Parallax and Tax[*]

Etan A. Green[1], Haksoo Lee[2], and David Rothschild[3]

[1]University of Pennsylvania
[2]Stanford University
[3]Microsoft Research

May 13, 2019

Abstract

Common valuations pose an obstacle to trade and, hence, an existential threat to brokers, who profit from taxing trade. We write down a model in which a broker drives a wedge between valuations by deceiving gullible traders. Deception facilitates arbitrage, and arbitrage generates brokerage fees. We then show that this process, which we call *parallax and tax*, explains a classic case of market inefficiency: the favorite-longshot bias in horserace parimutuel markets.

Keywords: deception, arbitrage, parallax and tax, favorite-longshot bias
JEL Classification: D22, D82, D83, L83

Profit is sweet, even if it comes from deception.

<div align="right">Sophocles</div>

1   Introduction

Inefficiencies characterize important financial markets. In the stock market, for instance, excess returns are persistent, rather than fleeting (Mehra and Prescott, 1985, 2003); prices swing wildly, even in the absence of news (LeRoy and Porter, 1981; Shiller, 1981, 1992; Shleifer, 1986); and price movements are predictable over short and long horizons (Fama and French, 1988; Jegadeesh and Titman, 1993, 2011).[1]

Market inefficiencies stem in part from the presence of irrational traders (Shiller, 1984; Shleifer and Summers, 1990; De Long, Shleifer, Summers and Waldmann, 1991). These "noise traders" rely on faulty heuristics and take advice credulously, and they misprice securities by trading on their systematically biased beliefs. Although arbitrage by more sophisticated investors attenuates the mispricing, various limits, such as transaction costs, illiquidity, and risk, preserve some degree of inefficiency (for a review, see Gromb and Vayanos, 2010).

A defining trait of noise traders is their gullibility (Akerlof and Shiller, 2015). In pump-and-dump schemes, for example, conspirators deceive gullible traders into overvaluing an asset (i.e., pump) so as to sell their own holdings at higher prices (i.e., dump). This paper characterizes a "parallax-and-tax" scheme in which a third-party broker, rather than an investor, deceives noise traders and then profits by taxing trade, rather than trading assets. In a visual parallax, the same object appears in different locations when viewed from different positions; here, a broker induces different valuations of the same asset from different investors. Common valuations pose an obstacle to trade (Tirole, 1982) and, hence, an existential threat to brokers. Deception drives a wedge between valuations, increasing trade and brokerage fees.

This paper models a parallax-and-tax scheme and uses the model to rationalize a classic case of market inefficiency. The model applies to a brokered market with informed arbitrageurs and credulous noise traders. We begin with an illustrative version in which the broker offers two securities: one that pays out if a given event occurs, and one that pays out if the event does not occur; later we extend the model to an arbitrary number of securities. The true probability of the event, p, is known to the broker and the arbitrageurs but not to the noise traders. The broker announces a prediction, q, which noise traders adopt as their beliefs. Noise traders trade in the first period. If $q \neq p$, they trade on mistaken beliefs and

---

[1]For reviews, see Barberis and Thaler (2003) and Shiller (2003).

misprice the securities. In the second period, arbitrageurs partially correct the mispricing. Some bias remains, however, because arbitrageurs pay brokerage fees.

In choosing q, the broker balances brokerage fees against the long-run costs of deception, both of which increase in the divergence between p and q. More deception increases arbitrage but also becomes more apparent *ex post*, decreasing future investment by noise traders. We show that if deception has sufficiently convex costs, the broker will deceive by overestimating unlikely outcomes (longshots) and underestimating likely outcomes (favorites).

In light of this result, we revisit the favorite-longshot bias in horserace parimutuel markets (Griffith, 1949; Thaler and Ziemba, 1988). In parimutuel markets, bettors place wagers on outcomes (e.g., a given horse winning the race), the broker 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. Consistent with previous literature, we find that wagers on longshots return far less in expectation than wagers on favorites. 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 (i.e., that pays a \$30 dividend on a winning \$1 wager) returns just 63 cents on the dollar.

This mispricing is commonly cited as real-world evidence for one of two tenets of prospect theory—i.e., of risk-loving preferences in the domain of losses (Thaler, 2015) or a tendency to overweight small probabilities (Barberis, 2013). We show that neither need be true. In our rendering, some bettors believe the track's predictions, and this is sufficient to generate a favorite-longshot bias.

Tracks provide bettors with a prediction for each horse, in the form of odds. These *morning-line odds* have no formal bearing on the parimutuel odds. Instead, their ostensible purpose, according to oddsmakers, "is to predict, as accurately as possible, how the betting public will wager on each race"—i.e., to predict the final parimutuel odds.[2]

Past studies of the favorite-longshot bias in horserace parimutuel markets ignore the morning-line odds, perhaps because they are difficult to obtain or because their formal irrelevance suggests economic irrelevance. For instance, Ottaviani and Sørensen (2009, 2129) justify studying racetrack parimutuel markets because of "the absence of bookmakers (who could induce biases)." We scrape the morning lines each morning for over a year, and we find that these predictions, made by the track oddsmaker, exhibit a favorite-longshot bias.[3]

---

[2]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`. Also see: Snyder (1978, 1117).

[3]To our knowledge, the only other researcher to collect the morning lines is Snyder (1978), who collects

The morning-line odds are miscalibrated: they are insufficiently short for favorites and insufficiently long for longshots. Horses with 1/1 morning-line odds begin the race with parimutuel odds of 1/2, on average, while horses with 30/1 morning-line odds end up with parimutuel odds longer than 50/1. As we show, the parimutuel odds separate from the morning lines at the end of the betting window, as late wagers concentrate on favorites. And as we discuss, journalistic accounts attribute this late wagering to professional gamblers.

The morning-line odds also imply biased win probabilities. If the track is to be believed, and if bettors are risk-neutral, then a longshot with 30/1 morning lines should win about 3% of the time. In practice, horses with morning-line odds of 30/1 win about 1% of the time. The morning-line odds provide useful ordinal information by correctly distinguishing between longshots and favorites. However, they provide biased cardinal information by overestimating the chances of longshots and underestimating the chances of favorites.

We find that across tracks, the extent of the miscalibration in the morning-line odds predicts the extent of the favorite-longshot bias in the parimutuel odds. Some tracks provide well-calibrated morning-line odds, and parimutuel odds at those tracks do not exhibit a favorite-longshot bias. Other tracks embed a severe favorite-longshot bias in the morning-line odds, and parimutuel odds at those tracks reflect a severe favorite-longshot bias.

These stylized facts and contextual details motivate an application of our two-period brokerage model. In the first period, noise traders infer beliefs, q, under the assumption that the morning-line odds reflect wagering by risk-neutral bettors. Noise traders then wager in proportion to these beliefs, generating odds that reproduce any bias inherent in the morning lines. In the second period, arbitrageurs 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 arbitrageurs to view favorites as underpriced. A large number of arbitrageurs place all wagers with positive expected value until no more exist. Late wagering concentrates on favorites, and the favorite-longshot bias moderates. The bias does not disappear, however, because arbitrageurs pay tax.

In our model, biased morning lines generate incremental income for the track. Whereas state-regulated tax rates cap the losses that noise traders should sustain in expectation, misinformed traders overbet losers and sustain excess losses. Because arbitrage is competitive, these excess losses flow in their entirety into the track's coffers, laundered through taxes on arbitrageurs. For tracks that embed a favorite-longshot bias in the morning-line odds, we estimate that misleading noise traders increases brokerage income by as much as 12%.[4]

---

them by hand and puzzles over the fact that they reflect a favorite-longshot bias.

[4]Other taxes on arbitrageurs, such as fees for data (e.g., Budish, Lee and Shim, 2019), are outside of our

3

Previous empirical studies of the favorite-longshot bias in horserace betting markets attempt to adjudicate between preference and belief-based explanations. In this literature, a model is proposed which allows for a downward-sloping relationship between the parimutuel odds and expected returns. At least one parameter is tuned to match the data. Parameter estimates are then interpreted as evidence of risk-loving preferences (Weitzman, 1965; Ali, 1977; Golec and Tamarkin, 1998), a bias towards overweighting small probabilities (Snowberg and Wolfers, 2010), or heterogeneity in beliefs (Gandhi and Serrano-Padial, 2014) or preferences (Chiappori, Salanié, Salanié and Gandhi, 2019).[5] Chiappori et al. (2019, 4) conclude that probability weighting is needed to match the data.

By comparison, our model is restricted in three ways. First, it predicts returns from the formally irrelevant morning-line odds, not the final parimutuel odds. Second, the only parameters in our estimation routine are the hyperparameters we use to (non-parametrically) estimate the beliefs of sophisticated bettors. Third, all agents are risk-neutral, and none transforms probabilities.

Despite these handicaps, our model closely predicts the observed relationship between odds and returns—and it predicts the observed variation in that relationship across tracks. In other models, this variation can only be rationalized by track-specific belief or preference parameters. In our model, variation in the favorite-longshot bias arises from variation in the extent to which tracks overestimate the chances of longshots.

Our model also explains two puzzling stylized facts. First, horserace parimutuel markets in Hong Kong and Japan do not exhibit a favorite-longshot bias (Busche and Hall, 1988; Busche, 1994). We observe that race cards in Hong Kong and Japan do not display morning lines.[6] Second, prediction markets do not reliably exhibit a favorite-longshot bias.[7] We observe that brokers do not make predictions in prediction markets.

A class of models show that heterogeneous beliefs can generate a favorite-longshot bias in parimutuel markets (Borel, 1938; Ali, 1977; Hurley and McDonough, 1995; Ottaviani and Sørensen, 2009, 2010).[8] Our work is most closely related to that of Gandhi and Serrano-

---

model.

[5] For a comprehensive review of the explanations offered for the favorite-longshot bias, see Ottaviani and Sørensen (2008).

[6] For Hong Kong race cards, see: `https://racing.hkjc.com/racing/Info/meeting/RaceCard/english/Local/`. For Japan race cards, see: `http://japanracing.jp/en/racing/result/`.

[7] In prediction markets, the favorite-longshot bias is sometimes present (Wolfers and Zitzewitz, 2004), sometimes absent (Servan-Schreiber, Wolfers, Pennock and Galebach, 2004; Cowgill and Zitzewitz, 2015), and sometimes reversed (Cowgill, Wolfers and Zitzewitz, 2009). Illiquidity and transaction costs are thought to explain inefficiencies in prediction markets (Wolfers and Zitzewitz, 2006).

[8] This result can also be generated in markets for Arrow-Debreu securities, which unlike parimutuel

4

Padial ([2014](#)), who empirically find that two types of bettors—those who know the true probabilities, and noise traders, whose beliefs follow an arbitrary distribution—are sufficient to generate the observed bias in horserace betting markets. By contrast, our approach endogenizes the beliefs of noise traders, rather than treating them as a primitive (cf. Kyle, 1985; Black, 1986). In our model, the broker induces heterogeneous beliefs by misleading noise traders.

Our setting is distinct from betting markets in which a bookmaker, rather than the market, sets the odds.[9] Bookmakers generate excess returns by taking risky positions against bettors with inaccurate beliefs, while using the tax to limit betting by "sharps" (Levitt, 2004). In our model, the track encourages the sharps to participate and generates a riskless profit by taxing them.

A similar distinction applies to models of deception by firms (e.g., Gabaix and Laibson, 2006; Heidhues, Kőszegi and Murooka, 2016). In these models, profits on naive customers offset losses on sophisticated customers.[10] Here, by contrast, the broker profits from both naifs and sophisticates.

Models of deception offer a new direction for structural work in behavioral economics, which has focused on models of behavioral biases, such as hyperbolic discounting and reference dependence (for a review, see DellaVigna, 2018). We provide a template for interactions between gullible and knowing agents that can be taken directly to data. Naive agents are assumed to believe what they are told, and sophisticated agents maximally exploit their gullibility. This approach generates parsimonious and estimable models—and, at least in our data, good predictions.

The remainder of the paper is organized as follows. Section 2 presents our model of a brokered market with two parimutuel securities. Section 3 describes the empirical context and the data. Section 4 illustrates the favorite-longshot bias and other stylized facts. Section 5 generalizes our model for an arbitrary number of securities, Section 6 compares theoretical predictions with the data, and Section 7 estimates each track's incremental income from misinformation. Section 8 concludes.

---

securities have a variable price and a fixed dividend (Manski, 2006; Ottaviani and Sørensen, 2015).

[9]In fixed-odds betting markets, returns for favorites outpace returns for longshots in horserace betting (Jullien and Salanié, 2000; Snowberg and Wolfers, 2010), but the reverse is generally true for professional American sports (e.g., Gandar et al., 1988; Woodland and Woodland, 1994; Levitt, 2004; Simmons et al., 2010). For models of pricing in fixed-odds markets, see: Shin (1991, 1992); Ottaviani and Sørensen (2005).

[10]For instance, Gabaix and Laibson (2006) model a firm that shrouds add-on costs, such as high-priced printer toner. It loses money from sophisticated customers, who only purchase the subsidized printers, and makes money from naive customers, who purchase the toner as well.

2 Theoretical framework

A broker offers two parimutuel securities whose returns are pegged to a binary event. One security pays off if the event occurs; the other pays off if the event does not occur. Let the probability of the event be $p \geq \frac{1}{2}$. Since the event is (weakly) more likely to occur than not, the security tied to the event is the favorite, and the other security is the longshot. The broker announces a prediction for the favorite, q, which some investors take at face value. We are interested in the broker's choice of q.[11]

2.1 Set-up

In a parimutuel market, investors purchase shares of a security for a set price (e.g., \$1), the broker takes a fee, σ, and the remaining investment is split among shareholders of the security that pays out. Let s be the share of investment in the favorite. If the event occurs, each dollar invested in the favorite returns $(1-\sigma)/s$, and investments in the longshot return 0. If the event does not occur, each dollar invested in the longshot returns $(1-\sigma)/(1-s)$, and investments in the favorite return 0. The more invested in a security, the lower its returns.

We consider a two-period model with two types of investors: noise traders and arbitrageurs. Noise traders invest in the first period; arbitrageurs invest in the second. Noise traders are gullible. Absent beliefs about p, they take the broker's prediction, q, at face value—i.e., they believe that the event will occur with probability q. Noise traders are also myopic: they fail to anticipate how the parimutuel returns will change in the second period.[12]

In the first period, noise traders invest in the security with the higher expected value according to their adopted beliefs. Given the brokerage fee, σ, both securities may portend losses. We rationalize the decision to invest with an additive utility shock, u, consistent with a direct utility from gambling (Conlisk, 1993), and we allow the shock to vary across noise traders. Thus, noise trader j values a \$1 investment in the favorite at $q(1-\sigma)/s + u_j$ and a

---

[11]The broker's prediction is cheap talk. We do not write down a cheap talk signaling game (e.g., Crawford and Sobel, 1982) because cheap talk games cannot rationalize biased signals, which we observe in our data. In a cheap talk signaling game, the receiver knows the bias and subtracts it out. Instead, we assume that some investors are gullible. This is similar in spirit to a cursed equilibrium (Eyster and Rabin, 2005), in which agents ignore the signal, though diamatrically opposed in practice—here, agents ignore their prior.

[12]It is straightforward to write down an alternative model in which noise traders correctly forecast second-period investment. In such a model, noise traders invest entirely in the longshot if $q < p$ and entirely in the favorite if $q > p$; arbitrageurs invest in the other security; arbitrage volume is higher; and broker income is higher as a result. We view this model as implausible because it requires noise traders to know how q compares to the truth—and to invest according to q anyway. In other words, one cannot be gullible without also being myopic.

$1 investment in the longshot at $(1-q)(1-\sigma)/(1-s)+u_j$. A large number of noise traders each decide whether to invest a \$1 endowment, and if so, which security to purchase. We are interested in s^*, the Walrasian equilibrium share of first-period investment in the favorite.

In the second period, arbitrageurs invest in the security with excess returns. Let x be the amount invested by arbitrageurs in the favorite, and let y be the amount invested by arbitrageurs in the longshot. Normalizing the total investment by noise traders to 1, total investment after arbitrage is $1+x+y$, and final parimutuel returns on a \$1 investment are $(1-\sigma)(1+x+y)/(s^*+x)$ for the favorite and $(1-\sigma)(1+x+y)/(1-s^*+y)$ for the longshot.

Unlike noise traders, arbitrageurs know p, and they do not receive direct utility from investing. The expected return on a \$1 investment in the favorite is:

$$\mathbb{E}\big[V(\$1 \text{ in favorite})\big] = p(1-\sigma)\frac{1+x+y}{s^*+x}$$

And the expected return on a \$1 investment in the longshot is:

$$\mathbb{E}\big[V(\$1 \text{ in longshot})\big] = (1-p)(1-\sigma)\frac{1+x+y}{1-s^*+y}$$

The broker price discriminates by offering arbitrageurs a rebate r on each dollar invested, regardless of the outcome. Hence, the broker's income is:

$$\pi(r) = \sigma \cdot 1 + (\sigma - r)(x+y)$$

For each dollar invested by noise traders, the broker takes a share σ and arbitrageurs invest $x+y$ dollars, of which the broker takes a share, $\sigma - r$. We assume that σ is fixed, as when brokers perfectly compete for noise traders or when its fees are regulated. However, the broker holds a monopoly over its own securities, and as such, chooses r to maximize π.

We assume that arbitrage is competitive. Hence, the second-period equilibrium is defined by arbitrage volumes $\{x^*, y^*\} \geq 0$ and a rebate r^* such that:

1. Arbitrageurs make zero profits in expectation.

   - $\mathbb{E}\big[V(\$1 \text{ in favorite})\big] = \$1 - r^*$ when $x > 0$
   - $\mathbb{E}\big[V(\$1 \text{ in longshot})\big] = \$1 - r^*$ when $y > 0$

2. Arbitrageurs leave no money on the table.

   - $\mathbb{E}\big[V(\$1 \text{ in favorite})\big] \leq \$1 - r^*$ when $x = 0$

7

- $\mathbb{E}\big[V(\$1 \text{ in longshot})\big] \leq \$1 - r^*$ when $y = 0$

3. The broker chooses the income-maximizing rebate: $r^* = \arg\max_r \pi(r)$

2.2 Equilibria

In the first-period Walrasian equilibrium, $s^* = q$, and only those for whom $u_j > \sigma$ invest. Consider an example with $\sigma = 0.2$ and $p = 0.6$. If the broker predicts $q = 0.5$, noise traders split their investment evenly among the securities. For a \$1 investment on either the favorite or the longshot, the parimutuel advertises a return of $(1 - \sigma)/q = (1 - \sigma)/(1 - q) = \$1.60$ were the security to pay out, and noise traders expect a return of $1 - \sigma = 80$ cents.

Noise traders are gullible, and deception by the broker induces them to misprice the securities. In the example above, the true expected returns are $p(1 - \sigma)/q = 96$ cents for the favorite and $(1 - p)(1 - \sigma)/(1 - q) = 64$ cents for the longshot. Let γ_1 be the ratio of first-period expected returns for the favorite and longshot under correct beliefs:

$$\gamma_1 = \frac{p(1 - q)}{q(1 - p)}$$

If $q < p$, the expected return on the favorite exceeds the expected return on the longshot—i.e., a favorite-longshot bias. And if $q > p$, the expected return on the longshot exceeds the expected return on the favorite—i.e., a reverse favorite-longshot bias.

We characterize equilibrium rebates, arbitrage volumes, parimutuel returns, and broker revenues. Let r^* be the income-maximizing rebate. The brokerage fee on arbitrage is:

$$\sigma - r^* = (1 - \sigma) \begin{cases} (1 - p)\big(\sqrt{\gamma_1} - 1\big), & q \leq p \\ p\big(\frac{1}{\sqrt{\gamma_1}} - 1\big), & q > p \end{cases}$$

When the broker tells the truth, $\gamma_1 = 1$, and arbitrage is free. Otherwise, fees on arbitrage are strictly positive.

Arbitrageurs invest in the security that promises excess returns. When the broker underestimates the favorite, arbitrageurs invest in the favorite:

$$\{x^*, y^*\} = \big\{q(\sqrt{\gamma_1} - 1), 0\big\}$$

8

When the broker overestimates the favorite, arbitrageurs invest in the longshot:

$$\{x^*, y^*\} = \left\{0, (1-q)\left(\frac{1}{\sqrt{\gamma_1}} - 1\right)\right\}$$

In both cases, the volume of arbitrage increases in the first-period bias. When the broker tells the truth, arbitrage volume is zero.

Arbitrage attenuates the bias. Arbitrageurs drive down expected returns on the overvalued security to $1 - r^*$ for every dollar invested, which increases returns on the undervalued security. Let γ_2 be the ratio of second-period expected returns for the favorite and longshot under correct beliefs. Regardless of q,

$$\gamma_2 = \sqrt{\gamma_1}$$

Arbitrage narrows, but does not eliminate, the gap in expected returns. When $q < p$, $1 < \sqrt{\gamma_1} < \gamma_1$; when $q > p$, $\gamma_1 < \sqrt{\gamma_1} < 1$. In other words, arbitrage brings the ratio of expected returns closer to one. Arbitrage does not equalize expected returns, however, because arbitrageurs pay some fees.

Consider the previous example in which $\sigma = 0.2$, $p = 0.6$, and $q = 0.5$. The broker sets a rebate of 13 cents, and arbitrageurs invest 11 cents in the favorite for every \$1 invested by noise traders. Parimutuel returns for the favorite decrease from \$1.60 to $(1-\sigma)(1+x^*)/(q+x^*) = \$1.46$, and parimutuel returns for the longshot increase from \$1.60 to $(1-\sigma)(1+x^*)/(q+x^*) = \$1.78$. Expected returns on a \$1 investment decrease from 96 cents to $1-r^* = 87$ cents for the favorite and increase from 64 cents to $(1-p)(1-\sigma)(1+x^*)/(1-q) = 71$ cents for the longshot. The gap in expected returns shrinks by half.

Given the optimal rebate, the broker's equilibrium income is:

$$\pi^*(q|p) = \sigma + (1-\sigma)\left[1 - \left(\sqrt{pq} + \sqrt{(1-p)(1-q)}\right)^2\right]$$

The first term is the broker's income from noise traders, and the second term is income from arbitrageurs. The broker's income from arbitrage is proportional to one minus the squared Bhattacharyya coefficient, which measures the similarity between two distributions. Hence, income from arbitrage, and thus total income, is increasing in the divergence between p and q. When the broker tells the truth, the Bhattacharyya coefficient is 1, and $\pi^* = \sigma$. When the broker lies, the coefficient is less than 1, and $\pi^* > \sigma$.

By assumption, arbitrageurs make exactly zero profits. Hence, the broker's income equals

9

the total losses sustained by noise traders, which exceed the tax, σ. Because arbitrage is competitive, excess losses incurred by noise traders flow in their entirety to the broker.

2.3 Deception

How does the broker choose q? Maximal deception, with $q = 0$, maximizes the broker's income but may not be advisable in the long run. If the event occurs, noise traders will know they were deceived. Even if it does not, large swings in second-period parimutuel returns will likely arise suspicion. In response, noise traders may seek out predictions from other sources, decide to invest with other brokers, quit investing altogether, or choose to become sophisticated—all of which decrease brokerage fees in the future. Less aggressive deception makes statistical inference more difficult and dampens swings in parimutuel returns, and is presumably less costly for the broker.

Rather than model these processes from first principles, we instead add a cost function to the broker's income that is increasing in the divergence between p and q—i.e., in the degree of deception. In particular, we assume that the cost of deception, $C_\alpha(q|p)$, is proportional to the Rényi (or α) divergence between p and q:

$$C_\alpha(q|p) \propto \frac{1}{\alpha - 1} \log \left(\frac{p^\alpha}{q^{\alpha-1}} + \frac{(1-p)^\alpha}{(1-q)^{\alpha-1}} \right),$$

We use the Rényi divergence because it has a single parameter, $\alpha \geq 0$, which modulates the function's convexity.[13] Its limiting case as $\alpha \to 1$ is the Kullback-Leibler divergence: $p \log[p/q] + (1-p) \log[(1-p)/(1-q)]$. Figure 1a shows Rényi divergences for different values of α. Regardless of α, the divergence is zero when $q = p$, positive when $q \neq p$ (strictly so when $\alpha > 0$), and increasing as q separates from p. For any α, the function is globally convex in q (Van Erven and Harremos, 2014). At $q = p$, the convexity in q is proportional to α.

Figure 1b shows the broker's income, cost, and profit when the cost of deception is proportional to the Kullback-Leibler divergence. Both income and cost increase with distance between p and q. When q is close to p, income increases faster than cost, implying that deception is profitable. In general, deception is profitable when the income function is more convex at $q = p$ than the cost function.[14]

When deception is profitable, the profit function has two local maxima—one when the

---

[13]Incidentally, the Rényi divergence is also commonly described as the most important divergence in information theory (e.g., Van Erven and Harremos, 2014).

[14]This implies an upper bound on the coefficient of proportionality of $(1 - \sigma)/2\alpha$.

Figure 1: Rényi divergence between q and $p = 0.6$ for different values of α (a), and broker financials when the cost of deception is proportional to the Kullback-Leibler divergence ($\alpha \to 1$).

<div style="display:flex">

(a) Rényi divergence

(b) Broker income, cost & profit

</div>
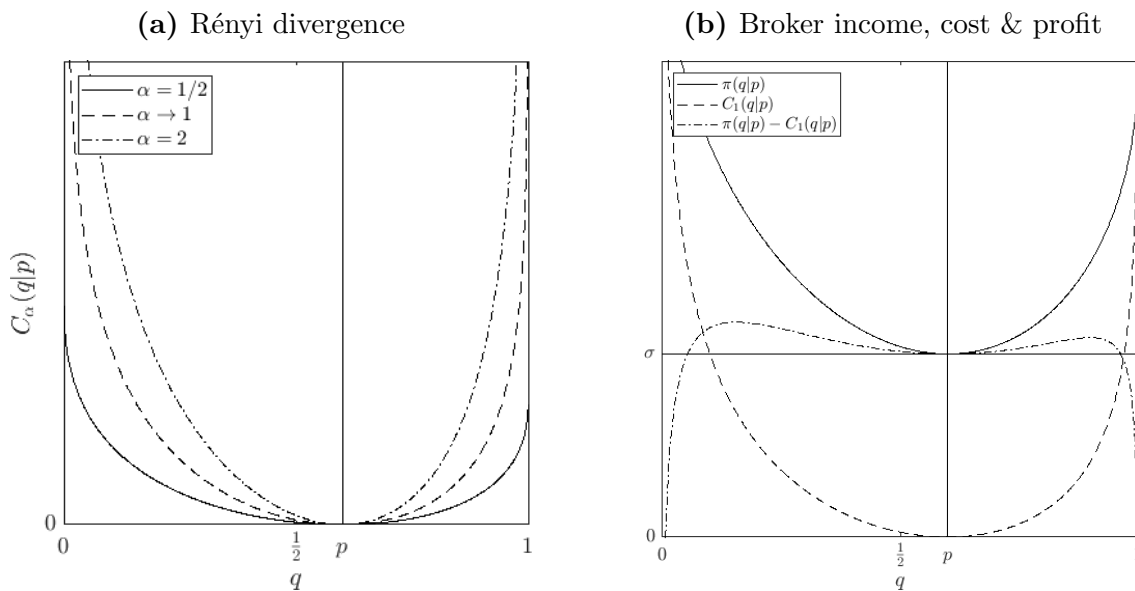
<u>Note</u>: In (b), $\sigma = 0.2$ and $C_1(q|p)$ includes a multiplicative constant of $(1 - \sigma)/3$.

broker underestimates the favorite and one when the broker overestimates the favorite. Which is preferable? If $\alpha = \frac{1}{2}$, the broker is indifferent between underestimation and overestimation.

Lemma 1. *If $\alpha = \frac{1}{2}$, then $\max_{q \geq p} \pi(q|p) - C_\alpha(q|p) = \max_{q \leq p} \pi(q|p) - C_\alpha(q|p)$.*

Proof. In Appendix A.1.

Using this lemma, we show that the broker overestimates the favorite when $\alpha < \frac{1}{2}$ and underestimates the favorite when $\alpha > \frac{1}{2}$.

Proposition 1. *Let $q^* = \arg\max_q \pi(q|p) - C_\alpha(q|p)$. If $\alpha < \frac{1}{2}$, then $q^* \geq p$. If $\alpha > \frac{1}{2}$, then $q^* \leq p$.*

Proof. In Appendix A.2.

The broker induces a favorite-longshot bias when α is sufficiently high—i.e., when deception has sufficiently convex costs. When the costs of deception increase rapidly, the broker underestimates the chances of the favorite and overestimates the chances of the longshot.

11

3 Context and data

We examine betting markets for North American horseraces, which are run as parimutuels. As in our model, bettors place wagers on outcomes (e.g., a certain horse winning the race), the track takes a state-regulated share σ 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 $(1 - \sigma)/s_i$ 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: $(1 - \sigma)/s_i - 1$.[15] For example, a winning bet at 2/1 odds pays a \$2 dividend for every \$1 wagered, along with the principal. Tracks update the parimutuel odds in real time on a central board called the tote board, as well as online, but only the final odds, or those the start of the race, are used to calculate payoffs.

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, 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 winners of two, three, four, five, or six races in a row, respectively.[16]

Popular accounts of horserace bettors support a simple typology of sophisticates and naifs. Some bettors receive a volume-based rebate. Others pay the full tax. Some bettors place their wagers electronically. Others line up at the track. Some bettors wait until the minutes or seconds before the race begins to place their wagers—i.e., when the live odds best predict the final odds (Ottaviani and Sørensen, 2006; Gramm and McKinney, 2009). Others place their wagers twenty minutes before the race begins. Some bettors train algorithms on expensive databases of race histories. Others do not have well-formed beliefs. Anecdotally,

---

[15]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.

[16]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. Parlays are inefficient. 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.

those who receive large rebates, who wager electronically, who time their bets strategically, and who analyze data tend to be the same bettors.[17]

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. Separately, we collect track- and pool-specific takeout rates from the Horseplayers Association of North America.[18]

Figure 2: An example race card entry, prominently displaying the morning-line odds (15-1).

Source: Equibase.com

We supplement these data with the morning-line odds, or the formally irrelevant win odds chosen by the oddsmaker at the track. Tracks publish morning lines in three places: 1) online; 2) on the tote board at the beginning of the betting window, when betting volume is low and parimutuel odds are unstable; and 3) in the race card, a printed pamphlet handed out at the track that provides information on each horse in each of the day's races. Figure 2 shows an example race card entry. Whereas historical race charts are freely available in an online repository, historical race *cards* are not. We scraped the morning-line odds each

---

[17]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. A *Bloomberg* profile on Bill Benter, whose algorithm made nearly $1 billion betting in horserace parimutuel markets, writes, "The odds change in the seconds before a race as all the computer players place their bets at the same time." https://www.bloomberg.com/news/features/2018-05-03/the-gambler-who-cracked-the-horse-racing-code.

[18]http://www.horseplayersassociation.org/

Table 1: Summary statistics by track.

Code	Name	State	Days	Races per day	Starts per race	Exotic share
ALB	Albuquerque	NM	97	9.6	8.0	.59
AP	Arlington	IL	94	8.1	7.3	.59
AQU	Aqueduct	NY	107	8.5	7.5	.67
BEL	Belmont	NY	90	9.0	7.8	.68
BTP	Belterra Park	OH	116	8.1	7.5	.65
CBY	Canterbury	MN	92	9.4	7.9	.58
CD	Churchill Downs	KY	83	9.6	8.5	.66
CT	Charles Town	WV	197	8.2	7.7	.64
DED	Delta Downs	LA	125	10.1	8.8	.68
DEL	Delaware	DE	121	8.8	7.1	.65
DMR	Del Mar	CA	97	8.4	8.4	.65
EVD	Evangeline	LA	136	9.0	7.8	.67
FG	Fair Grounds	LA	99	9.3	8.2	.65
FL	Finger Lakes	NY	185	8.6	6.6	.67
GG	Golden Gate	CA	144	8.2	7.0	.63
GP	Gulfstream	FL	171	10.3	8.2	.70
IND	Indiana	IN	109	9.0	8.4	.68
LA	Los Alamitos	CA	119	8.4	6.7	.76
LAD	Louisiana Downs	LA	104	7.3	7.4	.67
LRL	Laurel	MD	117	9.0	7.8	.66
MNR	Mountaineer	WV	146	8.9	7.1	.65
MVR	Mahoning Valley	OH	112	8.2	8.5	.68
PEN	Penn National	PA	129	8.1	7.4	.64
PID	Presque Isle	PA	101	7.9	7.2	.66
PRX	Parx	PA	144	9.0	8.1	.67
RP	Remington	OK	86	9.6	8.9	.67
SA	Santa Anita	CA	90	8.6	8.0	.67
SUN	Sunland	NM	70	9.2	8.7	.67
TAM	Tampa Bay Downs	FL	83	9.3	8.0	.67
TUP	Turf Paradise	AZ	125	7.9	7.8	.70

14

morning for 17 months from 2016 to 2018.[19] 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.[20]

Table 1 presents summary statistics by track. On average, each track hosts 7 to 10 races per day with 6 to 9 horses per race. Roughly two-thirds of wagers (in total dollars) are placed on exotic bets. Figure 17 in Section 7 shows takeout rates by track and pool. Gambling at the track is expensive. Takeouts range from 15% to 21% in win, place, and show pools, and from 12% to 31% in exotic pools. For comparison, bookmakers typically take 9% of winnings for common wagers on professional team sports.

All standard errors reported in 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.[21]

4   Stylized facts

Returns are systematically miscalibrated. Figure 3a shows observed returns in the win pool as a smooth function of log final odds. Wagers at all odds lose money on average, owing to the takeout. 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.[22] Bettors could reduce their losses by betting on favorites instead of longshots.

4.1   Heterogeneity

The extent of the favorite-longshot bias varies across tracks. Figure 3b shows expected returns in the win pool as a smooth function of log final odds, separately by track. (Individual plots are shown in Figure 10 in Section 6.) Some tracks exhibit little or no bias, whereas others exhibit severe bias.

---

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

[20]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.

[21]We resample entire pools, rather than individual outcomes within pools. This ensures that each resampled pool retains the same number of winning outcomes as in the original data (typically 1).

[22]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.

Figure 3: Expected returns for win bets.

(a) All tracks

(b) By track
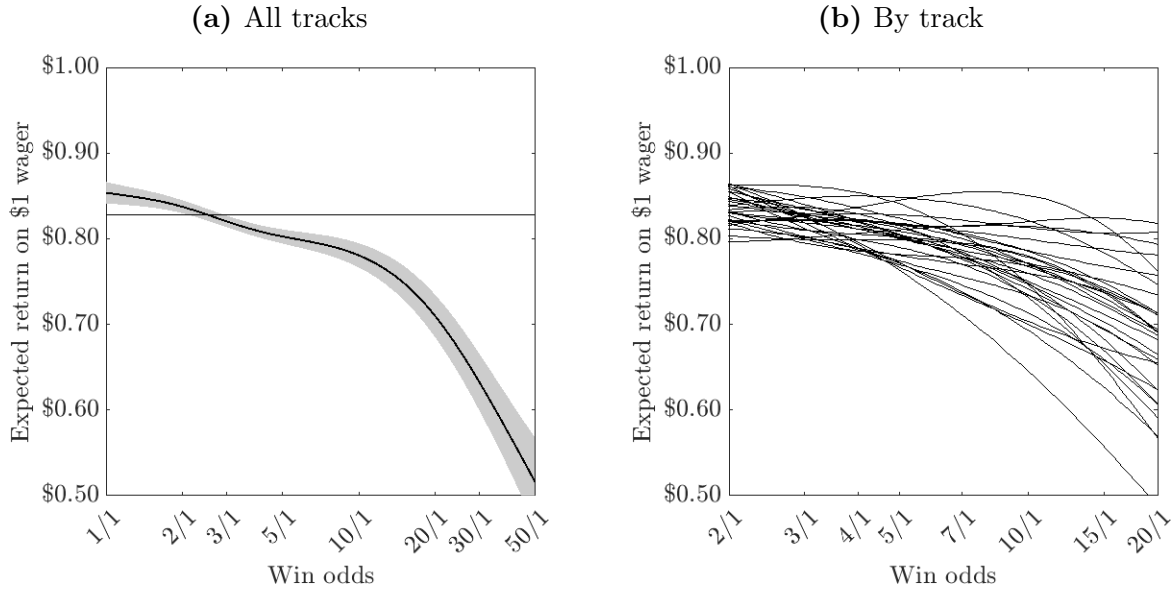
<u>Note</u>: All estimates are from a local linear regression with a Gaussian kernel. The bandwidth in (a), of 0.50 log odds, minimizes the leave-one-out mean-squared error. In (b), we use the same bandwidth for each track, of 1.0 log odds. The horizontal line in (a) marks 1 minus the average takeout.

We summarize the extent of the favorite-longshot bias at track t by its mean returns, μ_t, which we calculate by averaging observed returns in the win pool. Hence, μ_t measures the expected return on a \$1 wager for a bettor who picks horses uniformly at random. If horses are priced efficiently, random betting surrenders the takeout on average, and $\mu_t = 1 - \sigma_t$. When returns on longshots are low, however, $\mu_t < 1 - \sigma_t$. This occurs because random betting overbets longshots relative to the market.

Figure 4a shows mean returns by track, normalized by $1 - \sigma_t$.[23] The tracks are ordered by their returns at odds of 20/1 in Figure 3b, which roughly orders tracks by their normalized mean returns. 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. The variation across tracks is large. A gambler who chooses horses randomly will lose σ_t in expectation at some tracks and as much as

---

[23]In practice, the track's takeout rate is slightly higher than σ_t because of the breakage, the practice of rounding down parimutuel odds. In Figures 4a and 4b, we normalize mean returns by the true return rate, $1 / \sum_{i \in \mathcal{R}} 1/(O_i + 1)$.

16

Figure 4: Measures of the extent of the favorite-longshot bias, by track. Tracks are sorted by their expected returns at 20/1 odds in Figure 3b.

(a) Mean returns   **(b)** Weighted mean returns
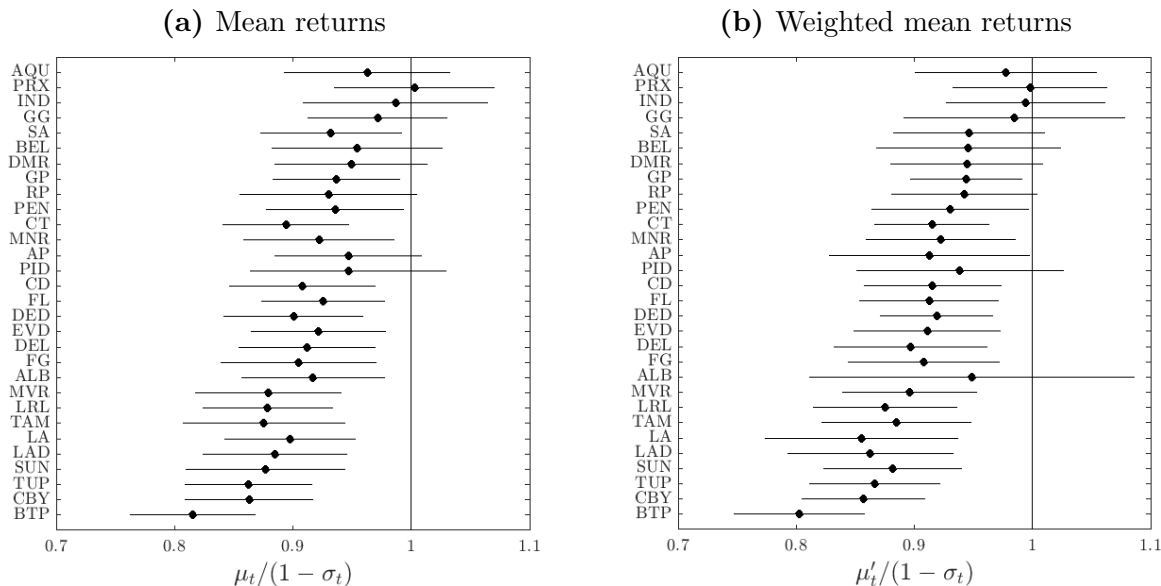
twice σ_t at others. This variation is also statistically significant: a Wald test rejects the null hypothesis that $\mu_t/(1 - \sigma_t)$ does not vary across tracks ($p < .001$).

The returns to random wagering depend on 1) mispricing in the parimutuel odds, as shown in Figure 3b, and 2) the distribution of those odds. For instance, a track with more longshots will have lower mean returns than a track with fewer longshots, even if longshots at both tracks are overbet to the same degree. To isolate the relationship between odds and returns, we re-weight observations such that the weighted distribution of win odds at each track converges to the distribution of win odds across all tracks.[24] Let μ'_t denote the weighted mean returns at track t, which we calculate by taking a weighted average of observed returns in the win pool. Hence, μ'_t measures expected returns when sampling horses according to the pooled distribution of odds, rather than uniformly at random. As with μ_t, $\mu'_t = 1 - \sigma$ when prices are efficient, and $\mu'_t < 1 - \sigma$ when longshots yield lower returns than favorites. Figure 4a shows weighted mean returns by track, normalized by $1 - \sigma_t$. Across tracks, weighted mean returns correlate with mean returns at 0.93. Again, variation across tracks is statistically significant ($p < .001$).

---

[24]For each track, we find weights, one for each observation and summing to 1, that minimizes the Kolmogorov-Smirnov test statistic—i.e., the maximal deviation—between the CDF of the weighted distribution of parimutuel odds and the CDF of the (unweighted) pooled distribution across a grid of (log) evenly spaced points. Without weighting, the KS statistic varies between 1.8pp and 7.8pp across tracks with a mean of 4.3pp. With optimal weights, it ranges from 0.1pp to 0.9pp, with a mean of 0.4pp.

4.2 Morning-line odds

The favorite-longshot bias appears to be related to the morning-line odds. One immediate observation is that the morning-line odds fail at their ostensible goal of predicting the parimutuel odds. Figure 5a 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 in excess of 50/1. The morning-line odds are truncated, as shown in Figure 16 in Section 6, rarely shorter than 1/1 and rarely longer than 30/1. If oddsmakers are trying to predict final odds, they could do better by predicting more extreme values.

Figure 5: Morning-line odds and parimutuel odds.

(a) Calibration plot

(b) Time series
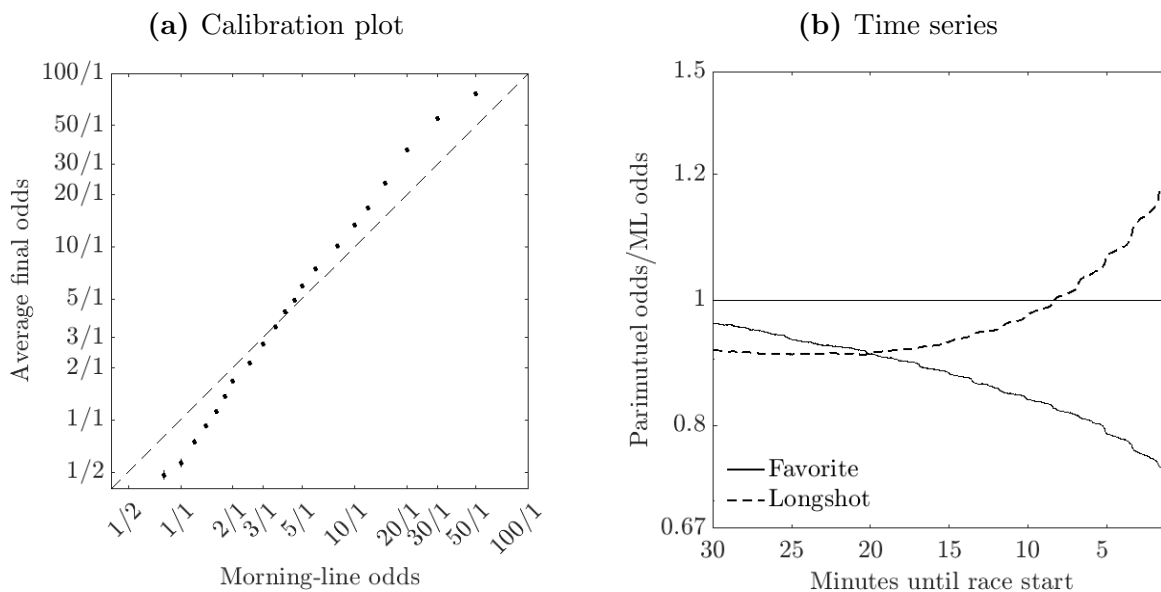
Note: In (a), the 95% confidence intervals are obscured by the point-estimate markers for all but the shortest odds. In (b), the favorite and longshot are defined by the morning-line odds; estimates reflect second-by-second averages interpolated from scrapes at roughly 2-minute intervals, for 6,500 races.

The separation between the morning lines and the parimutuel odds is a consequence of late wagering. Figure 5b plots a time series of live odds from a separate dataset of 6,500 US races. We plot the average log ratio of parimutuel odds to morning-line odds in the half-hour before the race, separately for the favorite and the longshot (as defined by the

18

morning-line odds). On average, the parimutuel odds for the favorite and longshot begin near their respective morning-line odds. Thereafter, the longshot's odds lengthen and the favorite's odds shorten, and this trend accelerates as the race nears (see also: Asch, Malkiel and Quandt, 1982; Camerer, 1998). This implies that late wagers are disproportionately placed on the favorite.

The morning-line odds not only mispredict the final odds—they also imply distorted beliefs about a horse's chances of winning. If bettors are risk-neutral, then the expected value of betting on any horse—given subjective beliefs q—must be the same for all horses in the race. That is, $q_i(O_i + 1) = q_j(O_j + 1) \ \forall i, j \in \mathcal{R}$, where $\mathcal{R}$ is the set of horses in a race, O denotes the final parimutuel odds, and $O + 1$ is the associated return on a winning \$1 wager. Further, if bettors place wagers such that the parimutuel odds converge to the morning lines, as the track oddsmaker predicts, then $q_i(l_i + 1) = q_j(l_j + 1) \ \forall i, j \in \mathcal{R}$, where l is the morning line. Given that $\sum q_j = 1$, the implied probability of horse i winning is:

$$q_i = \frac{1}{l_i + 1} \bigg/ \sum_{j \in \mathcal{R}} \frac{1}{l_j + 1} \tag{1}$$

The beliefs implied by the morning-line odds (under risk-neutrality) are the inverse of the associated returns on a \$1 wager, normalized such that these beliefs sum to 1. Typically, the normalization constant is roughly $1 - \sigma_t$, implying that the morning-line odds are odds that could occur in the parimutuel.

Figure 6: Observed win rate, p, as a function of the win probability implied by the morning-line odds, q, with a histogram of q. See Footnote 25 for details on the estimation.
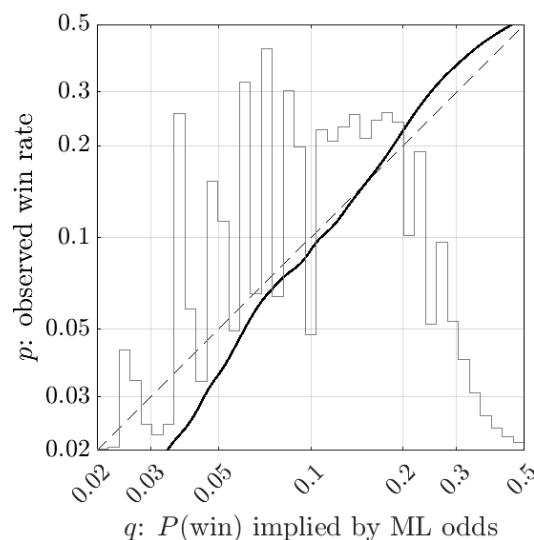
These beliefs are misleading. Figure 6 shows a calibration plot—the observed win rate (denoted by p) as a smooth function of q, the win probabilities implied by the morning-line odds.[25] Longshots underperform their implied win probabilities. At $q = 0.03$, for example, horses win just 1 in 100 races; at $q < 0.03$, horses effectively never win. Morning-line odds of 30/1 imply win probabilities of $q \approx 0.03$. 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 of $q \approx 0.02$. Of the 586 starts assigned morning-line odds of 50/1 or longer, just 2, or 0.3%, won the race. Symmetrically, favorites outperform their implied win probabilities. At $q = 0.3$, for example, $p = 0.37$.

The morning-line odds mislead at some tracks but not at others. Let $p_t(q_j)$ be the empirical rate at which a horse with implied probability q_j at track t wins, and let p_j be the normalized rate such that $\sum p_j = 1$. Figure 7 shows these normalized rates for each horse, separately by track.[26] Some tracks post well-calibrated morning-line odds (e.g., PRX), whereas others post miscalibrated morning-line odds (e.g., FG). Tracks that promulgate miscalibrated predictions mislead 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.

We summarize the extent of the miscalibration at a given track t by the Kullback-Leibler divergence between the implied winning probabilities, q_j, and the corresponding normalized winning rates, p_j, and we denote this measure δ_t. Specifically, we compute the divergence for each race as $\sum_{j \in \mathcal{R}} p_j \log(p_j/q_j)$, and we define δ_t as the average divergence across races. If $q = p$ for each entry in each race, $\delta_t = 0$; otherwise $\delta_t > 0$.

Across tracks, miscalibration in the morning-line odds predicts the favorite-longshot bias. Figure 8a shows a correlation of -0.61 ($p < .001$) between δ_t and mean returns, $\mu_t/(1 - \sigma_t)$,

---

[25]We regress an indicator for whether the observed horse won the race on $\log q$ using local constant regression with an adaptive bandwidth. For a given q, the bandwidth is $h\lambda(q)$, where h is 0.1, and λ is a local bandwidth factor that increases the bandwidth in regions of low density. This multiplicative factor is $\lambda(q) = \sqrt{\exp\left(\frac{1}{N}\sum_i \log \tilde{f}(\log q_i)\right)/\tilde{f}(\log q)}$, where $\tilde{f}$ is a kernel density estimate of $\log q$ using Silverman's rule-of-thumb bandwidth.

[26]We amend the estimation in four ways from the process described in Footnote 25. First, we estimate the function separately by track. Second, we predict win probabilities using cross validation, which ensures that the prediction for a given start is independent of whether that horse won the race. For each race, we estimate a smooth function using only a training sample of starts from *other* races at that track. Third, we ensure that each of these estimated functions is weakly increasing, as is true naturally in the pooled data in Figure 6. To do so, we follow Hall and Huang (2001) in assigning a weight to each start in the training sample and minimizing the Kullback–Leibler divergence of the weight vector via quadratic programming subject to the constraints that the function $p(q)$ is weakly increasing and that the weights sum to 1. Finally, we normalize the predicted win probabilities to sum to 1 in each race.

Figure 7: Predicted win rate (p) on the vertical axis as a function of the win probability implied by the morning-line odds (q) on the horizontal axis, for each entry at each track. See Footnote 26 for details on the estimation.
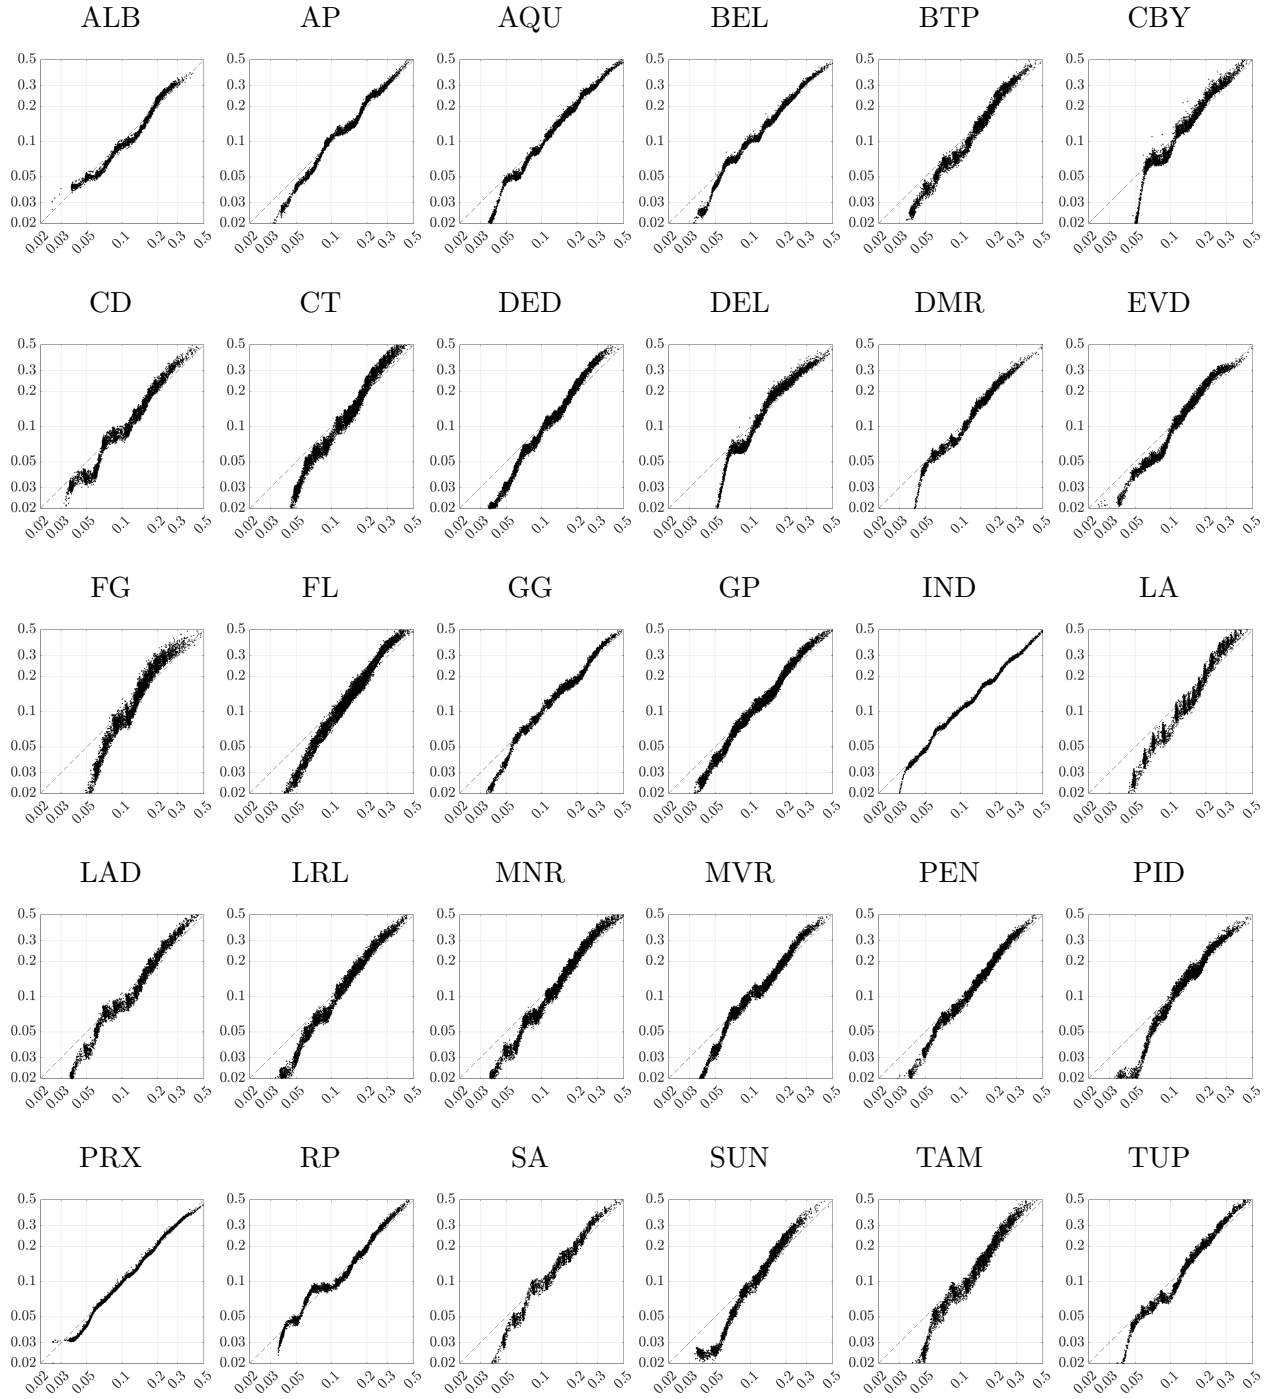
Figure 8: Scatter plot, with a regression fit line, of miscalibration in the morning-line odds, δ_t, and measures of the severity of the favorite-longshot bias, across tracks.

(a) Mean returns

(b) Weighted mean returns
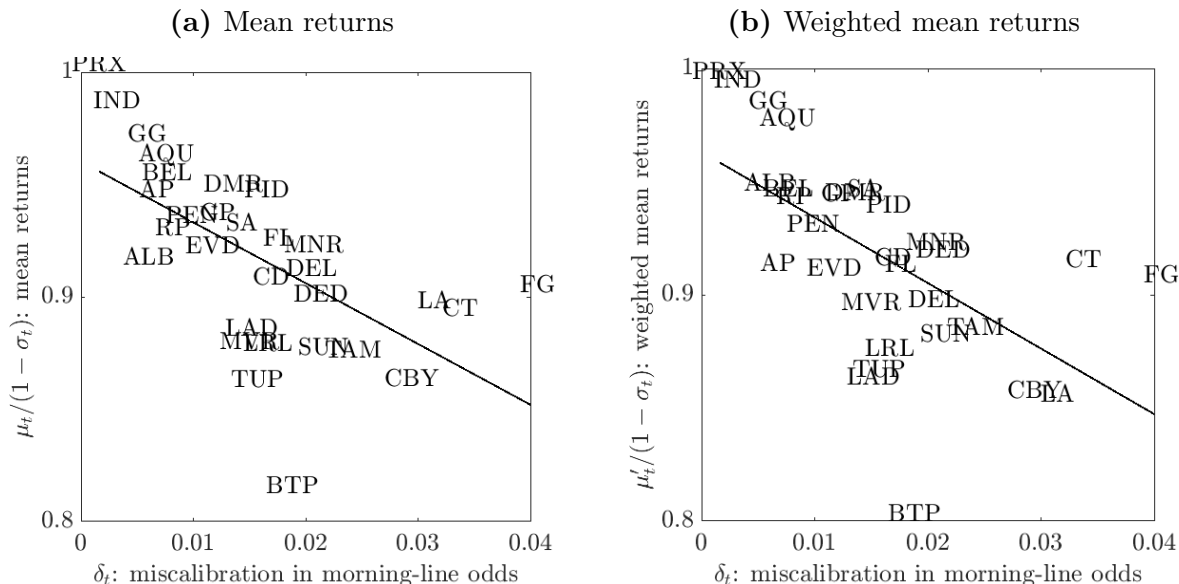
from Figure 4a. Figure 8b shows a correlation of -0.60 ($p < .001$) between δ_t and weighted mean returns, $\mu'_t/(1 - \sigma_t)$, from Figure 4b. The more severe the miscalibration in the morning-line odds, the more severe the favorite-longshot bias.

It is unclear why miscalibration in the morning-line odds varies across tracks. There are too few tracks in North America—let alone in our data—to identify correlates of miscalibration. Nonetheless, an anecdote suggests one predictor: falling income. Wagering at Canterbury Park (CBY) decreased from 2015 to 2016.[27] At the same time, Canterbury Park aggressively overestimated the chances of longshots, as seen in Figure 7.[28] Consistent with our model, a track with good reason to discount the future is among the most misleading.

5   Empirical model

These stylized facts motivate an application of our model from Section 2 to horserace parimutuel markets. In this section, we generalize the model to an arbitrary number of outcomes (e.g., horses in a race). This allows us to predict an array of quantities—arbitrage volume, optimal rebates, track income, parimutuel odds, and expected returns—solely from

---

[27]http://www.thoroughbreddailynews.com/oped-why-canterburys-takeout-reduction-didnt-work/
[28]At Canterbury, horses with 15-1 morning lines (average q of 5.6%) won 2.8% of races, horses with 20-1 morning lines (average q of 4.2%) won 1.1% of races, and horses with 30-1 morning lines (average q of 2.9%) did not win a race in 97 tries.

the beliefs of noise traders and arbitrageurs, which we infer from the morning-line odds.

5.1 Setup

The set-up is identical to Section 2, with the exception that there are now $N \geq 2$ outcomes, indexed by i. The track announces predictions q_i such that $\sum_{i=1}^{N} q_i = 1$.

In the first period, noise traders take the track's predictions at face value and wager on the outcome with the highest subjective expected value. Noise trader j values a \$1 investment in outcome i at $q_i(1-\sigma)/(1-s_i)+u_j$, where s_i is the share of first-period wagering on outcome i and u_j is the direct utility that noise trader j receives from wagering \$1. A large number of noise traders each decide whether to wager a \$1 endowment, and if so, which outcome to bet on. We are interested in s_i^*, the Walrasian equilibrium share of first-period wagering on each outcome.

Arbitrageurs wager in the second period. Unlike noise traders, arbitrageurs know p_i, and they do not receive direct utility from gambling. Let x_i be the amount wagered by arbitrageurs on outcome i. Normalizing total wagering by noise traders to 1, the expected return on a \$1 wager on outcome i is:

$$\mathbb{E}[V(\$1 \text{ on } i)] = p_i(1-\sigma)\frac{1+\sum_j x_j}{s_i^* + x_i}$$

The track price discriminates by offering arbitrageurs a rebate, r, on each dollar wagered regardless of the outcome. For every dollar wagered by noise traders, the track receives:

$$\pi(r) = \sigma + (\sigma - r)\sum_{i=1}^{N} x_i^*$$

We assume that arbitrage is competitive. Hence, the second-period equilibrium is defined by arbitrage volumes $\{x_i^*\} \geq 0$ and a rebate r^* such that:

1. Arbitrageurs make zero profits: $\mathbb{E}[V(\$1 \text{ on } i)] = \$1 - r^*$ when $x_i^* > 0$

2. Arbitrageurs leave no money on the table: $\mathbb{E}[V(\$1 \text{ on } i)] \leq \$1 - r^*$ when $x_i^* = 0$

3. The track chooses the income-maximizing rebate: $r^* = \arg\max_r \pi(r)$

5.2 Results

In the first-period Walrasian equilibrium, $s_i^* = q_i$, and only those for whom $u_j > \sigma$ invest. First-period parimutuel odds are $O_i^{(1)} = (1 - \sigma)/q_i - 1$.

For the second-period, we first derive the condition under which an equilibrium exists.

Lemma 2. *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 if and only if there exists an index m such that:

$$\frac{p_m}{q_m} > \sqrt{\frac{P_m(1 - P_m)}{Q_m(1 - Q_m)}} \geq \frac{p_{m+1}}{q_{m+1}}, \tag{2}$$

where $P_m \equiv \sum_{i=1}^m p_i$ and $Q_m \equiv \sum_{i=1}^m q_i$. In this equilibrium, $x_i > 0$ for $i \leq m$ and $x_i = 0$ for $i > m$.

Proof. In Appendix A.3.

The index m separates outcomes that arbitrageurs wager on from those they do not. P_m and Q_m are the cumulative subjective probabilities—held by arbitrageurs and noise traders, respectively—among outcomes wagered on by arbitrageurs.

Arbitrageurs wager on outcomes for which their own beliefs are high relative to those of noise traders.

Corollary 1. $P_m \geq Q_m$.

Proof. In Appendix A.4.

Under a mild condition, an equilibrium exists:

Proposition 2. *If $p_1 > q_1$, then there exists an $m \in \{1, \ldots, N-1\}$.*

Proof. In Appendix A.5.

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 arbitrageurs always wager on the outcome with the highest ratio of subjective beliefs. In addition, $m < N$, implying that arbitrageurs never wager on the outcome with the lowest ratio of subjective beliefs.

24

Equilibrium wagers by arbitrageurs are:

$$x_i^* = p_i\sqrt{\frac{Q_m(1-Q_m)}{P_m(1-P_m)}} - q_i \text{ for } i \in \{1,\dots,m\} \tag{3}$$

and $x_i^* = 0$ for $i \in \{m+1,\dots,N\}$. The amount wagered by arbitrageurs on outcome i is increasing in its true probability, p_i, and decreasing in the beliefs of noise traders, q_i, all else equal.

Other quantities all have analogous expressions to those in Section 2, with P_m replacing p and Q_m replacing q. Let:

$$\gamma_m \equiv \frac{P_m(1-Q_m)}{Q_m(1-P_m)}$$

Arbitrageurs collectively wager:

$$\sum_{i=1}^{N} x_i^* = Q_m(\sqrt{\gamma_m} - 1) \tag{4}$$

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

$$\sigma - r^* = (1-\sigma)(1-P_m)(\sqrt{\gamma_m} - 1) \tag{5}$$

And the track's income is:

$$\pi_t(r^*) = \sigma + (1-\sigma)\left(1 - \left(\sqrt{P_m Q_m} + \sqrt{(1-P_m)(1-Q_m)}\right)^2\right), \tag{6}$$

where σ is the track's income from noise traders, and the second term is track's income from arbitrageurs. The track's income from arbitrage is increasing in the divergence between P_m and Q_m. If $P_m = Q_m$, which occurs only when $p_i = q_i$ for all i, arbitrageurs place zero wagers, and the track generates zero income from arbitrage. Deviations between p_i and q_i, and hence between P_m and Q_m, generate wagers from arbitrageurs and income for the track.

The final parimutuel odds are:

$$O_i^{(2)} = (1-\sigma)\begin{cases} \frac{1}{p_i}\left[P_m + (1-P_m)\sqrt{\gamma_m}\right] - 1, & i \in \{1,\dots,m\} \\ \frac{1}{q_i}\left[Q_m\sqrt{\gamma_m} + (1-Q_m)\right] - 1, & i \in \{m+1,\dots,N\} \end{cases} \tag{7}$$

25

Before the rebate, the true expected value of a \$1 wager on outcome i is:

$$\mathbb{E}[V(\$1 \text{ on } i)] = (1 - \sigma) \begin{cases} P_m + (1 - P_m)\sqrt{\gamma_m}, & i \in \{1, \dots, m\} \\ \frac{p_i}{q_i}\left[Q_m\sqrt{\gamma_m} + (1 - Q_m)\right], & i \in \{m+1, \dots, N\} \end{cases} \tag{8}$$

Arbitrage equalizes returns among outcomes on which arbitrageurs wager. A \$1 wager on any of the first m outcomes has an expected return of $\$1 - r^*$. For $i > m$, wagers perform worse in expectation, following from (2), and they perform progressively worse as i increases, given that p_i/q_i decreases in i.

It is straightforward to see how a favorite-longshot bias in the track's predictions manifests in the parimutuel odds. A favorite-longshot bias characterizes the track's predictions if $p_1 \geq p_2 \geq \cdots \geq p_N$, as this implies that the track underestimates the favorite ($p_1 > q_1$) and overestimates the longshot ($p_N < q_N$). Further assume that $q_1 \geq q_2 \geq \cdots \geq q_N$—i.e., the track's predictions correctly order the outcomes by their true probabilities. Given that the parimutuel odds in (7) are inversely proportional to p_i, this indexing orders outcomes by their odds, from short to long. For outcomes with shorter odds (i.e., $i \leq m$), expected returns are flat. For outcomes with longer odds (i.e., $i > m$), expected returns are decreasing. Hence, the relationship between odds and expected returns is piecewise linear, similar to that observed in Figure 3a, with the kink located at the odds of outcome m.

5.3   Estimation

We infer beliefs from the morning lines. Noise traders take the morning-line odds at face value as in (1). In contrast, we imbue arbitrageurs with the well-calibrated beliefs, $p_{i,\text{win}}$, shown in Figure 7.[29] Noise traders assume that the morning lines imply well-calibrated beliefs. Arbitrageurs ensure that their beliefs are well calibrated.

We construct beliefs for exotic wagers from beliefs in the win pool using the Harville formula. Assuming that outcomes are independent, the probability of finishing in n^{th} place is the probability of winning a race without the first $n - 1$ finishers (Harville, 1973).

We first consider place and show wagers, which pay out if the chosen horse finishes in the

---

[29]Because this routine is computationally intensive, we do not repeat it during bootstrapping. In other words, we treat p_i as data rather an estimate.

top 2 or 3 positions, respectively. The probability of a place wager on horse i paying out is:

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

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

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

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 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:

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

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:

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

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 satisfies the relationship in (2); if more than one such index exists, we choose the one that maximizes

track income. For each outcome, we calculate the equilibrium odds in (7) and expected returns in (8). In Section 6, we compare the relationship between equilibrium odds and expected returns predicted by our model to that observed in the data. We then calculate the optimal rebate from (5) and track revenues under the optimal rebate (6), which we report in Section 7.[30]

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 tote board 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.

Name	Odds		Beliefs			Model predictions in win pool				
	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	.49	1.20	1.0	0.99	0.5	0.7	0.85
Tribal Heat	2/1	4.0	.27	.32	1.17	2.0	0.97	0.3	1.6	0.85
Dandy Candy	6/1	6.9	.12	.09	0.80	6.0	0.66	0.05	8.1	0.85
Click and Roll	8/1	15.6	.09	.06	0.68	8.0	0.57	0.02	12.6	0.85
Eveatetheapple	20/1	69.3	.04	.01	0.32	20.1	0.27	0	38.8	0.50
Movie Starlet	20/1	25.5	.04	.01	0.32	20.1	0.27	0	38.8	0.50
Boston Banshee	30/1	49.7	.03	.00	0.13	30.2	0.11	0	57.8	0.20

The beliefs ascribed to noise traders, q_i, are those of a risk-neutral bettor who believes that the final odds will coincide 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 arbitrageurs, p_i, are well calibrated—i.e.,

---

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

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

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

proportional to the track-specific rates at which horses with implied beliefs q_i actually win, as shown in Figure 7. For Charles Town, beliefs of $q_i > 0.2$ tend to be too pessimistic, and beliefs of $q_i < 0.2$ tend to be too optimistic. As a result, $p_i > q_i$ for 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 noise traders, 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, the practice of rounding down parimutuel odds. 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 $p = 0.49$ favorite is almost an even money bet. By contrast, at odds of 20.1, the $p = 0.01$ longshots return just 27 cents on every dollar wagered.

The second-period equilibrium consists of a set of wagers by arbitrageurs, x_i^*, such that they make zero profits and leave zero profits on the table. To find this equilibrium, we sort the outcomes in descending order of p_i/q_i and find the equilibrium index m. 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 income-maximizing equilibrium is $m = 4$, implying that $x_i > 0$ for the first 4 horses and $x_i = 0$ for the remaining 3. 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 return on a \$1 wager, $\mathbb{E}_i^{(2)}$, is 85 cents, with the 15-cent expected loss equaling the optimal rebate. For the longshots, expected losses attenuate in the second period. A \$1 wager on either of the horses with 20/1 morning lines, for instance, returns 50 cents in expectation at odds of 38.8, or 23 cents more than at first-period odds of 20.1.

In the win pool, arbitrageurs wager 89 cents for every dollar wagered by the noise traders. However, income from arbitrageurs comprises just 11% of the track's total income in the win pool, as the track offers arbitrageurs a 15-cent rebate. Arbitrage accounts for larger shares of the track's income in exotic pools. In the exacta pool, for instance, that share is 20%; in the superfecta pool, it is 30%. The compound beliefs attached to outcomes in exotic pools amplify differences in beliefs about win probabilities between noise traders and arbitrageurs and hence, taxes from arbitrageurs. This logic implies larger divergences between the cumulative beliefs P_m and Q_m. In the win pool, for example, $P_m = 0.97$ and $Q_m = 0.90$, whereas in the superfecta pool, $P_m = 0.81$ and $Q_m = 0.48$.

6 Model predictions

We begin by comparing predictions from the first- and second-period equilibria. Figure 9a shows a smoothed estimate of the relationship between log predicted odds and predicted returns for win bets, separately by period. (All predicted returns reflect expected returns absent the rebate and under well-calibrated beliefs.) In the first-period equilibrium, predicted odds approximate the morning-line odds, which embed a favorite-longshot bias that is more severe than that observed in the data.

Figure 9: Predicted and observed returns for win bets. The horizontal line marks the average takeout.

(a) Our model   **(b)** Representative-agent models
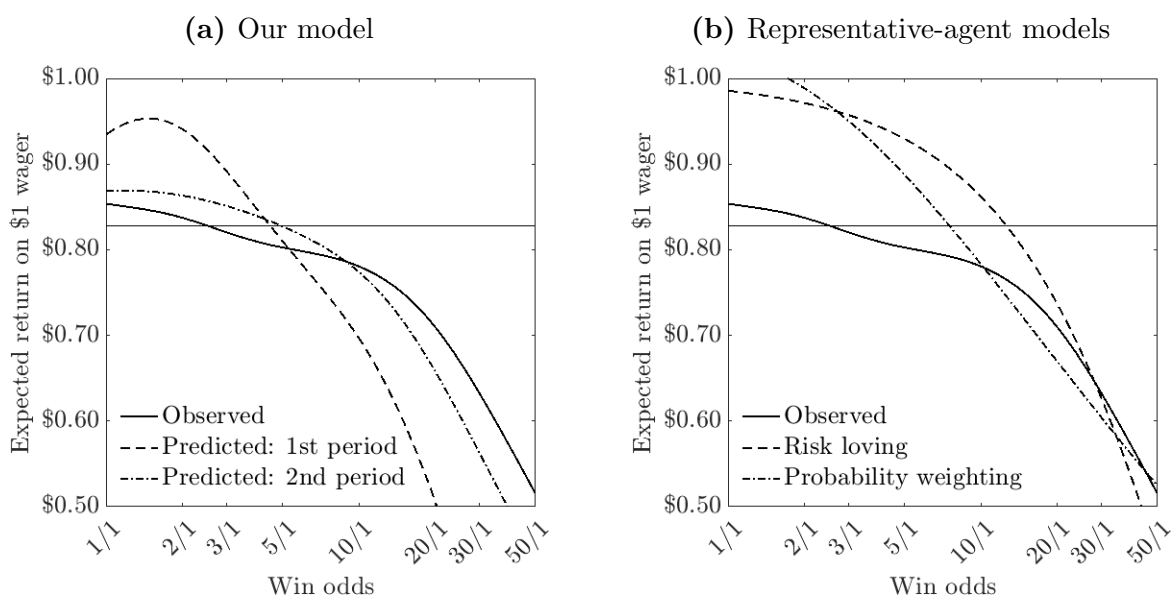
Note: Estimated using local linear regression with a Gaussian kernel. For comparability, we use the same bandwidth for the predictions as for the observed estimate, of 0.50 log odds.

The inclusion of arbitrageurs in the second period moderates the favorite-longshot bias, and the resulting predictions more closely follow observed returns. For a 1/1 favorite, first-period expected returns exceed observed returns by 8 cents; after arbitrage, the difference is 1.5 cents. For a 30/1 longshot, observed returns exceed first-period expected returns by 28 cents; after arbitrage, the difference is 7 cents.

We quantify the goodness of our model's predictions by measuring the average deviation: the mean absolute distance between predicted and observed returns, with the expectation

30

taken over the distribution of observed odds.[31] For a bettor who wagers randomly, the mean difference between her returns and those predicted by the model converges asymptotically to our average deviation measure. The average deviation is 3.4 (se: 0.6) cents for the model's second-period predictions, compared to 13.5 (0.6) cents before arbitrage.

Representative-agent models perform worse. Previous work has rationalized the favorite-longshot bias with risk-loving preferences (e.g., Weitzman, 1965; Ali, 1977) or a tendency to overweight small probabilities (e.g., Snowberg and Wolfers, 2010). Figure 9b shows predicted returns from each model.[32] For both models, predicted returns for favorites exceed observed returns, and the overall fit is poor. Average deviations, of 9.2 (0.5) cents for the risk-loving model and 7.0 (0.4) cents for the probability-weighting model, exceed the 3.4-cent deviation from our model with arbitrage. The difficulty is that neither model can rationalize highly negative returns for probable events. Risk-loving agents pay a large premium for lottery tickets, but they pay a small premium for gambles with little upside. Prospect-theory agents overweight large probabilities and thus need to *receive* a premium in order to wager on likely events. Neither explanation can account for a willingness to lose considerable sums, in expectation, when wagering on favorites.[33]

Our model with arbitrage also captures differences across tracks in the extent of the bias. Figure 10 shows estimates of the observed and predicted relationships between log odds and

---

[31]We estimate this distribution using a kernel density estimator with a Gaussian kernel and Silverman's rule-of-thumb bandwidth.

[32]We estimate these models following Snowberg and Wolfers (2010). The equilibrium conditions are $p_i U(O_i + 1) = 1 \ \forall i \in \mathcal{R}$ for the risk-loving model and $\pi(p_i)(O_i + 1) = 1 \ \forall i \in \mathcal{R}$ for the probability-weighting model. Each model is governed by one parameter. For the risk-loving model, Snowberg and Wolfers (2010) 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, 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.

[33]These explanations provided a better fit decades ago, when favorites were close to even-money bets (Thaler and Ziemba, 1988). See the end of Section 7 for a discussion of why the favorite-longshot bias has moderated over time.

Figure 10: Observed (solid line) and predicted (dashed line) returns for win bets with 95% confidence intervals, by track. The horizontal line marks $1 - \sigma_t$.
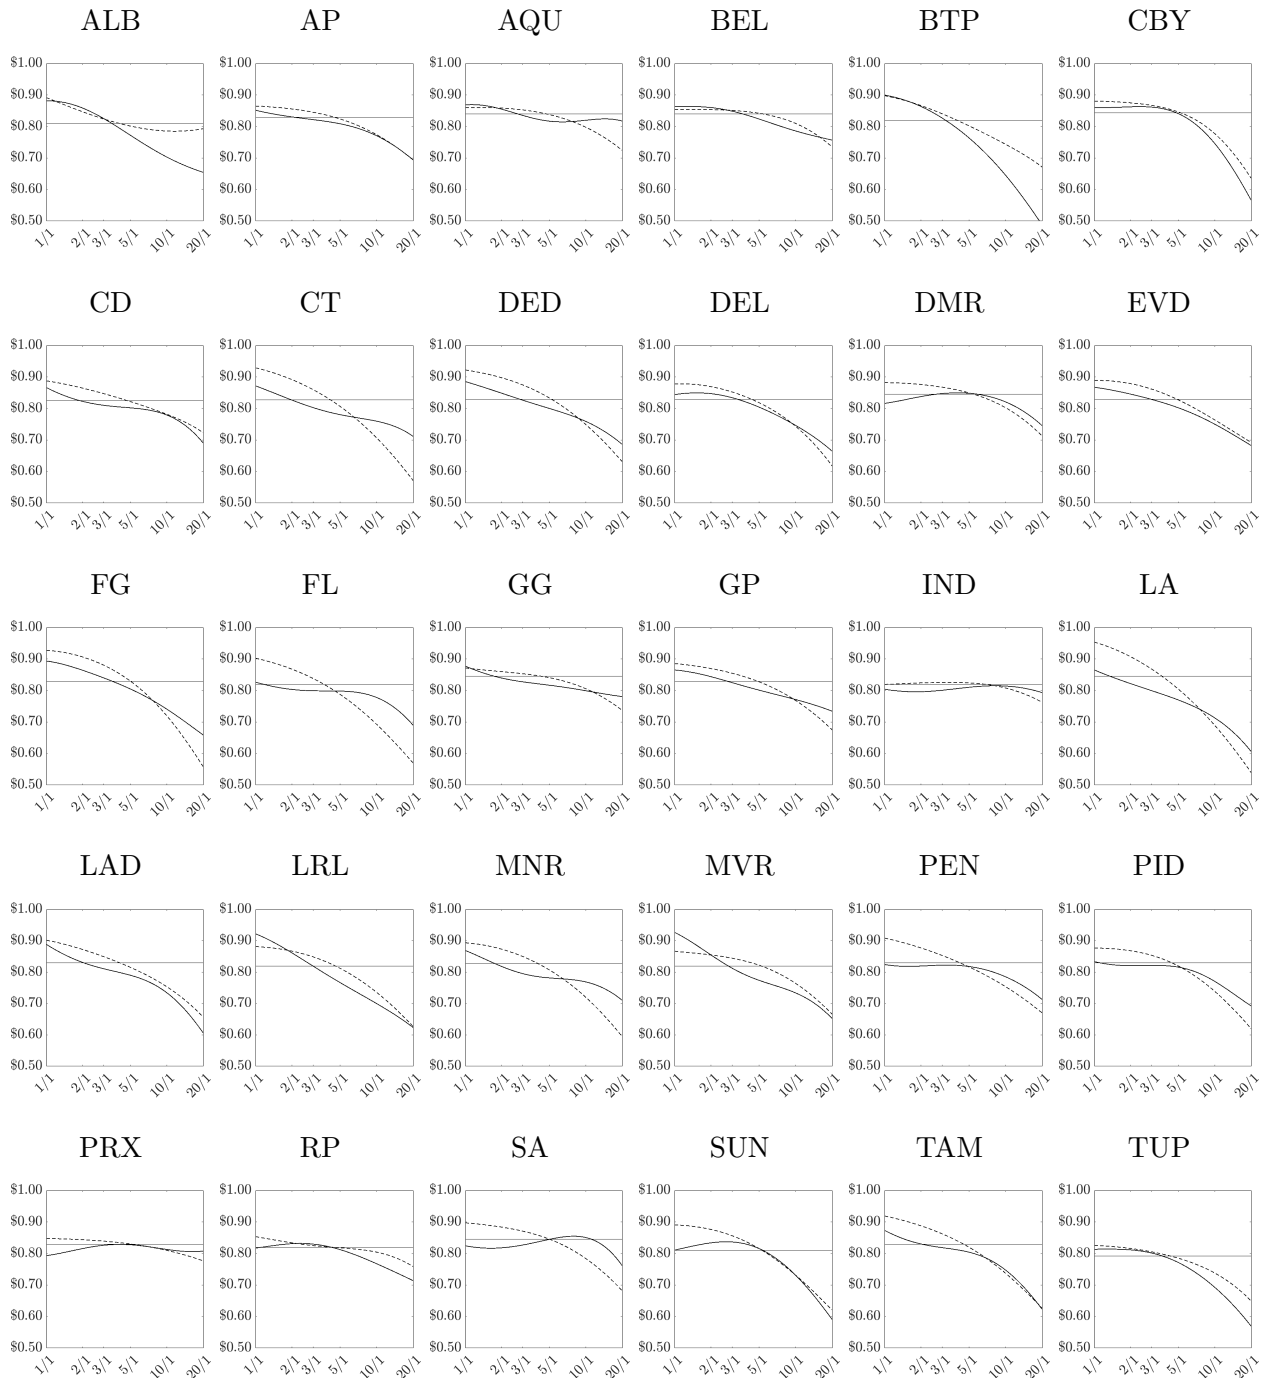
<u>Note</u>: Estimated using a local linear regression with a Gaussian kernel. For comparability, we use a common bandwidth for all estimates, of 1.0 log odds.

Figure 11: Measures of the severity of the favorite-longshot bias. The prediction from the model is on the horizontal axis, and the estimate from the data is on the vertical axis.

(a) Mean returns   **(b)** Weighted mean returns
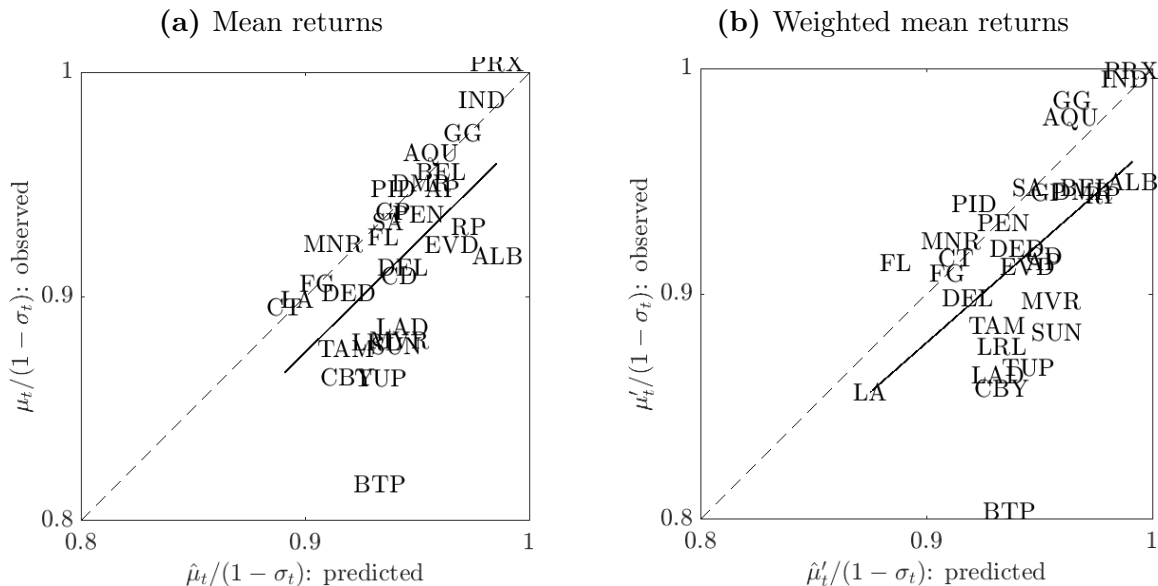
expected returns, separately by track. For most tracks, the predicted relationship follows the observed relationship. Figure 11a summarizes the model's fit at each track by comparing observed and predicted mean returns—i.e., the expected return on a randomly placed \$1 win bet, normalized by the state-sanctioned return, $1 - \sigma_t$. Observed and predicted values are correlated at 0.60 ($p < .001$). Figure 11b performs the same comparison for weighted mean returns, for which observed and predicted values are correlated at 0.56 ($p < .01$).

6.1 Other pools

We evaluate our second-period predictions in other pools as well. Since the race chart only lists final odds 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.

Figure 12 shows observed and predicted returns in place (12a) and show (12b) pools. 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 overestimation for favorites and an underestimation for longshots. In place pools, the average deviation is 5.6 cents; in show pools, it is 4.5 cents.

The final three sets of figures show expected returns on wagers involving two horses—

Figure 12: Observed and predicted returns for place and show bets, with 95% confidence intervals. The horizontal line marks the average takeout.

(a) Place

(b) Show
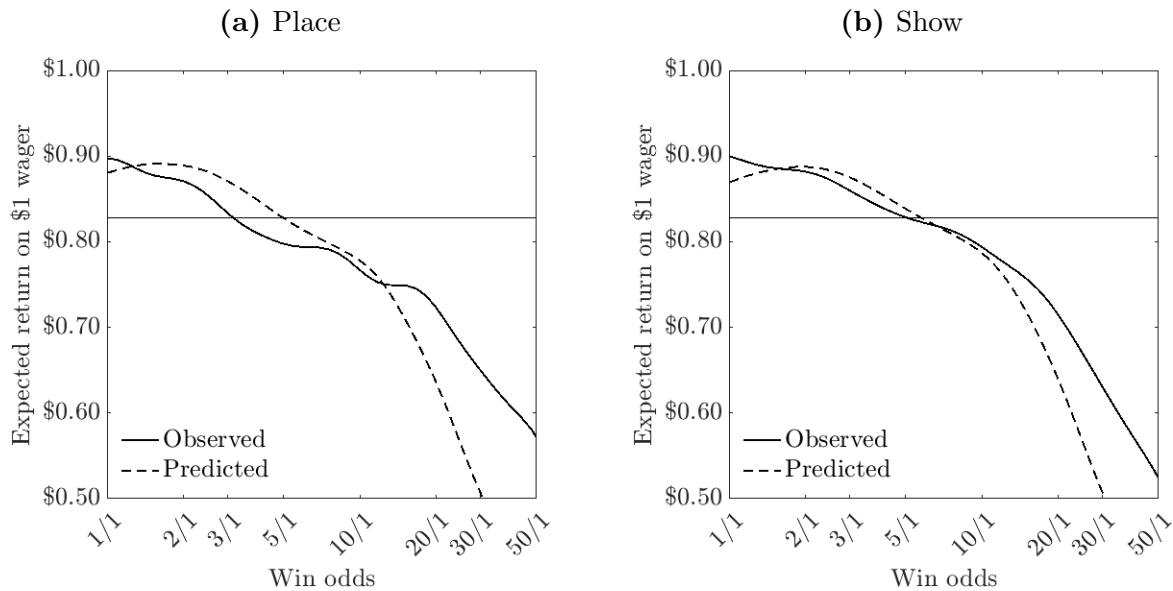
<u>Note</u>: Estimated using a local linear regression with a Gaussian kernel. The bandwidth, of 0.21 log odds for place bets and 0.26 log odds for show bets, minimizes the leave-one-out mean-squared error in the observed data.

those in exacta (13), quinella (14), and daily-double (15) pools. In each set of pool-specific figures, (a) shows observed returns, (b) shows predicted returns, and (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.)[34]

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 average deviations of 5.4 cents for exacta bets, 3.1 cents for quinella bets, and 6.1 cents for daily-double bets.

---

[34]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 13: Exacta: expected return on $1 wager.

(a) Observed   **(b)** Predicted   **(c)** Observed − predicted
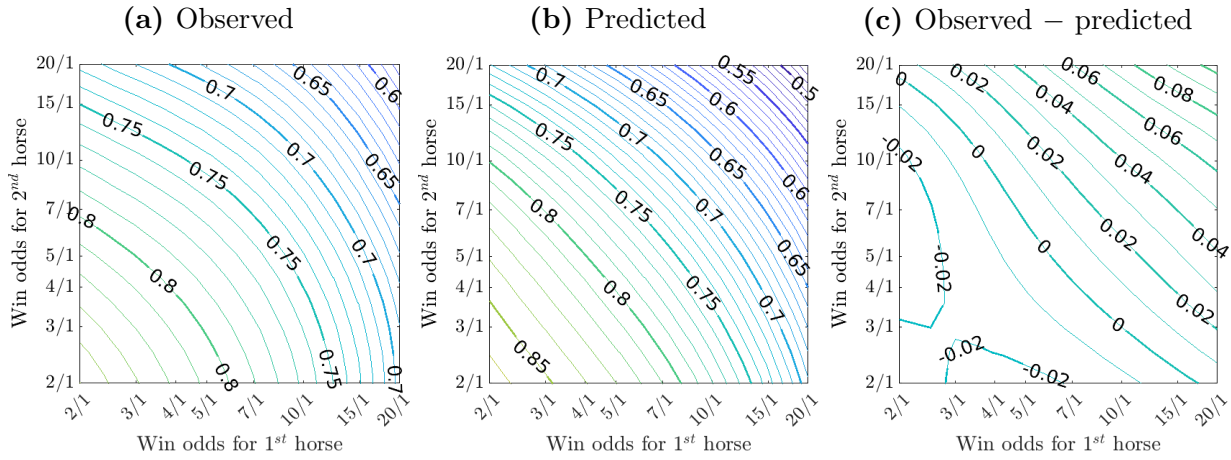
Figure 14: Quinella: expected return on $1 wager.

(a) Observed   **(b)** Predicted   **(c)** Observed − predicted
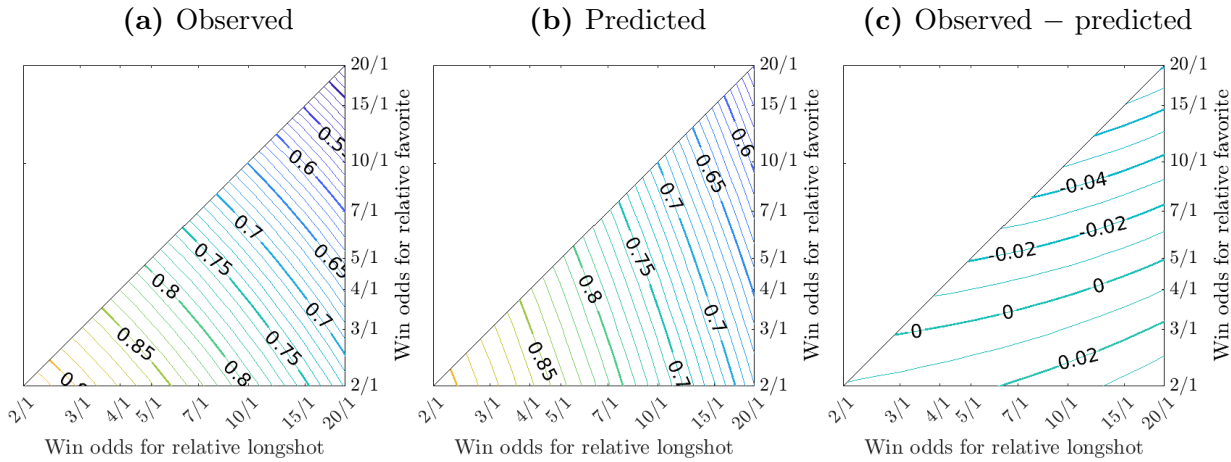
Figure 15: Daily double: expected return on $1 wager.

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

35
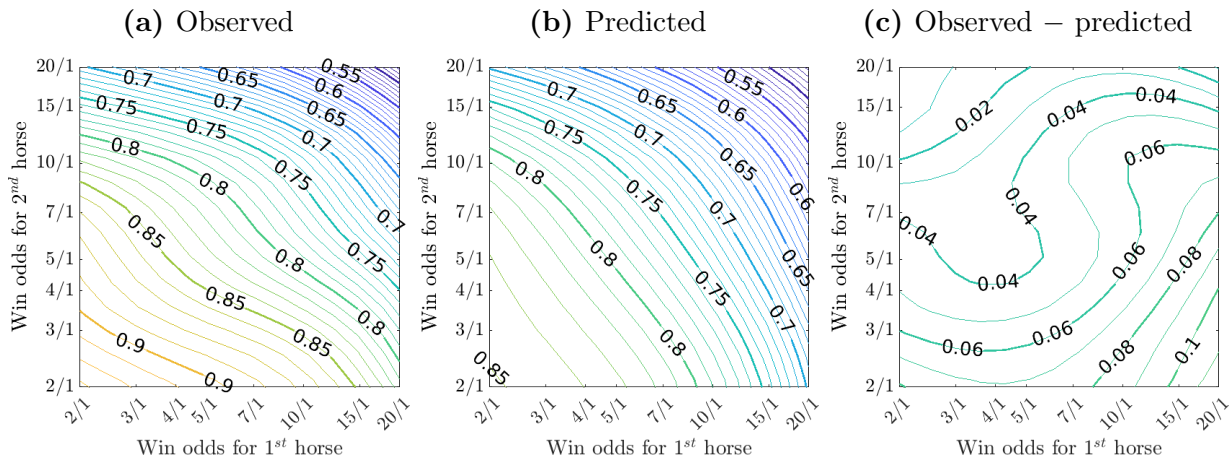
6.2 Rebate structure

We model the track as finding the optimal rebate in each betting pool in each race. In practice, it is unclear whether rebates are set for each race. Here, we consider a less flexible rebate structure, in which the track sets the profit-maximizing rebate in each betting pool. Specifically, we solve the model in Section 5 for a fixed rebate, which allows us to calculate the track's income in each race for a given rebate. We then use an optimizer to find the income-maximizing rebate in each pool. Appendix B shows versions of Figures 9a and 11 with the fixed rebate, which are appreciably identical to those presented above.

6.3 Limitations

Our model largely hits the mark, but its misses are generally of the same pattern: overestimating returns for favorites and underestimating returns for longshots. We suspect that this stems from the coarseness of the beliefs assigned to arbitrageurs. In our model, the beliefs of arbitrageurs derive from a single attribute: the morning-line odds. In practice, arbitrageurs use other information, such as past performance, to form higher resolution beliefs. This disconnect is most consequential for favorites with morning-line odds at the censoring threshold. From the morning lines, it is unclear whether such a favorite should be favored by a small margin or by an overwhelming margin. Presumably, sophisticated bettors know the difference. If the arbitrageurs 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.[35]

This limitation is evident in the predicted odds. Figure 16 compares the distributions of morning-line odds (bars), observed parimutuel odds (solid line), and predicted parimutuel odds (dashed line). The predicted odds are more dispersed than the morning-line odds, owing to the arbitrageurs in our model. However, the predicted odds are less dispersed than the observed odds—presumably due to the low resolution beliefs we assign to arbitrageurs.

One corrective would be to use a larger set of variables, along with a more complicated estimation routine, to assign arbitrageurs 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, as we see it is

---

[35]This issue is exacerbated by the non-parametric procedure we use to estimate the beliefs of arbitrageurs, which we describe in Figure 6. Extreme favorites are relatively rare; hence, their estimated win rates are biased towards those of less extreme favorites, which win less often.

Figure 16: Distributions of morning-line odds (bars), final odds (solid line), and predicted odds (dashed line) for win bets.
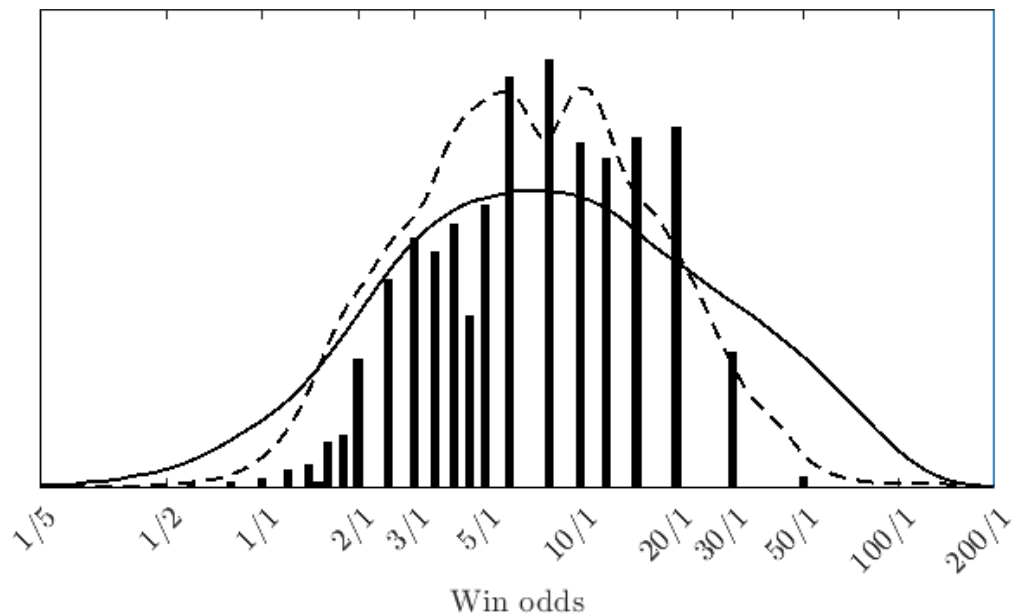
not that it makes perfect predictions—but that it gets close without sacrificing parsimony.

7 Track revenues

Our model also predicts the track's income from noise traders and arbitrageurs. The line in each subplot of Figure 17 shows the ratio of income from arbitrageurs to income from noise traders across betting pools at a given track. The error bars span the state-regulated takeout paid by noise traders (at the top) and the effective takeout paid by arbitrageurs (at the bottom) as estimated when the rebate is fixed by pool. Tracks that promulgate well-calibrated morning-line odds (e.g., PRX) generate minimal income from arbitrage. Other tracks (e.g., FG) levy higher tax rates on arbitrageurs and generate more income from arbitrage. Arbitrage volumes tend to be higher in exotic pools. This occurs because exotic wagers are tied to compound outcomes, which amplify differences in beliefs between noise traders and arbitrageurs, as discussed at the end of Section 5.

Figure 17: Ratio of track income from arbitrageurs to income from noise traders (line), the takeout (top of error bar), and the effective takeout paid by arbitrageurs when the rebate is fixed by pool (bottom of error bar), separately by track and pool.
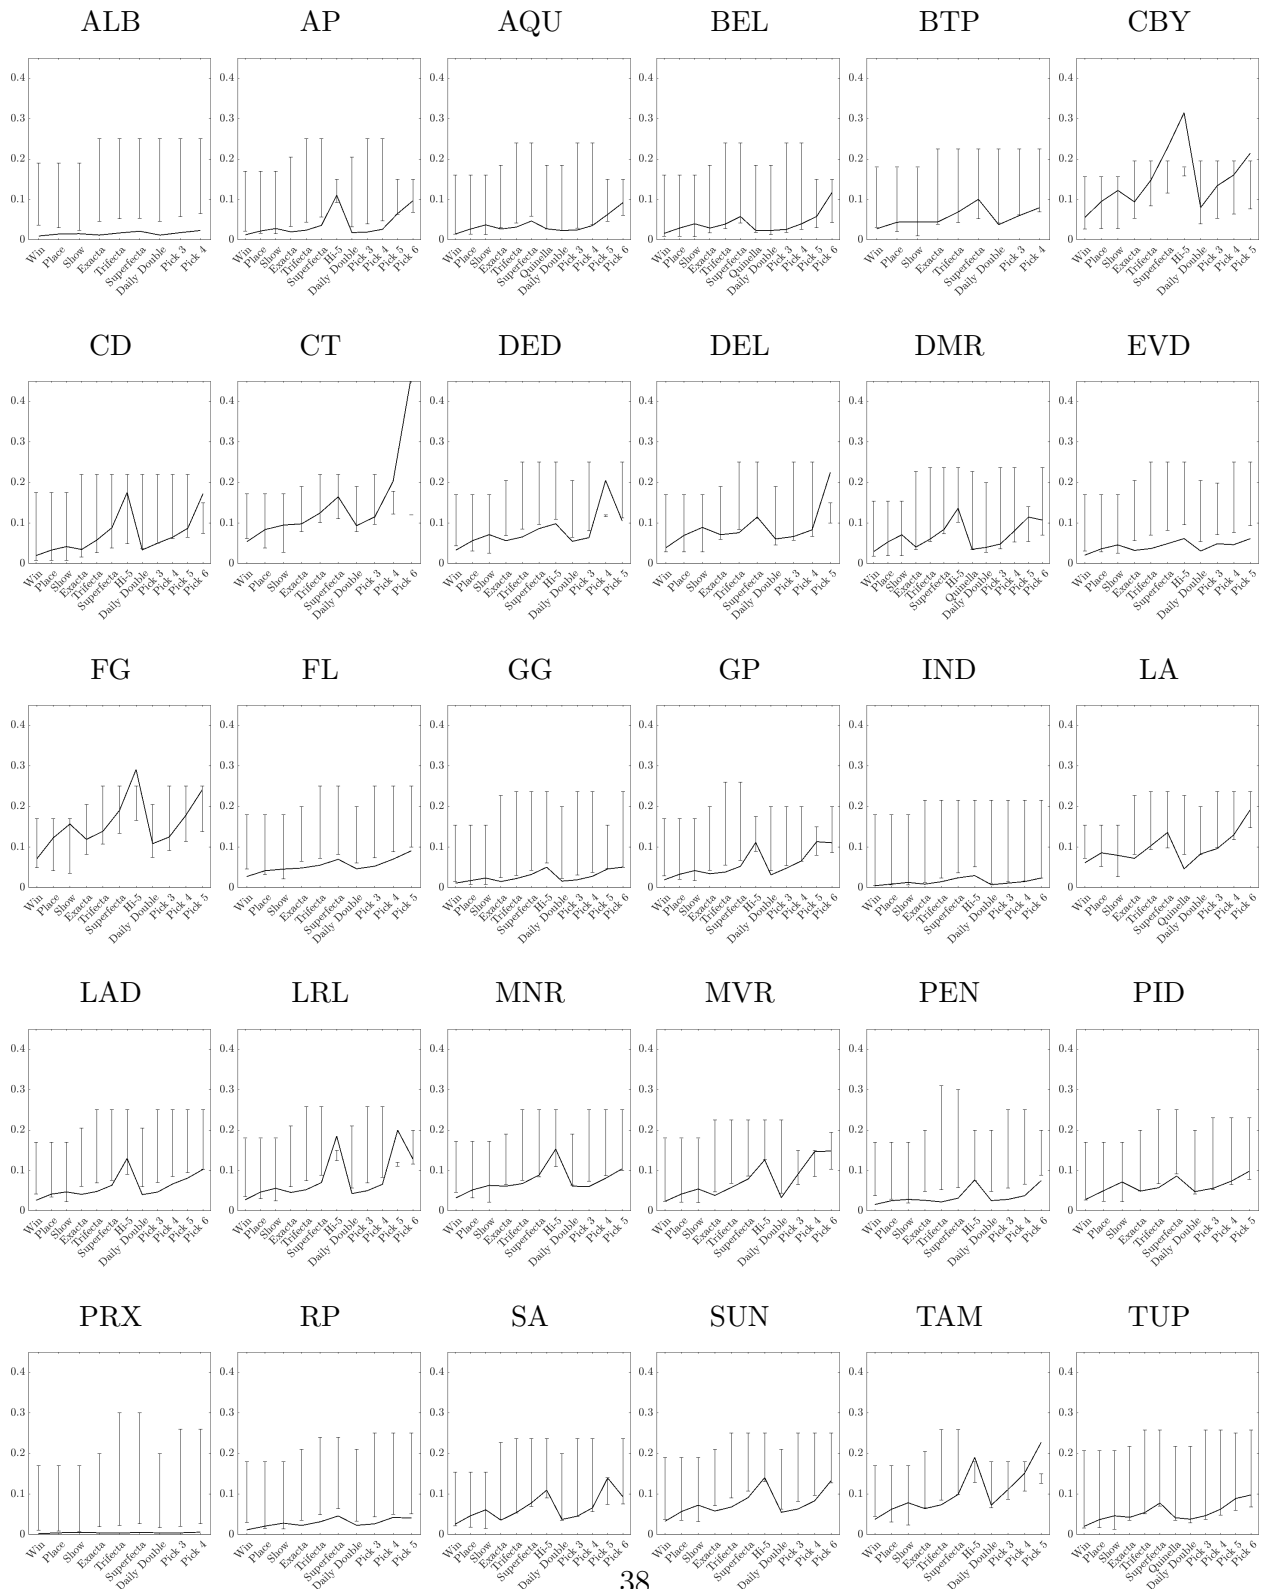
Figure 18: Ratio of track income from arbitrageurs to income from noise traders, by track.
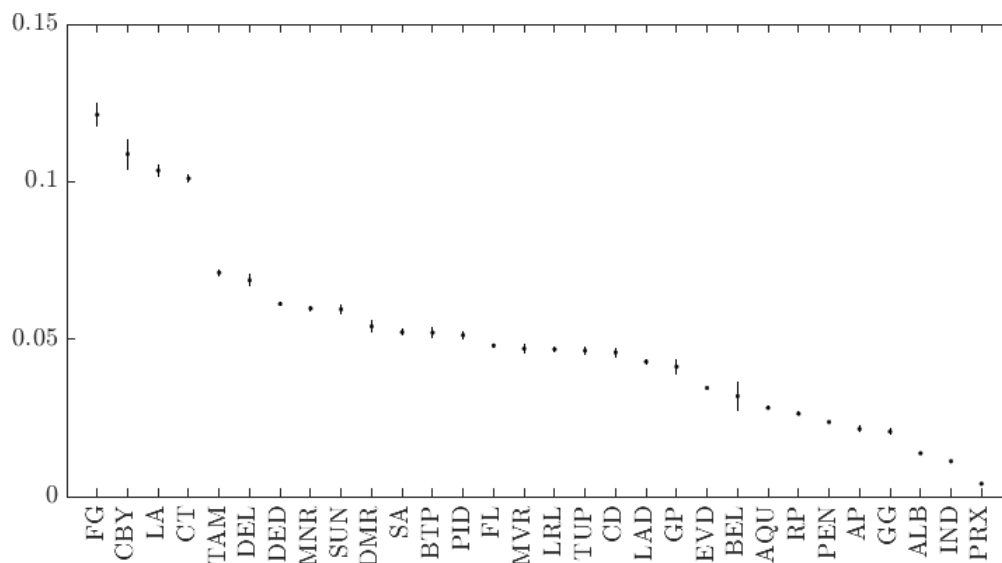
Figure 18 shows each track's incremental income from arbitrage.[36] We estimate that deceptive morning lines increase track income by as much as 12%. Cross-track variation is large, however, with a handful of tracks profiting minimally, if at all, from deception.

Readers should interpret these estimates with caution. For instance, arbitrageurs likely have more refined beliefs than those we assign them, as discussed in Section 6.3. If so, the predictions in Figure 18 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 arbitrageurs is imperfect, or if we model them as risk-averse, then they will wager less, and the track will make less money. (It is also possible that arbitrageurs collectively make negative profits, in which case our assumption of perfect competition leads to underestimation.)

Nevertheless, our estimates of arbitrage volume and rebates are consistent with those from industry sources. At an industry conference in 2012, the National Horsemen's Benevolent and Protective Association estimated that "just less than 20% of total pari-mutuel

---

[36]We weight each betting pool by the total amount wagered. 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 17, and the estimates in Figure 18 are meaningfully unchanged when dividing the pot equally among the win, place, and show pools.

Electronic copy available at: https://ssrn.com/abstract=3271248

handle...come[s] from high-volume shops,"[37] or the Advance Deposit Wagering companies that administer rebates to volume bettors. Using our model, we estimate that arbitrageurs account for 17% of wagering volume. At the same conference, one ADW cited an average rebate for its clients of 16.7% during a five-day period at Beulah Park in Ohio. That track has since closed. However, the owners subsequently opened the Mahoning Valley (MVR) racetrack under the same license, where we estimate an average rebate of 15.6%.[38]

The emergence of ADWs, and thus rebates, provides an explanation for why the favorite-longshot bias has moderated over time. In the 1980s, expected returns on a 1/1 favorite and a 50/1 longshot differed by 50 cents on the dollar (Thaler and Ziemba, 1988); they now differ by 35 cents. Rebates increase arbitrage volume, which moderates the bias. They also reduce the optimal amount of deception. Absent rebates, arbitrageurs participate only when the track skews the odds such that some wagers promise excess returns, as was true of extreme favorites in the 1980s (Thaler and Ziemba, 1988). When arbitrageurs receive rebates, arbitrage volume can be high even when all wagers lose money in expectation.

8 Discussion

This paper shows that the favorite-longshot bias in horserace parimutuel betting markets is consistent with a parallax-and-tax scheme and requires no primitive biases other than credulity. We write down a model in which the track deceives noise traders so as to profit by taxing arbitrageurs. In the data, 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. In our model, noise traders take the track's predictions at face value and overbet longshots, thereby making favorites attractive to informed arbitrageurs. Just before the race begins, arbitrageurs bet on favorites, which transforms excess losses sustained by noise traders into additional income for the track.

As with most work in forensic economics (for a review, see Zitzewitz, 2012), our evidence points to malfeasance but not necessarily to intentionality. In particular, we cannot say why tracks promulgate miscalibrated predictions. Some oddsmakers admit that the morning-line odds are intentionally biased, describing an "unwritten rule never to quote odds larger than 30 to 1," or an unwillingness to "point the finger too clearly at a horse's winning chances."[39]

---

[37]https://www.bloodhorse.com/horse-racing/articles/132355.

[38]We calculate these figures when the rebate is fixed within each pool. With a variable rebate, we estimate that arbitrage accounts for 19% of wagering volume and an average rebate at MVR of 15.5%.

[39]Quotes from Snyder (1978, 1117).

Others offer different motives, such as an aversion to extreme errors. "I take a somewhat conservative approach with my morning line," writes the oddsmaker at Arlington Park. "No odds-maker wants to see one of their 6-1 shots sent off as an 8-5 favorite, their 2-1 favorite sent off at 9-2, or a 30-1 morning line outsider at 5-1."[40] Whether or not such motives are true, a parallax-and-tax scheme makes them appear more credible. When a trader manipulates the price of an asset she holds, the motive is obvious. When a broker instead launders profits through taxes on arbitrage, alibis are more persuasive.

Parallax-and-tax schemes appear to characterize trade in other markets. In the lead-up to the financial crisis, investment banks and credit-rating agencies underestimated the risk in mortgage-backed securities. Unsuspecting 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 described in a Senate report on the financial crisis, "Goldman [Sachs] marketed Abacus securities to its clients, knowing the CDO was designed to lose value," and then "allowed a hedge fund, Paulson & Co. Inc., that planned on shorting the CDO to play a major but hidden role in selecting its assets...without disclosing the hedge fund's asset selection role or investment objective to potential investors" (Levin et al., 2011, 10). Investment banks created arbitrage opportunities by underestimating risk, and they profited by taxing the arbitrageurs.

Our results also inform a policy debate about gambling, one that has seen renewed interest following the Supreme Court's 2018 ruling that allows states to regulate sports betting. 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 suggest that gambling is better rationalized by simple enjoyment (Conlisk, 1993).[41] In our model, noise traders happily incur losses equal to the state-regulated tax. Presumably, they would be unhappy to learn that they are being misled, and that they incur additional losses as a result. Under this logic, welfare-enhancing regulations target misinformation by brokers.

---

[40] https://www.arlingtonpark.com/node/10222.

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

41

References

Akerlof, George A and Robert J Shiller (2015) *Phishing for phools: The economics of manipulation and deception*: Princeton University Press.

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

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.

Black, Fischer (1986) "Noise," *The journal of finance*, Vol. 41, pp. 528–543.

Borel, Émile (1938) "Sur le pari mutuel," *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences*, Vol. 207, pp. 197–200.

Budish, Eric, Robin S Lee, and John J Shim (2019) "Will the Market Fix the Market? A Theory of Stock Exchange Competition and Innovation."

Busche, Kelly (1994) "Efficient market results in an Asian setting," in *Efficiency of racetrack betting markets*: World Scientific, pp. 615–616.

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.

Chiappori, Pierre-André, Bernard Salanié, François Salanié, and Amit Gandhi (2019) "From Aggregate Betting Data to Individual Risk Preferences," *Econometrica: Journal of the Econometric Society*, pp. 1–36.

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

Cowgill, Bo, Justin Wolfers, and Eric Zitzewitz (2009) "Using Prediction Markets to Track Information Flows: Evidence from Google.," in *AMMA*, p. 3.

Cowgill, Bo and Eric Zitzewitz (2015) "Corporate prediction markets: Evidence from google, ford, and firm x," *The Review of Economic Studies*, Vol. 82, pp. 1309–1341.

Crawford, Vincent P and Joel Sobel (1982) "Strategic information transmission," *Econometrica: Journal of the Econometric Society*, pp. 1431–1451.

De Long, J Bradford, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann (1991) "The Survival of Noise Traders in Financial Markets," *Journal of Business*, pp. 1–19.

DellaVigna, Stefano (2018) "Structural behavioral economics," Technical report, National Bureau of Economic Research.

Eyster, Erik and Matthew Rabin (2005) "Cursed equilibrium," *Econometrica*, Vol. 73, pp. 1623–1672.

Fama, Eugene F and Kenneth R French (1988) "Dividend yields and expected stock returns," *Journal of Financial Economics*, Vol. 22, pp. 3–25.

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.

Gandar, John, Richard Zuber, Thomas O'brien, and Ben Russo (1988) "Testing rationality in the point spread betting market," *The Journal of Finance*, Vol. 43, pp. 995–1008.

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.

Griffith, RM (1949) "Odds Adjustments by American Horse-Race Bettors.," *American Journal of Psychology*, Vol. 62, p. 290.

Gromb, Denis and Dimitri Vayanos (2010) "Limits of arbitrage," *Annu. Rev. Financ. Econ.*, Vol. 2, pp. 251–275.

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, 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.

Jegadeesh, Narasimhan and Sheridan Titman (1993) "Returns to buying winners and selling losers: Implications for stock market efficiency," *The Journal of Finance*, Vol. 48, pp. 65–91.

———— (2011) "Momentum," *Annu. Rev. Financ. Econ.*, Vol. 3, pp. 493–509.

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

Kyle, Albert S (1985) "Continuous auctions and insider trading," *Econometrica: Journal of the Econometric Society*, pp. 1315–1335.

LeRoy, Stephen F and Richard D Porter (1981) "The present-value relation: Tests based on implied variance bounds," *Econometrica: Journal of the Econometric Society*, pp. 555–574.

Levin, Carl et al. (2011) *Wall Street and the Financial Crisis: Anatomy of a Financial Collapse*: United States Senate.

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.

Manski, Charles (2006) "Interpreting the predictions of prediction markets," *Economics Letters*, Vol. 91, pp. 425–429.

Mehra, Rajnish and Edward C Prescott (1985) "The equity premium: A puzzle," *Journal of monetary Economics*, Vol. 15, pp. 145–161.

———— (2003) "The equity premium in retrospect," *Handbook of the Economics of Finance*, Vol. 1, pp. 889–938.

Ottaviani, Marco and Peter N Sørensen (2005) "Parimutuel versus fixed-odds markets."

———— (2006) "The Timing of Parimutuel Bets."

———— (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.

————— (2015) "Price reaction to information with heterogeneous beliefs and wealth effects: Underreaction, momentum, and reversal," *American Economic Review*, Vol. 105, pp. 1–34.

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

Servan-Schreiber, Emile, Justin Wolfers, David M Pennock, and Brian Galebach (2004) "Prediction markets: Does money matter?" *Electronic markets*, Vol. 14, pp. 243–251.

Shiller, Robert J (1981) "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *The American Economic Review*, Vol. 71, pp. 421–436.

————— (1984) "Stock Prices and Social Dynamics," *Brookings Papers on Economic Activity*, Vol. 15, pp. 457–510.

————— (1992) *Market volatility*: MIT press.

————— (2003) "From efficient markets theory to behavioral finance," *Journal of economic perspectives*, Vol. 17, pp. 83–104.

Shin, Hyun Song (1991) "Optimal betting odds against insider traders," *The Economic Journal*, Vol. 101, pp. 1179–1185.

————— (1992) "Prices of state contingent claims with insider traders, and the favourite-longshot bias," *The Economic Journal*, Vol. 102, pp. 426–435.

Shleifer, Andrei (1986) "Do demand curves for stocks slope down?" *The Journal of Finance*, Vol. 41, pp. 579–590.

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.

Snyder, Wayne W (1978) "Horse racing: Testing the efficient markets model," *The Journal of finance*, Vol. 33, pp. 1109–1118.

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.

Tirole, Jean (1982) "On the possibility of speculation under rational expectations," *Econometrica: Journal of the Econometric Society*, pp. 1163–1181.

Van Erven, Tim and Peter Harremos (2014) "Rényi divergence and Kullback-Leibler divergence," *IEEE Transactions on Information Theory*, Vol. 60, pp. 3797–3820.

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

Wolfers, Justin and Eric Zitzewitz (2004) "Prediction markets," *Journal of economic perspectives*, Vol. 18, pp. 107–126.

——— (2006) "Interpreting prediction market prices as probabilities," Technical report, National Bureau of Economic Research.

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.

Zitzewitz, Eric (2012) "Forensic economics," *Journal of Economic Literature*, Vol. 50, pp. 731–69.

A   Proofs

A.1   Proof of Lemma 1

The first-order condition for the broker's profit, given $\alpha = 1/2$, is:

$$\frac{\partial \pi(q|p) - C_{\frac{1}{2}}(q|p)}{\partial q} = 0 \iff \left(\sqrt{pq} + \sqrt{(1-p)(1-q)}\right)^2 = \frac{c}{1-\sigma},$$

where $0 < c < 1 - \sigma$ is the coefficient of proportionality in the cost function. Substituting the above expression into the profit function yields the maximal profit:

$$\pi(q|p) - C_{\frac{1}{2}}(q|p) = \sigma + (1-\sigma)\left(1 - \frac{c}{1-\sigma}\right) + 2c\log\left(\sqrt{\frac{c}{1-\sigma}}\right),$$

which does not depend on q. $\square$

A.2   Proof of Proposition 1

Consider a prediction $q = p + \epsilon$ for $\epsilon \in (0, 1-p]$. Let $\epsilon^*(\alpha)$ be such that $C_\alpha(p+\epsilon|p) = C_\alpha(p - \epsilon^*(\alpha)|p)$ for some ϵ. We show that as α increases, so too does $p - \epsilon^*(\alpha)$. Hence, an increase in α increases the cost of underestimation more slowly than the cost of overestimation. Since under- and overestimation are equally profitable when $\alpha = \frac{1}{2}$, it must be the case that underestimation is more profitable when $\alpha > \frac{1}{2}$.

We begin by deriving $\epsilon^*(\alpha)$ when $\alpha = 1/2$, or ϵ^* for short, by solving $C_{\frac{1}{2}}(p + \epsilon|p) = C_{\frac{1}{2}}(p - \epsilon^*|p)$:

$$\epsilon^* = \epsilon + 4(2p-1)\left(p(1-p) - \sqrt{p(1-p)(p+\epsilon)(1-p-\epsilon)} - \frac{\epsilon}{2}(2p-1)\right) \quad (9)$$

Given that $\epsilon > 0$, it follows that $\epsilon^* > \epsilon$.

Now consider predictions that generate equivalent costs for a generic α:

$$\frac{c}{\alpha - 1}\log\left(\frac{p^\alpha}{(p+\epsilon)^{a-1}} + \frac{(1-p)^\alpha}{(1-p-\epsilon)^{a-1}}\right)$$
$$= \frac{c}{\alpha - 1}\log\left(\frac{p^\alpha}{\left(p - \epsilon^*(\alpha)\right)^{a-1}} + \frac{(1-p)^\alpha}{\left(1 - p + \epsilon^*(\alpha)\right)^{a-1}}\right),$$

or equivalently:

$$\log\left(\frac{p}{1-p}\right) = \frac{1}{\alpha}\Bigg(\log\left[\left(1-p+\epsilon^*(\alpha)\right)^{1-\alpha} - (1-p-\epsilon)^{1-\alpha}\right]$$

$$-\log\left[(p+\epsilon)^{1-\alpha} - \left(p-\epsilon^*(\alpha)\right)^{1-\alpha}\right]\Bigg) \quad (10)$$

Define the right hand side as $G\left(\alpha, \epsilon^*(\alpha)\right)$. By the Implicit Function Theorem,

$$\frac{\partial \epsilon^*(\alpha)}{\partial \alpha} = -\frac{\partial G\left(\alpha, \epsilon^*(\alpha)\right)}{\partial \alpha} \bigg/ \frac{\partial G\left(\alpha, \epsilon^*(\alpha)\right)}{\partial \epsilon^*(\alpha)}$$

We evaluate this derivative at $\alpha = 1/2$ and ϵ^*. Begin with the denominator:

$$\frac{\partial G\left(\alpha, \epsilon^*(\alpha)\right)}{\partial \epsilon^*(\alpha)} = \frac{1}{\sqrt{1-p+\epsilon^*}\left(\sqrt{1-p+\epsilon^*} - \sqrt{1-p-\epsilon}\right)} - \frac{1}{\sqrt{p-\epsilon^*}\left(\sqrt{p+\epsilon} - \sqrt{p-\epsilon^*}\right)}$$

It follows from $\epsilon^* > 0$ that $\partial G\left(\alpha, \epsilon^*(\alpha)\right)/\partial \epsilon^*(\alpha) < 0$. Hence, $\epsilon^*(\alpha)$ is increasing in α if the numerator is positive. The numerator is:

$$\frac{\partial G\left(\alpha, \epsilon^*(\alpha)\right)}{\partial \alpha} = 2\Bigg[\frac{\sqrt{p+\epsilon}\log(p+\epsilon) - \sqrt{p-\epsilon^*}\log(p-\epsilon^*)}{\sqrt{p+\epsilon} - \sqrt{p-\epsilon^*}}$$

$$- \frac{\sqrt{1-p+\epsilon^*}\log(1-p+\epsilon^*) - \sqrt{1-p-\epsilon}\log(1-p-\epsilon)}{\sqrt{1-p+\epsilon^*} - \sqrt{1-p-\epsilon}}$$

$$+ \log(p+\epsilon) + \log(1-p-\epsilon) - \log(p-\epsilon^*) - \log(1-p+\epsilon^*)\Bigg]$$

Since the natural log is a concave function, and since $\epsilon^* > \epsilon$, it follows that $\log(p+\epsilon) + \log(1-p-\epsilon) \geq \log(p-\epsilon^*) + \log(1-p+\epsilon^*)$. Hence, $\partial G\left(\alpha, \epsilon^*(\alpha)\right)/\partial \alpha > 0$ if:

$$\frac{\sqrt{1-p+\epsilon^*} - \sqrt{1-p-\epsilon}}{\sqrt{p+\epsilon} - \sqrt{p-\epsilon^*}} > \frac{\sqrt{1-p+\epsilon^*}\log(1-p+\epsilon^*) - \sqrt{1-p-\epsilon}\log(1-p-\epsilon)}{\sqrt{p+\epsilon}\log(p+\epsilon) - \sqrt{p-\epsilon^*}\log(p-\epsilon^*)}$$

The left-hand side equals $\sqrt{p/(1-p)}$, from (10). Since $\sqrt{p/(1-p)} \geq 1$, it follows that $\partial G\left(\alpha, \epsilon^*(\alpha)\right)/\partial \alpha > 0$ if:

$$\sqrt{p+\epsilon}\log(p+\epsilon) + \sqrt{1-p-\epsilon}\log(1-p-\epsilon) >$$

$$\sqrt{p-\epsilon^*}\log(p-\epsilon^*) + \sqrt{1-p+\epsilon^*}\log(1-p+\epsilon^*),$$

48

or if $y(\epsilon) > y(-\epsilon^*)$, where $y(x) = \log(p+x)\sqrt{p+x} + \log(1-p-x)\sqrt{1-p-x}$ for $x \in (0, 1-p]$. Observe that y is continuous and has the following properties:

1. $y''(x) > 0$

2. $y'(x) = 0$ when $x = 1/2 - p$

3. $y(x) = y(1 - 2p - x)$

The first two properties imply that $y'(x) > 0$ when $x > 1/2 - p$, and $y'(x) < 0$ when $x < 1/2 - p$. The third property implies that $y(\epsilon) > y(x)$ if $\epsilon > x > 1 - 2p - \epsilon$. Hence, $y(\epsilon) > y(-\epsilon^*)$ if $\epsilon > -\epsilon^* > 1 - 2p - \epsilon$. The first inequality is true because $\epsilon^* > 0$, and the second inequality follows from (9). $\qquad\square$

A.3   Proof of Lemma 2

The second-period equilibrium conditions on the expected returns to arbitrageurs imply an expression for $\sum_{i=1}^{N} x_i$, the total amount wagered by arbitragers. Let $Q_+ \equiv \sum_{i=1}^{N} q_i \mathbb{1}\{x_i > 0\}$ and $P_+ = \sum_{i=1}^{N} p_i \mathbb{1}\{x_i > 0\}$, the subjective probability—held by noise traders and arbitrageurs, respectively—of an outcome that attracts positive bets by arbitrageurs. Then,

$$\sum_{i=1}^{N} x_i = \frac{(1-\sigma)P_+ - (1-r)Q_+}{(1-r) - (1-\sigma)P_+}$$

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

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

Substituting $\sum x$ and r^* into the first equilibrium condition and solving for x_i yields:

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

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

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

49

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:

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

A.4 Proof of Corollary 1

We prove that $P_i \geq Q_i$ for all $i \in \{1, \ldots, N\}$. Since $P_N = Q_N = 1$, the result obtains if $P_i \geq Q_i$ is weakly decreasing in i—i.e., if:

$$\frac{P_i}{Q_i} \geq \frac{P_{i+1}}{Q_{i+1}} = \frac{P_i + p_{i+1}}{Q_i + q_{i+1}},$$

or equivalently, if:

$$\frac{P_i}{Q_i} \geq \frac{p_{i+1}}{q_{i+1}}$$

Since $p_i/q_i \geq p_{i+1}/q_{i+1}$, it is sufficient to show that:

$$\frac{P_i}{Q_i} \geq \frac{p_i}{q_i}$$

We prove this by induction. Observe that $P_1/Q_1 = p_1/q_1$ by definition. Assume that $P_i/Q_i \geq p_i/q_i$. It remains to be shown that:

$$\frac{P_{i+1}}{Q_{i+1}} = \frac{P_i + p_{i+1}}{Q_i + q_{i+1}} \geq \frac{p_{i+1}}{q_{i+1}},$$

or equivalently, that:

$$\frac{P_i}{Q_i} \geq \frac{p_{i+1}}{q_{i+1}},$$

50

which follows from the induction assumption and the ordering of the indices. $\qquad\square$

A.5   Proof of Proposition 2

We prove that if $p_1 > q_1$, then the equilbrium condition in (11) holds. Namely, there exists an index $m \in \{1, \ldots, N-1\}$ such that $(p_m/q_m)^2 > \lambda_m \geq (p_{m+1}/q_{m+1})^2$, where:

$$\lambda_i \equiv \frac{P_i(1 - P_i)}{Q_i(1 - Q_i)}$$

Observe that $(p_i/q_i)^2$ is weakly decreasing in i, which follows from the fact that p_i/q_i is weakly decreasing in i by construction. The following conditions on λ_i complete the proof:

1. λ_i is weakly decreasing in i.

2. $\lambda_i < (p_1/q_1)^2$.

3. $\lambda_i > (p_N/q_N)^2$.

The logic is that $(p_i/q_i)^2$ is collapsing on λ_i, which creates a crossing point. We offer a proof by contradiction. Assume that there is no index $i \in \{1, \ldots, N-1\}$ such that $(p_i/q_i)^2 > \lambda_i \geq (p_{i+1}/q_{i+1})^2$. Given condition 3, it then must be that $\lambda_{N-1} \geq (p_{N-1}/q_{N-1})^2$; otherwise, $m = N - 1$ would satisfy (11). And given condition 1, this inequality must hold generally— i.e., $\lambda_i \geq (p_i/q_i)^2$; otherwise, $m = i$ would satisfy (11). But this violates condition 2.

If condition 1 holds, then condition 2 follows from the assumption that $p_1 > q_1$, and condition 3 follows from $p_N < q_N$, which is implied by $p_1 > q_1$. Specifically, $p_1 > q_1$ implies:

$$\lambda_1 = \frac{p_1(1 - p_1)}{q_1(1 - q_1)} < \left(\frac{p_1}{q_1}\right)^2$$

And $p_N < q_N$ implies:

$$\lambda_{N-1} = \frac{(1 - p_N)p_N}{(1 - q_N)q_N} > \left(\frac{p_N}{q_N}\right)^2$$

We now prove condition 1. λ_i is weakly decreasing in i if $\lambda_i \geq \lambda_{i+1}$, or if:

$$\frac{P_i(1 - P_i)}{Q_i(1 - Q_i)} \geq \frac{P_{i+1}(1 - P_{i+1})}{Q_{i+1}(1 - Q_{i+1})},$$

<div align="center">51</div>

Observe that:

$$\frac{P_{i+1}(1 - P_{i+1})}{Q_{i+1}(1 - Q_{i+1})} = \frac{(P_i + p_{i+1})(1 - P_i - p_{i+1})}{(Q_i + q_{i+1})(1 - Q_i - q_{i+1})}$$
$$= \frac{P_i(1 - P_i) + p_{i+1}(1 - P_i - P_{i+1})}{Q_i(1 - Q_i) + q_{i+1}(1 - Q_i - Q_{i+1})}$$

Hence, λ_i is weakly decreasing in i if:

$$\frac{P_i(1 - P_i)}{Q_i(1 - Q_i)} \geq \frac{p_{i+1}(1 - P_i - P_{i+1})}{q_{i+1}(1 - Q_i - Q_{i+1})},$$

or equivalently, if:

$$P_i\left(1 - \frac{Q_{i+1}}{1 - Q_i}\right)q_{i+1} \geq Q_i\left(1 - \frac{P_{i+1}}{1 - P_i}\right)p_{i+1} \tag{12}$$

If $Q_{i+1} < 1 - Q_i$, then (12) can be rewritten as:

$$\frac{P_i}{Q_i} \geq \frac{p_{i+1}}{q_{i+1}}\left(1 - \frac{P_{i+1}}{1 - P_i}\right)\bigg/\left(1 - \frac{Q_{i+1}}{1 - Q_i}\right)$$

Since $P_i/Q_i \geq p_{i+1}/q_{i+1}$ (see Section A.4), the above inequality holds if:

$$1 - \frac{P_{i+1}}{1 - P_i} \leq 1 - \frac{Q_{i+1}}{1 - Q_i},$$

or if $P_{i+1}/Q_{i+1} \geq (1 - P_i)/(1 - Q_i)$. From Corollary 1, $P_{i+1}/Q_{i+1} \geq 1 \geq (1 - P_i)/(1 - Q_i)$. Thus, the inequality holds.

If $Q_{i+1} > 1 - Q_i$, then (12) can be rewritten as:

$$\frac{1 - P_i}{1 - Q_i} \leq \frac{p_{i+1}}{q_{i+1}}\left(1 - \frac{1 - P_{i+1}}{P_i}\right)\bigg/\left(1 - \frac{1 - Q_{i+1}}{Q_i}\right)$$

Since $(1 - P_i)/(1 - Q_i) \leq p_{i+1}/q_{i+1}$ (see Section A.4), the above inequality holds if:

$$1 - \frac{1 - P_{i+1}}{P_i} \geq 1 - \frac{1 - Q_{i+1}}{Q_i},$$

or if $P_i/Q_i \geq (1 - P_{i+1})/(1 - Q_{i+1})$. From Corollary 1, $P_i/Q_i \geq 1 \geq (1 - P_{i+1})/(1 - Q_{i+1})$. Thus, the inequality holds.

Finally, if $Q_{i+1} = 1 - Q_i$, then the inequality in (12) holds if $P_{i+1} \geq 1 - P_i$. From

Corollary 1, $P_{i+1} \geq Q_{i+1}$ and $1 - Q_i \geq 1 - P_i$. Thus, the inequality holds. $\qquad\square$

B   Estimates under a fixed rebate

Figure B1: Replications of Figures 9a (a), 11a (b), and 11b (c) when the rebate is fixed within pool.

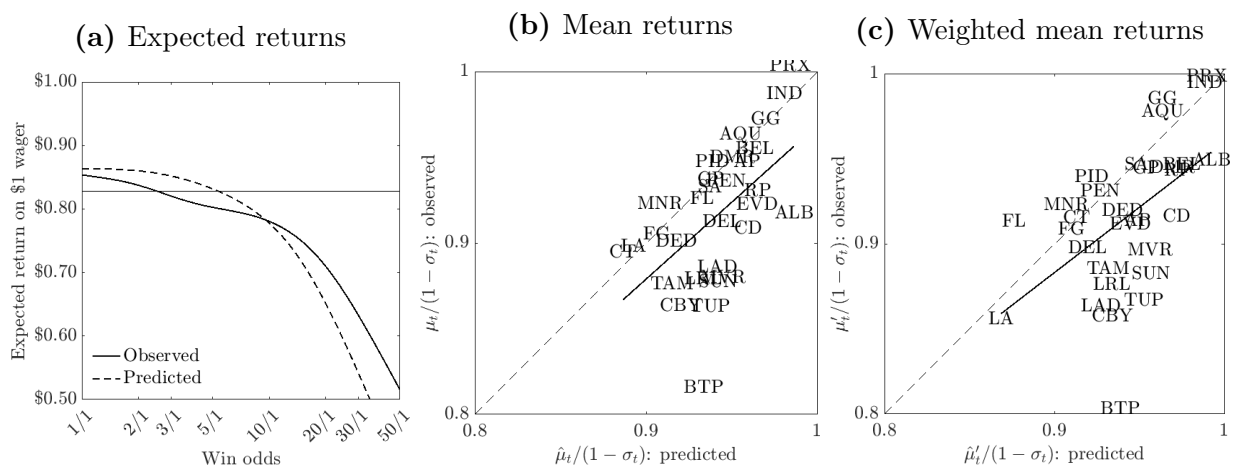**(a)** Expected returns   **(b)** Mean returns   **(c)** Weighted mean returns

Figure B1a shows a smoothed estimate of the relationship between log predicted odds and predicted returns for win bets; the average deviation between predicted and observed returns is 4.3 (se: 0.6) cents. Observed and predicted mean returns (B1b) are correlated at 0.56 ($p < .01$). Observed and predicted weighted mean returns (B1c) are correlated at 0.52 ($p < .01$).