The Favorite-Longshot Midas*

Etan A. Green[1], Haksoo Lee[1], and David Rothschild[2]

[1]University of Pennsylvania
[2]Microsoft Research

June 1, 2018

Abstract

The favorite-longshot bias in horse race betting markets, previously attributed to speculation or irrationality, instead results from profitable deception. Racetracks provide bettors with predictions that overestimate the chances of longshots and underestimate the chances of favorites. This deception creates arbitrage opportunities, which the track monetizes by taxing the arbitrageurs. Similar schemes enrich other market makers, such as investment banks during the boom in mortgage-backed securities.

*Please direct correspondence to: etangr@wharton.upenn.edu. We thank Bosko Blagojev and Will Cai for help with scraping and parsing the data; Joe Appelbaum, Dorothy Kronick, Cade Massey, Dave Pennock, Alex Rees-Jones, and Joe Simmons for helpful conversations and comments on previous drafts; and Wharton and Microsoft Research for generous financial support.

Profit is sweet, even if it comes from deception.

Sophocles

1 Introduction

We revisit a classic case of market inefficiency: the favorite-longshot bias in horse-race betting markets (Thaler and Ziemba, 1988). Wagers on longshots return far less in expectation than wagers on favorites. Two explanations have been offered: first, that bettors have non-standard preferences, and second, that bettors form biased beliefs (for reviews, see Sauer, 1998; Ottaviani and Sørensen, 2008). We show that neither need be true. Instead, the favorite-longshot bias arises from an indirect pump-and-dump scheme, in which the track deceives uninformed bettors so as to profit by taxing arbitrageurs.

Horse race betting markets in the United States are run as parimutuels. Bettors place wagers on outcomes (e.g., Secretariat winning the race), the track takes a tax, and the remainder is divided among winning wagers. If bettors are risk neutral and have correct beliefs, then expected returns should be the same for favorites and longshots. In practice, favorites are underbet, and longshots are overbet. Whereas a favorite with 1/1 odds (i.e., that pays a $1 dividend on a winning $1 wager) returns 85 cents on the dollar on average, a longshot with 30/1 odds returns just 63 cents on the dollar.

Previous explanations for the favorite-longshot bias mirror trends in economics. Early studies used parimutuel odds to estimate the preferences of a representative agent in an expected utility framework. Under this approach, a preference for risk rationalizes the favorite-longshot bias (Weitzman, 1965; Ali, 1977; Golec and Tamarkin, 1998). More recent work has favored psychological explanations. Here, biased beliefs explain the favorite-longshot bias—specifically, a tendency to overweight small probabilities in the manner of prospect theory (Jullien and Salanié, 2000; Snowberg and Wolfers, 2010).[1]

We show that the favorite-longshot bias is driven by profitable deception, rather than risk-seeking preferences or psychological biases. The track deceives uninformed bettors by suggesting inflated probabilities for longshots and depressed probabilities for favorites. This deception is strategic. By tricking uninformed bettors into underbetting favorites, the track

[1]Separately, a series of papers show that speculative or irrational bettors need only place a small share of the wagers in order to generate the observed bias (Hurley and McDonough, 1995; Ottaviani and Sørensen, 2009, 2010; Gandhi and Serrano-Padial, 2014). Because the tax makes it inefficient to bet against longshots, *any* overbetting of longshots—by virtue of risk-loving preferences, biased beliefs, or other anomalies—cannot be fully arbitraged away.

creates arbitrage opportunities for informed bettors and, from that arbitrage, additional tax revenue for itself. This rendering requires only that some bettors be gullible, rather than speculative or irrational.

There are two ways to bet on horse races: at the track or through off-track outlets such as betting websites. Whereas track bettors pay the full tax, off-track bettors receive a rebate—a portion of their wagers back, win or lose—that lowers their effective tax rate. This price discrimination sorts the informed from the uninformed. Off-track bettors tend to wager at the close of the betting window, when live odds are most informative. Many of these bettors are thought to handicap horses using algorithms trained on databases of race histories. Bettors at the track instead receive race cards, or curated pamphlets, featuring the morning-line odds—a set of odds, one for each horse, that has no formal bearing on the parimutuel odds.

We show that the morning lines embed a favorite-longshot bias. They are insufficiently short for favorites and insufficiently long for longshots. If bettors are risk neutral, favorites with 1/1 odds should win about half the time. In practice, horses with morning-line odds of 1/1 win nearly two in three races, yet morning lines are rarely shorter than 1/1. Similarly, longshots with 30/1 odds should win about 3% of the time. In practice, horses with 30/1 morning-line odds win about 1% of the time, yet morning lines are rarely longer than 30/1. Despite their formal irrelevance, the morning-line odds predict the parimutuel odds until minutes before the start of the race, after which the parimutuel odds shorten for favorites and lengthen for longshots.

These stylized facts motivate a two-period model with two risk-neutral representative agents. In the first period, track bettors infer beliefs, q, under the assumption that the morning-line odds reflect wagering by risk-neutral bettors with accurate beliefs. As expected-value maximizers, track bettors wager in proportion to their beliefs, generating odds that reproduce any bias inherent in the morning-line odds. (We rationalize their decision to gamble with a direct utility from gambling that exceeds the tax, as in Conlisk (1993).) Track bettors are unsophisticated in the sense of Milgrom and Roberts (1986): they fail to appreciate that the track may distort the morning lines for strategic reasons.[2]

In the second period, off-track bettors infer beliefs, $p(q)$, from the observed rates at which horses with implied probability q actually win. Given that the morning-line odds embed a favorite-longshot bias, this correction leads off-track bettors to view favorites as underpriced,

[2] Milgrom and Roberts (1986) show that competition among senders (i.e., tracks) or skepticism among receivers (i.e., track bettors) mitigate mistakes by receivers. The track is a monopolist in our model, and track bettors, without historical data on morning lines and outcomes, have little empirical basis for skepticism.

and with the rebate, perhaps profitably so. The track sets the profit-maximizing rebate as a monopolist, and a large number of off-track bettors place all wagers with positive expected value until no more exist. In equilibrium, late wagering concentrates on favorites, and the favorite-longshot bias moderates. The bias does not disappear, however, because off-track bettors pay some tax.

We estimate our model using a unique dataset. Past studies of the favorite-longshot bias in parimutuel betting markets ignore the morning-line odds, perhaps because they are difficult to obtain or because they have no formal bearing on the final odds.[3] We scrape morning-line odds each morning. Separately, we scrape minute-by-minute live odds, which we use to show how late wagering moves the parimutuel odds away from the morning lines.

Past studies engage in a curve-fitting exercise. A model is proposed which allows for a downward sloping relationship between the parimutuel odds and expected returns. Using the observed parimutuel odds, at least one parameter is tuned to match the slope in the data. Estimates are then interpreted as evidence of risk-loving preferences (e.g., Ali, 1977) or probability weighting (e.g., Snowberg and Wolfers, 2010).

By comparison, our model is restricted in two ways. First, it predicts returns by observing the formally irrelevant morning-line odds, not the parimutuel odds. And the only parameters in our estimation routine are hyper-parameters: the bandwidths used to (non-parametrically) estimate the beliefs of off-track bettors, which we select using cross validation. Despite these handicaps, our model more closely predicts the observed relationship between odds and returns than competing models with risk-loving or biased agents.[4]

Our model predicts differences across tracks in the extent of the favorite-longshot bias, which is nonexistent at some tracks and severe at others. In competing models, it is not obvious why preferences or beliefs would vary by track. In our model, deception in the morning-line odds generates the favorite-longshot bias. Some tracks do not distort the morning lines, and odds observed at those tracks are generally unbiased.[5] Other researchers have

[3]For instance, Ottaviani and Sørensen (2009) justify studying racetrack parimutuel markets because of "the absence of bookmakers (who could induce biases)." We show that the morning-line odds, devised by the track oddsmaker, induce a favorite-longshot bias.

[4]Gandhi and Serrano-Padial (2014) show that a model with two types of bettors—one with accurate beliefs and the other with random beliefs—predicts the favorite-longshot bias for win bets. They fit their model with three parameters: two that describe the distribution of beliefs among uninformed bettors, and a third measuring the share of uninformed bettors. We achieve comparable accuracy while specifying the beliefs of uninformed bettors and deriving the share of wagers from each type. In their model, the beliefs of uninformed bettors follow an arbitrary distribution. In our model, uninformed bettors take their beliefs from a manipulative agent. This implies divergent interpretations: the longshot-favorite bias arises because some bettors are deceived, not because some bettors bet randomly.

[5]Unfortunately, there are too few horse track parimutuel markets—in the world, let alone in our data—to

noted that parimutuel betting markets in Hong Kong do not exhibit a favorite-longshot bias (Busche and Hall, 1988). We observe that race cards in Hong Kong do not display morning lines.[6]

Deceptive morning lines generate excess profit for the track. State-regulated tax rates determine the losses that track bettors would sustain if they were well informed. Deceived track bettors sustain excess losses. Because arbitrage is competitive, these excess losses flow in their entirety into the track's coffers, laundered through taxes on arbitrageurs. Using our model, we estimate each track's incremental profit from deception. For tracks that embed a favorite-longshot bias in the morning-line odds, we estimate that as much as 10% of their gambling revenues come from taxing arbitrageurs.

In pump-and-dump schemes, traders distort asset prices (i.e., pump) so as to sell their holdings at higher prices (i.e., dump). Such schemes require the perpetrator to hold assets, entailing financial risk. They are also illegal. We illustrate an indirect pump-and-dump scheme, in which the market maker distorts asset prices so as to profit by taxing arbitrageurs. As in conventional pump-and-dump schemes, the market maker profits from its deception, but it does so without financial risk and in a manner that is apparently legal.

Other market makers perpetrate similar schemes. In the lead-up to the financial crisis, investment banks, along with rating agencies, misled unsuspecting investors about the risk in mortgage-backed securities. Deceived investors overpriced these securities, generating demand from savvy investors for ways to short them, which the same banks engineered and sold for large fees (Lewis, 2015). As with racetracks, investment banks created arbitrage opportunities by deceiving uninformed investors and profited by taxing the arbitrageurs.

Our model of deception by market makers is distinct from models of deception by firms (Gabaix and Laibson, 2006; Spiegler, 2011; Heidhues, Kőszegi and Murooka, 2016; Heidhues and Kőszegi, 2017). For instance, Gabaix and Laibson (2006) model a firm that shrouds add-on costs, such as high-priced printer toner. Because the firm cannot price discriminate, it loses money from sophisticated customers, who only purchase the subsidized printers, and makes money from naive customers, who purchase the toner as well. In our model, the market maker taxes informed and uninformed agents at different rates and profits from both.

This rendering is also distinct from betting markets in which a bookmaker, rather than a parimutuel, sets the betting odds. Bookmakers generate excess returns by taking risky positions against bettors with inaccurate beliefs, while using the tax to limit betting by

say with any confidence which factors correlate with deception.

[6] Hong Kong race cards can be accessed at: `http://racing.hkjc.com/racing/content/PDF/RaceCard/`.

"sharps" (Levitt, 2004). In our model, the track encourages the sharps to participate and generates a riskless profit by taxing them.

Considerable evidence shows that financial markets systematically misprice securities (for a review, see Barberis and Thaler, 2003). Leading explanations for these anomalies focus on noise traders who either possess non-standard preferences (Benartzi and Thaler, 1995) or form irrational beliefs (Shleifer and Summers, 1990). We show that a longstanding market anomaly, often rationalized by a preference for risk or a tendency to overweight small probabilities, is better explained by a model with zero noise traders. When some traders are gullible, deception by interested parties can separate prices from fundamental values.

The remainder of the paper is organized as follows. Section 2 describes the context, and Section 3 describes the data. Section 4 illustrates the favorite-longshot bias and other stylized facts. Section 5 details our model, Section 6 compares its predictions with the data, and Section 7 presents its estimates of each track's profit from deception. Section 8 concludes.

2 Context

In parimutuel markets, bettors place wagers on outcomes (e.g., a certain horse winning the race), the market maker takes a state-regulated share t known as the takeout, and the remainder is split among wagers on the winning outcome. Let s_i denote the share of wagers placed on outcome i. A \$1 wager on i returns $(1-t)/s_i$ if i occurs, and it returns 0 if i does not occur. Prospective returns are represented as odds, or as the dividend paid on a winning \$1 wager: $(1-t)/s_i - 1$.[7] For example, a winning bet at 2/1 odds pays a \$2 dividend for every \$1 wagered, along with the principal.

Tracks list the parimutuel odds in real time on a central board called the totalizator, or tote board, but only the final odds, or those the start of the race, are used to calculate payoffs. Live odds are also listed online.

Wagers on different types of outcomes are collected in separate pools, and odds are calculated within each pool. Almost all tracks offer win, place, and show pools, in which bets are placed on individual horses, and wagers pay out if the horse finishes first (win), first or second (place), or in the top 3 (show). Tracks also offer a selection of "exotic" pools, in which payoffs are contingent on multiple outcomes. Exacta, trifecta, superfecta, and hi-5 pools pay out if the bettor correctly predicts the first two, three, four, or five horses,

[7] In practice, tracks round odds down (to multiples of 5 cents for small odds and to larger multiples for higher odds), a practice known as breakage that slightly increases the track's effective takeout.

respectively, in order; quinella pools pay out if the bettor correctly predicts the first two horses in any order; and daily-double, pick-3, pick-4, pick-5, and pick-6 pools pay out if the bettor correctly predicts the the winners of two, three, four, five, or six races in a row, respectively.[8]

Tracks price discriminate between track bettors and off-track bettors, who wager either at off-track betting locations or online. Whereas track bettors pay the full takeout, off-track bettors receive a rebate—a proportion of the principal in return, win or lose—which lowers their effective tax rate. Rebates reflect a nuanced price differentiation strategy, varying by track, by pool, by off-track intermediary, by payout method (cash vs. redeemable points), and by bettor characteristics, including past wagering activity.

Figure 1: An example race card entry, prominently displaying the morning-line odds (15-1).

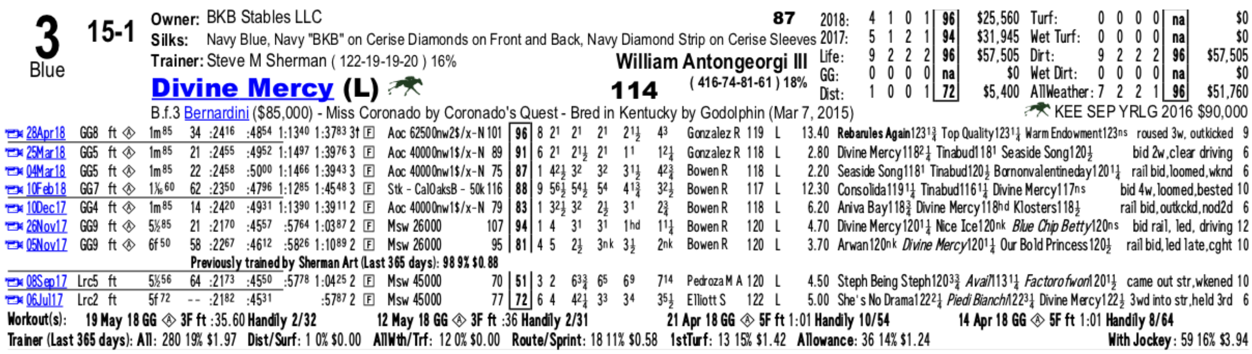

Source: Equibase.com

The rebate helps sort bettors by their level of sophistication. Informed bettors wager in the minutes or seconds before the race begins—i.e., when the live odds best predict the final odds (Gramm and McKinney, 2009). These late wagers are thought to be made disproportionately by off-track bettors, sometimes aided by algorithms.[9] Track bettors, by

[8]Many exotic wagers can also be constructed as a parlay, or contingent series, of wagers in other pools. For instance, a daily-double wager is logically equivalent to placing the full wager on a win bet in the first race and then placing any winnings on a win bet in the second race. Previous work finds that daily-double returns are less than those of the equivalent parlay, though not by a statistically significant margin—and claims that this constitutes evidence of efficient markets (Ali, 1979). We note that parlays are inefficient and hence, impose limits on arbitrage. Though the takeout is generally higher in exotic pools than in win, place, and show pools, the takeout in any exotic pool is always less than the takeout for the equivalent parlay. For example, paying the daily-double takeout once is less expensive than paying the win takeout twice.

[9]From twinspires.com, the online betting platform owned by Churchill Downs, "Late odds changes continue to confuse and confound horseplayers...There are big players in the pools and they are called 'whales.' Some whales, not all, use computer software to handicap and make their bets. Because of the volume of the whales' play, they are given rebates by advance deposit wagering companies to stimulate more betting." `https://www.twinspires.com/blog/2018/2/26/powell-explanation-behind-late-odds-changes`.

contrast, receive a race card, or a pamphlet with curated information about each horse. Figure 1 shows an example race card. Prominently displayed information include the horse's name, number, and color—and of greatest consequence for this paper, its morning-line odds.

The morning-line odds are win odds chosen by the oddsmaker at the track. Unlike in gambling markets with bookmakers, in which the bookmaker sets the final odds, the odds chosen by the track oddsmaker have no formal bearing on the odds that the track pays out. Nonetheless, the morning-line odds may influence the parimutuel odds. Publicly, oddsmakers argue that the morning-line odds attempt to "predict, as accurately as possible, how the betting public will wager on each race."[10] We propose an alternative explanation: distortion of the morning-line odds creates arbitrage opportunities for off-track bettors and, in so doing, additional revenue for the track.

3 Data

Most empirical studies of the favorite-longshot bias in horse-racing markets analyze data from the race chart, which summarizes the outcome of a race. We compile a similar dataset by collecting: 1) the final odds for each horse in the win pool, 2) the winning outcomes in each pool, 3) the returns to wagers on those outcomes, and 4) the total amount wagered in each pool.[11] Separately, we collect track- and pool-specific takeout rates from the Horseplayers Association of North America.

We supplement this standard dataset with the morning-line odds. Whereas historical race charts for US races are freely available, historical race *cards* are not. However, the morning-line odds are published online on the morning of the race, and we scraped these odds each morning for 17 months from 2016 to 2018.[12] In order to analyze the data within track, we restrict our sample to the 30 tracks for which we observe at least 5,000 horse starts. Our final dataset comprises 238,297 starts in 30,631 races.[13]

Table 1 presents summary statistics by track. Takeouts range from 12% to 31%. For comparison, US bookmakers typically take 9% of winnings on common wagers, such as spread

[10]Quoted in *Getting a Line on the Early Odds*, by Joe Kristufek. `https://www.arlingtonpark.com/node/10222`. Similarly, the horse racing columnist Bob Ehalt writes, "The intent [of the morning lines] is to frame how the betting on the race will unfold." `https://www.americasbestracing.net/gambling/2018-morning-line-odds-betting-tool`.

[11]One limitation of these data is that only in the win pool are final odds listed for all outcomes (i.e., horses); in other pools, only the odds for the winning outcome are listed.

[12]The scraper went offline from December 2, 2016 to July 23, 2017.

[13]This reflects the discarding of 91,727 starts in 12,164 races at 68 excluded tracks. Scraping or parsing issues led to the exclusion of an additional 4,147 races at both included and excluded tracks.

Table 1: Summary statistics by track.

Track			Counts			Takeout rates											Wager %	
Code	Name	State	Days	Races	Starts	WPS	Quinella	Exacta	Trifecta	Superfecta	Hi-5	Daily Double	Pick 3	Pick 4	Pick 5	Pick 6	WPS	Exotic
ALB	Albuquerque	NM	97	927	7454	.19	-	.25	.25	.25	-	.25	.25	.25	-	-	.41	.59
AP	Arlington	IL	94	758	5551	.17	-	.20	.25	.25	.15	.20	.25	.25	.15	.15	.41	.59
AQU	Aqueduct	NY	107	914	6836	.16	.18	.18	.24	.24	-	.18	.24	.24	.15	.15	.33	.67
BEL	Belmont	NY	90	814	6372	.16	.18	.18	.24	.24	-	.18	.24	.24	.15	.15	.32	.68
BTP	Belterra Park	OH	116	938	7004	.18	-	.23	.23	.23	-	.23	.23	.23	-	-	.35	.65
CBY	Canterbury	MN	92	869	6825	.16	-	.20	.20	.20	.18	.20	.20	.20	.20	-	.42	.58
CD	Churchill Downs	KY	83	793	6773	.17	-	.22	.22	.22	.22	.22	.22	.22	.22	.15	.34	.66
CT	Charles Town	WV	197	1619	12547	.17	-	.19	.22	.22	-	.19	.22	.18	-	.12	.36	.64
DED	Delta Downs	LA	125	1268	11126	.17	-	.20	.25	.25	.25	.20	.25	.12	.25	-	.32	.68
DEL	Delaware	DE	121	1066	7608	.17	-	.19	.25	.25	-	.19	.25	.25	.15	-	.35	.65
DMR	Del Mar	CA	97	813	6853	.15	.23	.23	.24	.24	.24	.20	.24	.24	.14	.24	.35	.65
EVD	Evangeline	LA	136	1221	9515	.17	-	.20	.25	.25	.25	.20	.20	.25	.25	-	.33	.67
FG	Fair Grounds	LA	99	924	7605	.17	-	.20	.25	.25	.25	.20	.25	.25	.25	-	.35	.65
FL	Finger Lakes	NY	185	1596	10554	.18	-	.20	.25	.25	-	.20	.25	.25	.25	-	.33	.67
GG	Golden Gate	CA	144	1188	8298	.15	-	.23	.24	.24	.24	.20	.24	.24	.15	.24	.37	.63
GP	Gulfstream	FL	171	1754	14371	.17	-	.20	.26	.26	.17	.20	.20	.20	.15	.20	.30	.70
IND	Indiana	IN	109	981	8203	.18	-	.21	.21	.21	.21	.21	.21	.21	-	.21	.32	.68
LA	Los Alamitos	CA	119	996	6631	.15	.23	.23	.24	.24	-	.20	.24	.24	-	.24	.24	.76
LAD	Louisiana Downs	LA	104	758	5617	.17	-	.20	.25	.25	.25	.20	.25	.25	.25	.25	.33	.67
LRL	Laurel	MD	117	1057	8266	.18	-	.21	.26	.26	.15	.21	.26	.26	.12	.20	.34	.66
MNR	Mountaineer	WV	146	1304	9299	.17	-	.19	.25	.25	.25	.19	.25	.25	.25	-	.35	.65
MVR	Mahoning Valley	OH	112	920	7792	.18	-	.23	.23	.23	.23	.23	.15	.15	-	.19	.32	.68
PEN	Penn National	PA	129	1041	7716	.17	-	.20	.31	.30	.20	.20	.25	.25	-	.20	.36	.64
PID	Presque Isle	PA	101	800	5764	.17	-	.20	.25	.25	-	.20	.23	.23	.23	-	.34	.66
PRX	Parx	PA	144	1303	10567	.17	-	.20	.30	.30	-	.20	.26	.26	-	-	.33	.67
RP	Remington	OK	86	826	7392	.18	-	.21	.24	.24	-	.21	.25	.25	.25	-	.33	.67
SA	Santa Anita	CA	90	775	6228	.15	-	.23	.24	.24	.24	.20	.24	.24	.14	.24	.33	.67
SUN	Sunland	NM	70	642	5598	.19	-	.21	.25	.25	.25	.21	.25	.25	-	.25	.33	.67
TAM	Tampa Bay Downs	FL	83	774	6213	.17	-	.20	.26	.26	.18	.18	.18	.18	.15	-	.33	.67
TUP	Turf Paradise	AZ	125	992	7719	.21	.22	.22	.26	.26	-	.22	.26	.26	.25	.26	.30	.70

or over-under bets, on other sports. The takeout is typically smaller in the win, place, and show (WPS) pools than in the exotic pools, yet as the final two columns show, more money is wagered in exotic pools at every track.

All standard errors reported in the remainder of the paper—either directly in the text or in tables, or as 95% confidence intervals in figures—are estimated from 10,000 bootstrap samples blocked by track.[14]

4 Stylized facts

Returns are systematically miscalibrated. Figure 2 shows observed returns in the win pool as a smooth function of log final odds. Wagers at all odds lose money on average, given the large takeouts. However, wagers on horses with longer odds lose more money. A $1 bet on a 1/1 favorite loses 15 cents in expectation, a smaller loss than the average takeout of 17 cents. By contrast, a 20/1 longshot loses 29 cents, and a 50/1 longshot loses 47 cents.[15] Bettors could reduce their losses by betting on favorites instead of longshots.

We show that models with non-standard agents do a poor job predicting the favorite-longshot bias. A long line of papers use final odds in win pools, along with observed returns, to estimate the preferences of a representative agent with accurate beliefs (e.g., Weitzman, 1965; Ali, 1977). Under this model, the favorite-longshot bias is rationalized by a taste for risk. Others observed that the same data can instead be used to estimate the beliefs of a risk-neutral representative agent who transforms probabilities in the manner predicted by prospect theory (Snowberg and Wolfers, 2010). Under this model, the favorite-longshot bias is rationalized by overweighting small probabilities and underweighting large probabilities.

Figure 3 shows predicted returns from each model.[16] For both models, predicted returns

[14] The standard bootstrapping approach is to resample observations, or in our case, an outcome in a given pool. However, doing so would create bootstrap samples for which the number of winning outcomes in each pool need not be one. As a result, we resample entire pools, rather than individual outcomes. The lone exception is Figure 5, for which the confidence intervals are derived from homoskedastic standard errors.

[15] These estimates are consistent with those from other analyses of parimutuel odds. For instance, Snowberg and Wolfers (2010) estimate expected returns of about 85% for a 1/1 favorite, 75% for a 20/1 longshot, and 55% for a 50/1 longshot. The favorite-longshot bias tends to be more severe in markets where bookmakers set the odds. For instance, Jullien and Salanié (2000) estimate expected returns of close to 100% for 1/1 favorites and 50% for 20/1 longshots using data from the United Kingdom, where all odds are set by bookmakers.

[16] We estimate these models following Snowberg and Wolfers (2010). In equilibrium, agents wager until the expected utility of a $1 bet equals $1, or until $p_i U(O_i + 1) = 1 \ \forall i \in \mathcal{R}$ in the risk-loving model. For the probability-weighting model, the equilibrium condition is $\pi(p_i)(O_i + 1) = 1 \ \forall i \in \mathcal{R}$. Each model is governed by one parameter. For the risk-loving model, they use the CARA utility function $U(x) = \big(1 - \exp(-\alpha x)\big)/\alpha$, where α modulates the agent's risk tolerance. For the probability-weighting model,

Figure 2: Expected returns for win bets, with 95% confidence intervals. The horizontal line marks the average takeout.

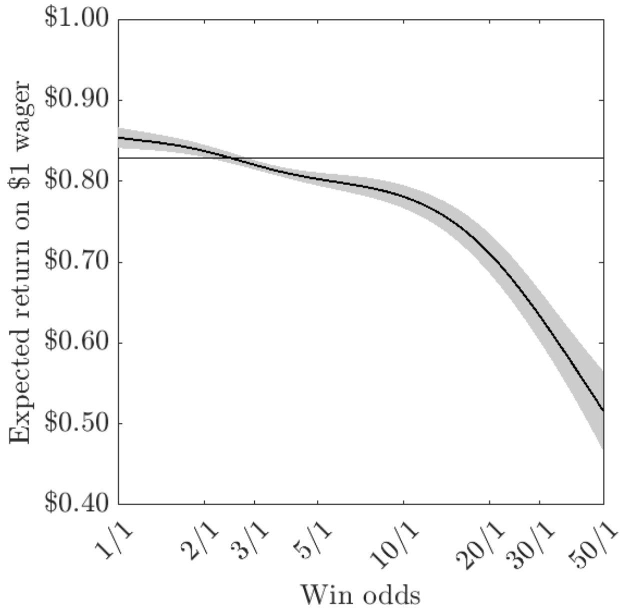

Note: Estimated using a local linear regression with a Gaussian kernel. The bandwidth, of log(1.65), minimizes the leave-one-out mean-squared error.

for favorites exceed observed returns. The difficulty is that neither class of models can rationalize highly negative returns for probable events. Risk-loving agents pay a large premium for lottery tickets, but they pay a small premium to gamble on likely events. Prospect-theory agents overweight large probabilities and thus need to *receive* a premium in order to wager on likely events. The decision to gamble is commonly rationalized by locally risk-loving preferences (e.g., Thaler and Ziemba, 1988) or probability weighting (e.g., Barberis, 2012), but

they use the weighting function $\pi(p) = \exp\big[-\big(-\log(p)\big)^{\beta}\big]$, where β modulates the degree to which the agent overweights small probabilities and underweights large ones (Prelec, 1998). We estimate each model by minimizing the squared distance between observed returns and expected returns, $p_i(O_i + 1)$, where O_i are the observed parimutuel odds, and p_i can be found by solving the appropriate equilibrium condition. In the risk-loving model, we estimate $\hat{\alpha} = -0.029$ (0.001), implying an extreme taste for risk. This representative agent is indifferent between \$100 for sure and a gamble that pays \$123 with 50% probability and \$0 otherwise. In the probability-weighting model, we estimate $\hat{\beta} = 0.900$ (0.003), implying overweighting of small probabilities. This representative agent behaves as if an event with 1.0% probability occurs 1.9% of the time. These estimates are more slightly more extreme than those of Snowberg and Wolfers (2010), who estimate $\hat{\alpha} = -0.017$ and $\hat{\beta} = 0.928$. Snowberg and Wolfers (2010) use these estimates to predict the odds for exotic bets, which they compare to the observed odds in the Jockey Club database. They find that both models have large prediction errors, though the errors are smaller for the probability-weighting model. Unfortunately, we cannot replicate this analysis using our data, as odds for exotic bets are not listed on the race chart.

Figure 3: Predicted returns for win bets from models with non-standard agents.

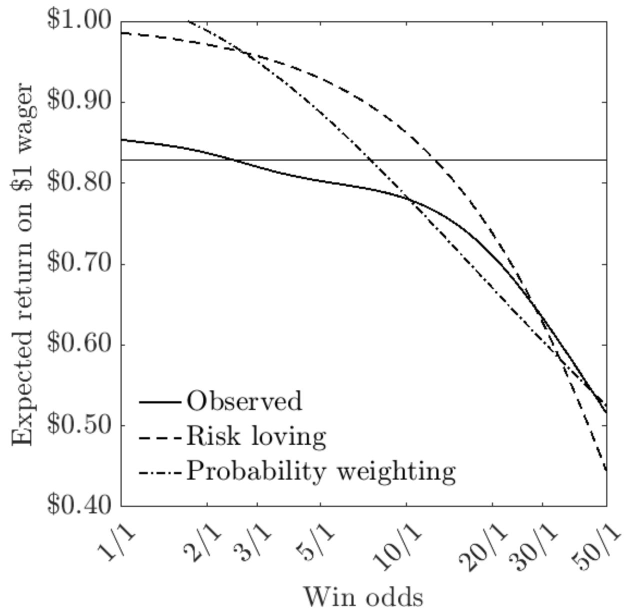

Note: Estimated using a local linear regression with a Gaussian kernel. For comparability, we use the same bandwidth for the predicted estimates as for the observed estimates: the bandwidth used in Figure 2, of log(1.65).

neither explanation can account for a willingness to lose considerable sums, in expectation, on gambles that are likely to pay out.[17]

These models also fail to predict heterogeneity in the favorite-longshot bias across tracks. Figure 4a shows expected returns in the win pool as a smooth function of log final odds, separately by track. Some tracks exhibit little or no bias, whereas others exhibit severe bias. We summarize the extent of the favorite-longshot bias at a given track by its mean returns, μ, or the expected return on a \$1 wager for a bettor who picks horses at random. A random betting strategy overbets longshots relative to the market. Were horses priced efficiently, overbetting longshots would not incur a penalty. Instead, a random betting strategy would surrender the takeout on average, and μ would equal $1 - t$. When longshots yield lower returns than favorites, however, overbetting longshots generates excess losses. Hence, a steeper favorite-longshot bias implies lower mean returns, μ. Figure 4b shows normalized mean returns, $\mu/(1 - t)$, by track. The tracks are ordered by their estimated

[17]These explanations were more sensible decades ago, when favorites were even-money bets, or close to it (Thaler and Ziemba, 1988). Our model suggests an explanation for why the favorite-longshot bias has moderated over time: the emergence of off-track betting.

Figure 4: Returns for win bets, by track. In (b), tracks are sorted by the estimated expected returns at 30/1 odds in (a).

(a) Expected returns **(b)** Normalized mean returns

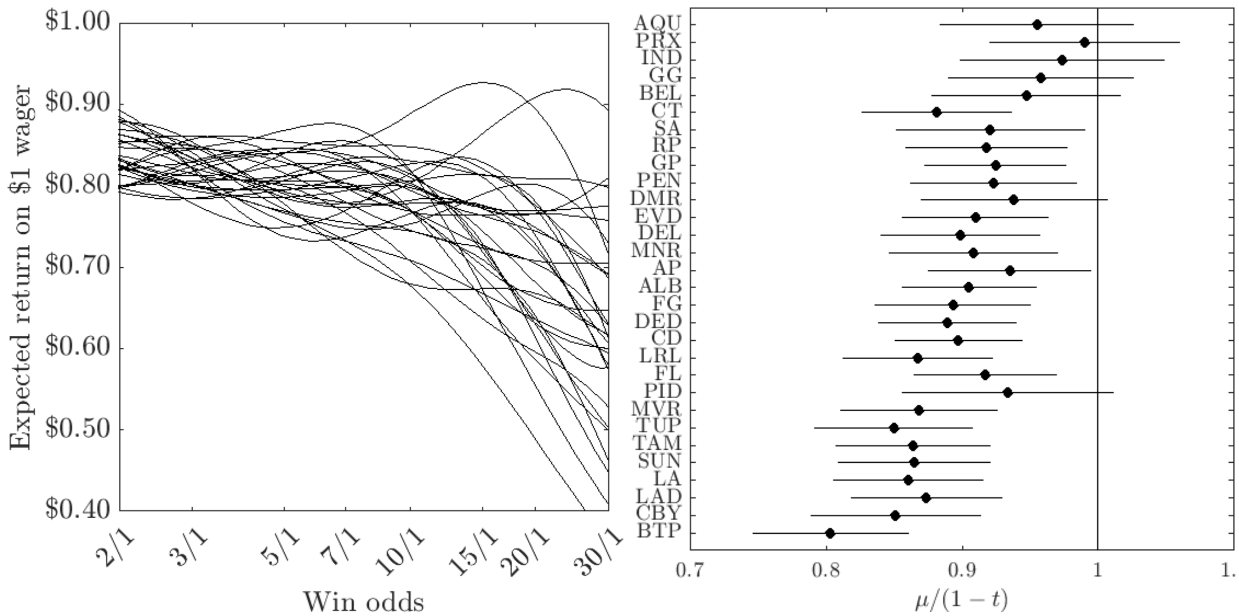

Note: Track-specific expected returns in (a) estimated as in Figure 2. For comparability, we use the same bandwidth as in the full sample, of log(1.65).

returns at win odds of 30/1, as shown in Figure 4a. For some tracks, expected returns are approximately flat and normalized mean returns are close to 1, implying that wagering randomly does not generate excess losses. For the remaining tracks, estimated returns are greater for favorites than for longshots, and random wagering generates excess losses. These differences across tracks cannot be explained solely by sampling variation. A Wald test rejects the null hypothesis that normalized mean returns, $\mu/(1-t)$, are equivalent at all tracks ($p < .001$).

The favorite-longshot bias appears to be related to the morning-line odds. One immediate observation is that morning-line odds fail at their ostensible goal of predicting final odds. Figure 5 shows average final odds at each observed morning-line odds. Morning-line odds are compressed—on average, those shorter than 4/1 are insufficiently short, and those longer than 4/1 are insufficiently long. For example, favorites assigned morning-line odds of 1/1 finish with final odds of 1/2 on average, and longshots assigned morning-line odds of 30/1 finish with final odds of 50/1 on average. Were oddsmakers truly trying to predict final odds,

Figure 5: Average log final odds at each morning-line odds, with 95% confidence intervals.

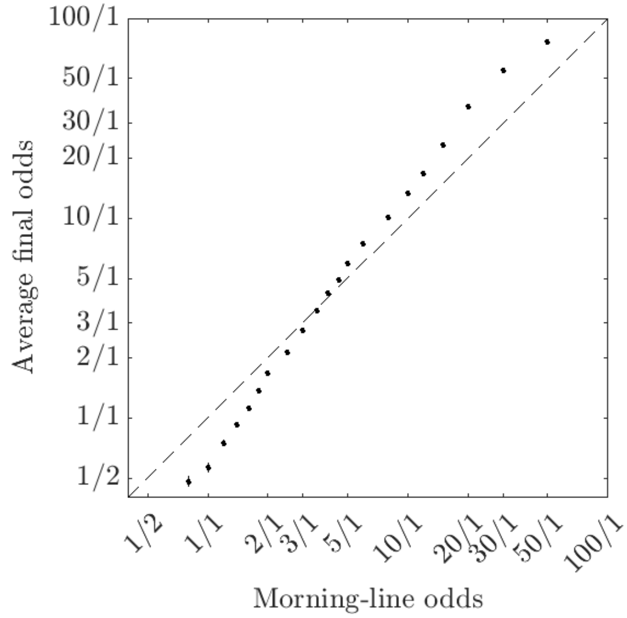

they could do better by predicting more extreme values.

Final odds diverge from morning-line odds just before the start of the race. Figure 6 shows a time series of live odds from a separate dataset of 1,362 US races, depicting the average log ratio of parimutuel odds to morning-line odds in the hour before the race, separately for the favorite and the longshot (as defined by the morning-line odds). On average, the parimutuel odds for the favorite and longshot hover around their respective morning-line odds until about 15 minutes before the start of the race. Thereafter, the longshot's odds lengthen and the favorite's odds shorten, implying late wagering on favorites (Asch, Malkiel and Quandt, 1982; Camerer, 1998).

The morning-line odds not only mispredict the final odds—they also imply distorted beliefs about a horse's chances of winning. Consider a win pool in which risk-neutral bettors wager until the final odds converge to the morning-line odds (denoted by l). In equilibrium, the expected value of betting on any horse—given subjective beliefs q—must be the same for all horses in the race. That is, $q_i(l_i + 1) = q_j(l_j + 1)\ \forall i, j \in \mathcal{R}$, where $\mathcal{R}$ is the set of horses in a race, and $l + 1$ is the return on a winning \$1 wager. Thus,

$$q_i = \frac{1}{l_i + 1} \sum_j \frac{1}{l_j + 1} \tag{1}$$

The subjective beliefs implied by the morning-line odds are simply the inverse of the associ-

Figure 6: Time series of the log ratio of live parimutuel odds to morning-line odds, separately for the favorite and longshot (as defined by the morning lines).

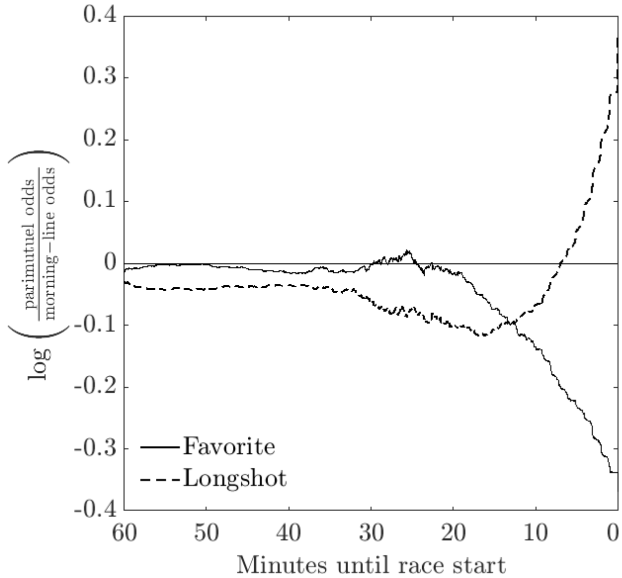

Note: These estimates reflect second-by-second averages over 1,362 races in 2017 and 2018 for which at least 20 time-stamps were observed during the last 60 minutes of the betting window, and for which the morning-line odds uniquely define a favorite and a longshot. We interpolate missing time-stamps by assigning the most recently observed live odds.

ated returns on a \$1 wager, normalized such that these beliefs sum to 1.[18]

These beliefs are misleading. Figure 7 shows a calibration plot—the observed win rate (denoted by p) as a smooth function of q, the win probabilities implied by the morning-line odds.[19] Longshots underperform their implied win probabilities, q. At $q = 0.03$, for example,

[18] In general, the inverse implied returns do not sum to 1. This happens because some horses are assigned morning-line odds but do not start the race (i.e., scratches and alternate entries); because morning-line odds are chosen from discrete bins; and because oddsmakers do not perfectly account for the takeout and breakage.

[19] We are interested in estimating a smooth function, $p(q)$, that measures the rate at which horses with an implied probability of winning, q, actually win the race. To do so, we regress an indicator for whether the observed horse won the race on $\log q$, its log probability of winning, as implied by the morning-line odds and as calculated in (1). In particular, we use the Nadaraya-Watson estimator, otherwise known as local constant regression, which unlike a local linear or local polynomial regression, ensures that the predicted values stay between the bounds of the outcome variable—i.e., between 0 and 1. The challenge is that we would like $p(q)$ to be weakly increasing, reflecting an assumption that the morning lines correctly order the horses by their chances of winning, even if they imply poorly calibrated beliefs. To do so, we follow Hall and Huang (2001) in assigning each observation a weight between 0 and 1, and, via quadratic programming, minimizing the Kullback–Leibler divergence of the weight vector, subject to the constraints that 1) the function $p(q)$, estimated at every half-percentage-point value of q, is weakly increasing; and 2) the weights sum to 1. This

Figure 7: Observed win rate, p, as a function of the win probability implied by the morning-line odds, q, with a histogram of q.

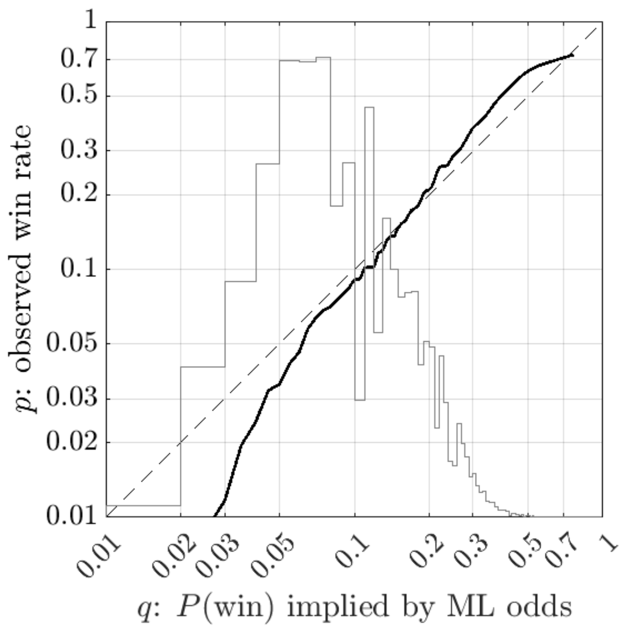

horses win just 1 in 100 races; at $q < 0.03$, horses effectively never win. Morning-line odds of 30/1 typically generate subjective beliefs around $q = 0.03$. Of the 7.869 starts assigned 30/1 morning-line odds, just 83, or 1.1%, won the race. Morning-line odds of 50/1 imply win probabilities around $q = 0.02$. Of the 588 starts assigned morning-line odds in excess of 30/1, just 2, or 0.3%, won the race. Symmetrically, favorites outperform their implied win probabilities. At $q = 0.5$, for example, $p = 0.63$.

The morning-line odds mislead at some tracks but not at others. Figure 8 replicates the estimate of p separately by track. Some tracks post well-calibrated morning-line odds (e.g., PRX), whereas others post misleading morning-line odds (e.g., LA). Tracks that promulgate distorted predictions deceive bettors in the same manner—by assigning insufficiently long morning lines to longshots and insufficiently short morning lines to favorites. In other words, these tracks embed a favorite-longshot bias in the morning-line odds.

routine does not guarantee convergence. For instance, if a lone horse with a high implied win probability q did not win its race, the monotonicity constraint will force the weight on this observation to 0, sending the Kullback–Leibler divergence to infinity. To avoid this pitfall, we employ a bandwidth $h \cdot \lambda(x)$, where h is a scalar bandwith, and λ is a local bandwidth factor that increases the bandwidth in regions of low density. This multiplicative factor is $\lambda(x) = \sqrt{\exp\left(\frac{1}{N}\sum_i \log \tilde{f}(\log q_i)\right)/\tilde{f}(x)}$, where $\tilde{f}$ is a kernel density estimate of $\log q$ using Silverman's rule-of-thumb bandwidth. The steps are as follows. We find the bandwidth, h, that minimizes the leave-one-out mean-squared error. Fixing the bandwidth, we then find the vector of minimally divergent weights that satisfy the weak monotonicity condition.

Figure 8: Observed win rate (p) as a function of the win probability implied by the morning-line odds (q), by track. The gray outline shows a histogram of q.

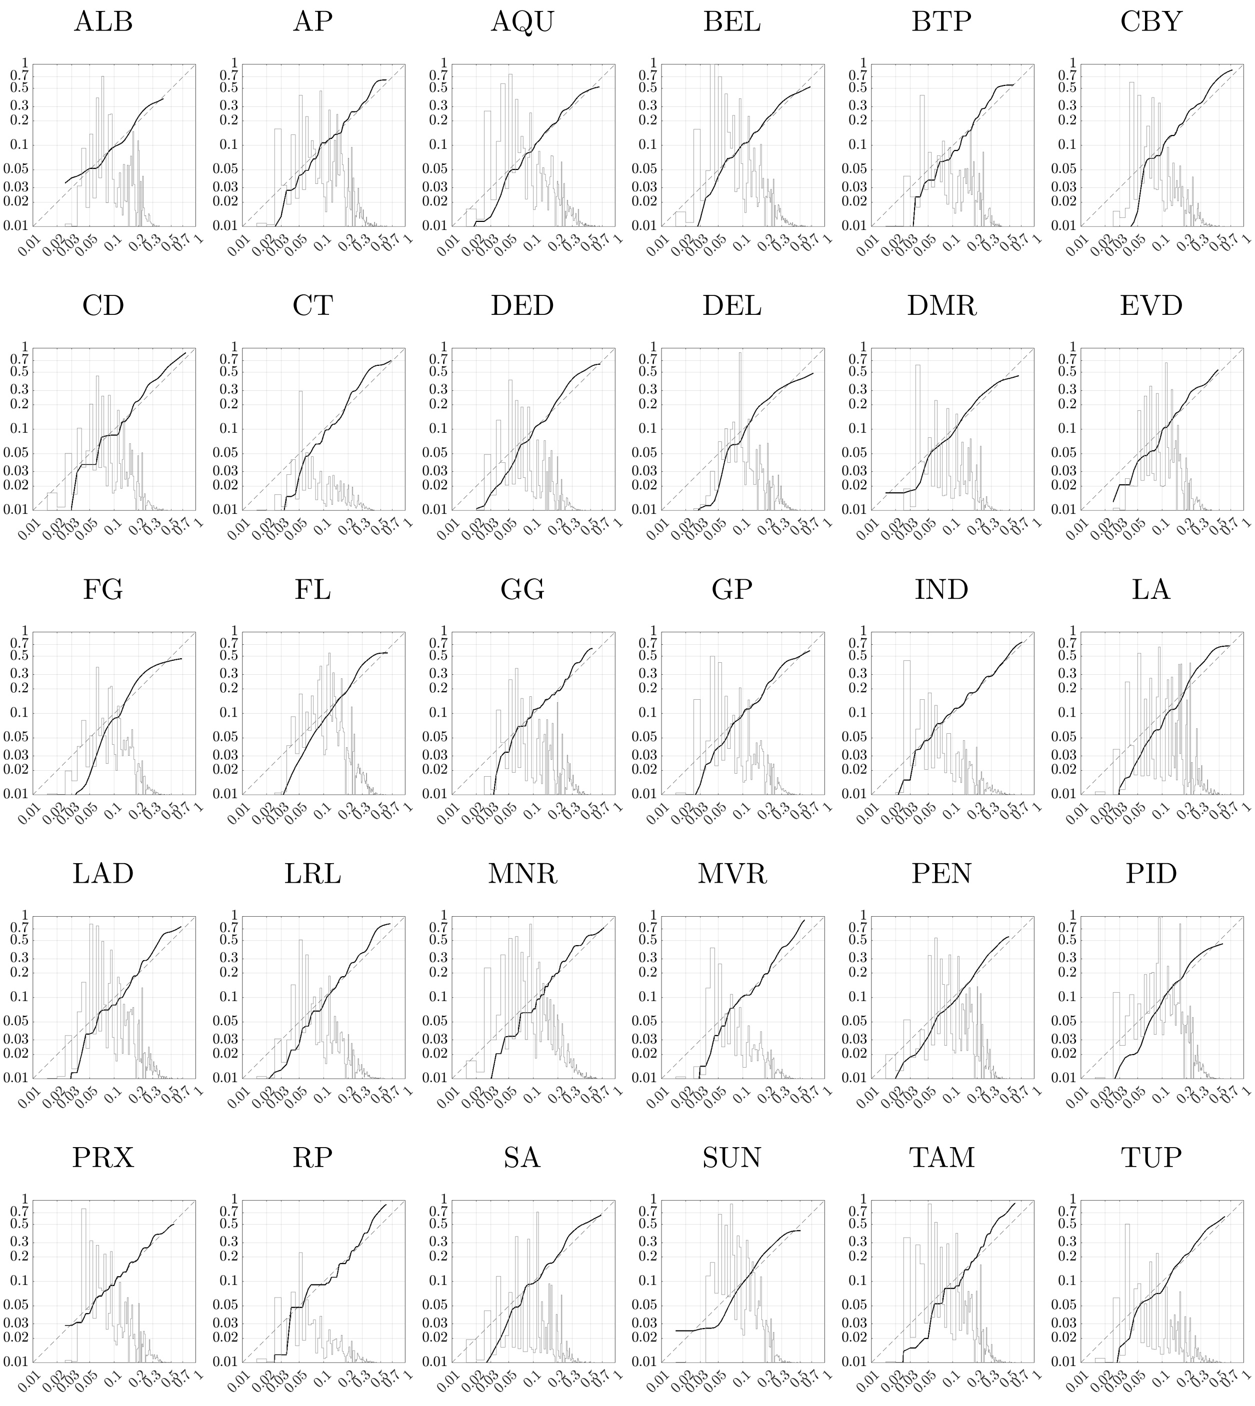

Figure 9: Scatter plot, with a regression fit line, of miscalibration in the morning-line odds, δ, and normalized mean returns for win bets, $\mu/(1-t)$.

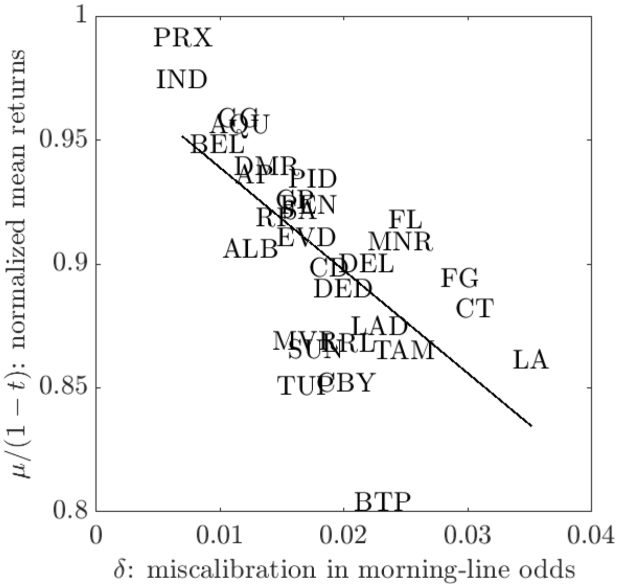

Across tracks, miscalibration in the morning-line odds predicts the favorite-longshot bias. To see this, define δ as the expected absolute difference between p and q, or $\int_0^1 |p-q|dF(q)$. In other words, δ measures miscalibration in the morning-line odds. On average, a randomly sampled horse will have an implied win rate, q, that differs by δ from the observed win rate, p, of horses with similar implied win rates. Second, let the normalized mean returns, or $\mu/(1-t)$ in Figure 4b measure the extent of the favorite-longshot bias. Figure 9 shows a strong correlation between δ and $\mu/(1-t)$ across tracks. The more severe the miscalibration in the morning-line odds, the more severe the favorite-longshot bias.

5 Theoretical framework

These stylized facts, along with the contextual background, motivate a new explanation for the favorite-longshot bias: it is an artifact of profitable deception by the track. We write down a model in which the track's gambling revenue is increasing in the disagreement among bettors, and we imbue those bettors with beliefs estimated from the morning-line odds. In contrast with past approaches, bettors in our model are neither speculative nor irrational.

5.1 Setup

Consider a two-period model with two risk-neutral representative agents: track bettors and off-track bettors. Track bettors gamble in the first period; off-track bettors gamble in the second. Track bettors gain direct utility from gambling; off-track bettors do not. Track bettors pay the full tax; off-track bettors receive a rebate that defrays part of the tax. Track bettors and off-track bettors differ in their beliefs, but neither type suffers from psychological biases in interpreting probabilities.

We analyze a generic betting pool with N outcomes, indexed by i (e.g., a win pool with N horses). A large number of (potential) track bettors each consider placing a \$1 wager on one outcome (e.g., Gandhi and Serrano-Padial, 2014). These track bettors share the same beliefs about the probability of each outcome obtaining, denoted q_i, but they receive different amounts of direct utility from gambling, u. Let $O_i^{(1)}$ denote the equilibrium odds for outcome i in the first period. The utility of a \$1 wager on outcome i is $q_i(O_i^{(1)} + 1) + u$—i.e., the expected value of the gamble under their beliefs, plus the utility shock. In equilibrium, only track bettors with $u \geq t$ choose to gamble, and the parimutuel odds are:

$$O_i^{(1)} = (1 - t) \cdot \frac{1}{q_i} - 1$$

The share of bets placed on outcome i is q_i. Hence, track bettors wager in proportion to their beliefs. We normalize the total amount wagered by track bettors to 1, which implies that track bettors wager q_i on outcome i.

In the second period, off-track bettors place bets of size x_i. The total amount bet is $1 + \sum x$, and the second-period parimutuel odds are:

$$O_i^{(2)} = (1 - t) \cdot \frac{1 + \sum x}{q_i + x_i} - 1$$

Off-track bettors possess subjective beliefs, p_i. They also receive a rebate, r, on each dollar they bet, win or lose. They do not receive any direct utility from gambling. Hence, the expected value of wagering x_i under their beliefs is:

$$\mathbb{E}_{\text{off-track}}[x_i] = x_i \cdot \left(p_i \cdot (1 - t) \cdot \frac{1 + \sum x}{q_i + x_i} + r \right) \tag{2}$$

As a monopolist, the track chooses the rebate, r, that maximizes its gambling revenue:

$$\pi = t \cdot 1 + (t - r) \cdot \sum x, \tag{3}$$

where the first and second terms are the track's revenues from track and off-track bettors, respectively. A higher rebate increases wagering by off-track bettors ($\sum x$), at the expense of lowering the track's effective tax rate on those wagers ($t - r$).

5.2 Results

We consider a competitive equilibrium among off-track bettors. This equilibrium consists of a set of wagers by off-track bettors, x_i, such that 1) $\mathbb{E}_{\text{off-track}}[x_i] = x_i$ for $x_i > 0$—i.e., off-track bettors make zero profits; and 2) $\mathbb{E}_{\text{off-track}}[x_i] \leq x_i$ for $x_i = 0$—i.e., no excess profits are left on the table. We first show that an equilibrium exists.

Lemma. *Reorder the outcome indices i such that:*

$$\frac{p_1}{q_1} \geq \frac{p_2}{q_2} \geq \cdots \geq \frac{p_i}{q_i} \geq \frac{p_{i+1}}{q_{i+1}} \geq \cdots \geq \frac{p_N}{q_N}$$

An equilibrium exists iff $\exists m$ such that:

$$\frac{p_m}{q_m} > \lambda_m \geq \frac{p_{m+1}}{q_{m+1}}, \tag{4}$$

where $\lambda_m \equiv \sqrt{\frac{P_m(1-P_m)}{Q_m(1-Q_m)}}$, $P_m \equiv \sum_i^m p_i$, and $Q_m \equiv \sum_i^m q_i$. Under this equilibrium, $x_i > 0$ for $i \leq m$ and $x_i = 0$ for $i > m$.

Proof. In Appendix A.

The index m separates outcomes that off-track bettors wager on from those they do not. P_m and Q_m are the cumulative subjective probabilities—held by off-track and track bettors, respectively—among outcomes wagered by off-track bettors. In equilibrium, off-track bettors wager on outcomes for which their own beliefs are high relative to those of track bettors.

Proposition. *If $p_1 > q_1$, then $\exists m \in \{1, \ldots, N-1\}$.*

Proof. In Appendix A.

An equilibrium exists so long as subjective beliefs do not coincide for every outcome—i.e., if $p_1 > q_1$ (and hence, $p_N < q_N$). If so, $m \geq 1$, implying that off-track bettors always wager on the outcome with the highest ratio of subjective beliefs. In addition, $m < N$, implying that off-track bettors never wager on the outcome with the lowest ratio of subjective beliefs.

Equilibrium wagers by off-track bettors are:

$$x_i^* = \begin{cases} p_i/\lambda_m - q_i, & i \in \{1, \ldots, m\} \\ 0, & i \in \{m+1, \ldots, N\} \end{cases} \tag{5}$$

The amount off-track bettors wager is increasing in their own subjective beliefs, p_i, and decreasing in those of track bettors, q_i, all else equal. Off-track bettors collectively wager:

$$\sum_i x_i^* = P_m/\lambda_m - Q_m \tag{6}$$

This implies:

Corollary. $P_m > Q_m$.

Proof. Since $x_1^* > 0$, it follows that $\sum_i x_i^* > 0$, and from (6) that $P_m > Q_m$.

In equilibrium, off-track bettors wager on a set of outcomes which they believe to be collectively more likely than do track bettors.

Let r^* be the revenue-maximizing rebate. Then,

$$\frac{t - r^*}{1 - t} = (1 - Q_m) \cdot \lambda_m - (1 - P_m) \tag{7}$$

The effective tax paid by off-track bettors, $t - r^*$, is a fixed fraction of $1 - t$, the share of the pot split among winning wagers.

The track's gambling revenue is:

$$\pi(r^*) = t + (1 - t) \cdot \left[\sqrt{P_m \cdot \big(1 - Q_m\big)} - \sqrt{Q_m \cdot \big(1 - P_m\big)}\right]^2, \tag{8}$$

where t is the track's revenue from track bettors, and the second term is the track's revenue from off-track bettors, equivalent to $(t - r^*) \sum_i x_i^*$. If off-track bettors hold accurate beliefs, these two terms can instead be interpreted, respectively, as the state-regulated losses sustained by track bettors and their losses in excess of the tax. By assumption, off-track bettors make zero profits. Hence, the gambling revenues collected by the track equal the total losses sustained by track bettors, which exceed t. Because arbitrage is competitive, excess losses incurred by track bettors flow in their entirety to the track.

The track's revenues from off-track bettors are increasing in the distance between P_m and Q_m. If $P_m = Q_m$, which occurs only when $p_i = q_i$ for all i, off-track bettors place zero wagers, the track offers no rebate, and the track makes no money from off-track bettors. Deviations between p_i and q_i, and hence between P_m and Q_m, generate bets from arbitrageurs and profits for the track.

Our model is agnostic about how p_i and q_i diverge. Nonetheless, it is straightforward to see how a favorite-longshot bias could arise. Imagine that $p_1 \geq p_2 \geq \cdots \geq p_N$—i.e., relative to track bettors, off-track bettors are bullish on favorites and bearish on longshots. Further, assume that track bettors and off-track bettors agree on how to order the outcomes by their likelihood, even if they disagree on the exact probabilities. That is, $q_1 \geq q_2 \geq \cdots \geq q_N$. Finally, assume that the beliefs of off-track bettors, p_i, are accurate.

First-period wagering generates a favorite-longshot bias. Expected returns, which are proportional to p_i/q_i, are largest for favorites and smallest for longshots. Second-period wagering concentrates on favorites, moderating the favorite-longshot bias. The final odds are:

$$O_i^{(2)} = \begin{cases} \frac{1-t}{p_i} \cdot \left[P_m + (1 - Q_m) \cdot \lambda_m\right] - 1, & i \in \{1, \ldots, m\} \\ \frac{1-t}{q_i} \cdot \left[P_m/\lambda_m + (1 - Q_m)\right] - 1, & i \in \{m+1, \ldots, N\} \end{cases} \tag{9}$$

Wagering by off-track bettors on favorites ($i \leq m$) lengthens odds on longshots ($i > m$) and, in turn, shortens odds for some favorites—in particular, those for which $p_i/q_i > P_m + (1 - Q_m)\lambda_m$. At the track, the expected value of a \$1 bet, given accurate beliefs p_i, is:

$$\mathbb{E}_{\text{track}}[x_i = 1] = (1 - t) \cdot \begin{cases} P_m + (1 - Q_m) \cdot \lambda_m, & i \in \{1, \ldots, m\} \\ \frac{p_i}{q_i} \cdot \left[P_m/\lambda_m + (1 - Q_m)\right], & i \in \{m+1, \ldots, N\} \end{cases} \tag{10}$$

Second-period betting equalizes returns for the first m favorites. A \$1 wager on any of the first m outcomes returns \$1 to off-track bettors and $\$1 - r^*$ at the track. Wagers on less likely outcomes perform worse, which follows from (4). The favorite-longshot bias manifests as a piecewise relationship similar to that observed in Figure 2, with the kink located at the odds of outcome m. For outcomes with shorter odds (i.e., $i < m$), expected returns are flat. For outcomes with longer odds (i.e., $i > m$), expected returns are decreasing.

5.3 Estimation

Beliefs in the win pool are derived from the morning-line odds. We infer the beliefs of track bettors by assuming that they take the morning-line odds at face value—i.e, as a well-calibrated, or unbiased, prediction of the final parimutuel odds in the win pool were all bettors risk neutral like themselves. Specifically, track bettors form beliefs $q_{i,\text{win}}$ as in (1)—i.e., by inverting the morning lines and then normalizing by race such that $\sum_{i\in\mathcal{R}} q_{i,\text{win}} = 1$.[20] In contrast, we imbue off-track bettors with well-calibrated beliefs, $p_{i,\text{win}}$. Specifically, we non-parametrically estimate the track-specific rates at which horses associated with similar naive beliefs, $q_{i,\text{win}}$, actually win the race, as shown in Figure 8 and described in Footnote 19. We then normalize these estimates by race such that $\sum_{i\in\mathcal{R}} p_{i,\text{win}} = 1$.[21] Whereas track bettors assume that the morning lines imply well-calibrated beliefs, off-track bettors ensure that their beliefs are well calibrated.

We use beliefs in the win pool to construct beliefs in other pools. To do so, we assume that the probability of a horse finishing in n^{th} place is simply its probability of winning the race, divided by the cumulative winning probability among horses that did not finish in the first $n-1$ positions (Harville, 1973). The race for second place, for instance, can be thought of as a race within a race in which the first-place finisher does not participate.

We first consider place and show wagers, which pay out if the chosen horse finishes in the top 2 or 3 positions, respectively. The probability of a place wager on horse i paying out is:

$$p_{i,\text{place}} = p_{i,\text{win}} + \sum_{j\neq i} p_{j,\text{win}} \cdot \frac{p_{i,\text{win}}}{1 - p_{j,\text{win}}}$$

where the second term is the probability of horse i finishing second. Beliefs for track bettors, $q_{i,\text{place}}$, can be calculated in a corresponding manner. Similarly, the probability of a show wager on horse i paying out is:

$$p_{i,\text{show}} = p_{i,\text{place}} + \sum_{j\neq i} \sum_{k\neq(i,j)} p_{j,\text{win}} \cdot \frac{p_{k,\text{win}}}{1 - p_{j,\text{win}}} \cdot \frac{p_{i,\text{win}}}{1 - p_{j,\text{win}} - p_{k,\text{win}}}$$

where the second term is the probability of horse i finishing third.

We also consider a range of exotic wagers. Bettors may wager on the order of the first

[20] In the win pool, track bettors need not form these beliefs explicitly. They behave as if maximizing expected value given $q_{i,\text{win}}$ by simply wagering on the horse with the highest ratio of morning-line odds to current parimutuel odds.

[21] Because this routine is computationally intensive, we do not repeat it during bootstrapping. Instead, we use the same estimates of p_i (and q_i) in every resampled race.

n horses in a single race, as in exacta and trifecta pools, which we denote by order-n. The probability of a sequence, $\vec{v}$, is:

$$p_{\text{order-}n}(\vec{v}) = p_{\vec{v}_1,\text{win}} \times \frac{p_{\vec{v}_2,\text{win}}}{1 - p_{\vec{v}_1,\text{win}}} \times \cdots \times \frac{p_{\vec{v}_n,\text{win}}}{1 - p_{\vec{v}_1,\text{win}} - \cdots - p_{\vec{v}_{n-1},\text{win}}}$$

A variant of the exacta (i.e., order-2) is the quinella, in which the bettor predicts the first two horses regardless of order. Assuming conditional independence, the probability of a quinella bet on (i, j) paying out is:

$$p_{\text{quin}}(i, j) = p_{i,\text{win}} \cdot \frac{p_{j,\text{win}}}{1 - p_{i,\text{win}}} + p_{j,\text{win}} \cdot \frac{p_{i,\text{win}}}{1 - p_{j,\text{win}}}$$

This is equivalent to the probability that an exacta bet on either (i, j) or (j, i) pays out.

Bettors may also wager on the winner of n consecutive races, as in daily-double and pick-3 pools, which we denote by pick-n. Assuming independence between races, the probability of a sequence, $\vec{v}$, is:

$$p_{\text{pick-}n}(\vec{v}) = \prod_{i=1}^{n} p^{(i)}_{\vec{v}_i,\text{win}}$$

where the superscript indexes the race.

Finding the equilibrium in each pool proceeds according to the model. We order outcomes by p_i/q_i, from greatest to smallest. Using a grid search, we find the index m that maximizes the track's gambling revenue in (8).[22] For each outcome, we calculate the second-period equilibrium odds (9) and corresponding expected returns (10). In Section 6, we compare the relationship between equilibrium odds and expected returns predicted by our model to that observed in the data. For reference, we also depict the favorite-longshot bias implied by first-period equilibrium odds and expected returns. We then calculate the optimal rebate from (7), which unlike the takeout is unobserved, and track revenues under the optimal rebate (8), which we report in Section 7.[23]

[22] In other words, we consider the equilibrium under the globally revenue-maximizing rebate. In theory, many equilibria may exist, one for every locally revenue-maximizing rebate. In the data, 37% of pools have more than one equilibrium. Multiple equilibria are most common for pick-n bets, occurring in 50% of such pools. They are least common for win, place, and show bets, occurring in 32% of those pools.

[23] The expressions for equilibrium quantities are slightly different in place and show pools, where more than one outcome obtains. Let k denote the number of outcomes that pay out—i.e., $k = 2$ in place pools and $k = 3$ in show pools—with $\sum_i q_i = \sum_i p_i = k$. Observe that the parimutuel pays out $1/k^{th}$ of the post-tax

5.4 Example

Table 2 illustrates the estimation routine for an example race at Charles Town. T Rex Express, the favorite with 1/1 morning-line odds, finished with final odds of 3/10 on the totalizer and in first place on the track. As a result, win bets on T Rex Express paid out a divided of 30 cents on every dollar wagered. Relative to the morning lines, parimutuel odds lengthened for the other six horses in the race.

Table 2: Example race at Charles Town.

	Odds		Beliefs			Model predictions in win pool				
Name	M/L	Final	q_i	p_i	p_i/q_i	$O_i^{(1)}$	$\mathbb{E}_i^{(1)}$	x_i	$O_i^{(2)}$	$\mathbb{E}_i^{(2)}$
T Rex Express	1/1	0.3	.41	.52	1.26	1.0	1.05	0.6	0.6	0.85
Tribal Heat	2/1	4.0	.27	.30	1.11	2.0	0.91	0.3	1.8	0.85
Dandy Candy	6/1	6.9	.12	.09	0.78	6.0	0.64	0.06	8.3	0.85
Click and Roll	8/1	15.6	.09	.06	0.64	8.0	0.53	0.02	13.5	0.85
Eveatetheapple	20/1	69.3	.04	.01	0.34	20.1	0.28	0	40.6	0.55
Movie Starlet	20/1	25.5	.04	.01	0.34	20.1	0.28	0	40.6	0.55
Boston Banshee	30/1	49.7	.03	.00	0.00	30.2	0.00	0	60.4	0.00

The beliefs ascribed to track bettors, q_i, are those that a representative risk-neutral bettor would hold if the final odds coincided with the morning-line odds. Specifically, they are inversely proportional to the returns implied by the morning-line odds, up to a normalizing constant, as in (1). By contrast, the beliefs ascribed to off-track bettors, p_i, are well calibrated. In particular, they are proportional to the track-specific rates at which horses with implied beliefs q_i actually win, as shown in Figure 8, up to a normalizing constant. For Charles Town, beliefs of $q_i > 0.2$ are too pessimistic on average, and beliefs of $q_i < 0.2$ are too optimistic on average. As a result, $p_i > q_i$ for the T Rex Express ($q = 0.41$) and for Tribal Heat ($q = 0.27$), and $p_i < q_i$ for the other 5 horses.

The first-period odds in the win pool, $O_i^{(1)}$, reflect betting by risk-neutral track bettors, given beliefs q_i. They diverge from the morning-line odds only because the morning-line odds do not precisely reflect the takeout, of 17.25%, and the breakage, or the practice of rounding

pot to winning wagers. Hence, second-period odds can be written as:

$$O_i^{(2)} = \frac{1-t}{k} \cdot \frac{1+\sum x}{q_i/k + x_i} - 1,$$

Modified expressions for the equilibrium odds, expected returns, optimal rebate, and track revenues can be derived in the same manner.

down odds to the nearest 10 cents. The first-period odds generate a severe favorite-longshot bias. The first-period expected returns on a \$1 wager—denoted $\mathbb{E}_i^{(1)}$ and taken over p_i—are sharply decreasing in the odds. At first-period odds of 1.0, the $p = 0.52$ favorite is better than an even money bet. By contrast, at odds of 20.1, the $p = 0.01$ longshots return just 28 cents on every dollar wagered—and the $p = 0$ longshot returns 0 at any odds.

The second-period equilibrium consists of a set of wagers by off-track bettors, x_i, such that these arbitrageurs make zero profits and leave zero profits on the table. In our model, the track finds the revenue-maximizing equilibrium by setting the optimal rebate, which determines wagering by off-track bettors. If rebates and wagers by off-track bettors were observable, we could assess whether these quantities are consistent with our model. Given that they are unobservable, we instead infer them from the model. Specifically, we sort the outcomes in descending order of p_i/q_i, and we find the index m that maximizes the track's gambling revenues in (8). In the example above, this ordering coincides with sorting the horses in ascending order of the morning-line odds—a product of the favorite-longshot bias embedded in the morning lines. The index $m = 4$ is optimal, implying that $x_i > 0$ for the first 4 horses and $x_i = 0$ for the remaining 3.[24] Second-period wagers concentrate on the favorite, and the predicted odds, $O_i^{(2)}$, shorten for the first two horses and lengthen for all others. A less severe favorite-longshot bias remains. For any of the first three horses, the expected returns on a \$1 wager, $\mathbb{E}_i^{(2)}$, is 85 cents, with the 15-cent expected loss equaling the optimal rebate. For the longshots, losses are lower than under the first-period odds. A \$1 wager on either of the horses with 20/1 morning lines, for instance, returns 55 cents at odds of 40.6, or 21 cents more than at odds of 20.1.

In total, off-track bettors wager nearly as much as track bettors in the win pool (i.e., $\sum_i x_i = 0.97$). Whereas the track collects 17 cents from every dollar wagered at the track, off-track bettors receive a 15-cent rebate, leaving the track with just 2 cents. As a result, revenues from off-track bettors comprise 12% of the track's gambling revenues in the win pool. Arbitrage accounts for larger shares of the track's gambling revenues in exotic pools. In the exacta pool, for instance, that share is 21%; in the superfecta pool, it is 41%. The compound beliefs attached to outcomes in exotic pools amplify differences in beliefs about win probabilities between track and off-track bettors and hence, taxes from arbitrageurs. If off-track bettors are more optimistic than track bettors about two favorites, they will be even more optimistic about an exacta wager on those horses. This logic implies larger

[24] This set includes 3 horses for which off-track bettors are more pessimistic than track bettors about its chances. Wagers on the favorite lengthen the odds on other horses, which makes them more appealing.

divergences between the cumulative beliefs P_m and Q_m, as measured by (8). In the win pool, for example, $P_m = 0.97$ and $Q_m = 0.90$, whereas in the superfecta pool, $P_m > 0.99$ and $Q_m = 0.79$.

6 Model predictions

We first evaluate predictions of a simplified version of our model with only track bettors. Thereafter, we consider predictions from the complete model with off-track bettors. Figure 10 shows a smoothed estimate of the relationship between log predicted odds and predicted returns for win bets, separately for the simplified version and for the complete model. (All predicted returns reflect expected returns at the track—i.e., before any rebate.) With only track bettors, predicted odds approximate the morning-line odds, which embed a favorite-longshot bias that is more severe than that observed in the data. Hence, the simplified model overpredicts returns for favorites and underpredicts returns for longshots.

Figure 10: Predicted returns for win bets, with only track bettors and with track and off-track bettors. The horizontal line marks the average takeout.

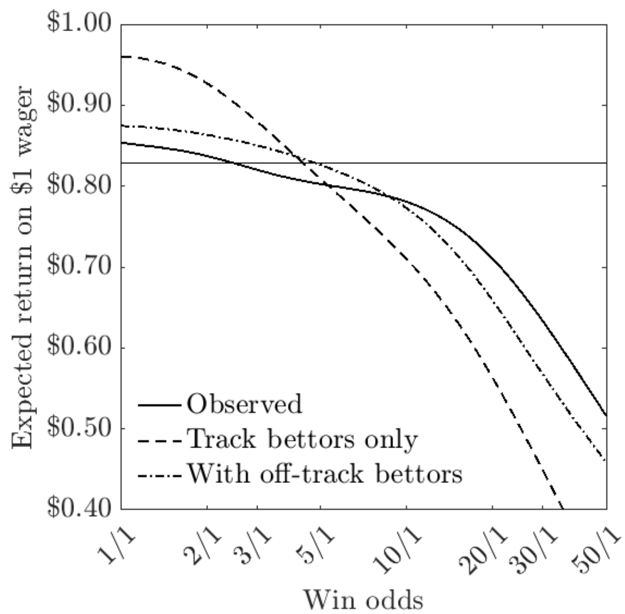

<u>Note</u>: Estimated using a local linear regression with a Gaussian kernel. For comparability, we use the same bandwidth for the predicted estimates as for the observed estimates: the bandwidth used in Figure 2, of log(1.65).

Figure 11: Observed (solid line) and predicted (dashed line) returns for win bets with 95% confidence intervals, by track. The horizontal line marks the state-sanctioned return, or $1-t$.

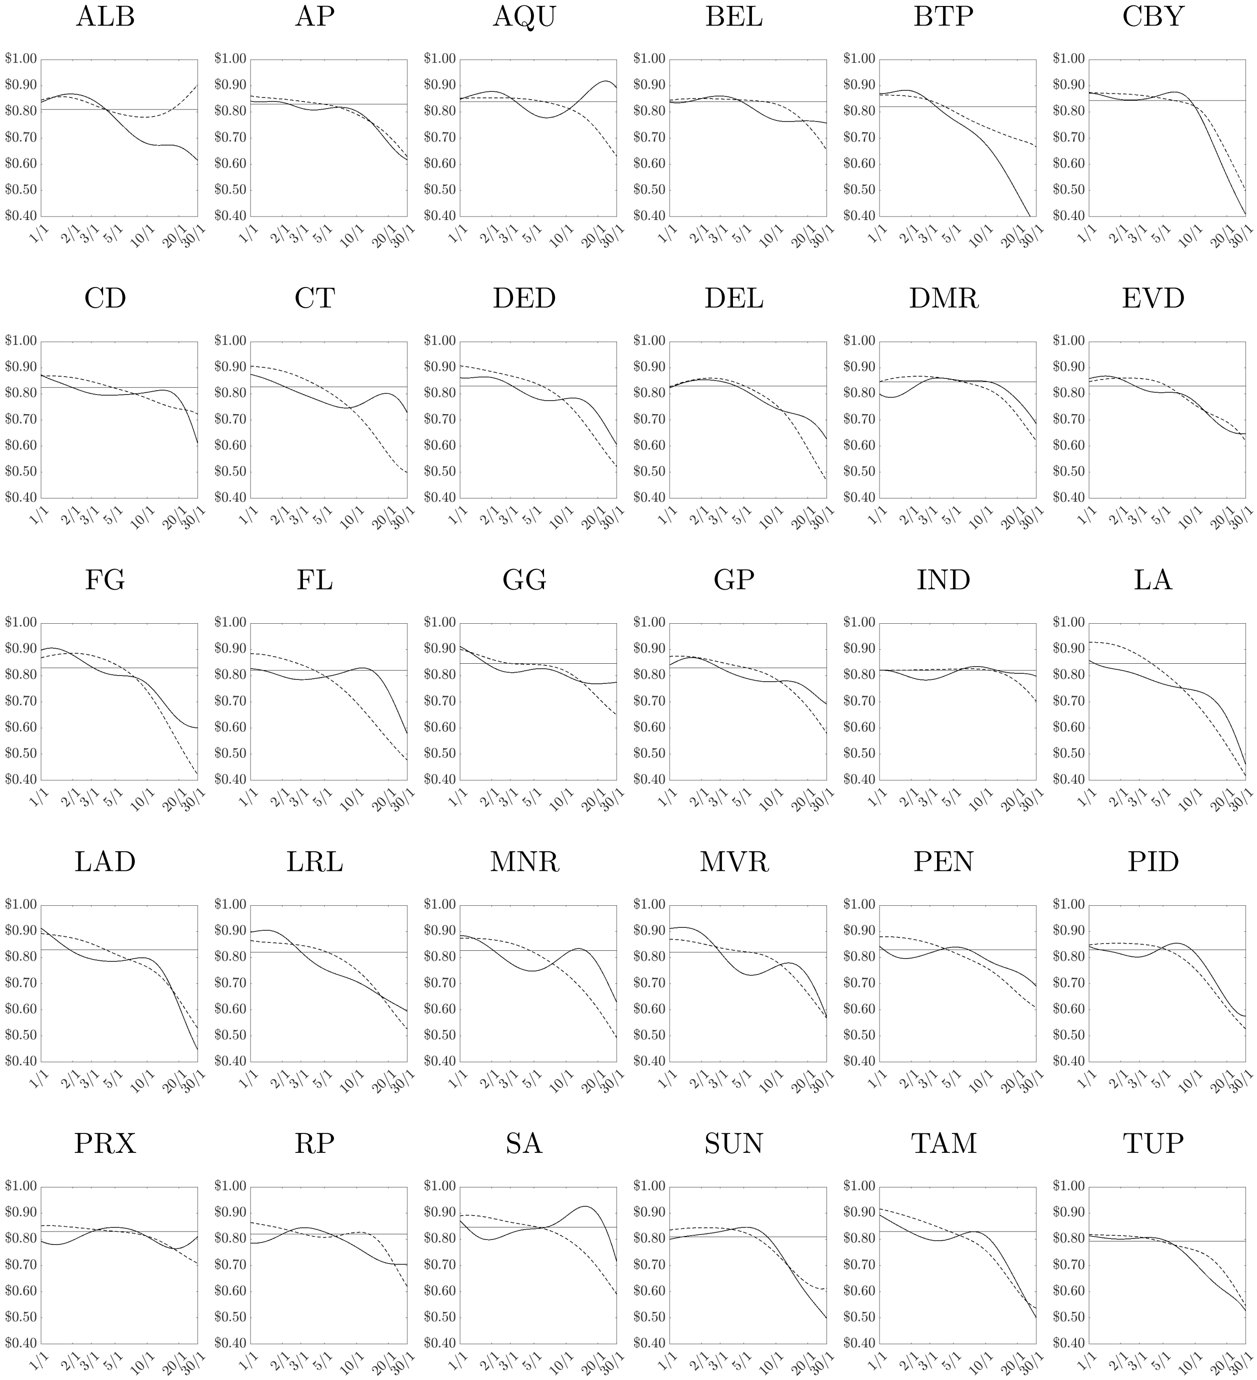

Note: Estimated using a local linear regression with a Gaussian kernel. For comparability, we use a common bandwidth for all estimates: the bandwidth used in Figure 2, of log(1.65).

The inclusion of off-track bettors moderates the favorite-longshot bias, and the resulting predictions more closely follow observed returns. For a 1/1 favorite, predicted returns exceed observed returns by just 2 cents, compared to 11 cents in the simplified model. For a 30/1 longshot, observed returns exceed predicted returns by 7 cents, compared to 19 cents in the simplified model. (All predictions that follow pertain to the complete model.)

We quantify the goodness of the model's predictions by measuring the expected absolute deviation, or the average distance between the predicted and observed lines in Figure 10, with the expectation taken over the distribution of observed odds.[25] For a bettor who wagers randomly, the absolute difference between her returns and those predicted by the model will, in expectation, equal our expected absolute deviation measure. This measure is just 3.1 (se: 0.5) cents for the complete model. Comparable measures for the representative-agent models in Figure 3 are larger: 9.2 (0.5) cents for the risk-loving model and 7.0 (0.4) cents for the probability-weighting model.

Figure 12: Normalized mean returns, observed and predicted.

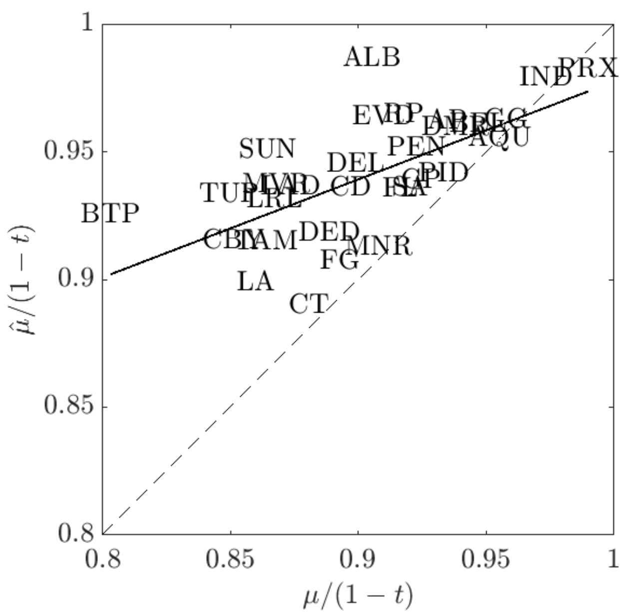

Our model also captures differences across tracks in the extent of the bias. Figure 11 shows smoothed estimates of the observed and predicted relationships between log odds and expected returns, separately by track. For most tracks, the predicted relationship approximates the observed relationship. Figure 12 summarizes the model's fit at each track by

[25]We estimate this distribution using a kernel density estimator with Silverman's rule-of-thumb bandwidth, and we approximate the integral by evaluating the absolute deviation, along with the density of observed log odds, along a grid of 100 equally spaced points spanning the entire range of the distribution.

comparing observed and predicted normalized mean returns—i.e., the expected return on a randomly placed $1 win bet, normalized by the state-sanctioned return, $1-t$. Observed and predicted values are correlated across tracks, at $\rho = 0.65$, implying that the model predicts variation in the extent of the favorite-longshot bias across tracks. However, the model generally overpredicts the extent of the bias, for two reasons. First, our model ignores breakage, the practice of rounding down odds to the nearest 10 cents, which increases the track's effective takeout. Second, the model overpredicts returns for favorites at many tracks, as in the pooled comparison in Figure 10, for reasons that we discuss at the end of this section.

Figure 13: Observed and predicted returns for place and show bets, with 95% confidence intervals. The horizontal line marks the average takeout.

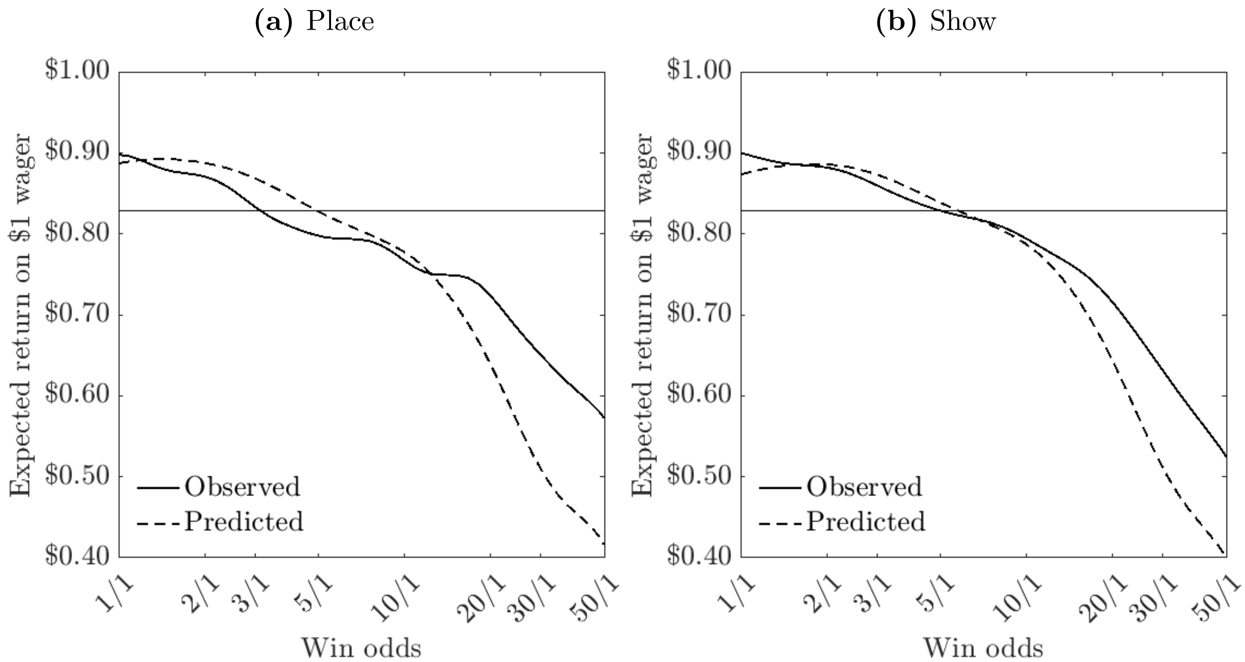

Note: Estimated using a local linear regression with a Gaussian kernel. The bandwidth, of log(1.23) for place bets and log(1.29) for show bets, minimizes the leave-one-out mean-squared error in the observed data.

We evaluate the predictions of our model in other pools as well. Figure 13 shows observed and predicted returns in place (13a) and show (13b) pools. (Since we observe final odds for all outcomes only in the win pool, in other pools, we estimate the relationship between log odds in the win pool and expected returns in the given pool.) A favorite-longshot bias characterizes returns for both place and show bets. In both pools, favorites with 1/1 win odds return 90 cents on the dollar, and longshots with 30/1 win odds return about 65 cents on the dollar. As in the win pool, the model fit is generally close, with a slight overprediction for favorites and an underprediction for longshots. In place pools, the expected absolute deviation is 5.2 (0.3) cents; in show pools, it is 4.1 (0.2) cents.

The final three sets of figures show expected returns on wagers involving two horses—those in exacta (14), quinella (15), and daily-double (16) pools. In each set of pool-specific figures, the first figure (a) shows observed returns, the second (b) shows predicted returns, and the third (c) shows the difference between observed and predicted returns. In all figures, expected returns are expressed in terms of the log win odds for the first and second horses chosen. (Since finishing order is irrelevant for quinella wagers, we show expected returns in terms of log win odds for the relative longshot and relative favorite.)[26]

For exotic bets, the favorite-longshot bias is severe. In expectation, wagering $1 on two randomly selected horses returns 67 cents in exacta pools, 65 cents in quinella pools, and 68 cents in daily-double pools—incurring far larger losses than the average takeouts of 21, 22, and 20 cents, respectively. Predicted returns from our model approximate observed returns in each pool, with expected absolute deviations of 4.9 (0.8) cents for exacta bets, 3.0 (1.6) cents for quinella bets, and 6.2 (0.7) cents for daily-double bets. For exacta bets, the deviation pattern is similar to that for bets on a single horse, with the model overpredicting returns for favorites and underpredicting returns for longshots; for quinella and daily-double bets, the deviation patterns are more idiosyncratic.

Our model largely hits the mark, but its misses are generally of the same pattern—overpredicting returns for favorites and underpredicting returns for longshots. We suspect that this stems from the coarseness of the beliefs assigned to off-track bettors. In our model, the beliefs of off-track bettors derive from a single attribute: the morning-line odds. Hence, every horse at a given track with the same morning lines is allocated the same probability of winning the race, before normalization. In practice, sophisticated bettors use other information, such as past performance, to form higher resolution beliefs. This disconnect is most consequential for favorites with morning-line odds at a censoring threshold. From the morning lines, it is unclear whether such a favorite is merely favored, or is Secretariat. Presumably, sophisticated bettors know the difference. If off-track bettors in our model were similarly discerning, they would wager even greater sums on favorites, and the model would predict lower returns for favorites and higher returns for longshots.

The poor resolution of the beliefs held by off-track bettors is evident in the predicted odds.

[26]Each figure is estimated using a local linear regression with a bivariate Gaussian kernel. In each pool, the bandwidth pair minimizes the leave-one-out mean-squared error in the observed data (a).

Figure 14: Exacta: expected return on \$1 wager.

(a) Observed **(b)** Predicted **(c)** Observed − predicted

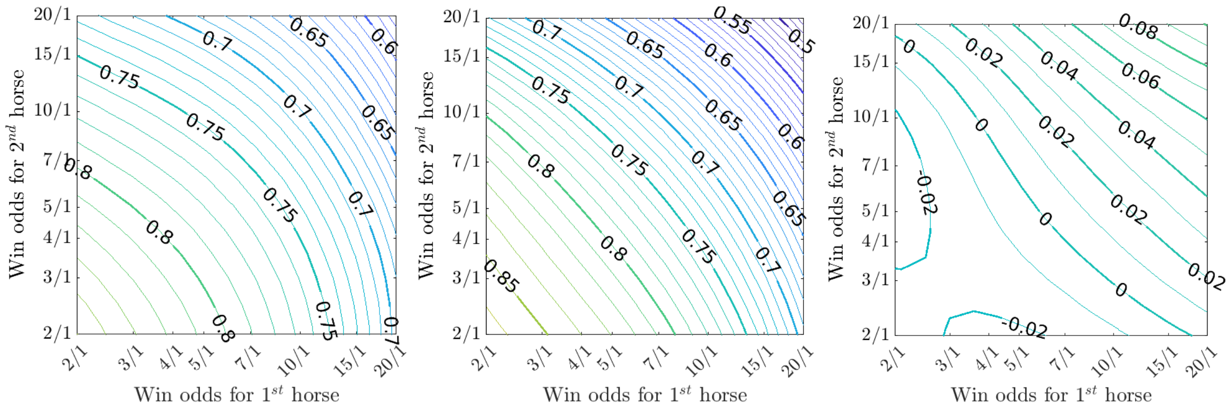

Figure 15: Quinella: expected return on \$1 wager.

(a) Observed **(b)** Predicted **(c)** Observed − predicted

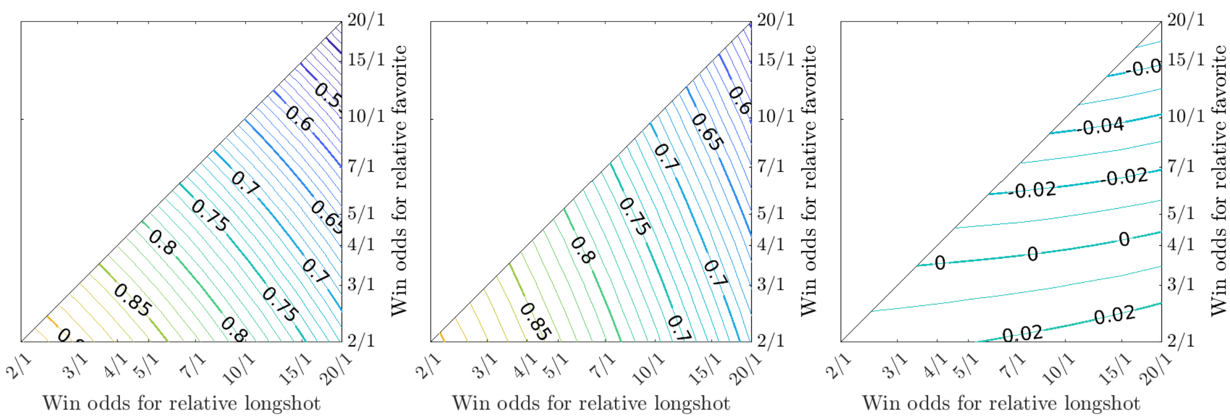

Figure 16: Daily double: expected return on \$1 wager.

(a) Observed **(b)** Predicted **(c)** Observed − predicted

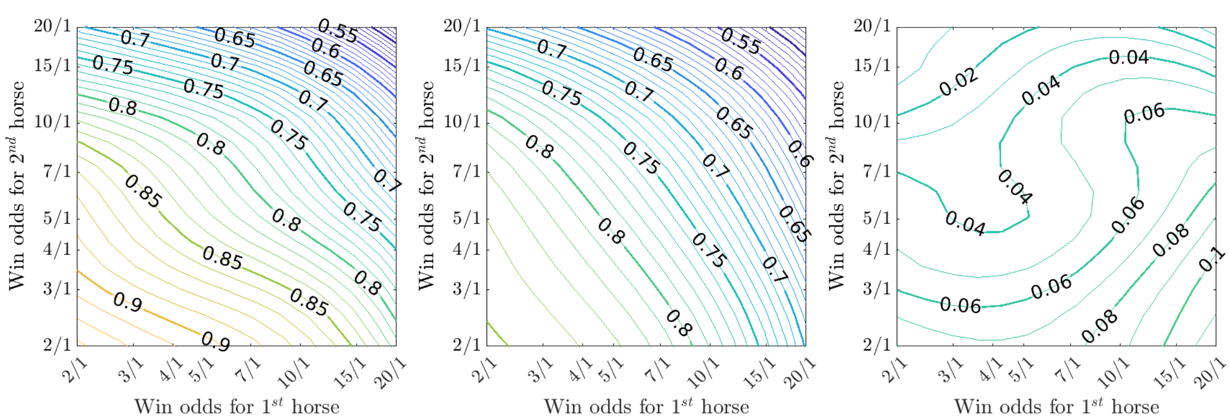

Figure 17: Distributions of morning-line odds (bars), final odds (solid line), and predicted odds (dashed line) for win bets.

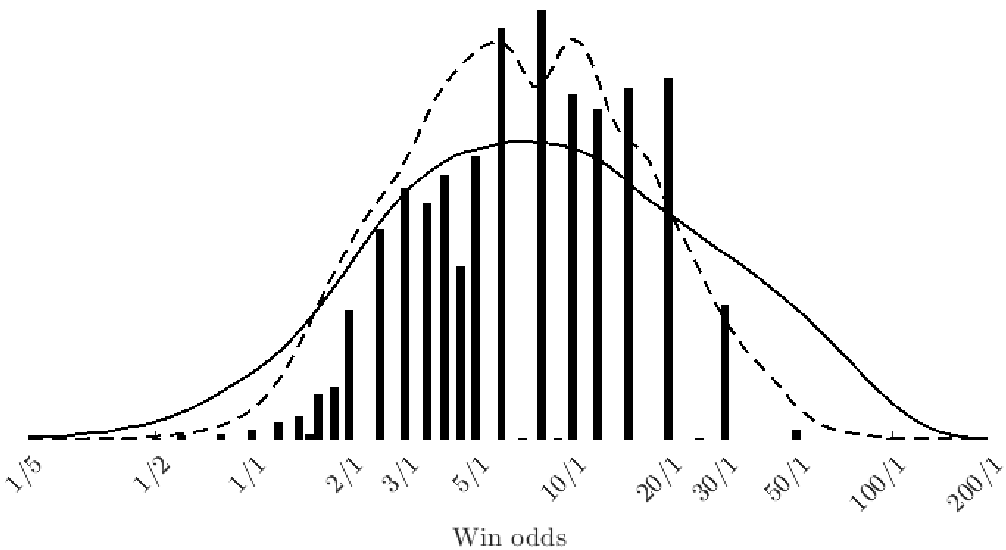

Note: Distributions of (log) parimutuel odds—both observed and predicted—estimated using a kernel density estimator with a Gaussian kernel and Silverman's rule of thumb bandwidth as calculated using the observed data.

Figure 17 compares the distributions of morning-line odds (bars), observed parimutuel odds (solid line), and predicted parimutuel odds (dashed line). The predicted odds exhibit more variance than the morning-line odds but less variance than the observed odds. One corrective would be to use a larger set of variables, along with a more complicated estimation routine, to assign off-track bettors more refined beliefs. Doing so would likely make the predicted odds more extreme, as in the observed distribution, and the predicted favorite-longshot bias less steep, as in the observed relationship. But we are loathe to complicate matters in search of a marginal increase in accuracy. The virtue of our model is not that it makes perfect predictions, but that it gets close without sacrificing parsimony.

7 Track revenues

Our model also predicts the track's gambling revenues, separately from track and off-track bettors. Predicted revenues from off-track bettors reflect arbitrage opportunities created by distortion in the morning-line odds. Importantly, they do not reflect revenues generated by

Figure 18: The distribution of optimal rebates across races (box plot) and the share of gambling revenues from off-track bettors (line), along with the takeout (star), by track and wager type. In the box plots, the box represents the interquartile range, and plus signs denote outliers.

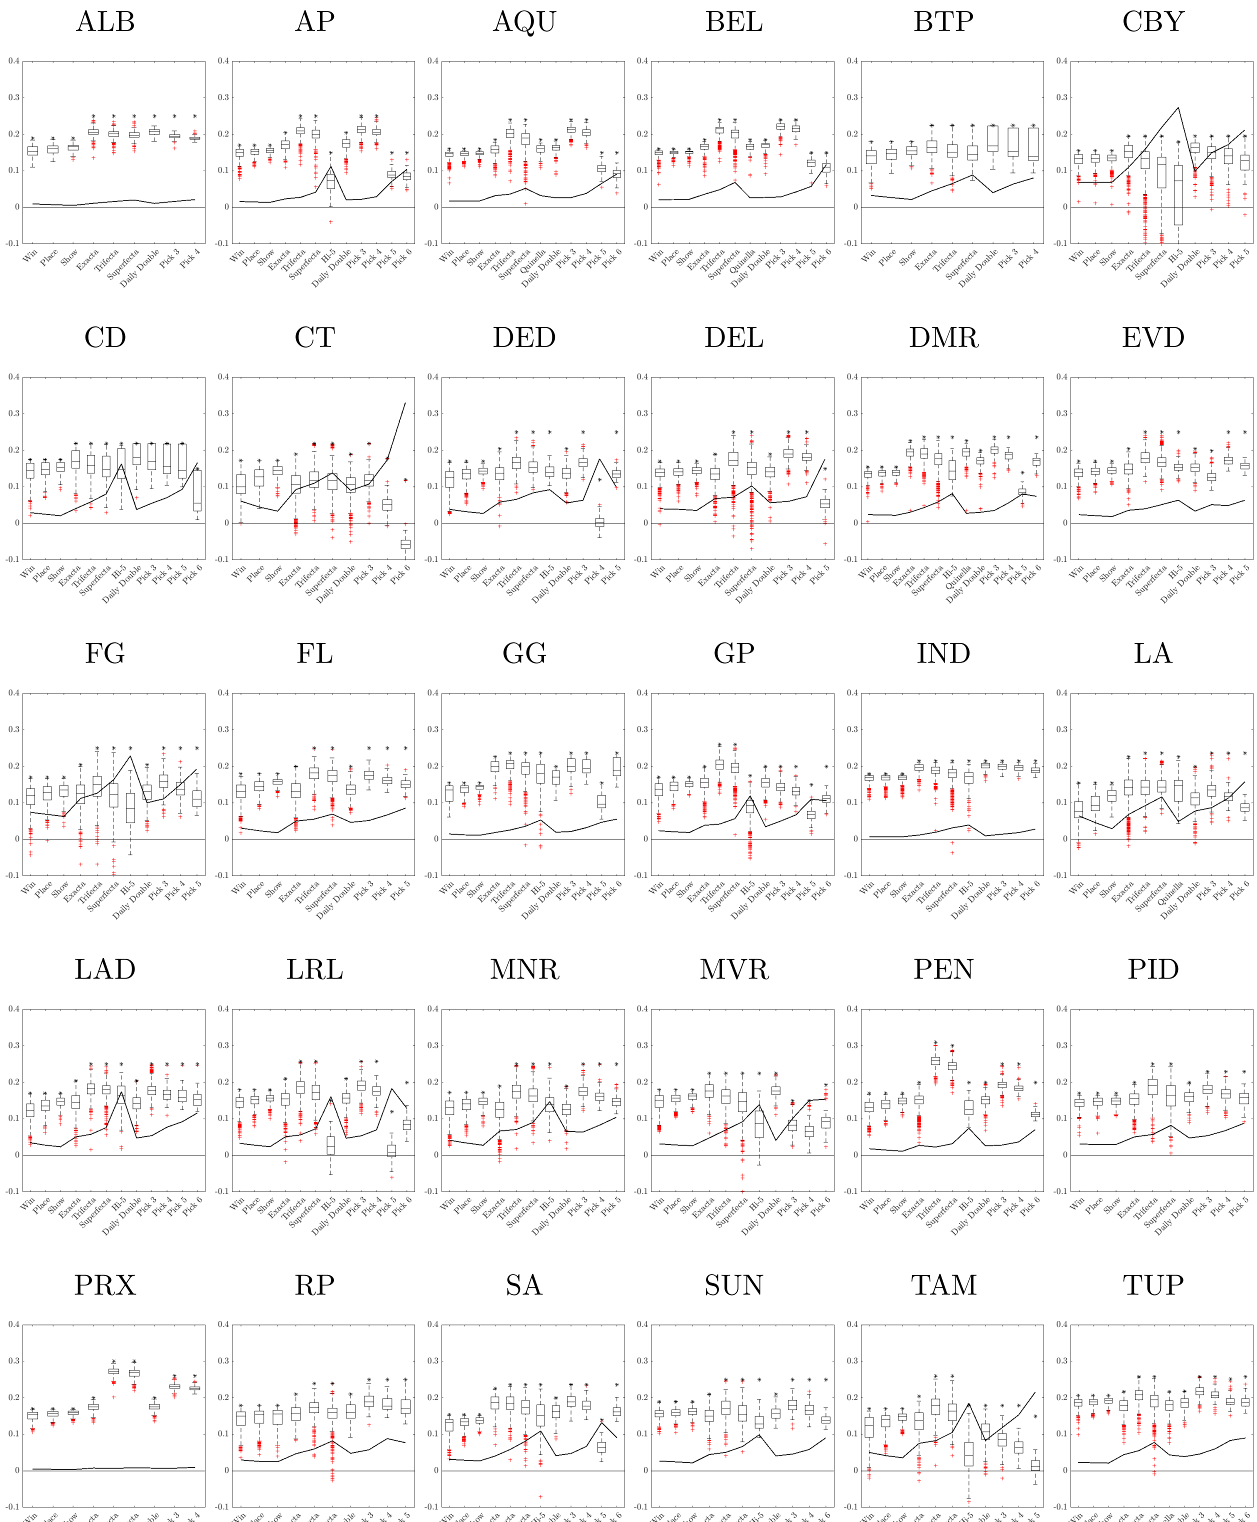

other means of deception, by other manners of arbitrage, or from off-track bettors that engage in behaviors other than arbitrage, such as speculation. In this sense, they are estimates of a quantity that cannot be directly observed.

How profitable is deception? Figure 18 shows the model's predictions for each wager type at each track, displaying the distribution of optimal rebates across races (box plot) and the share of gambling revenues from off-track bettors (line), along with the takeout (star). At tracks that promulgate well-calibrated morning-line odds (e.g., PRX), rebates concentrate near the takeout, and the track profits minimally, if at all, from taxing arbitrage. At other tracks, tax rates for off-track bettors are larger, and arbitrage generates more revenue for the track. Miscalibration in the morning-line odds creates arbitrage opportunities, reducing the rebate required to entice off-track bettors into the market. In some cases, arbitrage opportunities are so enticing that optimal rebates are negative—i.e., off-track bettors pay a surcharge.

Figure 18 also shows that exotic wagers are relatively more profitable for the track. This occurs because compound wagers, in which the bettor predicts the finishing order or the winners of consecutive races, amplify differences between track and off-track bettors in their beliefs for win wagers, as illustrated at the end of Section 5. Given that exotic pools comprise the majority of wagers at every track (Table 1), the vast majority of track revenues from off-track bettors come from exotic pools.

Figure 19: Estimated share of gambling revenues from off-track bettors, by track.

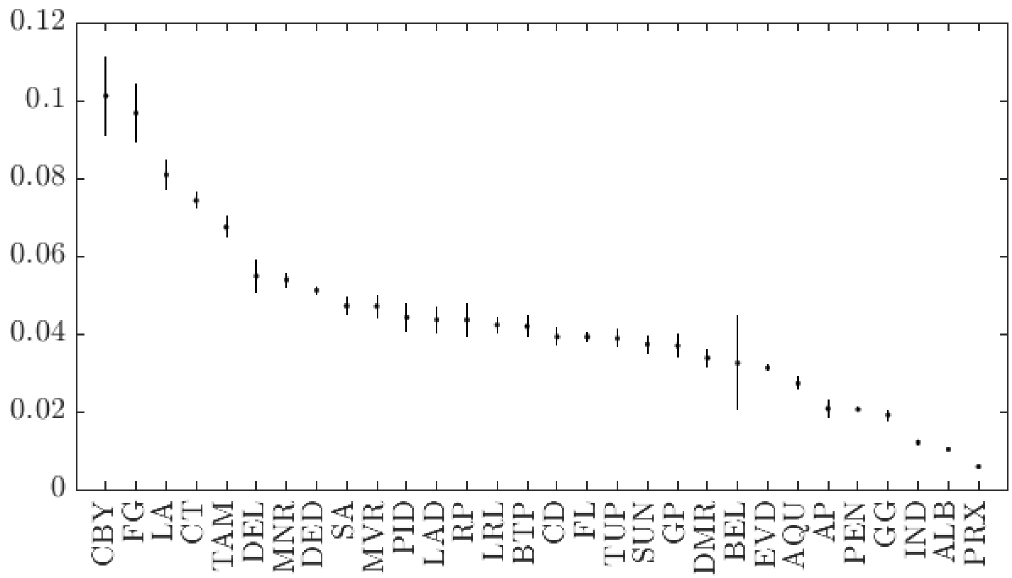

Figure 19 aggregates the predicted share of revenues from off-track bettors by track, weighting each pool by the total amount wagered.[27] We estimate that as much as 10% of gambling revenues are generated by deception in the morning-line odds. Cross-track variation is large, however, with a handful of tracks profiting minimally, if at all, from deception.

These estimates may underestimate or overestimate their true values. For instance, off-track bettors likely have more refined beliefs than those we assign them, as discussed at the end of Section 6. If so, the predictions in Figure 19 are probably underestimates, as greater resolution in p_i generally increases the deviation between P_m and Q_m. On the other hand, our assumption of perfect competition among arbitrageurs, if wrong, may lead to overestimation. If we instead assume that competition among off-track bettors is imperfect, or if we model them as risk averse, then off-track bettors will wager less, and the track will make less money. (It is also possible that off-track bettors collectively make negative profits, in which case our assumption of perfect competition leads to underestimation.) A second reason for overestimation is that we model the track as finding the optimal rebate in each betting pool, even when the optimal rebate is negative. In practice, tracks may not be so sophisticated, nor are they allowed to charge off-track bettors more than punters pay at the track.

8 Discussion

In recent years, a number of phenomena with long-standing behavioral explanations have been shown to be consistent with standard models (e.g., Farber, 2015; Miller and Sanjurjo, Forthcoming). This paper shows that the favorite-longshot bias—commonly cited as evidence of irrational behavior in markets (Barberis, 2013; Thaler, 2015)—is better explained by a model with rational, if uninformed, agents.

In our model, the track deceives uninformed bettors so as to profit by taxing arbitrageurs. The morning-line odds facilitate this deception. Billed as the track's prediction of the final parimutuel odds, the morning lines at most tracks instead embed a favorite-longshot bias. Uninformed bettors take the morning lines at face value and overbet longshots, thereby making favorites attractive to informed arbitrageurs. At the last minute, arbitrageurs bet on favorites, which moderates the favorite-longshot bias and transforms excess losses by

[27]For win, place, and show pools, the race chart lists the combined amount wagered rather than the pool-specific amounts. We assume that the listed total was wagered entirely in the win pool, reflecting the unpopularity of place and show bets. At the Santa Anita race track, for example, nearly two in three dollars wagered in the win, place, or show pools is wagered in the win pool (`https://lat.ms/2il84x7`). The model's predictions are generally similar across the three pools, as shown in Figure 18, and the estimates in Figure 19 are meaningfully unchanged when dividing the pot equally among the win, place, and show pools.

uninformed bettors into additional profit for the track.

Critics have noted that behavioral economics focuses on making economic models more psychologically realistic, rather than making them more predictive (e.g., Pesendorfer, 2006). Our approach prioritizes predictive accuracy over descriptive accuracy. We include features of the market that help make good predictions, namely the morning-line odds and the distinction between track and off-track bettors, and we exclude features that might make the model more realistic though not necessarily more predictive, such as non-standard preferences and psychological biases. The merits are twofold. First, modeling agents as gullible rather than behavioral makes our model parsimonious. Whereas alternatives to expected utility maximization typically involve parametric assumptions, gullible agents make straightforward inferences from observable information. Second, the model predicts the favorite-longshot bias observed in the data. We run a figurative horse race between speculative, biased, and gullible agents, and we find that gullibility wins out.

Our model is designed to show how deceptive morning-line odds can generate a favorite-longshot bias, not how the morning lines are themselves generated. As a result, the model cannot say why the morning lines embed a favorite-longshot bias, rather than the reverse or some other pattern.[28] An extension might allow for the direct utility that track bettors receive from gambling to vary with their beliefs. Promulgating optimistic predictions for longshots could thus be rationalized by a desire among track bettors to believe, for instance, that the underdog can prevail or that the race will be competitive. If so, one implication is that tracks may distort the morning lines for reasons other than—and perhaps without awareness of—the resulting profits from taxing arbitrage.

Finally, our results inform a policy debate about how best to regulate gambling, one that has seen renewed interest following the Supreme Court's 2018 ruling allowing states to legalize wagering on sports. Gambling is commonly rationalized by way of prospect theory. Wagers that lose money on average are appealing to agents who are risk loving in the domain of losses (Thaler and Ziemba, 1988; Thaler and Johnson, 1990) or who overweight small probabilities (Barberis, 2012). Both of these explanations imply that gambling is a mistake. The gambler will later regret the wagers she placed with the hope of getting back to even, or those she placed with a biased perception of the odds. Under this logic, welfare-enhancing interventions target gamblers, for instance by "nudging" them to gamble less. Our results

[28] A reverse favorite-longshot bias, in which returns for longshots exceed those for favorites, characterizes bookmaker odds for spread bets in Major League Baseball (Woodland and Woodland, 1994) and the National Football League (Levitt, 2004). NFL bettors overbet favorites even when they are told (truthfully) that the lines have been exaggerated, and the favorite is unlikely to cover the spread (Simmons et al., 2010).

show that gambling at the track is better rationalized by the enjoyment of gambling (Conlisk, 1993).[29] In our model, track bettors happily incur losses equal to the state-regulated tax. Presumably, they would be unhappy to learn that they are incurring additional losses as a result of being deceived by the track. Under this logic, welfare-enhancing regulations target deception by market makers.

References

Ali, Mukhtar M (1977) "Probability and Utility Estimates for Racetrack Bettors," *Journal of Political Economy*, Vol. 85, pp. 803–815.

——— (1979) "Some Evidence of the Efficiency of a Speculative Market," *Econometrica*, pp. 387–392.

Asch, Peter, Burton G Malkiel, and Richard E Quandt (1982) "Racetrack betting and informed behavior," *Journal of Financial Economics*, Vol. 10, pp. 187–194.

Barberis, Nicholas (2012) "A model of casino gambling," *Management Science*, Vol. 58, pp. 35–51.

——— (2013) "The psychology of tail events: Progress and challenges," *American Economic Review*, Vol. 103, pp. 611–16.

Barberis, Nicholas and Richard Thaler (2003) "A survey of behavioral finance," *Handbook of the Economics of Finance*, Vol. 1, pp. 1053–1128.

Benartzi, Shlomo and Richard H Thaler (1995) "Myopic loss aversion and the equity premium puzzle," *The Quarterly Journal of Economics*, Vol. 110, pp. 73–92.

Busche, Kelly and Christopher D Hall (1988) "An Exception to the Risk Preference Anomaly," *The Journal of Business*, Vol. 61, pp. 337–346.

Camerer, Colin F (1998) "Can asset markets be manipulated? A field experiment with racetrack betting," *Journal of Political Economy*, Vol. 106, pp. 457–482.

Conlisk, John (1993) "The utility of gambling," *Journal of Risk and Uncertainty*, Vol. 6, pp. 255–275.

Farber, Henry S (2015) "Why you can't find a taxi in the rain and other labor supply lessons from cab drivers," *The Quarterly Journal of Economics*, Vol. 130, pp. 1975–2026.

Gabaix, Xavier and David Laibson (2006) "Shrouded attributes, consumer myopia, and information suppression in competitive markets," *The Quarterly Journal of Economics*, Vol. 121, pp. 505–540.

[29] There are other reasons that people gamble, such as addiction, that call for different policy responses.

Gandhi, Amit and Ricardo Serrano-Padial (2014) “Does belief heterogeneity explain asset prices: The case of the longshot bias,” *The Review of Economic Studies*, Vol. 82, pp. 156–186.

Golec, Joseph and Maurry Tamarkin (1998) “Bettors love skewness, not risk, at the horse track,” *Journal of Political Economy*, Vol. 106, pp. 205–225.

Gramm, Marshall and C Nicholas McKinney (2009) “The effect of late money on betting market efficiency,” *Applied Economics Letters*, Vol. 16, pp. 369–372.

Hall, Peter and Li-Shan Huang (2001) “Nonparametric kernel regression subject to monotonicity constraints,” *Annals of Statistics*, pp. 624–647.

Harville, David A (1973) “Assigning probabilities to the outcomes of multi-entry competitions,” *Journal of the American Statistical Association*, Vol. 68, pp. 312–316.

Heidhues, Paul and Botond Kőszegi (2017) “Naivete-based discrimination,” *The Quarterly Journal of Economics*, Vol. 132, pp. 1019–1054.

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka (2016) “Inferior products and profitable deception,” *The Review of Economic Studies*, Vol. 84, pp. 323–356.

Hurley, William and Lawrence McDonough (1995) “A note on the Hayek hypothesis and the favorite-longshot bias in parimutuel betting,” *The American Economic Review*, pp. 949–955.

Jullien, Bruno and Bernard Salanié (2000) “Estimating preferences under risk: The case of racetrack bettors,” *Journal of Political Economy*, Vol. 108, pp. 503–530.

Levitt, Steven D (2004) “Why are gambling markets organised so differently from financial markets?” *The Economic Journal*, Vol. 114, pp. 223–246.

Lewis, Michael (2015) *The Big Short: Inside the Doomsday Machine*: WW Norton & Company.

Milgrom, Paul and John Roberts (1986) “Relying on the information of interested parties,” *The RAND Journal of Economics*, pp. 18–32.

Miller, Joshua B and Adam Sanjurjo (Forthcoming) “Surprised by the Gambler’s and Hot Hand Fallacies? A Truth in the Law of Small Numbers,” *Econometrica*.

Ottaviani, Marco and Peter N Sørensen (2008) “The favorite-longshot bias: An overview of the main explanations,” *Handbook of Sports and Lottery Markets*, pp. 83–101.

——— (2009) “Surprised by the parimutuel odds?” *American Economic Review*, Vol. 99, pp. 2129–34.

——— (2010) “Noise, information, and the favorite-longshot bias in parimutuel predictions,” *American Economic Journal: Microeconomics*, Vol. 2, pp. 58–85.

Pesendorfer, Wolfgang (2006) “Behavioral economics comes of age: A review essay on advances in behavioral economics,” *Journal of Economic Literature*, Vol. 44, pp. 712–721.

Prelec, Drazen (1998) “The probability weighting function,” *Econometrica*, pp. 497–527.

Sauer, Raymond D (1998) “The economics of wagering markets,” *Journal of economic Literature*, Vol. 36, pp. 2021–2064.

Shleifer, Andrei and Lawrence H Summers (1990) “The noise trader approach to finance,” *Journal of Economic Perspectives*, Vol. 4, pp. 19–33.

Simmons, Joseph P, Leif D Nelson, Jeff Galak, and Shane Frederick (2010) “Intuitive biases in choice versus estimation: Implications for the wisdom of crowds,” *Journal of Consumer Research*, Vol. 38, pp. 1–15.

Snowberg, Erik and Justin Wolfers (2010) “Explaining the favorite–long shot bias: Is it risk-love or misperceptions?” *Journal of Political Economy*, Vol. 118, pp. 723–746.

Spiegler, Ran (2011) *Bounded Rationality and Industrial Organization*: Oxford University Press.

Starr, Ross M (1969) “Quasi-equilibria in markets with non-convex preferences,” *Econometrica*, pp. 25–38.

Thaler, Richard H (2015) *Misbehaving: The Making of Behavioral Economics*: WW Norton & Company.

Thaler, Richard H and Eric J Johnson (1990) “Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice,” *Management science*, Vol. 36, pp. 643–660.

Thaler, Richard H and William T Ziemba (1988) “Anomalies: Parimutuel betting markets: Racetracks and lotteries,” *Journal of Economic perspectives*, Vol. 2, pp. 161–174.

Weitzman, Martin (1965) “Utility analysis and group behavior: An empirical study,” *The Journal of Political Economy*, Vol. 73, pp. 18–26.

Woodland, Linda M and Bill M Woodland (1994) “Market Efficiency and the Favorite-Longshot Bias: The Baseball Betting Market,” *The Journal of Finance*, Vol. 49, pp. 269–279.

A Proofs

The equilibrium conditions on the expected returns to off-track bettors (2)—namely that $\mathbb{E}_{\text{off-track}}[x_i] = x_i$ for $x_i > 0$, and $\mathbb{E}_{\text{off-track}}[x_i] \leq x_i$ for $x_i = 0$—imply an experession for $\sum x$, the total amount wagered by off-track bettors. Let $Q_+ \equiv \sum_i q_i \cdot \mathbb{1}\{x_i > 0\}$ and $P_+ = \sum_i p_i \cdot \mathbb{1}\{x_i > 0\}$, the subjective probability—held by track and off-track bettors, respectively—of an outcome that attracts positive bets by off-track bettors. Then,

$$\sum x = \frac{(1-t) \cdot P_+ - (1-r) \cdot Q_+}{(1-r) - (1-t) \cdot P_+}$$

Substituting this expression into the track's profit function (3) and solving the first-order condition yields the profit-maximizing rebate, r^*:

$$r^* = 1 - (1-t) \cdot \left[P_+ + \sqrt{\frac{P_+ \cdot (1-P_+) \cdot (1-Q_+)}{Q_+}} \right]$$

Substituting $\sum x$ and r^* into the first equilibrium condition—i.e., $\mathbb{E}_{\text{off-track}}[x_i] = x_i$ for $x_i > 0$—and solving for x_i yields:

$$x_i = p_i \cdot \sqrt{\frac{Q_+ \cdot (1-Q_+)}{P_+ \cdot (1-P_+)}} - q_i, \quad \forall i \text{ s.t. } x_i > 0$$

Hence, $x_i > 0$ if and only if:

$$\left(\frac{p_i}{q_i}\right)^2 > \frac{P_+ \cdot \left(1 - P_+\right)}{Q_+ \cdot \left(1 - Q_+\right)}$$

Without loss of generality, reorder the outcome indices i such that:

$$\frac{p_1}{q_1} \geq \frac{p_2}{q_2} \geq \cdots \geq \frac{p_i}{q_i} \geq \frac{p_{i+1}}{q_{i+1}} \geq \cdots \geq \frac{p_N}{q_N}$$

As a result, $Q_+ \equiv Q_m = \sum_{i=1}^m q_i$, where $x_i > 0$ for $i \leq m$ and $x_i = 0$ for $m < i \leq N$. Similarly, $P_+ \equiv P_m = \sum_{i=1}^m p_i$.

An equilibrium exists if there exists an index m such that $x_i > 0$ for $i \in \{1, \ldots, m\}$, and $x_i = 0$ for $i \in \{m+1, \ldots, N\}$. Given that p_i/q_i is decreasing in i by construction, this equilibrium condition can be written as:

$$\exists m \in \{1, \ldots, N-1\} \text{ s.t. } \left(\frac{p_m}{q_m}\right)^2 > \frac{P_m \cdot (1-P_m)}{Q_m \cdot (1-Q_m)} \geq \left(\frac{p_{m+1}}{q_{m+1}}\right)^2 \tag{11}$$

We show the existence of an equilibrium when N is large and (11) has a continuous representation. Let $p(m)$ and $q(m)$ denote continuous probability density functions with support on $[1, N]$ such that $\frac{d}{dm}\frac{p(m)}{q(m)} \leq 0$ and $\frac{p(1)}{q(1)} > 1$ (and hence, $\frac{p(N)}{q(N)} < 1$). In addition, let $P(m)$ and $Q(m)$ denote the cumulative density functions, $\int_1^m p(u)du$ and $\int_1^m q(u)du$, respectively, where $P(1) = Q(1) = 0$ and $P(N) = Q(N) = 1$. The continuous representation of the (11) is:

$$\exists m \in (1, N) \text{ s.t. } \left(\frac{p(m)}{q(m)}\right)^2 = \frac{P(m) \cdot \big(1-P(m)\big)}{Q(m) \cdot \big(1-Q(m)\big)} \tag{12}$$

Note that because $p(m)$ and $q(m)$ are continuous, both sides of the equality in (12) are also continuous functions. For brevity, let $\lambda(m)^2$ equal the right-hand expression in (12).

To show the existence of an equilibrium, it is sufficient to observe that $\lim_{m \to 1} \lambda(m)^2 < \big(p(1)/q(1)\big)^2$ and $\lim_{m \to N} \lambda(m)^2 > \big(p(N)/q(N)\big)^2$. First, consider the limit as $m \to 1$:

$$\lim_{m \to 1} \lambda(m)^2 = \lim_{m \to 1} \frac{p(m) \cdot \big(1-2P(m)\big)}{q(m) \cdot \big(1-2Q(m)\big)} = \frac{p(1)}{q(1)} < \left(\frac{p(1)}{q(1)}\right)^2,$$

where the first equality follows from L'Hospital's rule, and the inequality follows from the ordering of $p(m)/q(m)$. Now consider the limit as $m \to N$:

$$\lim_{m \to N} \lambda(m)^2 = \lim_{m \to N} \frac{p(m) \cdot \big(1-2P(m)\big)}{q(m) \cdot \big(1-2Q(m)\big)} = \frac{p(N)}{q(N)} > \left(\frac{p(N)}{q(N)}\right)^2,$$

where the inequality again follows from the ordering of $p(m)/q(m)$. These opposing inequalities, along with the continuousness of $\big(p(m)/q(m)\big)^2$ and $\lambda(m)^2$, imply the existence of at least one fixed point $m \in (1, N)$. Thus in the discrete case, an equilibrium exists if p_m and q_m are approximately continuous in the sense of Starr (1969). Of practical importance, we note that (11) is satisfied for every observed pair of p and q vectors.