

The Favorite-Longshot Midas*

Etan A. Green¹, Haksoo Lee², and David Rothschild³

¹University of Pennsylvania

²Stanford University

³Microsoft Research

February 15, 2019

Abstract

Deception underlies a classic case of market inefficiency: the favorite-longshot bias in horserace parimutuel markets. Racetracks provide bettors with predictions. At some tracks, these predictions are well calibrated, and markets are efficient. Other tracks overestimate the chances of longshots; at these tracks, longshots are overbet. The deception is strategic. Deception facilitates arbitrage, and arbitrage generates additional income for the track.

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

*Please direct correspondence to: etangr@wharton.upenn.edu. We thank Bosko Blagojević and Will Cai for help with scraping and parsing the data; Joe Appelbaum, Hamsa Bastani, Gerard Cachon, Botond Köszegi, Dorothy Kronick, David Laibson, Simone Marinesi, Cade Massey, Ken Moon, Dave Pennock, Alex Rees-Jones, Joe Simmons, and Ran Spiegler for helpful conversations and comments on previous drafts; and Wharton and Microsoft Research for generous financial support.

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, often 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

The consensus explanation is that many traders are irrational (Shiller, 1984; Black, 1986; 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 present an obstacle to trade (Tirole, 1982) and, hence, an existential threat to brokers. Deception drives a wedge between valuations, increasing trade and brokerage fees.

Pump-and-dump schemes require the perpetrator to take a position in the market, entailing financial risk. They are also illegal. Parallax-and-tax schemes, by contrast, shield the broker from financial risk. If they are illegal, they appear to go unprosecuted. In the lead-up to the financial crisis, investment banks and credit-rating agencies misled unsuspecting investors about the risk in mortgage-backed securities. Deceived investors overpriced these securities, generating demand from savvy investors for ways to short them, which the same

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

banks engineered and sold for large fees (Lewis, 2015). Investment banks created arbitrage opportunities by deceiving noise traders, and they profited by taxing the arbitrageurs.

We begin with an illustrative model of a brokered market with informed arbitrageurs and credulous noise traders. 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. 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 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-term 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, generating a favorite-longshot bias.²

In light of this result, we revisit a classic case of market inefficiency: 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.

We show that the favorite-longshot bias arises from a parallax-and-tax scheme. The track deceives naive bettors by suggesting inflated probabilities for longshots and depressed probabilities for favorites. This deception induces naive bettors to underbet favorites, which creates arbitrage opportunities for sophisticated bettors and, from that arbitrage, incremental tax revenue for the track.

²This model is most closely related to a series of papers showing that heterogeneous beliefs can generate a favorite-longshot bias in parimutuel markets (Borel, 1938; Hurley and McDonough, 1995; Ottaviani and Sørensen, 2009, 2010; Gandhi and Serrano-Padial, 2014; Chiappori, Salanié, Salanié and Gandhi, 2019) and in markets for Arrow-Debreu securities (Ottaviani and Sørensen, 2015). Our model incorporates the broker, who induces heterogeneous beliefs through deception.

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, those who receive the rebate, who wager electronically, who time their bets strategically, and who analyze data tend to be the same bettors.³

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.⁴

We show that 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; horses with 30/1 morning-line odds end up with parimutuel odds longer than 50/1. If the track is to be believed, and if bettors are risk neutral, then a 1/1 favorite should win about half the time. In practice, horses with morning-line odds of 1/1 win nearly two in three races, yet morning lines are rarely shorter than 1/1. Similarly, a 30/1 longshot should win about 3% of the time. In practice, horses with 30/1 morning-line odds win about 1% of the time, yet morning lines are rarely longer than 30/1. The morning-line odds underestimate the chances of favorites and overestimate the chances of longshots. In other words, they embed a favorite-longshot bias.

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

³From [twinspires.com](https://www.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>.

⁴Quoted 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, p. 1117).

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 contextual details and stylized facts 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, and with the rebate, perhaps profitably so. The track sets the income-maximizing rebate, and 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 some tax.

Deceptive morning lines generate excess profit for the track. State-regulated tax rates cap the losses that noise traders should sustain. Deceived traders overbet losers and sustain losses in excess of these caps. Because arbitrage is competitive, these excess losses flow in their entirety into the track's coffers, laundered through taxes on arbitrageurs. We use our model to estimate each track's incremental income from deception. For tracks that embed a favorite-longshot bias in the morning-line odds, we find that as much as 12% of their gambling income comes from taxing arbitrage.

A large empirical literature on the favorite-longshot bias in horserace betting markets uses parimutuel odds to estimate the preferences or beliefs of bettors. In this literature, a model is proposed which allows for a downward sloping relationship between the final 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) conclude that probability weighting is needed to match the data.⁶

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

⁶From page 4: “Models relying on an expected utility (EU) framework do not perform well; their fit is quite poor, even for heterogeneous versions of the model...Our preferred models are relatively parsimonious versions of homogeneous rank-dependent expected utility (RDEU) preferences, and of homogeneous NEU preferences. In both cases, the main role is played by distortions of probabilities.”

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, outperforming representative-agent models with risk-loving preferences or probability weighting. Moreover, our model predicts the observed variation in the favorite-longshot bias 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.

The favorite-longshot bias in horserace parimutuel markets is commonly cited as real-world evidence of prospect theory—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. Our rendering requires only that some bettors be gullible, rather than speculative or innumerate.

Our model also explains three stylized facts that have puzzled researchers for decades. First, prediction markets do not reliably exhibit a favorite-longshot bias.⁷ We observe that brokers do not mediate trade in prediction markets. Second, 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.⁸ Third, the bias at North American racetracks 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. We suspect that this is due to technologies that facilitate arbitrage—in particular, rebates to online bettors. 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

⁷In 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).

⁸Hong Kong race cards can be accessed at: <https://racing.hkjc.com/racing/Info/meeting/RaceCard/english/Local/>. Japan race cards can be accessed at <http://japanracing.jp/en/racing/result/>.

high even when all wagers lose money in expectation.

Past studies of the favorite-longshot bias in horserace parimutuel markets ignore the morning-line odds.⁹ One reason is that historical morning lines are difficult to obtain. Whereas data on race outcomes, such as the final parimutuel odds, are freely available on the web, historical morning-line odds are only available in an expensive database that is marketed to sophisticated bettors. We collect the morning-line odds by scraping them each morning for over a year. A second reason is that the formal irrelevance of the morning lines appears to have persuaded some scholars of their economic irrelevance. For instance, [Ottaviani and Sørensen \(2009, p. 2129\)](#) justify studying racetrack parimutuel markets because of “the absence of bookmakers (who could induce biases).” We show that the morning-line odds, devised by the track oddsmaker, induce a favorite-longshot bias.

Our model of deception by brokers is distinct from models of deception by firms ([Gabaix and Laibson, 2006](#); [Carlin and Manso, 2010](#); [Spiegler, 2011](#); [Heidhues, Kőszegi and Murooka, 2016](#); [Heidhues and Kőszegi, 2017](#); [Kondor and Kőszegi, 2017](#)). In these models, profits on naive customers typically offset losses on sophisticated customers.¹⁰ In our model, the broker profits from both naive and sophisticated investors. A similar distinction applies to fixed-odds betting markets, in which a bookmaker, rather than the market, sets the odds.¹¹ 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.

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 income from deception at each track. Section 8 concludes.

⁹To our knowledge, the lone exception is [Snyder \(1978\)](#), who collects the morning-line odds by hand and puzzles over the fact that they reflect a favorite-longshot bias.

¹⁰For 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.

¹¹In 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 \(2005\)](#).

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 a gullible subset of investors take at face value. We are interested in the broker’s choice of q .¹²

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.¹³

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

¹²The 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 diametrically opposed in practice—here, agents ignore their prior.

¹³It 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:

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

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

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

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.
 - $\mathbb{E}[V(\$1 \text{ in favorite})] = \$1 - r^*$ when $x > 0$
 - $\mathbb{E}[V(\$1 \text{ in longshot})] = \$1 - r^*$ when $y > 0$
2. Arbitrageurs leave no money on the table.
 - $\mathbb{E}[V(\$1 \text{ in favorite})] \leq \$1 - r^*$ when $x = 0$

- $\mathbb{E}[V(\text{\$1 in longshot})] \leq \text{\$1} - r^*$ 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 now characterize equilibrium rebates, arbitrage volumes, parimutuel returns, and broker revenues. Let r^* be the income-maximizing rebate. When $q < p$, the brokerage fee on arbitrage is:

$$\sigma - r^* = (1 - \sigma)(1 - p)(\sqrt{\gamma_1} - 1)$$

When $q > p$,

$$\sigma - r^* = (1 - \sigma)p\left(\frac{1}{\sqrt{\gamma_1}} - 1\right)$$

In both cases, fees on arbitrage are positive (i.e., $r^* < \sigma$) and proportional to the square root of the first-period bias, γ_1 . When the broker tells the truth, arbitrage is free.

Arbitrageurs invest in the security that promises excess returns. When the broker un-

derestimates the favorite, arbitrageurs invest in the favorite:

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

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 \$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 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 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 :

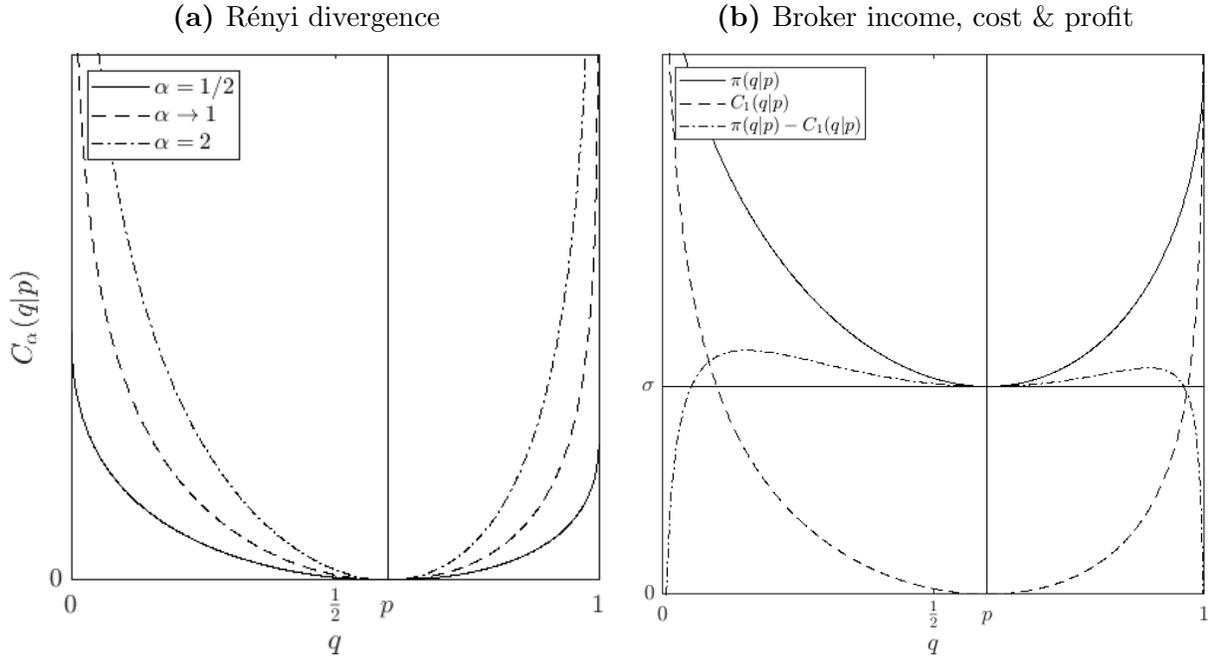
$$C_\alpha(q|p) \propto \frac{1}{\alpha - 1} \log \left(\frac{p^\alpha}{q^{\alpha-1}} + \frac{(1-p)^\alpha}{(1-q)^{\alpha-1}} \right),$$

We use the Rényi divergence because it has a single parameter, $\alpha \geq 0$, which modulates the function’s convexity.¹⁴ Its limiting case as $\alpha \rightarrow 1$ 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

¹⁴Incidentally, the Rényi divergence is also commonly described as the most important divergence in information theory (e.g., Van Erven and Harremos, 2014).

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 (i.e., when $\alpha \rightarrow 1$).



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

deception is profitable. In general, deception is profitable when the income function is more convex at $q = p$ than the cost function.¹⁵

When deception is profitable, the profit function has two local maxima—one when the 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 $q^* \leq p$.*

¹⁵This implies an upper bound on the coefficient of proportionality of $(1 - \sigma)/2\alpha$.

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.

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: $(1 - \sigma)/s_i - 1$.¹⁶ 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.¹⁷

¹⁶In 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.

¹⁷Many exotic wagers can also be constructed as a parlay, or contingent series, of wagers in other pools. For instance, a daily-double wager is logically equivalent to placing the full wager on a win bet in the first race and then placing any winnings on a win bet in the second race. Previous work finds that daily-double returns are less than those of the equivalent parlay, though not by a statistically significant margin—and claims that this constitutes evidence of efficient markets (Ali, 1979). We note that parlays are inefficient and hence, impose limits on arbitrage. Though the takeout is generally higher in exotic pools than in win, place, and show pools, the takeout in any exotic pool is always less than the takeout for the equivalent parlay. For example, paying the daily-double takeout once is less expensive than paying the win takeout twice.

Table 1: Summary statistics by track.

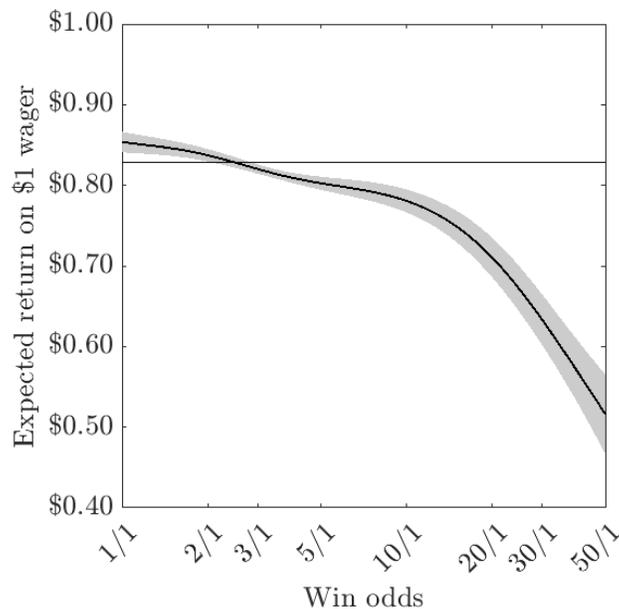
Track		Counts					Takeout rates						Wager %					
Code	Name	State	Days	Races	Starts	WPS	Quinella	Exacta	Trifecta	Superfecta	HR-5	Daily Double	Pick 3	Pick 4	Pick 5	Pick 6	WPS	Exotic
ALB	Albuquerque	NM	97	927	7454	.19	-.25	.25	.25	-.25	.25	.25	-.25	.25	-.25	.25	.41	.59
AP	Arlington	IL	94	758	5551	.17	-.20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

in the exotic pools, yet as the final two columns show, more money is wagered in exotic pools at every track.

All standard errors reported in the 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.²²

4 Stylized facts

Figure 3: Expected returns for win bets, with 95% confidence intervals. The horizontal line marks the average takeout.



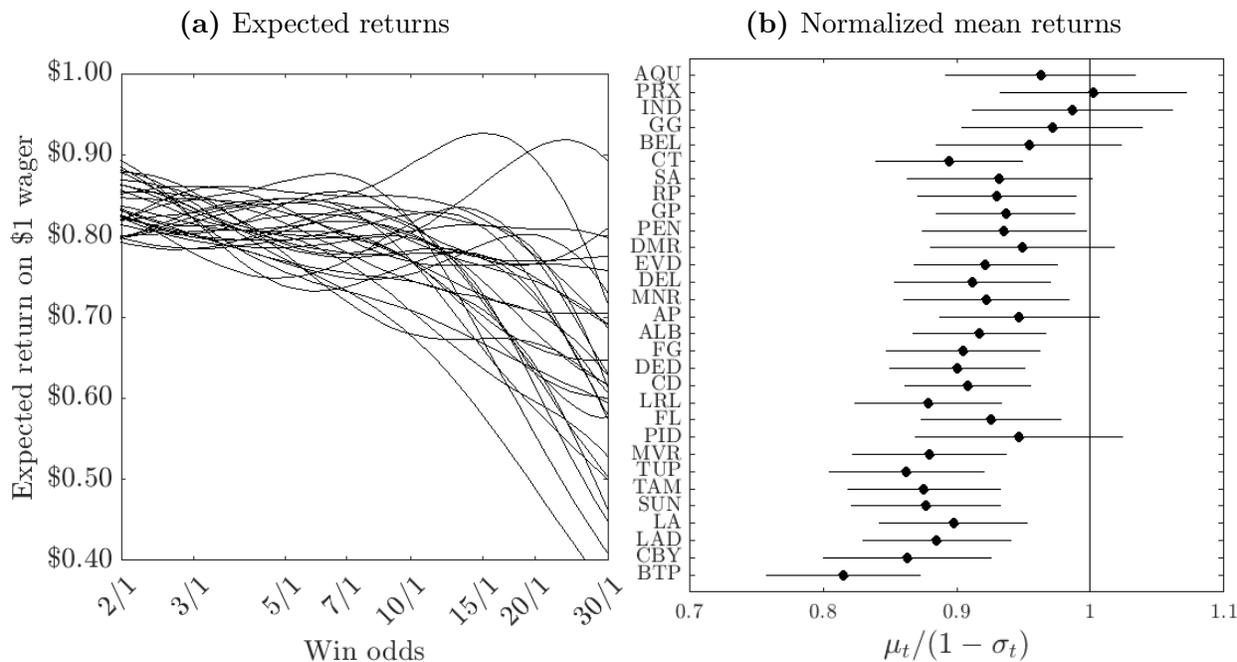
Note: Estimated using a local linear regression with a Gaussian kernel. The bandwidth, of $\log(1.65)$, minimizes the leave-one-out mean-squared error.

Returns are systematically miscalibrated. Figure 3 shows observed returns in the win pool as a smooth function of log final odds. Wagers at all odds lose money on average, given the large takeouts. However, wagers on horses with longer odds lose more money. A \$1 bet on a 1/1 favorite loses 15 cents in expectation, a smaller loss than the average takeout of

²²The standard bootstrapping approach is to resample observations, or in our case, an outcome in a given pool. However, doing so would create bootstrap samples for which the number of winning outcomes in each pool need not be one. As a result, we resample entire pools, rather than individual outcomes. The lone exception is Figure 5, for which the confidence intervals are calculated from homoskedastic standard errors.

17 cents. By contrast, a 20/1 longshot loses 29 cents, and a 50/1 longshot loses 47 cents.²³ Bettors could reduce their losses by betting on favorites instead of longshots.

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



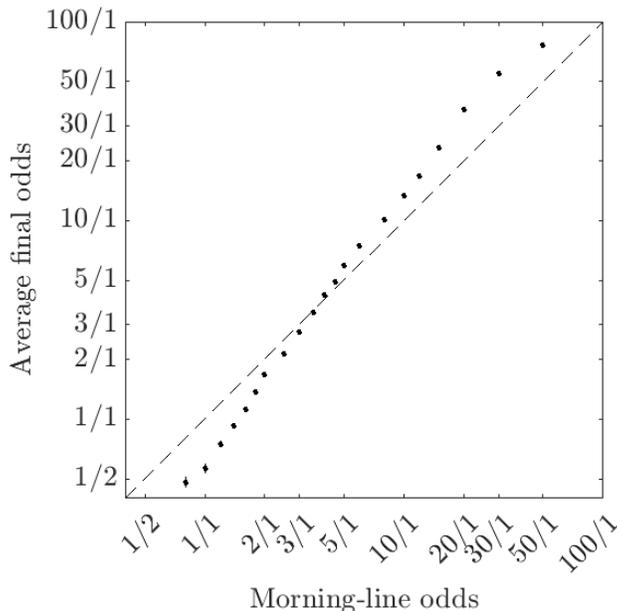
Note: Track-specific expected returns in (a) estimated as in Figure 3. For comparability, we use the same bandwidth as in the full sample, of $\log(1.65)$.

The extent of the favorite-longshot bias varies across tracks. Figure 4a shows expected returns in the win pool as a smooth function of log final odds, separately by track. (Individual plots are shown in Figure 4a in Section 6.) Some tracks exhibit little or no bias, whereas others exhibit severe bias. We summarize the extent of the favorite-longshot bias at track t by its mean returns, μ_t , or the expected return on a \$1 wager for a bettor who picks horses at random. A random betting strategy overbets longshots relative to the market. Were horses priced efficiently, overbetting longshots would not incur a penalty. Instead, a random betting strategy would surrender the takeout on average, and μ_t would equal $1 - \sigma_t$. When

²³These estimates are consistent with those from other analyses of parimutuel odds. For instance, [Snowberg and Wolfers \(2010\)](#) estimate expected returns of about 85% for a 1/1 favorite, 75% for a 20/1 longshot, and 55% for a 50/1 longshot. The favorite-longshot bias tends to be more severe in markets where bookmakers set the odds. For instance, [Jullien and Salanié \(2000\)](#) estimate expected returns of close to 100% for 1/1 favorites and 50% for 20/1 longshots using data from the United Kingdom, where all odds are set by bookmakers.

longshots yield lower returns than favorites, however, overbetting longshots generates excess losses. Hence, a steeper favorite-longshot bias implies lower mean returns, μ_t . Figure 4b shows normalized mean returns, $\mu_t/(1 - \sigma_t)$, by track.²⁴ The tracks are ordered by their estimated returns at win odds of 30/1, as shown in Figure 4a. For some tracks, expected returns are approximately flat and normalized mean returns are close to 1, implying that wagering randomly does not generate excess losses. For the remaining tracks, estimated returns are greater for favorites than for longshots, and random wagering generates excess losses. These differences across tracks cannot be explained by sampling variation alone. A Wald test rejects the null hypothesis that normalized mean returns, $\mu_t/(1 - \sigma_t)$, are equivalent at all tracks ($p < .001$).

Figure 5: Average log final odds at each morning-line odds. For all but the shortest odds, the 95% confidence intervals are obscured by the point estimates.

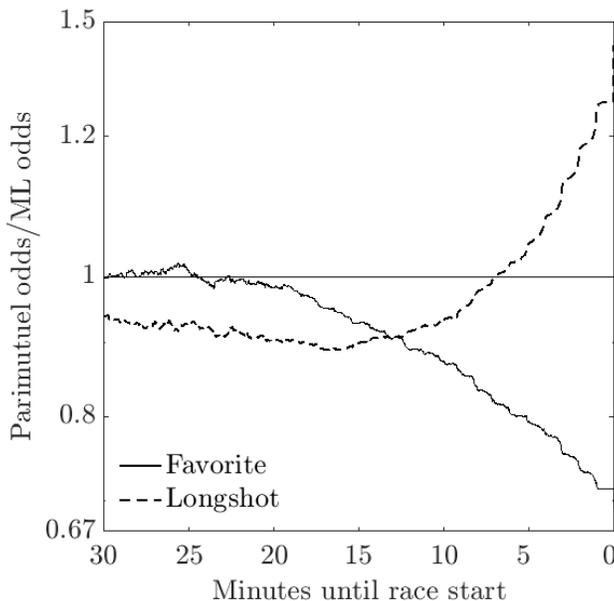


The favorite-longshot bias appears to be related to the morning-line odds. One immediate observation is that the morning-line odds fail at their ostensible goal of predicting the parimutuel odds. Figure 5 shows average final odds at each observed morning-line odds. Morning-line odds are compressed—on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

²⁴In practice, the track's takeout rate is slightly higher than σ_t because of the breakage, the practice of rounding down parimutuel odds, typically to the nearest 10 cents. Figure 4b normalizes mean returns by the true return rate, $1/\sum_{i \in \mathcal{R}} 1/(O_i + 1)$.

morning-line odds of 30/1 finish with final odds in excess of 50/1. If oddsmakers are trying to predict final odds, they could do better by predicting more extreme values.

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



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

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

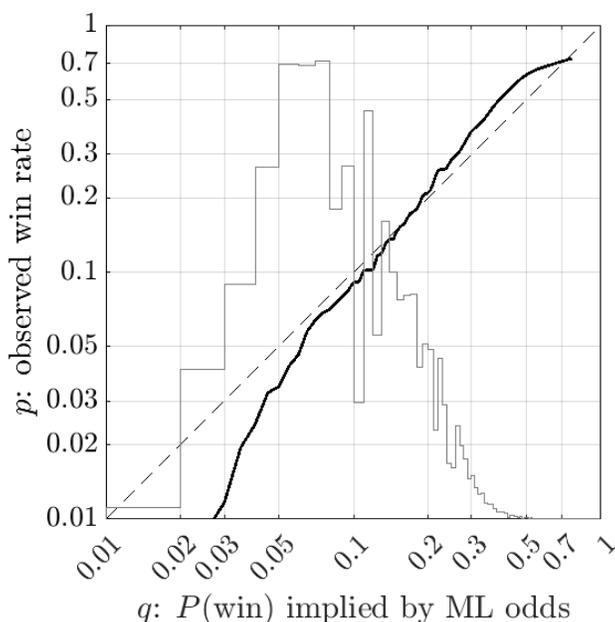
The morning-line odds not only mispredict the final odds—they also imply distorted beliefs about a horse’s chances of winning. Imagine that bettors place wagers such that the parimutuel odds converge to the morning lines, as the track oddsmaker predicts. If these bettors are risk neutral, then the expected value of betting on any horse—given subjective

beliefs q —must be the same for all horses in the race. That is, $q_i(l_i + 1) = q_j(l_j + 1) \forall i, j \in \mathcal{R}$, where \mathcal{R} is the set of horses in a race, and $l + 1$ is the return on a winning \$1 wager. Given that $\sum_{j \in \mathcal{R}} q_j = 1$, the subjective probability of horse i winning is:

$$q_i = \frac{1}{l_i + 1} / \sum_{j \in \mathcal{R}} \frac{1}{l_j + 1} \quad (1)$$

The beliefs implied by the morning-line odds are simply the inverse of the associated returns on a \$1 wager, normalized such that these beliefs sum to 1.²⁵

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



These beliefs are misleading. Figure 7 shows a calibration plot—the observed win rate (denoted by p) as a smooth function of q , the win probabilities implied by the morning-line odds.²⁶ Longshots underperform their implied win probabilities, q . At $q = 0.03$, for example,

²⁵Typically, the inverse implied returns sum to about $1/(1 - \sigma_t)$ —i.e., the morning-line odds are odds that could occur in the parimutuel.

²⁶We are interested in estimating a smooth function, $p(q)$, that measures the rate at which horses with an implied probability of winning, q , actually win the race. To do so, we regress an indicator for whether the observed horse won the race on $\log q$, its log probability of winning, as implied by the morning-line odds and as calculated in (1). In particular, we use the Nadaraya-Watson estimator, otherwise known as local constant regression, which unlike a local linear or local polynomial regression, ensures that the predicted values stay between the bounds of the outcome variable—i.e., between 0 and 1. The challenge is that we would like $p(q)$ to be weakly increasing, reflecting an assumption that the morning lines correctly order the horses by

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 longer than 50/1, just 2, or 0.3%, won the race. Symmetrically, favorites outperform their implied win probabilities. At $q = 0.5$, for example, $p = 0.63$.

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

Across tracks, miscalibration in the morning-line odds predicts the favorite-longshot bias. As before, we measure the extent of the favorite-longshot bias using the normalized mean returns, or $\mu_t/(1 - \sigma_t)$ in Figure 4b. Miscalibration is measured as follows. Let $p_t(q_j)$ be the empirical rate at which a horse with implied probability q_j at track t wins, as shown in Figure 8, and let p_j be the normalized rate such that $\sum_{j \in \mathcal{R}} p_j = 1$. We measure miscalibration in a given race as the Kullback-Leibler divergence between the vector of implied winning probabilities, q_j , and the corresponding vector of normalized winning rates, p_j : $\sum_{j \in \mathcal{R}} p_j \log(p_j/q_j)$. If $q = p$ for each horse in the race, the divergence is zero; otherwise it is positive. We then define the miscalibration at track t , denoted δ_t , as the average Kullback-Leibler divergence among races at that track. Figure 9 shows a strong correlation between δ_t and $\mu_t/(1 - \sigma_t)$. 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 mis-

their chances of winning, even if they imply poorly calibrated beliefs. To do so, we follow [Hall and Huang \(2001\)](#) in assigning each observation a weight between 0 and 1, and, via quadratic programming, minimizing the Kullback-Leibler divergence of the weight vector, subject to the constraints that 1) the function $p(q)$, estimated at every half-percentage-point value of q , is weakly increasing; and 2) the weights sum to 1. This routine does not guarantee convergence. For instance, if a lone horse with a high implied win probability q did not win its race, the monotonicity constraint will force the weight on this observation to 0, sending the Kullback-Leibler divergence to infinity. To avoid this pitfall, we employ a bandwidth $h\lambda(x)$, where h is a scalar bandwidth, and λ is a local bandwidth factor that increases the bandwidth in regions of low density. This multiplicative factor is $\lambda(x) = \sqrt{\exp(\frac{1}{N} \sum_i \log \tilde{f}(\log q_i)) / \tilde{f}(x)}$, where \tilde{f} is a kernel density estimate of $\log q$ using Silverman’s rule-of-thumb bandwidth. The steps are as follows. We find the bandwidth, h , that minimizes the leave-one-out mean-squared error. Fixing the bandwidth, we then find the vector of minimally divergent weights that satisfy the weak monotonicity condition.

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

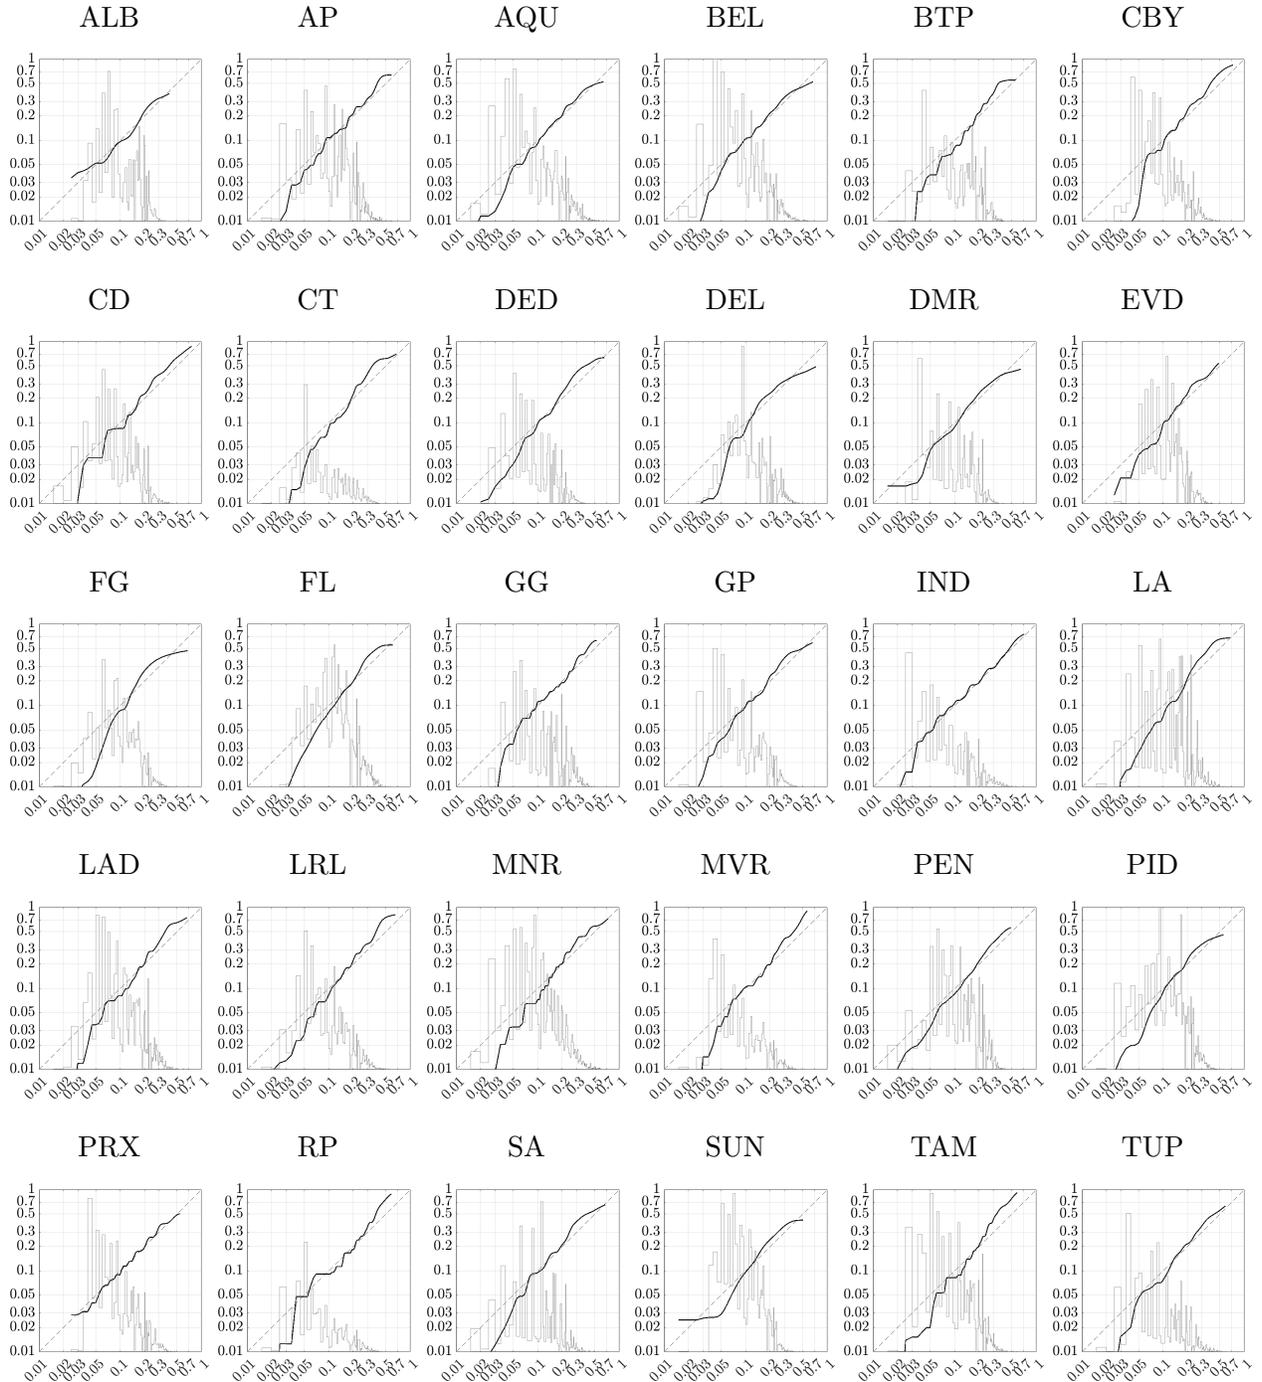
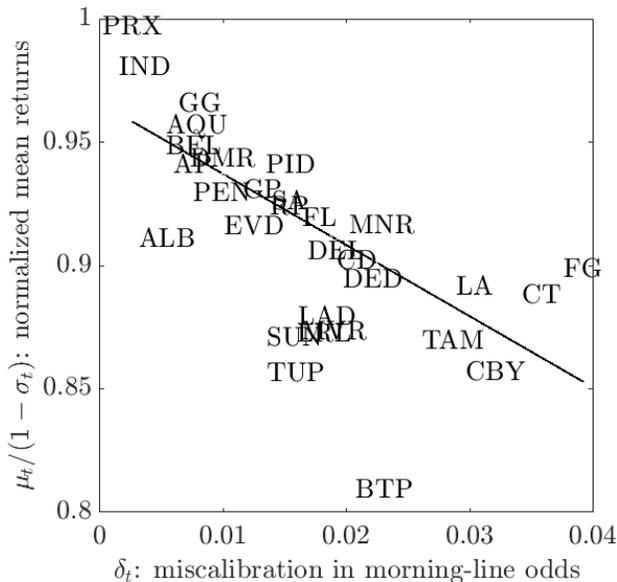


Figure 9: Scatter plot, with a regression fit line, of miscalibration in the morning-line odds, δ_t , and normalized mean returns for win bets, $\mu_t/(1 - \sigma_t)$, across tracks.



calibration. Nonetheless, an anecdote suggests one predictor: falling income. Wagering at Canterbury Park (CBY) decreased from 2015 to 2016 even though the track cut its takeouts below state-regulated levels.²⁷ At the same time, Canterbury Park aggressively overestimated the chances of longshots, as shown in Figure 9.²⁸ 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 the beliefs of noise traders and arbitrageurs, which we infer from the morning-line odds.

²⁷<http://www.thoroughbreddailynews.com/oped-why-canterburys-takeout-reduction-didnt-work/>

²⁸At 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.

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 equilibrium, we first derive the condition under which one exists.

Lemma 2. *Reorder the outcome indices i such that:*

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

An equilibrium exists if and only if there exists an index m such that:

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

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 $m \in \{1, \dots, N-1\}$.*

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.

Equilibrium wagers by arbitrageurs are:

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

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:

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

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

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

And the track's income is:

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

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 profits for the track.

The final parimutuel odds are:

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

Without a rebate, the true expected value of a \$1 wager on outcome i is:

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

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 $p_1 \geq p_2 \geq \dots \geq p_N$, as this implies that the track underestimates the favorite ($p_1 > q_1$) and overestimates the longshot ($p_N < q_N$). Further assume that $q_1 \geq q_2 \geq \dots \geq q_N$ —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 3, with the kink located at the odds of outcome m .

5.3 Estimation

We infer the beliefs of noise traders by assuming that they take the morning-line odds at face value—i.e, as a well-calibrated, or unbiased, prediction of the parimutuel odds in the win pool. Specifically, noise traders form beliefs $q_{i,\text{win}}$ as in (1)—i.e., by inverting the morning lines and then normalizing by race such that $\sum_{i \in \mathcal{R}} q_{i,\text{win}} = 1$.²⁹ In contrast, we imbue arbitrageurs with well-calibrated beliefs, $p_{i,\text{win}}$. Specifically, we non-parametrically estimate the track-specific rates at which horses associated with similar naive beliefs, $q_{i,\text{win}}$, actually win the race, as shown in Figure 8 and described in Footnote 26. We then normalize these estimates by race such that $\sum_{i \in \mathcal{R}} p_{i,\text{win}} = 1$.³⁰ Whereas noise traders assume that the morning lines imply well-calibrated beliefs, arbitrageurs ensure that their beliefs are well calibrated.

We use these beliefs to construct beliefs for exotic outcomes. In doing so, we adopt the Harville assumption: the probability of a horse finishing in n^{th} place is simply its probability of winning the race, divided by the cumulative winning probability among horses that did not finish in the first $n - 1$ positions (Harville, 1973). The race for second place, for instance,

²⁹In the win pool, noise traders need not form these beliefs explicitly. They behave as if maximizing expected value given $q_{i,\text{win}}$ by simply wagering on the horse with the highest ratio of morning-line odds to current parimutuel odds.

³⁰Because this routine is computationally intensive, we do not repeat it during bootstrapping. Instead, we use the same estimates of p_i (and q_i) in every resampled race.

can be thought of as a race within a race in which the first-place finisher does not participate.

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

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

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.³¹ 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), which unlike the takeout is unobserved, and track revenues under the optimal rebate (6), which we report in Section 7.³²

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	.52	1.26	1.0	1.05	0.6	0.6	0.85
Tribal Heat	2/1	4.0	.27	.30	1.11	2.0	0.91	0.3	1.8	0.85
Dandy Candy	6/1	6.9	.12	.09	0.78	6.0	0.64	0.06	8.3	0.85
Click and Roll	8/1	15.6	.09	.06	0.64	8.0	0.53	0.02	13.5	0.85
Eveatetheapple	20/1	69.3	.04	.01	0.34	20.1	0.28	0	40.6	0.55
Movie Starlet	20/1	25.5	.04	.01	0.34	20.1	0.28	0	40.6	0.55
Boston Banshee	30/1	49.7	.03	.00	0.00	30.2	0.00	0	60.4	0.00

³¹In other words, we consider the equilibrium under the income-maximizing rebate. In the data, 37% of pools have more than one equilibrium. Multiple equilibria are most common for pick- n bets, occurring in 50% of such pools. They are least common for win, place, and show bets, occurring in 32% of those pools.

³²The 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:

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

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

The beliefs ascribed to noise traders, q_i , are those that a representative risk-neutral bettor would hold if the final odds coincided with the morning-line odds. Specifically, they are inversely proportional to the returns implied by the morning-line odds, up to a normalizing constant, as in (1). By contrast, the beliefs ascribed to arbitrageurs, p_i , are well calibrated—i.e., proportional to the track-specific rates at which horses with implied beliefs q_i actually win, as shown in Figure 8, up to a normalizing constant. For Charles Town, beliefs of $q_i > 0.2$ are too pessimistic on average, and beliefs of $q_i < 0.2$ are too optimistic on average. As a result, $p_i > q_i$ for 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 odds to the nearest 10 cents. The first-period odds generate a severe favorite-longshot bias. The first-period expected returns on a \$1 wager—denoted $\mathbb{E}_i^{(1)}$ and taken over p_i —are sharply decreasing in the odds. At first-period odds of 1.0, the $p = 0.52$ favorite is better than an even money bet. By contrast, at odds of 20.1, the $p = 0.01$ longshots return just 28 cents on every dollar wagered—and the $p = 0$ longshot returns 0 at any odds.

The second-period equilibrium consists of a set of wagers by 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 equilibrium $m = 4$ is unique, 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 returns on a \$1 wager, $\mathbb{E}_i^{(2)}$, is 85 cents, with the 15-cent expected loss equaling the optimal rebate. For the longshots, expected losses attenuate in the second period. A \$1 wager on either of the horses with 20/1 morning lines, for instance, returns 55 cents in expectation at odds of 40.6, or 21 cents more than at first-period odds of 20.1.

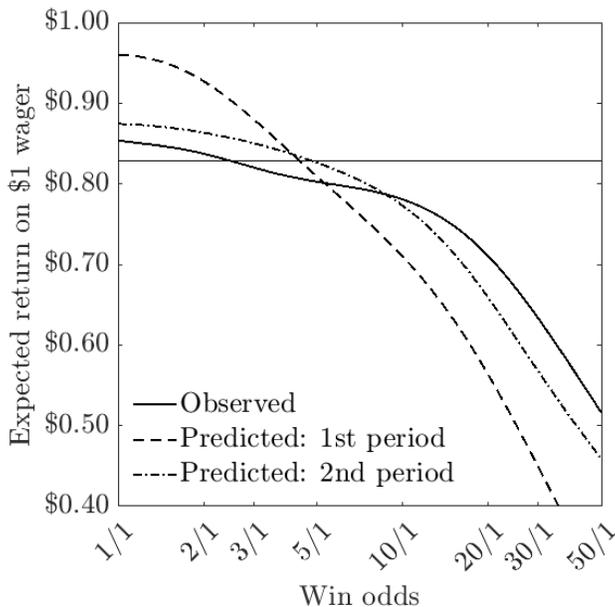
In the win pool, arbitrageurs wager 97 cents for every dollar wagered by the noise traders. However, income from arbitrageurs comprises just 12% of the track’s total income in the win pool, as the track offers arbitrageurs a 15-cent rebate. Arbitrage accounts for larger shares of the track’s income in exotic pools. In the exacta pool, for instance, that share

is 21%; in the superfecta pool, it is 41%. The compound beliefs attached to outcomes in exotic pools amplify differences in beliefs about win probabilities between 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, $P_m > 0.99$ and $Q_m = 0.79$.

6 Model predictions

We begin by comparing predictions from the first- and second-period equilibria. Figure 10 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 correct 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 10: Predicted and observed returns for win bets. The horizontal line marks the average takeout.



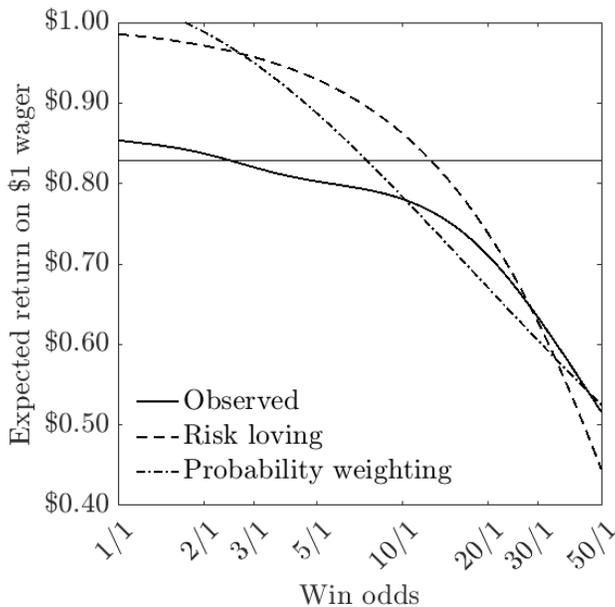
Note: Estimated using a local linear regression with a Gaussian kernel. For comparability, we use the same bandwidth for the predicted estimates as for the observed estimates: the bandwidth used in Figure 3, of $\log(1.65)$.

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 11 cents; after arbitrage, the difference is 2 cents. For a 30/1 longshot, observed returns exceed first-period expected returns by 19 cents; after arbitrage, the difference is 7 cents. Second-period predictions slightly exaggerate the extent of the bias for reasons that we discuss at the end of this section.

We quantify the goodness of our model’s predictions by measuring the expected absolute deviation: the average distance between predicted and observed returns, with the expectation taken over the distribution of observed odds.³³ For a bettor who wagers randomly, the absolute difference between her returns and those predicted by the model converges asymptotically to our expected absolute deviation measure. This measure is just 3.1 (se: 0.5) cents for the model’s second-period predictions, compared to 9.7 (se: 1.6) cents before arbitrage.

Figure 11: Predicted returns for win bets from models with non-standard agents.



Note: Estimated using a local linear regression with a Gaussian kernel. For comparability, we use the same bandwidth for the predicted estimates as for the observed estimates: the bandwidth used in Figure 3, of $\log(1.65)$.

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 11 shows predicted

³³We estimate this distribution using a kernel density estimator with Silverman’s rule-of-thumb bandwidth.

returns from each model.³⁴ For both models, predicted returns for favorites exceed observed returns, and the overall fit is poor. Expected absolute 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 3.1-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. The decision to gamble is commonly rationalized by risk-loving preferences (e.g., [Thaler and Johnson, 1990](#)) or probability weighting (e.g., [Barberis, 2012](#)), but neither explanation can account for a willingness to lose considerable sums, in expectation, when wagering on favorites.³⁵

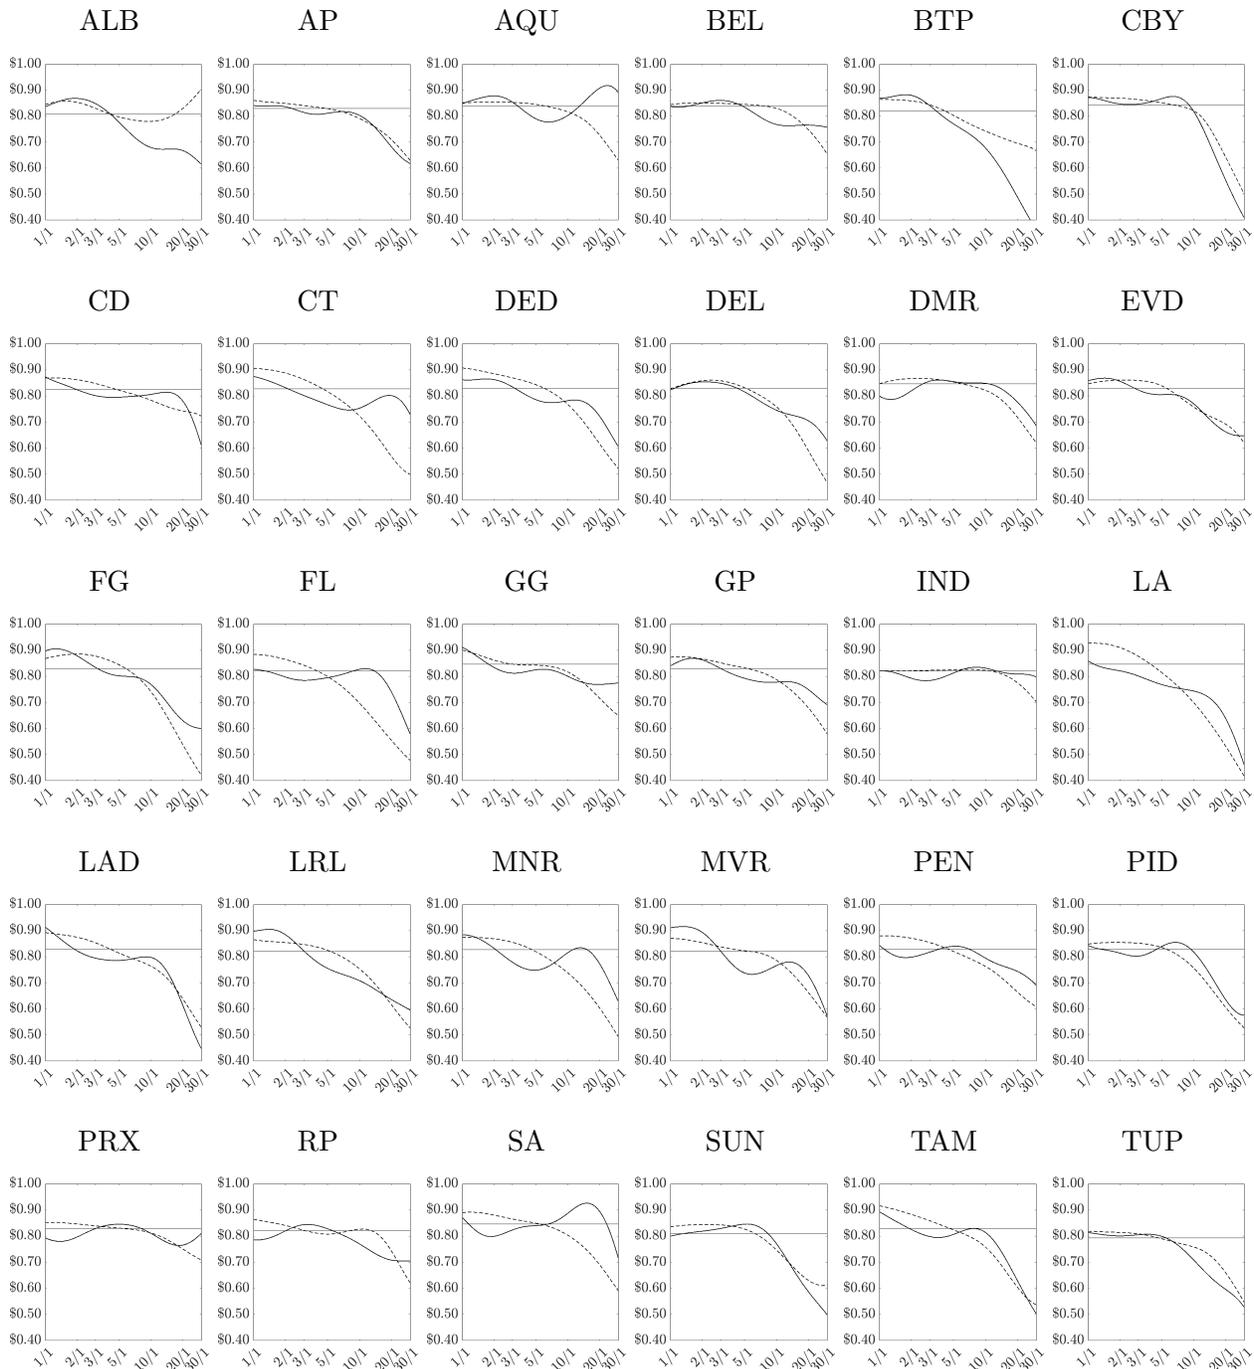
Our model with arbitrage also captures differences across tracks in the extent of the bias. Figure 12 shows smoothed estimates of the observed and predicted relationships between log odds and expected returns, separately by track. For most tracks, the predicted relationship closely follows the observed relationship. Figure 13 summarizes the model’s fit at each track by comparing observed and predicted normalized mean returns—i.e., the expected return on a randomly placed \$1 win bet, normalized by the state-sanctioned return, $1 - \sigma_t$. Observed and predicted values are correlated at $\rho = 0.63$, and predicted values explain 39% of the variation in the observed values across tracks.

We evaluate our second-period predictions in other pools as well. Figure 14 shows ob-

³⁴We estimate these models following [Snowberg and Wolfers \(2010\)](#). In equilibrium, agents wager until the expected utility of a \$1 bet equals \$1, or until $p_i U(O_i + 1) = 1 \forall i \in \mathcal{R}$ in the risk-loving model. For the probability-weighting model, the equilibrium condition is $\pi(p_i)(O_i + 1) = 1 \forall i \in \mathcal{R}$. Each model is governed by one parameter. For the risk-loving model, [Snowberg and Wolfers \(2010\)](#) use the CARA utility function $U(x) = (1 - \exp(-\alpha x))/\alpha$, where α modulates the agent’s risk tolerance. For the probability-weighting model, they use the weighting function $\pi(p) = \exp[-(-\log(p))^\beta]$, 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.

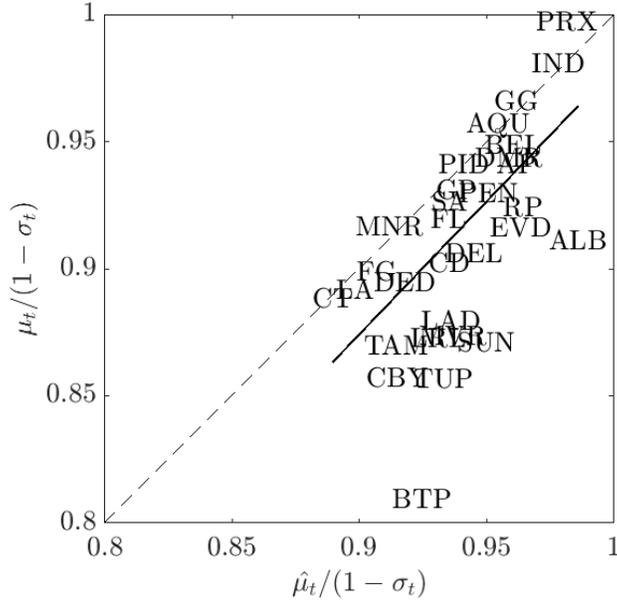
³⁵These explanations were more sensible decades ago, when favorites were close to even-money bets ([Thaler and Ziemba, 1988](#)).

Figure 12: Observed (solid line) and predicted (dashed line) returns for win bets with 95% confidence intervals, by track. The horizontal line marks the state-sanctioned return, or $1 - \sigma_t$.



Note: Estimated using a local linear regression with a Gaussian kernel. For comparability, we use a common bandwidth for all estimates: the bandwidth used in Figure 3, of $\log(1.65)$.

Figure 13: Normalized mean returns, observed and predicted.



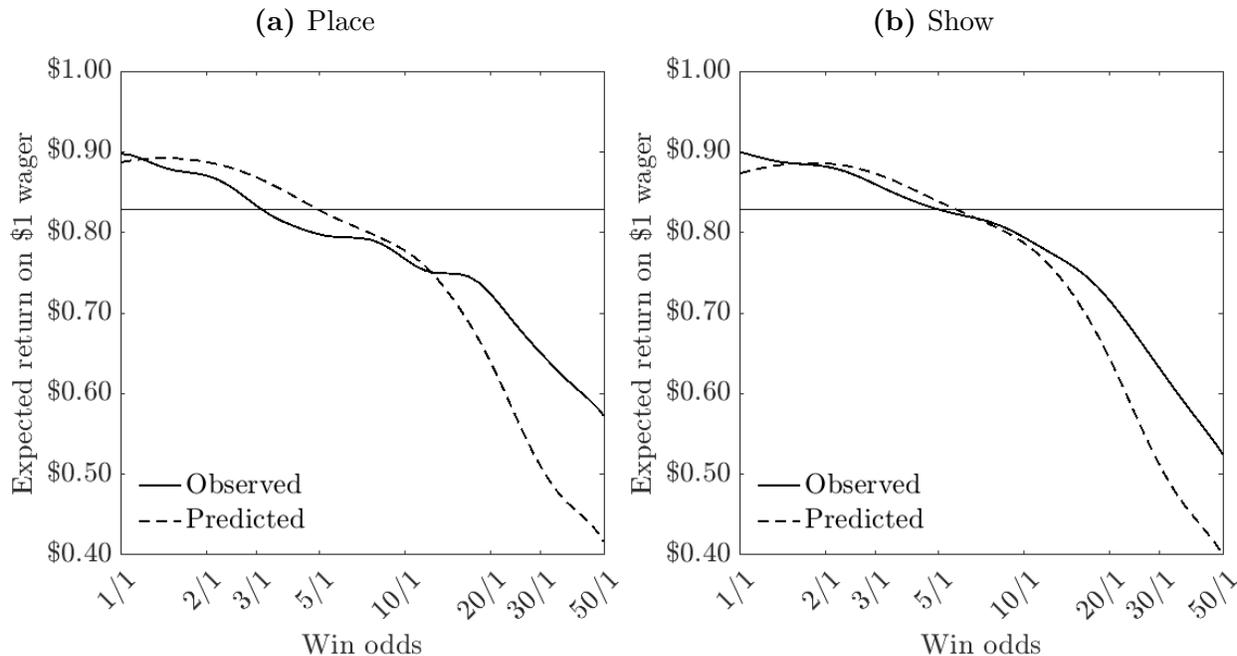
served and predicted returns in place (14a) and show (14b) pools. (Since we only observe final odds in the win pool, in other pools, we estimate the relationship between log odds in the *win* pool and expected returns in the given pool.) A favorite-longshot bias characterizes returns for both place and show bets. In both pools, favorites with 1/1 win odds return 90 cents on the dollar, and longshots with 30/1 win odds return about 65 cents on the dollar. As in the win pool, the model fit is generally close, with a slight overprediction for favorites and an underprediction for longshots. In place pools, the expected absolute deviation is 5.2 (0.3) cents; in show pools, it is 4.1 (0.2) cents.

The final three sets of figures show expected returns on wagers involving two horses—those in exacta (15), quinella (16), and daily-double (17) pools. In each set of pool-specific figures, the first figure (a) shows observed returns, the second (b) shows predicted returns, and the third (c) shows the difference between observed and predicted returns. In all figures, expected returns are expressed in terms of the log win odds for the first and second horses chosen. (Since finishing order is irrelevant for quinella wagers, we show expected returns in terms of log win odds for the relative longshot and relative favorite.)³⁶

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

³⁶Each figure is estimated using a local linear regression with a bivariate Gaussian kernel. In each pool, the bandwidth pair minimizes the leave-one-out mean-squared error in the observed data (a).

Figure 14: Observed and predicted returns for place and show bets, with 95% confidence intervals. The horizontal line marks the average takeout.



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

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

Our model largely hits the mark, but its misses are generally of the same pattern—overpredicting returns for favorites and underpredicting returns for longshots. We suspect that this stems from the coarseness of the beliefs assigned to arbitrageurs. In our model, the beliefs of arbitrageurs derive from a single attribute: the morning-line odds. Hence, every horse at a given track with the same morning lines is allocated the same probability of winning the race, before normalization. In practice, arbitrageurs use other information, such as past performance, to form higher resolution beliefs. This disconnect is most consequential

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

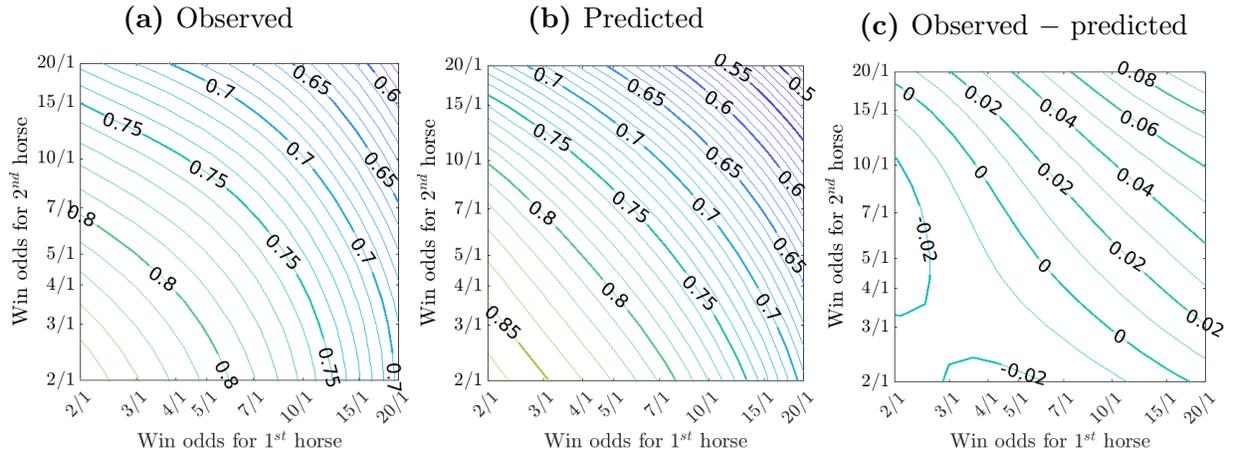


Figure 16: Quinella: expected return on \$1 wager.

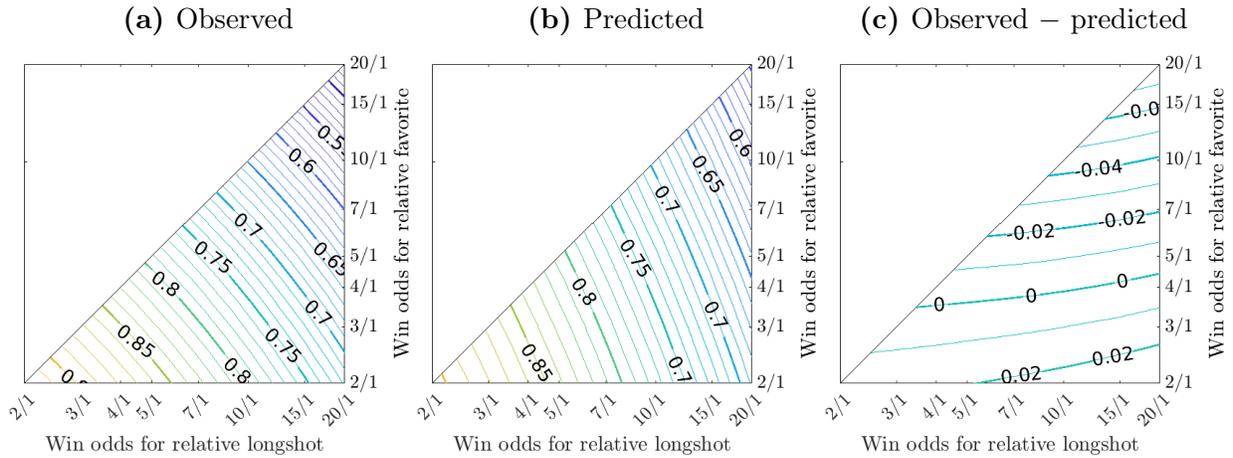
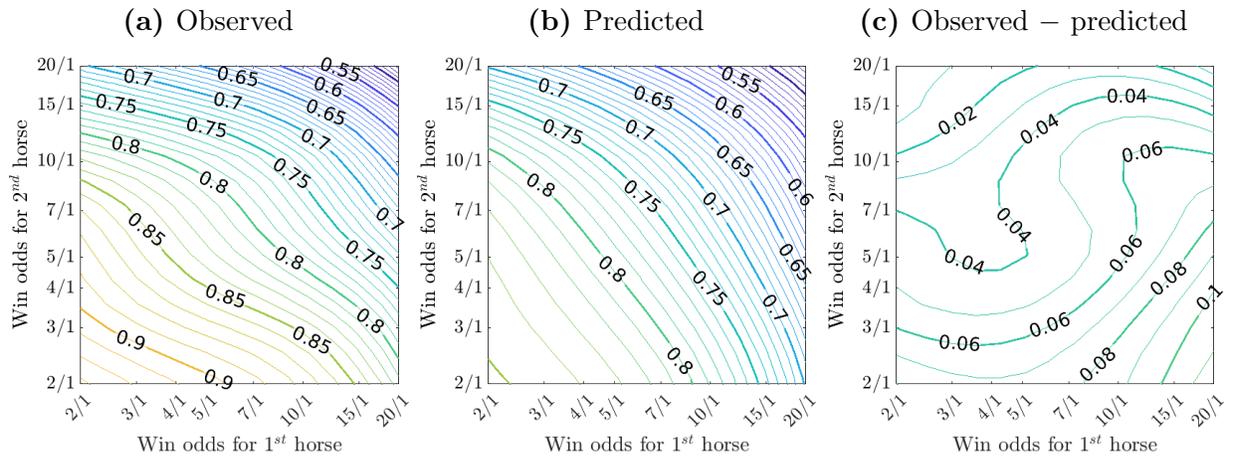
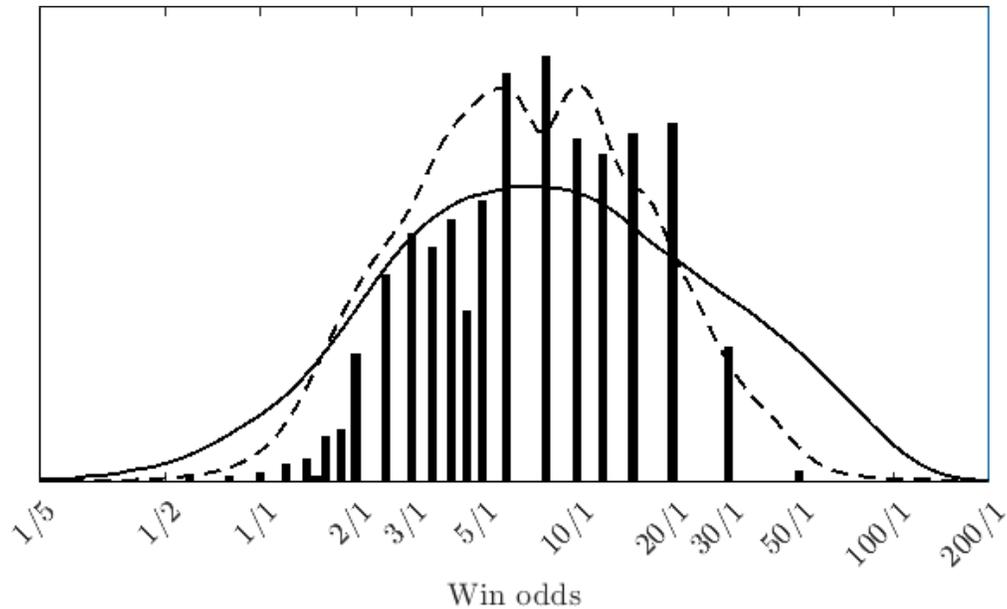


Figure 17: Daily double: expected return on \$1 wager.



for favorites with morning-line odds at the censoring threshold. From the morning lines, it is unclear whether such a favorite is merely more likely to win, or is Secretariat. 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.

Figure 18: Distributions of morning-line odds (bars), final odds (solid line), and predicted odds (dashed line) for win bets.



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

This limitation is evident in the predicted odds. Figure 18 compares the distributions of morning-line odds (bars), observed parimutuel odds (solid line), and predicted parimutuel odds (dashed line). The predicted odds exhibit more variance than the morning-line odds, owing to the arbitrageurs in our model. However, the predicted odds exhibit less variance 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 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. Figure 19 shows the predicted share of income from arbitrage (line) for each wager type at each track, as well as takeout rate (star) and the distribution of optimal rebates across races (box plot). At tracks that promulgate well-calibrated morning-line odds (e.g., PRX), rebates concentrate near the takeout, and the track generates minimal income from arbitrage. At other tracks, effective tax rates on arbitrage are higher, and arbitrage generates more income for the track. Miscalibration in the morning-line odds creates arbitrage opportunities, reducing the rebate required to entice arbitrageurs into the market. In some cases, arbitrage opportunities are so enticing that optimal rebates are negative—i.e., arbitrageurs pay a surcharge. Exotic wagers generate a disproportionate share of the track’s income from arbitrage. This occurs because compound outcomes, such as finishing order or the winners of consecutive races, amplify differences in beliefs between noise traders and arbitrageurs, as illustrated at the end of Section 5.

Figure 20 shows each track’s share of income from arbitrage.³⁷ We estimate that as much as 10% of gambling revenues are generated by deception in the morning-line odds. Cross-track variation is large, however, with a handful of tracks profiting minimally, if at all, from deception.

These estimates may underestimate or overestimate their true values. For instance, arbitrageurs likely have more refined beliefs than those we assign them, as discussed at the end of Section 6. If so, the predictions in Figure 20 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

³⁷We 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/2i184x7>). The model’s predictions are generally similar across the three pools, as shown in Figure 19, and the estimates in Figure 20 are meaningfully unchanged when dividing the pot equally among the win, place, and show pools.

Figure 19: The distribution of optimal rebates across races (box plot) and the share of track income from arbitrage (line), along with the takeout (star), by track and wager type. In the box plots, the box represents the interquartile range, and plus signs denote outliers.

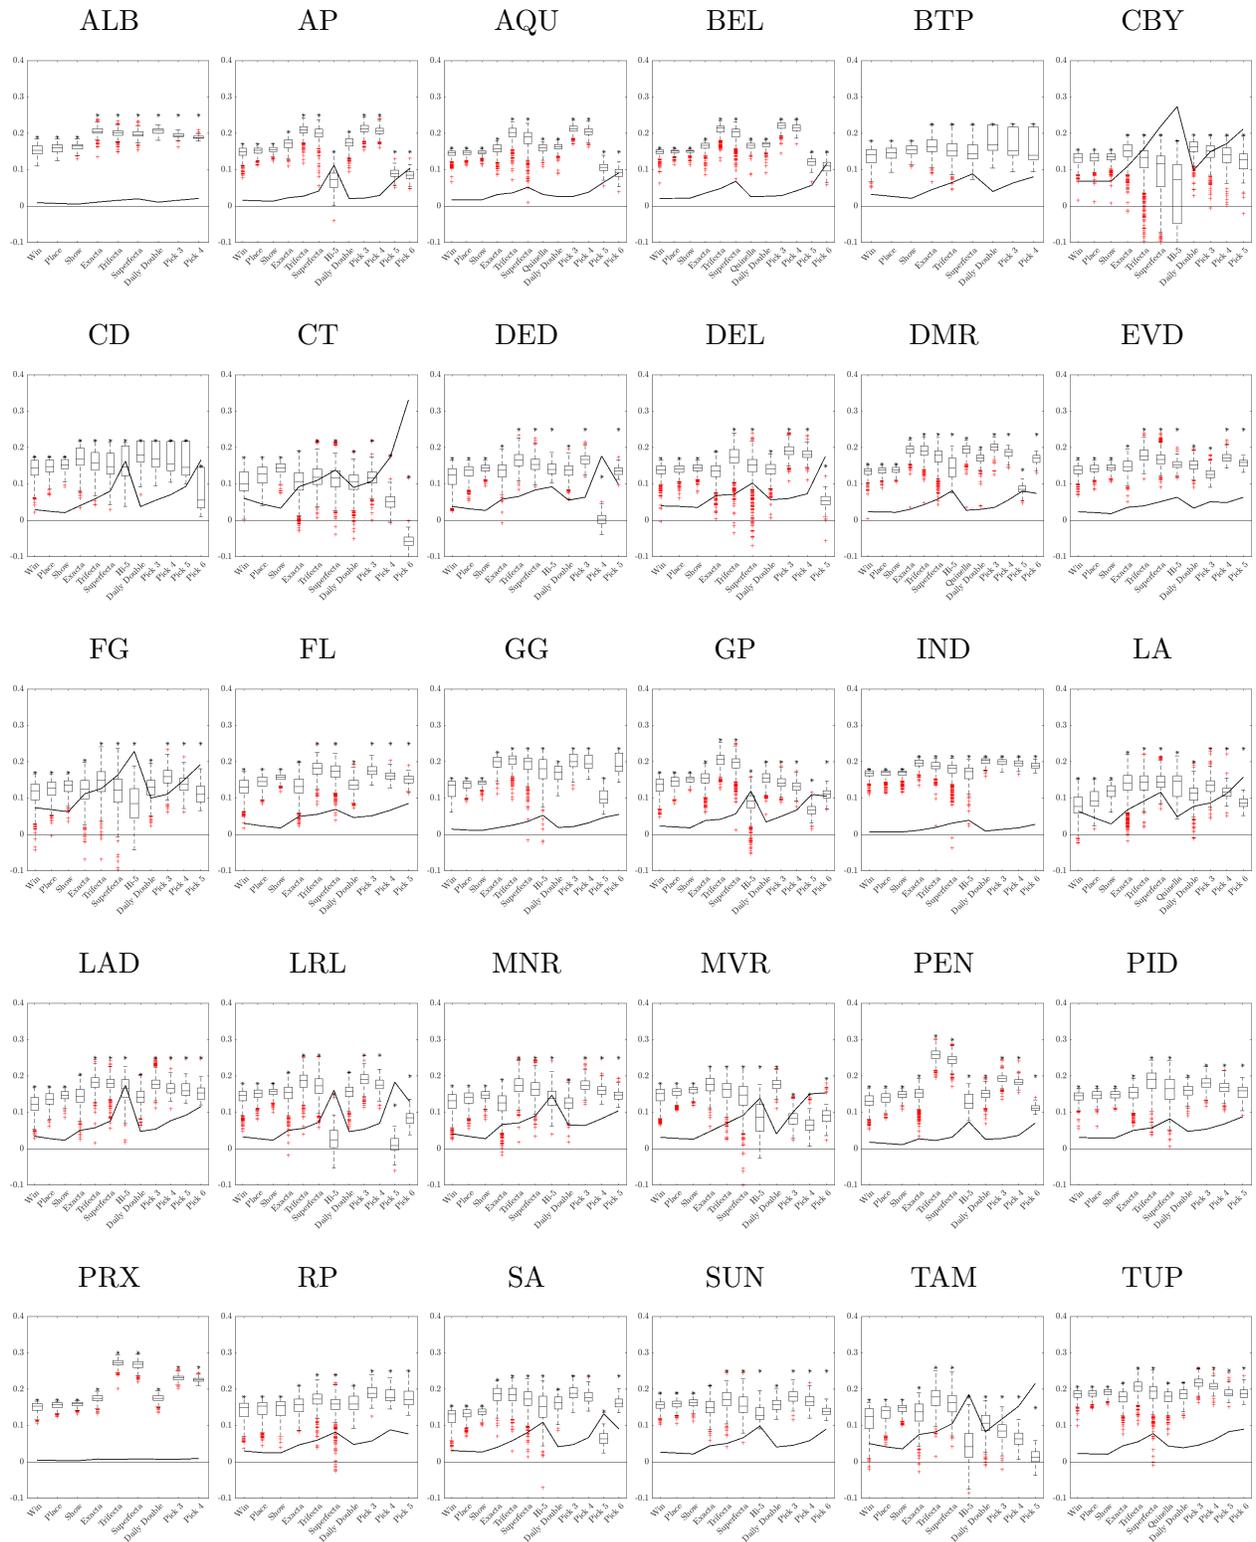
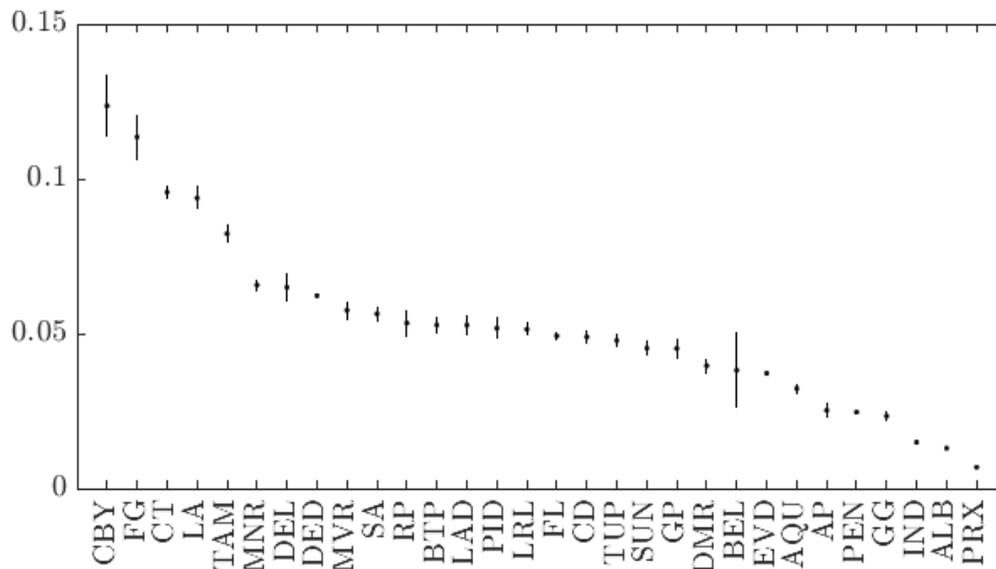


Figure 20: Estimated share of track income from arbitrage, by track.



possible that arbitrageurs collectively make negative profits, in which case our assumption of perfect competition leads to underestimation.) A separate issue concerns rebates. We model the track as finding the optimal rebate in each betting pool in each race. In practice, rebates differ across betting pools but not across races. They also differ across bettors. It is unclear whether these discrepancies bias our estimates upward or downward.

8 Discussion

The favorite-longshot bias in horserace parimutuel betting markets arises from a parallax-and-tax scheme. The track deceives uninformed bettors so as to profit by taxing arbitrageurs. The morning-line odds facilitate this deception. Billed as the track’s prediction of the final parimutuel odds, the morning lines at most tracks instead embed a favorite-longshot bias. Uninformed bettors take the 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 moderates the favorite-longshot bias and transforms excess losses sustained by uninformed bettors into additional income for the track. Consistent with our model, the favorite-longshot bias is more severe at tracks that tell larger lies.

One limitation of our analysis is that we cannot say whether the deception in the morning lines is *intended* to increase profits. Of course, track oddsmakers offer different reasons for

why the odds are miscalibrated, such as an aversion to extreme errors,³⁸ an “unwritten rule never to quote odds larger than 30 to 1,” and an unwillingness to “point the finger too clearly at a horse’s winning chances.”³⁹ A parallax-and-tax scheme makes these justifications appear credible. When a trader manipulates the price of an asset she holds, the conflict of interest is obvious. When a broker instead launders profits through taxes on arbitrage, she can more persuasively plead ignorance.

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).⁴⁰ 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 deceived, and that they incur additional losses as a result. Under this logic, welfare-enhancing regulations target deception by brokers.

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.
- (1979) “Some Evidence of the Efficiency of a Speculative Market,” *Econometrica*, pp. 387–392.
- Asch, Peter, Burton G Malkiel, and Richard E Quandt (1982) “Racetrack betting and informed behavior,” *Journal of Financial Economics*, Vol. 10, pp. 187–194.

³⁸In *Getting a Line on the Early Odds*, Joe Kristufek writes, “I take a somewhat conservative approach with my morning line. 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.” <https://www.arlingtonpark.com/node/10222>

³⁹Quotes from Snyder (1978, 1117).

⁴⁰There are other reasons that people gamble, such as addiction, that call for different policy responses.

- 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.
- 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.
- Carlin, Bruce Ian and Gustavo Manso (2010) “Obfuscation, learning, and the evolution of investor sophistication,” *The Review of Financial Studies*, Vol. 24, pp. 754–785.
- 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.
- 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 Maury 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 and Botond Köszegi (2017) “Naivete-based discrimination,” *The Quarterly Journal of Economics*, Vol. 132, pp. 1019–1054.
- Heidhues, Paul, Botond Köszegi, 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.
- Kondor, Péter and Botond Köszegi (2017) “Financial Choice and Financial Information,” Technical report.
- 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.
- 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.
- 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.
- Spiegler, Ran (2011) *Bounded Rationality and Industrial Organization*: Oxford University Press.
- 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 Harremoës (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.

A Proofs

A.1 Proof of Lemma 1

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

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

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 . □

A.2 Proof of Proposition 1

Consider a prediction $q = p + \epsilon$ for $\epsilon \in (0, 1 - p]$. Let $\epsilon^*(\alpha)$ be such that $C_\alpha(p + \epsilon|p) = C_\alpha(p - \epsilon^*(\alpha)|p)$. We show that as α increases, so too does $p - \epsilon^*(\alpha)$, for a fixed ϵ . 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 :

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

or equivalently:

$$\begin{aligned} \log \left(\frac{p}{1-p} \right) = \frac{1}{\alpha} \left(\log [(1-p + \epsilon^*(\alpha))^{1-\alpha} - (1-p-\epsilon)^{1-\alpha}] \right. \\ \left. - \log [(p + \epsilon)^{1-\alpha} - (p - \epsilon^*(\alpha))^{1-\alpha}] \right) \quad (10) \end{aligned}$$

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

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

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

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

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

$$\begin{aligned} \frac{\partial G(\alpha, \epsilon^*(\alpha))}{\partial \alpha} = 2 \left[\frac{\sqrt{p+\epsilon} \log(p+\epsilon) - \sqrt{p-\epsilon^*} \log(p-\epsilon^*)}{\sqrt{p+\epsilon} - \sqrt{p-\epsilon^*}} \right. \\ - \frac{\sqrt{1-p+\epsilon^*} \log(1-p+\epsilon^*) - \sqrt{1-p-\epsilon} \log(1-p-\epsilon)}{\sqrt{1-p+\epsilon^*} - \sqrt{1-p-\epsilon}} \\ \left. + \log(p+\epsilon) + \log(1-p-\epsilon) - \log(p-\epsilon^*) - \log(1-p+\epsilon^*) \right] \end{aligned}$$

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, $\partial G(\alpha, \epsilon^*(\alpha))/\partial \alpha > 0$ 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 $\partial G(\alpha, \epsilon^*(\alpha))/\partial \alpha > 0$ if:

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

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). \square

A.3 Proof of Lemma 2

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

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

Substituting this expression into the track's profit function and solving the first-order con-

dition yields the profit-maximizing rebate, r^* :

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

Substituting $\sum_i x$ 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:

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

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

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

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 $i \in \{1, \dots, m\}$, and $x_i = 0$ for $i \in \{m + 1, \dots, N\}$. Given that p_i/q_i is decreasing in i by construction, this equilibrium condition can be written as:

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

A.4 Proof of Corollary 1

We prove that $P_i \geq Q_i$ for all $i \in \{1, \dots, N\}$. 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:

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

which follows from the induction assumption and the ordering of the indices.

A.5 Proof of Proposition 2

We prove that if $p_1 > q_1$, then the equilibrium condition in (11) holds. Namely, there exists an index $m \in \{1, \dots, N-1\}$ 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 $i \in \{1, \dots, N-1\}$ 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:

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

Observe that:

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

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:

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

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

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

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, $P_{i+1}/Q_{i+1} \geq 1$ and $(1-P_i)/(1-Q_i) \leq 1$. 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, $P_i/Q_i \geq 1$ and $(1-P_{i+1})/(1-Q_{i+1}) \leq 1$. 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.