes more risk-averse, their threshold probabilities are higher and no longer rise monotonically with the true probability that the bet wins. Indeed, with log utility, the gambler only takes a competitively-priced bet if their subjective belief in its probability of winning is at least 0.5. The right panel illustrates the outcomes for betting volumes for uniformly distributed beliefs with $\sigma = 0.1$. For the near-risk-neutral case, betting volumes are larger on longshots but as risk-aversion rises, no bets are placed on extreme longshots.

¹²See for example Weitzman (1965) and Quandt (1986).

These calculations suggest the most realistic calibration of our model is one with risk-averse bookmakers but bettors that are at most very mildly averse to risk. Risk averse bettors with γ of 0.5 or above would generate a number of counter-factual predictions. This would predict a complete absence of bets on longshots which runs counter to the fact that real-world gamblers do take on these bets. Moreover, the zero variance option for a bookmaker faced with these betting volumes for risk-neutral odds would be to price teams that are very likely to win worse than competitive odds and teams less likely to win better than competitive odds, so a risk-averse bookmaker could set odds implying a counterfactual reverse favorite-longshot bias.


There are a number of explanations for why a configuration of risk-averse bookmakers but approximately risk-neutral bettors makes sense. First, there is the scale of risk involved. Most gamblers place small bets that do not have life-altering consequences if they win or lose. Rabin (2000) pointed out that people should be approximately risk-neutral when the stakes are small. If utility functions are already concave in the region of the changes in wealth generated by winning or losing small bets, then they should be even more risk-averse when it comes to making decisions with big financial consequences. Rabin shows that people turning down positive expected value risky small bets would imply counter-factual predictions for the pricing of insurance of various types and for decisions such as whether to invest in stocks.¹³ In contrast, bookmakers are taking a large volume of bets and can be at risk of losing a lot of money, so Rabin's arguments suggest they could be plausibly considered risk averse.

Second, the empirical studies used to estimate risk aversion tend to consider far wider variations in income or assets than would be generated by the betting decisions made by most gamblers. For example, Layard, Mayraz and Nickell (2008) and Gandelman and Hernandez-Murillo (2013) use self-reported happiness from surveys to assess the curvature of utility functions across a wide range of incomes. It is questionable whether these estimates are relevant for the risk assessment of a gambler placing a \$20 bet.

Third, it seems likely that people are heterogeneous in their attitudes to risk. Even if most of the population did have an aversion to small-scale risks, this does not mean the population of people who would consider betting on sport share this aversion. While we have shown that it is not necessary to assume that gamblers are risk lovers, we have seen that even modest levels of risk-aversion generate counterfactual predictions about betting volumes. Together, these considerations point to a benchmark of risk-averse bookmakers and risk-neutral gamblers as a reasonable one.

¹³Rabin and Thaler (2001) provide further discussion of these ideas.

Figure 13: Threshold betting probabilities for risk averse gamblers and amount of bets placed ($\mu = 0.01, \theta = 0.04, \sigma = 0.1$)



6. Conclusion

Sports betting markets have long attracted economists as a way of testing hypotheses about how people take decisions involving risk. While these betting markets themselves may not have been of much importance to the overall economy, they provided an environment well suited to direct testing of hypotheses about decision making with risky outcomes. In recent years, however, the sports betting industry has increased in size worldwide to a point where its importance for economists should no longer be just as a testing ground for theories about decision making under risk. This is particularly true in the US due to the legalization of sports betting in 2018. To give a potential benchmark, in the UK, which has a well-developed sports betting market, the total amount of betting on sports per year is currently about £36 billion, equivalent to £530 per person.¹⁴ An equivalent amount of sports betting per person in the US would see annual totals of around \$210 billion, equivalent to about 1 percent of current GDP.

Given the growing importance of sports betting and the widespread use of fixed odds rather than a pari-mutuel approach, it is important to understand how prices are set in fixed-odds markets. We have provided a model of how a fixed-odds betting market works when there is competition and free entry. The requirement for bookmakers to cover costs and earn a market-determined rate of return affects the demand for bets on favorites differently to the demand for bets on longshots once people disagree about the probability of teams winning. These margins discourage those who would bet on favorites more than those who would bet longshots. Because longshot bets generate more variance for their profits, risk averse bookmakers will respond to this pattern by setting odds to discourage bets on longshots and encourage them on favorites.

One way to interpret our results is that they are driven by the misplaced over-confidence of those who are more confident than the average bettor in the longshot's chances of winning. However, it is worth emphasizing that the gamblers in our model are not systematically poor at understanding low probability events in the manner suggested by prospect theory. For each person that is incorrectly confident in the longshot team, the model also contains a symmetric person who is over-confident in the favorite's prospects.

An implication of our model is that gamblers in the newly emerging US sports betting market are going to lose a lot of money betting on longshots. This suggests a possible role for government agencies to consider running information campaigns to explain that the likely outcome of betting on sports is that gamblers will lose and to specifically document the dangers of taking high odds bets that only pay off in unlikely outcomes.

The model could be extended in several directions. Bet size could be endogenized and the number of competitors in the sporting event could be increased. Perhaps more importantly, other market

¹⁴Sports betting revenue figures taken from <https://www.gamblingcommission.gov.uk/statistics-and-research/publication/industry-statistics-november-2021>

structures could be considered. While we have argued that competitive markets could generate the pattern of favorite-longshot bias seen in the data, that does not mean that all real-world betting markets meet the ideal of a competitive market. In particular, there are reasons to be concerned about the market structure of the emerging US legal sports betting market. Bloomberg have reported that the two leading firms (FanDuel and DraftKings) have 65% of the current market.¹⁵ In those circumstances, outcomes for gamblers may be worse than those described in this paper.

References

- [1] Ali, Mukhtar (1977). "Probability and Utility Estimates for Racetrack Bettors," *Journal of Political Economy*, Volume 85, pages 803-815.
- [2] Ang, Andrew (2014). *Asset Management: A Systematic Approach to Factor Investing*, Oxford University Press.
- [3] Berkowitz, Jason, Craig Depken and John Gandar (2017). "A Favorite-Longshot Bias in Fixed-Odds Betting Markets: Evidence from College Basketball and College Football," *Quarterly Review of Economics and Finance*. Volume 63, pages 233-239
- [4] Buchdahl, Joseph (2016). *Squares & Sharps, Suckers & Sharks: The Science, Psychology Philosophy of Gambling*. High Stakes Publishing.
- [5] Buhagiar, Ranier, Dominic Cortis, and Philip Newall (2015). "Why Do Some Soccer Bettors Lose More Money Than Others?" *Journal of Behavioural and Experimental Finance*, Volume 18, pages 85-93.
- [6] Cain, Michael, David Law and David Peel (2000). "The Favorite-Longshot Bias and Market Efficiency in UK Football Betting," *Scottish Journal of Political Economy*, Volume 47, pages 25-36.
- [7] Chetty, Raj (2006). "A New Method of Estimating Risk Aversion," *American Economic Review*, Volume 96, pages 1821-1834.
- [8] Gabriel, Paul and James Marsden (1990). "An Examination of Market Efficiency in British Race-track Betting," *Journal of Political Economy*, Volume 98, pages 874-885.
- [9] Gandelman, Nestor and Ruen Hernandez-Murillo (2013). "What Do Happiness and Health Satisfaction Data Tell Us About Relative Risk Aversion?" *Journal of Economic Psychology*, Volume 39, pages 301-312.

¹⁵<https://www.bloomberg.com/news/articles/2021-09-21/how-fanduel-gained-more-fantasy-sports-gamblers-than-draftkings-dkng>

- [10] Griffith, Richard (1949). "Odds Adjustment by American Horse Race Bettors," *American Journal of Psychology*, Volume 62, pages 290–294.
- [11] Forrest, David and Ian McHale (2007). "Anyone For Tennis (Betting)?" *European Journal of Finance*, Volume 13, pages 751-768.
- [12] Hurley, William and Lawrence McDonough (1995). "A Note on the Hayek Hypothesis and the Favorite-Longshot Bias in Parimutuel Betting," *American Economic Review*, Volume 85, pages 949-955.
- [13] Kahneman, Daniel and Amos Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, Volume 47, pages 263-291.
- [14] Layard, Richard, Guy Mayraz and Steve Nickell (2008). "The Marginal Utility of Income," *Journal of Public Economics*, Volume 92, pages 1846-1857.
- [15] Miller, Ed and Matthew Davidow (2019) *The Logic of Sports Betting*.
- [16] Ottaviani, Marco and Peter Norman Sørensen (2008). "The Favorite-Longshot Bias: An Overview of the Main Explanations," in *Handbook of Sports and Lottery Markets*, Elsevier.
- [17] Ottaviani Marco and Peter Norman Sørensen (2009). "Surprised By Pari-Mutuel Odds," *American Economic Review*, Volume 99, pages 2129-2134.
- [18] Ottaviani Marco and Peter Norman Sørensen (2010). "Noise, Information and the Favorite-Longshot Bias in Parimutuel Predictions," *American Economic Journal: Microeconomics*, Volume 2, pages 58-85.
- [19] Quandt, Richard (1986). "Betting and Equilibrium," *Quarterly Journal of Economics*, Volume 101, pages 201-208.
- [20] Rabin, Matthew (2000). "Risk-Aversion and Expected Utility Theory: A Calibration Theorem," *Econometrica*, Volume 68, pages 1281-1292.
- [21] Rabin, Matthew and Richard Thaler (2001). "Anomalies: Risk Aversion," *Journal of Economic Perspectives*, Volume 15, pages 219–232
- [22] Shin, Hyun Song (1991). "Optimal Betting Odds Against Insider Traders," *Economic Journal*, Volume 101, pages 1179-1185.
- [23] Shin, Hyun Song (1992). "Prices of State-Contingent Claims with Insider Traders and the Favorite-Longshot Bias," *Economic Journal*, Volume 102, pages 426-435.
- [24] Snowberg, Erik and Justin Wolfers (2010). "Explaining the Favorite-Longshot Bias: Is it Risk-Love or Misperceptions?" *Journal of Political Economy*, Volume 118, pages 723-746.

- [25] Weitzman, Martin (1965). "Utility Analysis and Group Behavior: An Empirical Study," *Journal of Political Economy*, Volume 73, pages 18-26.