Do Sports Bettors Need Consumer Protection? Evidence from a Field Experiment*

Link to latest version

Matthew Brown† Nick Grasley‡ Mariana Guido‡

April 17, 2026

Abstract

We conduct a field experiment with high-volume sports bettors to study how biases impact sports betting demand and evaluate corrective policy interventions. We find substantial overoptimism about financial returns: the average participant predicts that they will break even, but loses 7.5 cents for every dollar wagered. We also find evidence of self-control problems, though these are quantitatively smaller. Our estimates imply that surplus-maximizing sports betting taxes are about twice as large as U.S. sports betting taxes. Though targeted interventions can improve efficiency relative to uniform taxes, experimental evidence shows that two specific interventions favored by the gambling industry are ineffective.

*We thank Matthew Gentzkow, Hunt Allcott, and Doug Bernheim for their mentorship. We thank Adrian Blattner, Harsh Gupta, Caroline Hoxby, Zach Freitas-Groff, Lily Liu, Neale Mahoney, Muriel Niederle, Roger Noll, Sebastian Reidmiller, Pete Robertson, Adam Rosenberg, Daniel Sacks, Dimitry Taubinsky, Jeremiah Weinstock, Jason Weitze; as well as seminar participants at Stanford University, the University of Notre Dame, the Wisconsin School of Business, the Booth School of Business, MIT, and DePaul University for valuable comments and suggestions. We thank Ryan Murphy, Samuel Coolidge, and Quinn Koether at SharpSports for their assistance in creating the betting account syncing portal. We gratefully acknowledge financial support from the George P. Shultz Dissertation Support Fund and the Program in Regulatory Policy at the Stanford Institute for Economic Policy Research, and from the Bill Lane Center for the American West and the McCoy Family Center for Ethics and Society at Stanford. This experiment was approved by Stanford's IRB and pre-registered at the AEA RCT registry (AEARCTR-0013310). The pre-analysis plan is available at https://www.socialscienceregistry.org/trials/13310 and survey instruments are available at https://mattbrownecon.github.io/

†Department of Economics, University of Chicago. Email: mattbrownecon@uchicago.edu

‡Department of Economics, Stanford University

1 Introduction

It was mostly illegal to bet on sports in the United States until the Supreme Court overturned federal restrictions in 2018. In the following years, Americans became familiar with the low-friction, smartphone-based mobile sportsbooks that have revolutionized legal betting markets worldwide. By 2024, just six years after legalization, Americans wagered $149 billion on sports (American Gaming Association, 2025b). The U.S. is not a global outlier in this respect. Per capita online sports betting losses in the U.K., a more mature betting market, are even higher.[1]

These developments have sparked debates about how to regulate sports betting markets. Opponents of regulation argue that people derive entertainment value from betting on sports, and that imposing regulatory barriers would impede the consumption of an enjoyable product. Proponents of regulation argue that betting may be driven by behavioral biases, such as overoptimism and self-control problems. Motivated by these concerns, policymakers and firms have proposed a variety of interventions to protect consumers. Evaluating these proposals requires evidence on the nature and magnitude of the alleged biases, as well as on whether the proposed remedies effectively target them.

In this paper, we present experimental evidence on how behavioral biases affect sports betting consumption, and we evaluate the effectiveness of various policy interventions at addressing these biases. We conduct a field experiment that links observational data from the sportsbook accounts of high-volume sports bettors to a longitudinal survey on their beliefs and preferences. Using techniques from the behavioral economics literature, we estimate overoptimism about financial returns and self-control problems. Our unique data allow us to document heterogeneity in these biases across sports bettors. We also experimentally test the effectiveness of two bias-correction interventions that are popular among firms and regulators. Finally, we estimate a structural model of biased betting demand and simulate the welfare effects of proposed corrective policies.

Our field experiment is designed to measure policy-relevant parameters from a simple model of biased sports betting. The model allows for both a true preference for sports betting and the two biases of interest, overoptimism about financial returns and self-control problems. We show that in our model, both biases can be defined as price-metric

[1] We estimate that annual online betting losses (excluding online casino gambling) were $61 per U.K. adult and $48 per U.S. adult in 2024 dollars. To get these numbers, we begin with aggregate annual losses of £2.63 billion and $12.87 billion respectively, from government and industry reports (American Gaming Association, 2025a, 2024; Gambling Commission, 2025). These annual losses differ from annual dollars wagered because bettors sometimes win back their stake. We convert pounds into dollars using the average exchange rate (1.276, from FRED) and divide by the adult populations of each country (55 million for the U.K. and 266 million for the U.S.) to put losses per capita units.

uninternalized costs. As the literature in behavioral public finance has emphasized, these price-metric biases are key statistics for policy evaluation (Diamond, 1973; Allcott and Taubinsky, 2015). The experiment measures bias in these units.

To study these biases, we link survey data to observational data from betting accounts for a sample of high-volume sports bettors. We focus on high-volume bettors because the distribution of sports betting consumption is highly skewed. High-volume bettors account for most of the activity in this market, so they are the group that would be most impacted by proposed regulations. Compared to a representative sample of weekly sports bettors, our analysis sample was more educated and exhibited less bias according to qualitative survey measures, suggesting that our quantitative bias estimates are conservative for the population of interest. All participants in the study took three surveys over two months. They were also required to sync their sportsbook accounts to our study via an online portal, which allowed us to observe their consumption directly.

To measure overoptimism about the financial returns to betting, we elicited predictions of expected future financial returns. We compare average predicted returns to average realized returns and interpret the difference as a measure of average overoptimism. We find that the average participant predicted that they would roughly break even, but in fact lost 7.5¢ for every dollar wagered. This result implies that the average participant internalized *none* of the expected financial costs of betting. We also estimate overoptimism at the individual level, using Empirical Bayes methods to deal with noisy returns. We find a great deal of heterogeneity across bettors, estimating that 9% of participants are underoptimistic about financial returns, while 10% are overoptimistic by more than 20¢/dollar. Overoptimism is especially pronounced among those who bet on multi-leg wagers, or *parlays* – one of the fastest-growing segments of the sports betting industry.

To measure self-control problems, we elicited valuations of an experimental incentive to reduce future sports betting (the *Bet Less Bonus*) and predictions of future dollars wagered, following Carrera et al. (2022) and Augenblick and Rabin (2019). These data allow us to estimate perceived self-control problems and naivete about their self-control problems respectively. We find evidence of modest average perceived self-control problems: the average participant was willing to pay 12% more for the Bet Less Bonus than they would have been if they had no self-control problems. When translated into price units, though, perceived self-control problems are much smaller than overoptimism. We also do not find evidence of naivete.

We designed experimental treatments to study two behavioral interventions that allegedly reduce biased consumption. Such interventions are prominent in debates about sports betting regulation, and the interventions in our experiment mirrored real-world interventions implemented on U.S. sportsbooks. To address overoptimism about financial returns, some sportsbook apps and websites have increased the visibility of past winnings and losses. In our experiment, participants in the *history transparency treatment* viewed summary information about their past net returns. To address self-control problems, sportsbooks typically include an option for bettors to voluntarily set binding caps on their own future betting activity. In our experiment, participants in the *limits treatment* were prompted to make an active choice about whether they would like to use this tool.

We find these interventions do not eliminate biased consumption. The history transparency treatment did not reduce forward-looking overoptimism, because the average participant was not overoptimistic about their past returns – they were only overoptimistic about the future. This result is surprising given the prominence of selective memory accounts of overoptimism (Bénabou and Tirole, 2002); it is more consistent with models where agents remember the content of past signals, but misinterpret them (Thaler, 2024). In the limits treatment, some participants chose to set binding limits, partially mitigating overconsumption from self-control problems. However, many people who perceived themselves to have self-control problems chose not to restrict their future behavior through limits. Suggestive survey evidence indicates that uncertainty about future normative demand shocks may drive the low takeup of this binding commitment tool (Laibson, 2015).

Finally, we estimate a specialized version of our model and conduct counterfactual simulations to evaluate corrective policies. We use multiple strategies to estimate the price-sensitivity of demand. With our preferred estimates, the implied optimal uniform corrective tax on sports betting is 4.97¢ per dollar. This rate is more than twice as high as sports betting taxes in the U.S., but close to the sports betting tax in Germany. Because of heterogeneity in bias, though, the optimal uniform tax provides only about half the welfare gains of a first-best perfectly targeted policy.

The inefficiency of the uniform tax illustrates the importance of designing well-targeted policy interventions. Targeted interventions disproportionately reduce biased betting while allowing unbiased betting to continue. One approach is to implement bias-correction interventions, though our experimental results on sportsbooks' history transparency and limit interventions highlight that not all such interventions accomplish their stated goals. An alternative approach is to impose restrictions on the kinds of sports bets that are most likely to be driven by bias. Our results on overoptimism suggest that differential regulations of parlay wagers would be well-targeted.

This paper makes several contributions to the literature on the welfare effects of gambling, beginning with our reduced-form evidence on bias. Previous work studying financial misperceptions in gambling has relied on either choices from framed lab experiments (Chegere et al., 2022; Donkor et al., 2025), inferences from aggregate market outcomes (Snowberg and Wolfers, 2010; Levitt, 2004), or survey proxies of bias (Lockwood et al., 2025). Using our linked observational and survey data, we provide the first direct evidence that people overestimate the returns to their own sports bets.[2] We also provide the first evidence that gamblers are willing to pay for an incentive that reduces their future consumption, complementing a literature on compulsive betting from psychiatry.[3] Finally, we document that bias plays a much larger role in parlay betting than traditional sports betting, either because parlays cause overoptimism, because they attract more overoptimistic bettors, or both.

Our second set of contributions comes from connecting the reduced-form evidence on bias to a tractable economic model for quantitative policy analysis. Two influential recent papers Baker et al. (2024) and Hollenbeck et al. (2024) show that sports betting legalization caused adverse financial outcomes for households. While these results quantify some important costs of sports betting, they do not engage with whether consumers internalize those costs in their consumption choices.[4] By using our individual-level data on biases to discipline the model, we quantify the efficiency costs of uniform policies relative to alternatives that target distortions. We use this targeting framework to evaluate sportsbooks' bias-correcting interventions, applying insights from recent lab experiments on the welfare effects of nudges to a field setting (Ambuehl et al., 2022; Allcott et al., 2025; List et al., 2023).

The rest of the paper proceeds as follows. Sections 2-9 present the background, stylized model, experimental design, descriptive facts, reduced-form results, model estimation strategy, quantitative policy analysis, and conclusion, respectively.

[2] Heirene et al. (2022) and Braverman et al. (2014) also use linked observational and survey data, but they only study backward-looking recollections, not forward-looking predictions. We find that forward-looking predictions are biased and backward-looking recollections are not.

[3] See Potenza et al. (2019) for a review of the psychiatric literature on gambling disorders. The gambling literature has also studied the adoption and impacts of voluntary wager limits (Ivanova et al., 2019; Auer and Griffiths, 2022), but to our knowledge no one has quantified the demand for commitment with a willingness to pay measure. The existing literature also focuses exclusively on the takeup of binding limits. One result of our paper is that participants who are willing to pay to reduce future betting are still reluctant to take up binding commitments, consistent with the view that the lack of demand for binding commitment does not necessarily indicate a lack of perceived time-inconsistency (Laibson, 2015; Strack and Taubinsky, 2021).

[4] There is a broader reduced-form literature in economics on the impacts of legal gambling. See for example Guryan and Kearney (2010), Guryan and Kearney (2008), Muggleton et al. (2021), Evans and Topoleski (2002), Kearney (2005), Akee et al. (2015), and Gerstein et al. (1999). Much of this literature has focused on lottery and casino betting. Other papers specifically studying the rollout of legal sports betting in the U.S. include: Taylor et al. (2024); Couture et al. (2024); Matsuzawa and Arnesen (2024); Humphreys and Ruseski (2024).

2 Context on Sports Betting

Modern sports betting differs from other popular kinds of gambling in ways that make concerns about overoptimism and self-control problems particularly salient. Policymakers can address these concerns with a variety of instruments, including taxes and specialized behavioral interventions that attempt to directly correct bias.

2.1 How Modern Sports Betting Works

The vast majority of sports bets are placed online. Online sports betting is legal in all EU countries, the majority of U.S. states, and other countries including most of South America and Australia. In the U.S., revenue from mobile platforms accounted for 94% of legal sports betting revenues (American Gaming Association, 2024). The prominence of mobile platforms is a departure from traditional legal sports betting, which usually occurred at brick-and-mortar sportsbooks or casinos.

Mobile platforms could exacerbate self-control problems in gambling consumption. It is well-known that some gamblers report feeling urges to gamble that are challenging to resist.[5] Combining gambling with smartphones could make it harder to overcome these urges. A public relations chair from Gamblers Anonymous articulates the concern as follows: "They have access to it 24/7 in the palm of their hands. The temptation is always there. You can stay away from casinos and racetracks but you can't stop using your phone" (Vice, 2022).

The expected financial returns to sports betting depend on the odds set by sportsbooks and on a bettor's skill at picking advantageous bets. The *hold rate*, or the share of total dollars wagered that are not paid out as winnings, is one measure of the average financial cost of betting. For example, the average hold rate across U.S. sportsbooks in 2024 was 9% (American Gaming Association, 2025a). In principle, skilled bettors can overcome the hold rate and earn positive returns by identifying wagers with mispriced odds. In practice, it is challenging to consistently identify such opportunities, because sportsbooks use sophisticated methods to set odds. Professional sports bettors often emphasize that amateurs underestimate how challenging it is to beat the house.[6]

The relevance of skill in sports betting creates scope for biased beliefs. In games where

[5] See Feeney (2023) for survey evidence. The psychiatric literature classifies gambling disorders in the same diagnostic category as substance use disorders related to alcohol and other drugs, in part because of evidence the underlying psychological and neurological pathways driving these urges are similar across these contexts (American Psychiatric Association, 2013; Potenza et al., 2019; Goudriaan et al., 2019)

[6] The legendary professional sports bettor Billy Walters writes in his memoir: "For the average bettor, [breaking even is] like trying to swim the English Channel at night, doing the backstroke, without a wetsuit. Surrounded by sharks." (Tamny, 2023)

skill is not relevant, such as lotteries, there is evidence that consumers are sophisticated about negative expected returns (Lockwood et al., 2025). In such games, there are usually few plausible reasons for a consumer to believe that one combination of numbers is a better deal than any other. The psychology of sports betting may be different. Sports fans regularly consume media where commentators provide opinions and arguments about how teams and players will perform. If consumers overestimate the extent to which they can predict sports outcomes, they may overestimate the financial returns to betting.

Overoptimism may be especially relevant for an increasingly prominent kind of bet called *parlays*. Parlays are bets where multiple outcomes must occur jointly for the bet to pay off. Sportsbooks have aggressively promoted parlays, in part because hold rates for parlays are much larger than for regular bets, anecdotally ranging from 20% to 30% (Greenberg, 2022). Given these large financial costs, they are surprisingly popular among consumers: in 2024, parlays accounted for the majority of sportsbook revenues in the states that reported data separately by bet type (Simonetti, 2025). Critics worry that people do not understand how costly parlays are, possibly because it is challenging to aggregate independent probabilities of the component outcomes into the probability that the parlay pays off (Greenberg, 2022).

2.2 Existing Regulations and Taxes

In jurisdictions where it is legal, sports betting is typically a source of government tax revenue. The rate of taxation varies a great deal across jurisdictions. Figure 1 illustrates variation in sports betting taxes across U.S. states. We harmonize all taxes to commensurate units by reporting the ratio of tax revenue to dollars wagered in each state in 2023. At the high end, states like New York taxed sportsbook revenues at 50%, which means they received 4.5% of all dollars wagered (since the average hold rate in New York was 9%). Other states, like Nevada and Arizona, received less than 1% of dollars wagered. On average, sports betting in the U.S. was taxed at a rate of about 2¢ per dollar wagered.

Sports betting taxes are rapidly evolving. In the U.S., Ohio and Illinois doubled their sports betting taxes in 2024. In 2025, France, the U.K., and Brazil all passed regulations that increased taxes on online sports betting. One goal of this paper is to provide empirical evidence and a theoretical framework that can inform debates about the optimal sports betting tax.

Aside from taxes, regulators have considered non-price interventions to ensure that users gamble responsibly. These interventions have taken many forms. We focus here on two prominent interventions that are designed to correct our biases of interest. To correct overoptimism, some have called for sportsbooks to provide users with transparent data

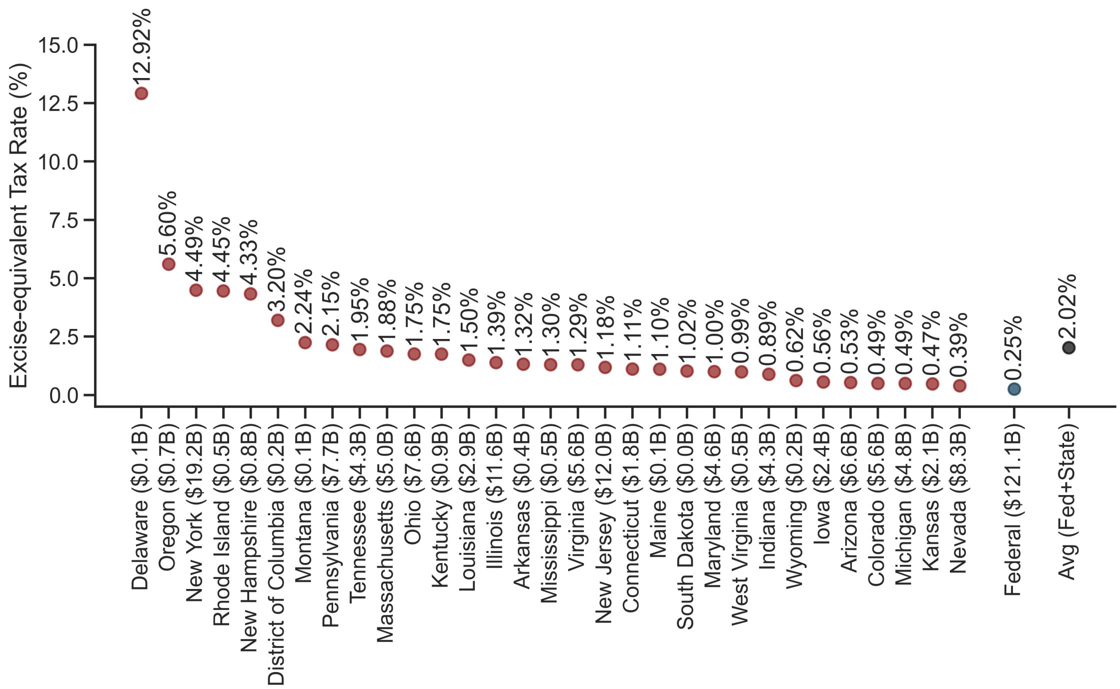

Figure 1: Summary of U.S. sports betting taxes in 2023

Notes: The figure summarizes sports betting activity and taxes in 2023. Labels report dollars wagered or "handle" in each jurisdiction for 2023. Points illustrate the share of dollars wagered that states retain as revenues, or "excise-equivalent" taxes. For state taxes, the excise-equivalent tax is typically the product of a tax levied on sports betting revenues and the hold rate. On the right, we also report the federal excise tax on gambling and the combined average excise-equivalent rate across all states, weighted by dollars wagered. Data on state-level wagers and revenues was compiled by Legal Sports Report https://www.legalsportsreport.com/sports-betting/revenue/ (accessed September 19, 2024).

on their past wins and losses. For example, in Australia, sportsbooks are required to email customers "meaningful statements on their wagering activity" every month (Commonwealth of Australia, 2022). To correct self-control problems, some states require sportsbooks to let users set binding limits on their wagering activity. As of 2022, more than two thirds of U.S. jurisdictions with legal sports betting required operators to provide a mechanism for users to self-limit.[7] Even when regulators do not legally mandate these kinds of bias-correction interventions, many sportsbooks voluntarily implement them as part of a public commitment to "responsible gaming."

The upsides of bias-correction interventions in sports betting mirror the benefits of nudges in other contexts (Thaler and Sunstein, 2008). They are light-touch, cheap to implement, politically feasible, and potentially well-targeted at heterogeneous bias. However, skeptics have argued that in practice, sportsbooks' responsible gaming interventions do not have material impacts and distract from the need for larger-scale interventions

[7] https://www.americangaming.org/resources/responsible-gaming-regulations-and-statutes-guide

(Rose-Berman, 2024; Chater and Loewenstein, 2023). Our study provides experimental evidence on the extent to which history transparency and voluntary limit tools accomplish their stated goals.

3 A Model of Biased Sports Betting

We present a simple model of sports betting where overoptimism about financial returns and self-control problems can cause overconsumption. The model defines the estimands that our experiment is designed to measure. These estimands, which amount to price-metric valuation errors, are key statistics for optimal corrective policy (Diamond, 1973; Allcott and Taubinsky, 2015).

3.1 Setup and Bias Definitions

We study a static sports betting consumption decision: an agent i chooses consumption x_i in a given period. We interpret x_i as the choice of how many dollars to wager. The model abstracts from the dynamics of consumption and from other dimensions of choice (such as choices about which individual bets to place), since these are not the focus of our empirical investigation.

The agent can derive financial and nonfinancial value from sports betting. The financial value of betting arises because if the agent wins, they can use their winnings to consume a numeraire good. The nonfinancial value of betting is a reduced-form object capturing all reasons for betting other than the agent expecting to consume their winnings. Overoptimism about financial returns and self-control problems cause the agent to choose as if they have higher financial and nonfinancial value respectively.

To model the financial value of betting, we define the *realized net return* as $a \in [-1, \infty)$. The realized net return is -1 if the agent loses all of his money, 0 if he breaks even, and greater than 0 if he earns money. The distribution of net returns is governed by a cumulative distribution function F_i. Indexing the returns distribution by i allows expected returns to vary across agents, for example, because some bettors are more skilled than others. After net returns are realized, the agent receives an exogenous endowment y_i plus his total returns $x_i \cdot a$ as numeraire consumption.

We allow for overoptimism by distinguishing between the true return distribution F_i and the perceived return distribution $\tilde{F}_i$. We focus throughout on overoptimism about expected returns, defining a bias parameter $\gamma_i^O = E_{\tilde{F}_i}[a] - E_{F_i}[a]$.

To model the nonfinancial value of betting, we define a reduced-form nonfinancial subutility function $z_i(x_i; \tilde{F}_i)$, which we assume is concave.[8] This function captures rea-

[8]Concavity captures the diminishing marginal nonfinancial value of betting. One interpretation is that

sons for enjoying betting other than the consumption of future winnings. For example, over 80% of participants in our experiment agreed that "betting makes watching sports games more enjoyable." We allow for heterogeneous nonfinancial utility: since z is indexed by i, different agents may convert dollars wagered into utils at different rates. In this sense, our framework allows for general differences in the extent to which betting provides "entertainment value." We also allow perceived financial returns to directly affect the nonfinancial utility of betting, since $\tilde{F}_i$ enters as an argument.

We incorporate self-control problems by adding a temptation parameter to the model, following Banerjee and Mullainathan (2010). In the moment, the agent chooses as if his nonfinancial utility is $z_i(x;\tilde{F}_i) + \gamma_i^{SC} x$ rather than $z_i(x;\tilde{F}_i)$. The temptation parameter γ_i^{SC} captures the extent to which compulsions cause the agent to bet more than would be best for himself, given his understanding of his net returns distribution.

We assume quasilinear utility for ease of exposition and empirical tractability.[9] Normalizing the marginal utility of numeraire consumption to one, we define a *decision utility function*, which the agent maximizes.

$$\begin{aligned} u_i^{decision}(x;\tilde{F}_i) = & \overbrace{y_i + E_{F_i}[a]\cdot x}^{\text{Expected utility from numeraire consumption}} + \overbrace{\gamma_i^{O}\cdot x}^{\text{Effect of overoptimism}} \\ & + \underbrace{z_i(x;\tilde{F}_i)}_{\text{Nonfinancial utility from betting}} + \underbrace{\gamma_i^{SC}\cdot x}_{\text{Effect of self-control problems}} \end{aligned} \tag{1}$$

We next define the *normative utility function*, which governs the agent's welfare. The normative utility function is what the agent would maximize if they were unbiased.

$$u_i^{normative}(x;\tilde{F}_i) = y_i + E_{F_i}[a]\cdot x + z_i(x;\tilde{F}_i) \tag{2}$$

Biased agents consume as if they overvalue gambling by $\gamma_i^{O} + \gamma_i^{SC}$ per dollar wagered. Since these parameters are defined in units of numeraire consumption (which we interpret as money) per consumption unit, we refer to them as *price-metric biases* following Bernheim and Taubinsky (2018).

3.2 Graphical illustration and connection to experimental design

Figure 2 illustrates the price-metric biases in a demand curve diagram. The demand for betting is decreasing in the price of wagering a dollar (i.e, expected losses). There are separate demand curves for in-the-moment choices, which are subject to temptation ("short-term demand") and choices made in advance, which are not ("long-term demand"). An

betting has time costs as well as financial costs, and the opportunity cost of time rises as the agent places more bets.

[9] In a previous version of this paper, we studied a more general version of the model with nonlinear numeraire consumption utility. Analysis with this more general model generates qualitatively similar results as the results in this paper, under plausible parameterizations (Brown et al., 2025).

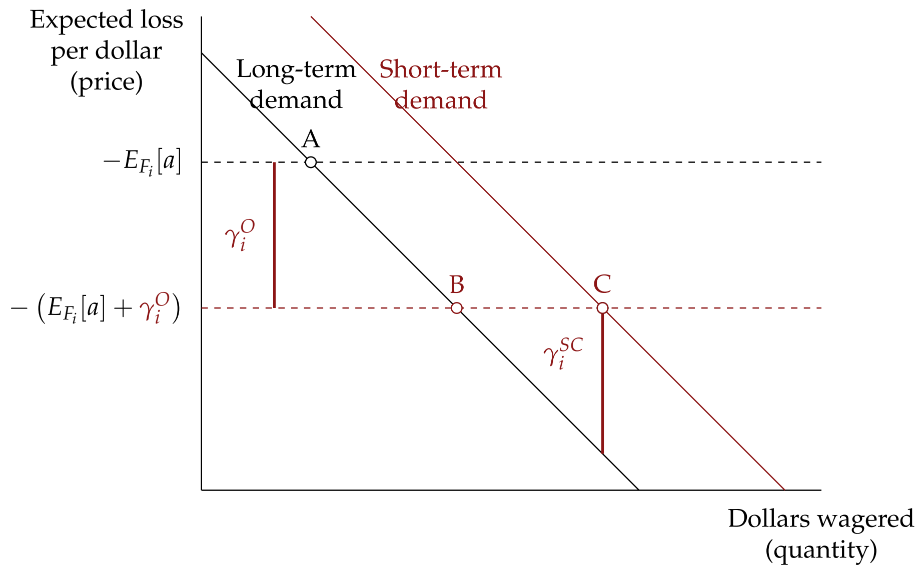

Figure 2: Overconfidence and self-control problems as price-metric biases

Notes: The figure illustrates the definitions of price-metric bias and the effects of bias on consumption. Long-term demand represents choices the agent would make if choosing in advance, and short-term demand represents choices they would make if choosing in the moment. Point A represents unbiased choices, point B represents choices with overoptimism only, and point C represents choices with self-control problems and overoptimism. The linear functional form of demand is only for illustration.

unbiased agent would consume where their long-term demand curve intersects the true price of betting (point A). Adding overoptimism causes the agent to underestimate the price of betting and overconsume (point B). The vertical distance between the perceived price and the true price is the overoptimism parameter γ_i^O. Adding self-control problems causes choices to be driven by short-term demand rather than long-term demand (point C). The vertical distance between the short-term and long-term demand curves is the self-control problems parameter γ_i^{SC}. Since both biases are defined in price units, they both appear on the graph as vertical distances.

These price-metric biases are crucial objects for policy design (Allcott and Taubinsky, 2015). They represent net costs of betting that are not internalized in consumption decisions. Policy interventions can improve welfare by internalizing these costs. For example, a corrective tax could induce a price change equal to the sum of the price-metric biases. Doing so would shift consumption from C to the socially efficient point A.

Motivated by the policy relevance of price-metric bias, we designed our experiment to measure overoptimism and self-control problems in these units. Our experiment is not designed to isolate the mechanisms driving either bias, and we do not take a stand on any particular microfoundations that could generate either bias. Instead, we view γ_i^O and γ_i^{SC} as reduced-form objects capturing the extent to which overoptimism and self-control problems cause net costs of betting to not be internalized in decisions, whatever the underlying behavioral model of bias.[10]

3.3 Discussion

While we allow for general heterogeneity in bias across people, we impose that within person, the bias parameters γ_i^O and γ_i^{SC} are constant. This assumption makes counterfactual simulations tractable by allowing us to hold price-metric bias fixed in scenarios where consumption changes. One violation of this assumption is if people are more or less biased when they are wagering fewer dollars. For readers concerned about such issues, we note that our empirical bias estimates have natural interpretations even when bias can vary with consumption.[11]

For normative analysis, our model assumes that overoptimism and self-control problems are indeed biases rather than nonstandard preferences. We briefly clarify the nature of these assumptions and justify them below.[12]

We only treat overoptimism as a bias to the extent that it causes misperceptions of average financial returns. Overoptimism may have other effects. For example, it may be more fun to bet if you think you are going to win. We incorporate these effects through the nonfinancial utility function z's dependence on $\tilde{F}_i$. There is no obvious reason for the planner to treat such preferences as normatively irrelevant, so we include them in the normative utility function (2). In contrast, the misperceptions of average financial returns captured by γ_i^O represent misunderstandings of the consequences of consumption. Such misperceptions are objective mistakes, and the case for intervening to correct them is strong. Therefore, they do not appear in the normative utility function (2).

Considering self-control problems, our normative assumption amounts to the view that long-run preferences are associated with true well-being rather than short-run preferences. This so-called *long-run criterion* is sometimes disputed on conceptual grounds (Bernheim, 2016). However, in this context, it is consistent with psychiatric and neurobiological evidence that problem gamblers are unable to appropriately weigh the consequences of their actions when they experience gambling compulsions (Potenza et al., 2019).

[10] In the language of Chetty et al. (2009) and Chetty (2015), these parameters are the sufficient statistics for welfare analysis that arise from the underlying behavioral model.

[11] Let $\gamma_i^O(x)$, $\gamma_i^{SC}(x)$ denote overoptimism and self-control problems that can vary with dollars wagered x. Our estimates of overoptimism correspond to estimates of average overoptimism between 0 and the chosen level of consumption, $\left(\int_0^{x_i} \gamma_i^O(x)dx\right)/x_i$. Our estimates of self-control problems correspond to average self-control problems at the chosen level of consumption, $\gamma^{SC}(x_i)$.

[12] See Bernheim (2016) and Bernheim and Taubinsky (2018) for a more detailed discussion of normative approaches in behavioral public economics.

Phase	Date	Sample Size
Recruitment and intake	March 13 - April 8	545,197 viewed social media ads 12,912 clicked on ads 6,155 satisfied initial eligibility criteria 2,062 consented and provided contact info 666 synced at least one account 555 synced all accounts
Survey 1	April 9	533 completed survey 1
Survey 2	May 10	486 completed surveys 1 and 2
Survey 3	June 10	472 completed surveys 1, 2, and 3, of which 447 provided data for all accounts through June 10 and 444 were also not in the MPL group (analysis sample)

Table 1: Experiment timeline and sample sizes

Notes: The table illustrates the number of unique participants who completed each step of our study.

4 Field Experiment

4.1 Overview

Participants in our experiment shared data on their sports betting activity and completed three surveys over a two-month period. Table 1 presents sample sizes over the course of the experiment.

Our intake procedures were designed to recruit a sample of high-volume U.S. sports bettors. We focus on high-volume bettors because they account for the vast majority of sports betting consumption.[13] We focus on the U.S. because it is one of the fastest-growing markets and because our tool for collecting observational data only works for U.S. sportsbooks. To reach this population, we conducted a targeted social media ad campaign. Appendix Figure A.1 presents a sample advertisement, and Appendix Table A.1 summarizes our targeting procedures. Upon clicking our ad, participants began an intake survey with an initial eligibility screening module. We excluded participants unless they self-reported betting on sports at least once a week and wagering at least $100 over the last 30 days. 6,155 people satisfied these eligibility criteria, and 2,062 agreed to participate after viewing the introductory materials.

Participants were required to share data on their sports betting activity. Our data collection procedure was designed to ensure that we observed the near-universe of sports wagers placed by our participants. Before we told participants about the data sharing re-

[13] Forrest and McHale (2024) show, using data from U.K.operators, that the top 5% of bettors accounted for 64% of sportsbook revenues in 2023.

quirements, we asked them whether they used each of six popular sportsbooks: DraftKings, FanDuel, BetMGM, ESPN BET, Caesars, and Hard Rock Bet. These sportsbooks account for the vast majority of legal sports bets in the U.S.[14] After we elicited this list, we asked participants to sync their accounts to the research study via a portal that we developed with an external company, SharpSports. We provide a screenshot of this portal in Appendix Figure A.2.

For participants who complied with syncing requirements, we observe detailed information about every bet placed between Jan 1, 2024 through the end of the study on June 10, 2024. This observational data has several advantages relative to self-reports. Most importantly, given that we are studying misperceptions, we did not want to rely on potentially misreported consumption and returns data as the ground truth. The data also contain detailed bet characteristics for every wager.

Participants received surveys 1, 2, and 3 on April 9, May 10, and June 10, respectively. Throughout this paper, we define period 1 to be the 30-day period from April 10 to May 9 and period 2 to be the 30-day period from May 11 to June 9, so that period t is always the period following survey t. We define pre-periods 0, -1, -2, and -3 similarly as 30-day periods counting backwards from April 8.

All randomized treatments were implemented in survey 1, so we undertook procedures to minimize attrition among the 533 participants who completed survey 1 (Appendix A.1). Overall, 516 (97%) completed at least one follow-up survey and provided follow-up data for at least one account, and 447 (84%) completed both surveys 2 and 3 while providing complete betting data. To define our final analysis sample of 444 participants, we exclude three participants who were randomized into a small treatment condition that ensured the incentivization of multiple price list elicitations. We discuss selection into the sample and differential attrition in Section 5.3.

We pre-registered our experiment at the AEA RCT registry (AEARCTR-0013310). Our pre-analysis plan specified the analysis sample, variable definitions, and regression specifications for estimation of randomized treatment effects. Appendix D summarizes differences between the body of the paper and the pre-registered analysis and reports all pre-registered estimates.

4.2 Measuring Overoptimism About Financial Returns

In the first module of every survey, we asked:

[14] DraftKings and FanDuel alone account for more than 70% of sports betting transactions (Baker et al., 2024).

How much do you expect to gain or lose for every \$100 that you wager on [synced apps] over the next thirty days?

We compare these predictions to true net returns to measure overoptimism γ^O. In surveys 1 and 2, we randomized whether these predictions were incentivized.

To ensure that responses coincided with the target beliefs, we guided participants through a structured elicitation procedure. For example, the procedure required participants to confirm their understanding that "breaking even" corresponds to winning \$0 per \$100. We also invited respondents to voluntarily revise their responses if they predicted expected net returns outside the range $[-0.4, 0.25]$. As we pre-registered, we truncate all net return predictions to lie within this range.

To enrich our understanding of perceptions, we asked analogous questions about past returns. In surveys 1, 2, and 3, participants were asked to recall their net winnings per \$100 wagered in the period from Jan 1 to April 7, $t = 1$, and $t = 2$, respectively. In survey 1, we also asked participants to estimate the average net returns of American sports bettors in 2023. These questions were incentivized for all participants.

4.3 Measuring Self-control Problems

We separately estimate perceptions of self-control problems and naivete about self-control problems. We define perceived self-control problems $\tilde{\gamma}_i^{SC}$ as the agent's prediction about how much their future self overvalues sports betting. We allow the agent to underestimate his own self-control problems, defining naivete as $\tilde{\gamma}^{SC} - \gamma^{SC}$, the difference between perceived and true self-control problems. Adding these two objects yields γ^{SC}, the parameter governing distortions in sports betting consumption.

Our design follows a recent literature in behavioral economics that uses similar methods to estimate models of time-inconsistency with partial sophistication (Acland and Levy, 2015; Chaloupka et al., 2019; Augenblick and Rabin, 2019; Carrera et al., 2022; Allcott et al., 2022b,a). This approach has a number of advantages relative to alternative popular designs (DellaVigna and Malmendier, 2006; Read and van Leeuwen, 1998). We briefly summarize these advantages in Appendix B.2.

Measuring Perceived Self-control Problems To measure perceived self-control problems, we created an incentive called the Bet Less Bonus. In survey 1, we introduced it to all participants with the following text:

You may have the opportunity to earn money by betting less on sports over the next 30 days!

If you are selected for the Bet Less Bonus, you will receive a $6 payment for every $10 that you reduce your average daily betting, up to a maximum bonus of $[M].

You'll only get paid if you wager less than $[B] per day, which is slightly more than how much you've been wagering recently.

The benchmark B and maximum payment M were replaced with personalized values based on past betting consumption.[15] For bettors who wagered less than $B·30 in the relevant period, the bonus corresponded to a 2¢ increase in the price of wagering $1.[16] The payment arrived at the end of survey 3.

Valuations of the Bet Less Bonus identify perceived self-control problems because if someone predicts that their future self will bet too much, they will place higher value on an incentive to bet less in the future. To elicit valuations of the Bonus, we used a multiple price list (MPL). Participants made a series of binary choices between receiving the Bonus and receiving a fixed payment of varying size.[17] We incentivized the MPL choices by informing participants that they may be randomly selected to have their choices "count," in which case one of the MPL options would determine their payment.

In practice, 0.5% of participants were randomized to have MPL choices count. For the other 99.5% of participants, we randomly assigned to receive the Bonus (with probability 0.35) or to a control arm (with probability 0.65). This random assignment induces exogenous variation in the price of betting, which we use to measure the price-sensitivity of demand. This and all other randomized treatments were independently cross-randomized, stratified by period-0 net returns and period-0 dollars wagered.

Since the reasoning that causes perceived self-control problems to increase bonus valuations is somewhat subtle, we took several steps to ensure that participants understood the decision. One simple but important choice was about framing: we called the incentive the "Bet Less Bonus" to highlight that the incentive would reduce future consumption.

[15] The benchmark B was the average daily wagers on baseline supported accounts in 2024, rounded up to the nearest $10. We capped bonus payments at $90: the maximum payment M was $\max\{B/10 \cdot 6, 90\}$.

[16] Since the bonus was active for 30 days, a $10 reduction in daily wagers corresponds to a reduction in wagers of $300 across the whole period. Therefore, the $6 payment corresponds to a subsidy of $0.02 per dollar of wager reductions. We framed the bonus as a subsidy for reducing daily betting rather than a subsidy per dollar reduced because participants were inattentive to the latter framing in pilots.

[17] We treat the lowest fixed payment where a participant chose the bonus as an upper bound on their bonus valuation, and the highest fixed payment where they chose the payment as a lower bound. We refined valuation estimates further by asking participants to state the exact fixed payment which would have made them indifferent between a fixed payment and the bonus. When this self-reported indifference point falls within our MPL bounds, we use it as the bonus valuation. When it does not, we use the midpoint of the MPL bounds.

We also asked participants to predict how the Bonus would affect their consumption using an interactive screen that transparently mapped consumption reductions into Bonus payouts. Finally, following Allcott et al. (2022b), we included the following text to bring relevant considerations to mind:

> *How might you decide?*
>
> *You told us that if you are selected for Bet Less Bonus, you expect to wager $[Prediction]. Thus, you expect that you would earn a Bet Less Bonus of $[ExpectedBonus].*
>
> *You might prefer $[ExpectedBonus] instead of the Bet Less Bonus if you don't want any pressure to bet less.*
>
> *You might prefer the Bet Less Bonus instead of $[ExpectedBonus] if you want to give yourself extra incentive to bet less.*

Measuring Naivete We measure naivete about self-control problems by comparing predictions of future dollars wagered to the truth. We infer naivete if people systematically underestimate future consumption. Intuitively, people who underestimate their future self-control problems will also underestimate their future consumption, while people who are sophisticated about self-control problems will not underestimate future consumption. As with the predictions of future returns, we used a structured elicitation procedure to minimize the influence of misunderstandings, and we randomized whether the prediction was incentivized.[18]

4.4 Bias-correction Interventions

We created experimental treatments to evaluate two bias-correcting interventions. Our experimental interventions mirror real interventions that sportsbooks have implemented with the express goal of reducing bias.

Addressing Overoptimism: History Transparency In the *history transparency treatment*, participants receive information about their past net returns. The treatment mimics interventions like DraftKings's *My Stat Sheet*, which provides similar information about historical activity at a single sportsbook.[19] Similar interventions that provide information

[18]To make sure people did not neglect changes in the sports schedule over time, we began by asking participants qualitative open-ended questions about which sports they might bet on over the next 30 days vs. the past 30 days. Then, to make sure participants understood the scale of their own past betting, we told them their own dollars wagered over the past 30 days. We next asked whether they expected to bet more, less, or about the same over the next 30 days compared to the past 30 days. Finally, we elicited a numerical prediction of future dollars wagered.

[19]The stated goal of *My Stat Sheet* is to "help customers evaluate their play and make informed choices." DraftKings press release: https://draftkings.gcs-web.com/news-releases/news-release-details/draftkings-launches-my-stat-sheet-new-tool-promote-responsible

about past performance have been studied in various lab (Möbius et al., 2022) and field (Carrera et al., 2022) settings.

We randomized 50% of participants to be treated. They viewed the following text:

> *You said you [won/lost] $[recollection] for every $100 that you wagered.*
>
> *In fact, you [won/lost] $[realization] for every $100 that you wagered.*
>
> *This calculation used data from [number] bets on [synced accounts] in 2024.*

After viewing the information, participants answered an open-ended question about whether their recollection was close to the truth. Then, participants were allowed (but not forced) to update their predictions about the future.

We study whether the intervention reduced bias on average, and whether the intervention's treatment effects were correlated with pre-intervention bias.

Addressing Self-control Problems: Voluntary Wager Limits In the *limits treatment*, participants were prompted to make an active choice about how to use sportsbooks' weekly wager limit tools. These tools allow users to cap the amount that they are allowed to wager in a given calendar week. These caps are binding, at least in the short term: if a user attempts to remove the cap, the change will not take effect until after a cooling-off period. Every sportsbook in our study has some version of a weekly limit tool.

We framed limit tools as a way for participants to control their own sports betting activity. In survey 1, we elicited a stated "ideal" weekly wager volume for all participants with the following question:

> *Ideally, in the next 30 days, how much would you wager per week on [synced accounts]?*

Then, the 50% of participants in the limits treatment viewed the following text:

> *For this study, you are required to choose a weekly wager limit for your [synced accounts].*
>
> *If you don't want to limit your betting, that's fine! In that case, you can set very high limits, like $9,999,999. Then you wouldn't be restricting your betting at all. The only requirement is: you must type some number into the weekly limit box for each of your synced accounts.*

We used comprehension questions to confirm understanding.

Next, we asked participants to plan the weekly wager limits across their synced accounts. We included a reminder of their previously stated ideal on the planning screen.[20] Then, we provided video instructions on how to set limits for each account in succession. After showing the video instructions, we asked participants which limits they actually chose. These final self-reports are our measure of chosen limits.

We study whether participants chose restrictive limits, and whether restrictiveness correlates with our measures of self-control problems. We interpret these choices as the restrictiveness that people would choose absent hassle costs and informational frictions.[21] Outside our experiment, most U.S. bettors do not use limits at all.[22] In other contexts, such as Norway (Auer and Griffiths, 2022) and Italy (Calvosa, 2017), players are required to make choices about limit choices upon account creation. Participant choices from the limits treatment are informative about the plausible impacts of such a regulation in the American market.

4.5 Qualitative Measures of Bias and Other Supplemental Variables

We collected several qualitative measures of overoptimism and self-control problems to supplement and validate our quantitative estimates. The *Gambling Literacy Index*, adapted from Wood et al. (2017), measures the extent to which gamblers understand that they cannot beat the house. We interpret it as a proxy for overoptimism. The *Problem Gambling Severity Index* (Holtgraves, 2009) is a prominent screening tool for problem gambling. We interpret it as a proxy for self-control problems. Appendix A.2 contains the pre-registered questions and definitions for both indices. In addition to these measures from the literature, we asked participants whether they thought they were sports betting "too little," "too much," or "the right amount." This question is another qualitative measure of (perceived) self-control problems.

We also asked participants to predict their responses to hypothetical changes in the price of betting. These questions provide supplemental evidence on the sensitivity to naturally framed price changes, as a complement to the randomized treatment effect of the Bet Less Bonus. In survey 3, we asked about the following two scenarios:

[20] The reminder read: *For reference, you told us earlier that in upcoming weeks, you'd ideally wager a total of $[ideal] per week on [synced apps].* If the chosen limits differed from the ideal, we displayed a message indicating this fact.

[21] For example, the screen where a user can set a limit is sometimes hidden behind opaque menus.

[22] DraftKings reported to the Massachusetts Gaming Commission that less than 5% of users in that jurisdiction used limits (Massachusetts Gaming Commission, 2023). In our sample, 93% of users had heard of self-imposed limit tools before, but only 7% of participants had used them.

Suppose that all sportsbooks made their odds 2% worse for an extended period of time. How much, if at all, would you reduce your betting?

Suppose you learned that you were 2% worse at making money betting than you'd thought. How much, if at all, would you reduce your betting?

We elicited answers in percent terms.

Another kind of hypothetical price change was a change in the Bet Less Bonus's payment rate. We asked participants to predict consumption under payment rates of $\{0.01, 0.02, 0.04, 0.083\}$¢ per dollar. We varied the payment rates within participant and on the same decision screen, so these responses tell us about how much the price change from the Bet Less Bonus affected predictions when that price change was salient. These questions also appeared on survey 3.

These prediction questions were not pre-registered. We added them to the survey after we observed the treatment effect of the bonus in survey 1, with the goal of improving our interpretation of that effect and improving our price-sensitivity analysis.

We also asked various other qualitative and open-ended questions, which we describe as they become relevant.

5 Sample Characteristics

5.1 Betting Activity

Figure 3 summarizes betting activity in the 30 days before the study began. The median participant wagered $153/week in the pre-study period, confirming that our recruitment procedures successfully targeted high-volume bettors. This wager volume would put the median member of our sample close to 95th percentile of wager volume among all sports bettors (Forrest and McHale, 2024).[23] Even within our sample, bet volume has a long right tail – more than 25% of our sample wagered more than $1000 per week. The distributions of number of bets per week (median = 17.5) and average bet size (median = $10.3) are similarly skewed. Finally, we present in the lower right panel the participants' average decimal odds, or the potential payout per dollar wagered. Long shot bets are popular: the median person's average bet has decimal odds of 4.6

[23] We draw this conclusion by comparing $153/week to results in Forrest and McHale (2024), who report on the distribution of dollars wagered on sports over a year-long period in the U.K. from 2018-2019. That paper uses a multi-operator dataset, which makes it comparable with our data. They find that the 95th percentile of wager volume was £65/week, or $102/week in 2024 dollars. This number is lower than our median dollars wagered, but the pre-study period was a relatively high-volume betting period, so it is reasonable to say that our median would be close to this 95th percentile.

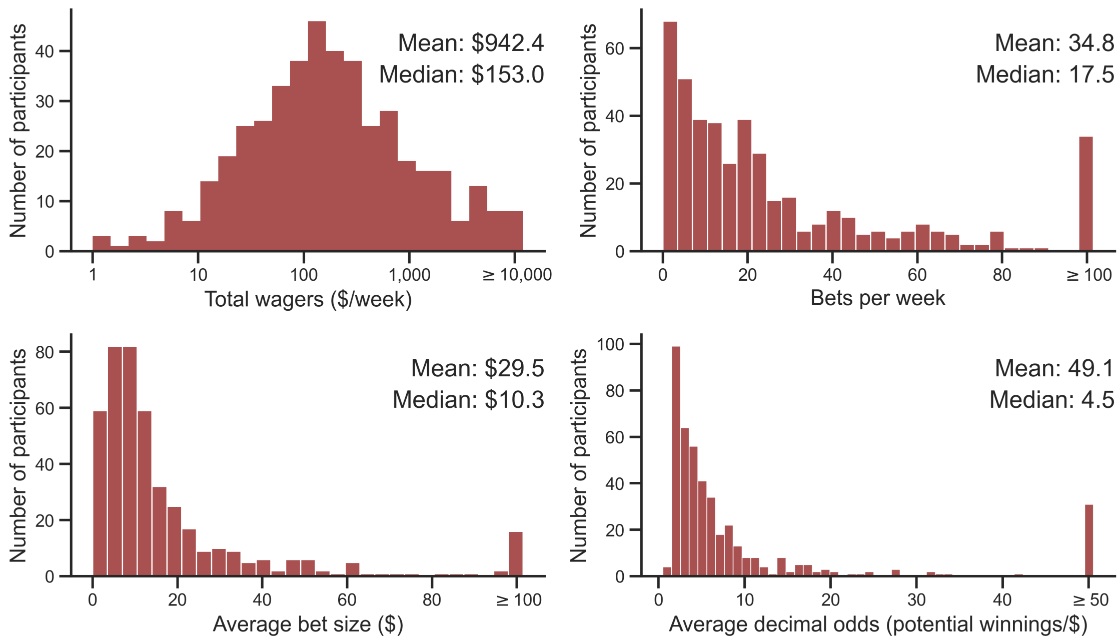

Figure 3: Pre-study betting activity across individuals

Notes: These plots illustrate patterns of betting activity in the analysis sample (N = 444) over the 30-day period prior to survey 1. The x-axis in the bottom right panel is a person's average decimal odds across all bets in this pre-period, weighting by dollars wagered. If a bet has decimal odds of 5, that means when a person wagers $1 on that bet, he stands to win $5.

5.2 Qualitative Evidence of Bias

We present qualitative evidence on overoptimism and self-control problems in Figure 4. The top left panel reports on responses to the question: "Which of the following statements are reasons that you bet on sports?" The most popular selections relate to the nonfinancial value of betting (e.g, 81% said sports betting "makes watching sports more fun"). We included two options that related to our biases of interest. 39% selected "On average, I win money by betting" as a reason for betting, but only 2% selected "I can't stop myself from betting, even though I wish I could bet less."

The top right panel provides further evidence on perceived self-control problems. 20% of the sample self-report betting on sports too much. For comparison, previous work shows that over 50% of respondents reported that they used social media too much on a question with analogous language (Allcott et al., 2022a).

The bottom two panels report on the qualitative bias measures that we borrow from the gambling studies literature. The bottom left panel shows responses to the questions that make up the Gambling Literacy Index. For each question, around 20% of participants gave responses that could plausibly be interpreted as factual misunderstandings of how gambling returns work. The bottom right panel illustrates Problem Gambling Severity

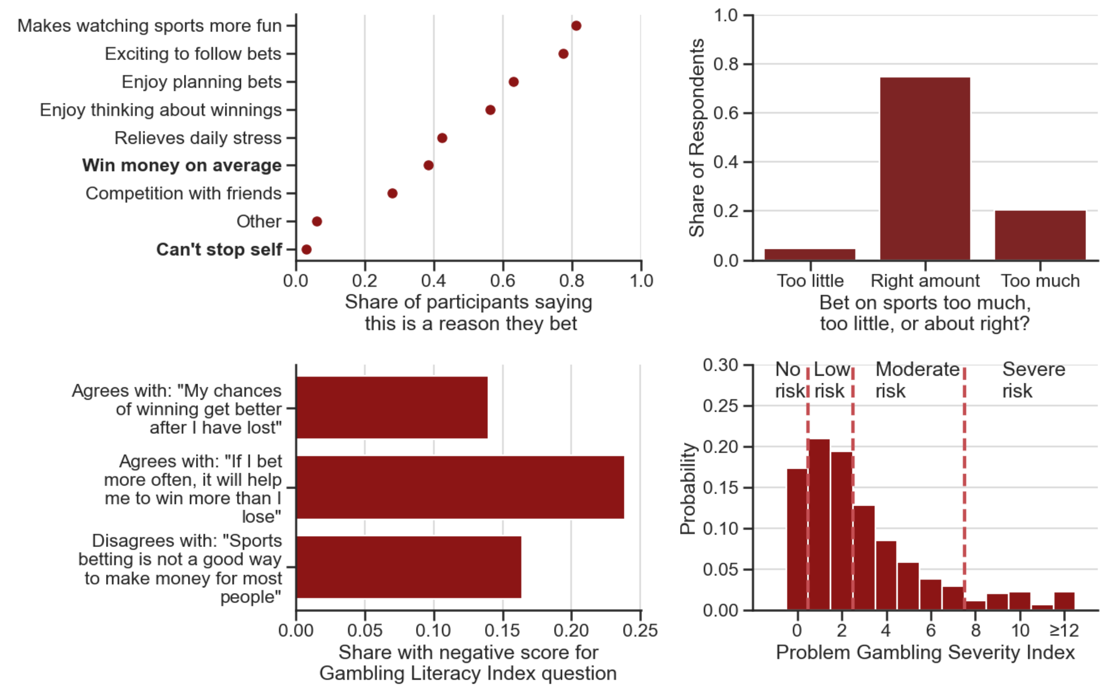

Figure 4: Qualitative evidence of bias

Notes: The top left panel reports on responses to the question: "Which of the following statements are reasons that you bet on sports?" The top right panel reports on responses to the question: "How do you feel about your own sports betting in a typical week?" The bottom left panel reports the share of respondents who give answers that receive negative scores for each of the three components of the Gambling Literacy Index. The bottom right panel shows the distribution of Problem Gambling Severity Index scores. Scores in buckets of $\{0, [1,2], [3,7], [7,21]\}$ are sometimes classified by practitioners as "no risk," "low risk," "moderate risk," and "high risk" for problem gambling, respectively.

Index scores. A majority of the sample records a score of 2 or lower, which is traditionally regarded by practitioners as low or no risk of problem gambling.

Overall, these findings suggest that overoptimism and self-control problems are relevant to the sports betting decisions of some, but not all, members of our sample. They also suggest that the vast majority of participants derive some form of normative entertainment value from betting. Regulators therefore face a tradeoff between reducing biased consumption and preserving entertainment value. The quantitative evidence presented in following sections clarifies the nature of this tradeoff more explicitly.

5.3 Demographics and Representativeness

We use external data from surveys conducted by Grubbs and Kraus (2023c) (henceforth "GK") to understand the sports bettor population. GK conduct surveys containing basic demographic questions, the Gambling Literacy Index, and the Problem Gambling Severity Index. They study nationally representative samples of the general population, self-

Variable	Grubbs and Krauss Census Matched	 Weekly Lottery Bettors	 Weekly Sports Bettors	Brown, Grasley, and Guido Analysis Sample
N	2806	406	517	444
Demographics				
Age	51.59	55.21	41.47	39.92
White	0.66	0.62	0.59	0.81
Male	0.46	0.53	0.69	0.96
Bachelor's degree or higher	0.34	0.25	0.50	0.82
Graduate degree	0.13	0.08	0.19	0.39
Household income ($000s)	68 (62)	67 (57)	101 (84)	156 (116)
Qualitative bias measures				
Gambling Literacy Index	4.00 (2.30)	3.12 (2.74)	1.53 (3.03)	3.55 (2.05)
Problem Gambling Severity Index	0.99 (2.69)	2.83 (4.21)	6.77 (5.06)	2.89 (2.85)

Table 2: Demographics, qualitative bias measures, and representativeness

Notes: Columns 1 through 3 present average demographics and qualitative bias measures for subsamples of the Grubbs and Kraus nationally representative study. They show averages for the census-matched sample, weekly lottery bettor sample, and weekly sports bettor samples respectively. Column 4 presents averages for our main analysis sample. Standard deviations are in parentheses.

identified weekly lottery bettors, and self-identified weekly sports bettors.[24] Columns 1, 2, and 3 of Table 2 report average demographics and qualitative bias measures for these three samples respectively. Unlike lottery bettors, sports bettors are younger, have higher education, and earn more than the average American. Sports bettors also have worse biases than lottery bettors according to qualitative measures.

We next compare our analysis sample to the GK weekly sports bettors sample. Our analysis sample was not designed to match the GK weekly sports bettor sample – recall that we intentionally oversampled high-volume bettors – but it is nevertheless a useful comparison group. Our analysis sample is whiter, more male, more highly educated and earns more than the GK weekly sports bettors. Our sample also has higher scores on the Gambling Literacy Index and lower scores on the Problem Gambling Severity Index. This result suggests that, if anything, bias estimates for our analysis sample may be conservative estimates of bias for frequent bettors more generally.

6 Reduced-form Evidence

This section summarizes our quantitative estimates of overoptimism, self-control problems, price-sensitivity, and the effects of bias-correcting interventions.

[24]Grubbs and Kraus (2023c) (henceforth GK) conducted two YouGov surveys to learn about the population of sports bettors. The first was a sample of 2,806 Americans defined to match the 2019 ACS on demographics. The second was a sample of 1,557 self-reported sports bettors, weighted to match the demographics of sports bettors in the first sample. We construct two subsamples of GK data. *GK weekly sports bettors* includes all participants from either survey who report sports betting weekly or more frequently (N = 406). *GK weekly lottery bettors* includes all participants who reported buying lottery tickets weekly or more, from the first survey only (N = 517). The authors and collaborators have published several other descriptive studies using this data, including Grubbs and Kraus (2023b, 2022, 2023a); Connolly et al. (2024)

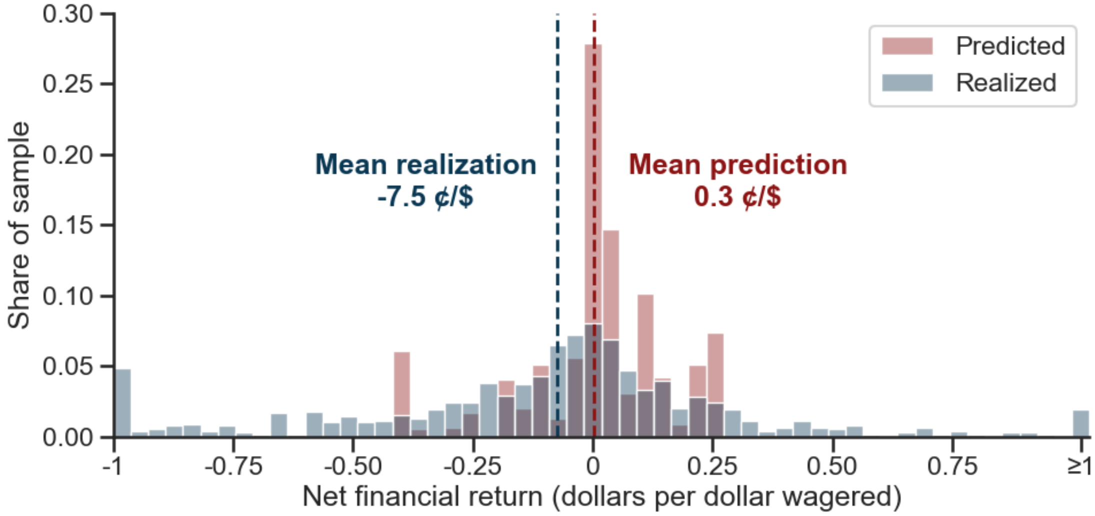

Figure 5: Evidence on overoptimism from predicted and realized returns

Notes: This figure shows the distribution of predicted and realized winnings, pooling across period 1 (April 9 to May 9) and period 2 (May 11 to June 10). Predictions are elicited in surveys 1 and 2 respectively. We restrict the sample of predictions to participants who placed wagers in the relevant period. Predictions are censored to lie within $[-0.4, 0.25]$. Dotted lines and annotations represent unweighted averages.

6.1 Overoptimism About Financial Returns

We find that participants overestimate the financial returns to betting. On average, participants predicted that they would roughly break even, but the average participant lost 7.5¢ per dollar. We illustrate the distributions of predicted and realized returns in Figure 5, pooling across surveys 1 and 2. We interpret the average difference between predictions and realizations as a transparent estimate of average overoptimism.

The observable feature of betting activity that is most strongly associated with overoptimistic predictions is the share of dollars wagered on parlays. Panel A of Figure 6, which plots the heterogeneity of prediction errors with respect to bettor characteristics, shows that bettors with high parlay shares are 18¢ per dollar more overoptimistic than low parlay share bettors. Panel B illustrates this relationship with further detail. The association arises because even though high parlay share bettors actually earn much less on average, their predictions do not reflect their worse prospects.[25] These results indicate that consumer protection concerns around parlays have some legitimacy.

We also show in Panel A of Figure 6 that low education is the demographic characteristic most strongly associated with overoptimistic prediction. Therefore, overoptimism is a regressive bias in the sense that it is more prevalent among disadvantaged populations.

[25]The association between high parlay share and misprediction persists even after controlling for other observables (Appendix Table B.1).

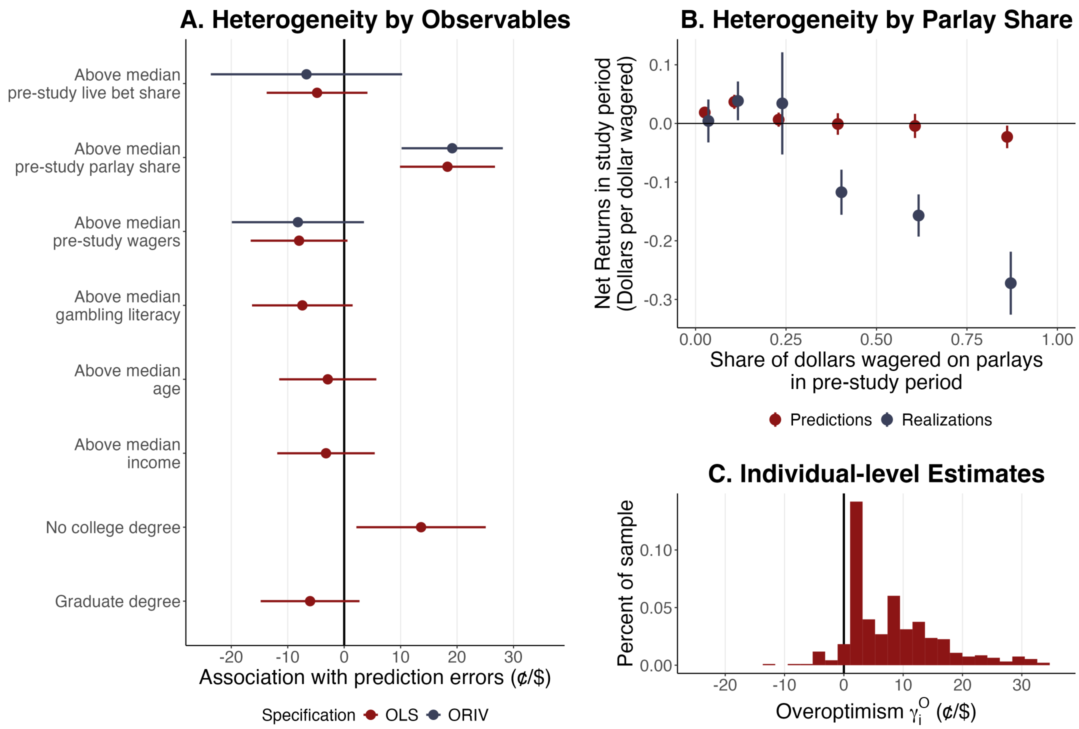

Figure 6: Heterogeneity in overoptimism

Notes: Panel A shows coefficient estimates from univariate regressions of signed prediction errors about net returns (pooled across periods 1 and 2 when possible) on binary explanatory variables. An observation is a member of the analysis sample who placed wagers both in the pre-study period ($t = -1, -2$) and the study period ($t = 1, 2$). Error bars represent 95% confidence intervals. Pre-study variables are measured as the average of the relevant quantities in periods -1 and -2. The "ORIV" points show estimates using the Obviously Related Instrumental Variables procedure of Gillen et al. (2019) to correct for measurement error in the pre-study variables. Panel B shows how predicted and realized net returns vary with dollars wagered on parlays in the pre-study period. Points are binned means and error bars are 95% confidence intervals. Panel C shows the histogram of individual-specific overoptimism estimates, computed via the shrinkage procedure described in Appendix B.3.

We next turn to individual-specific estimates of overoptimism. As we have emphasized, allowing for rich heterogeneity in overoptimism is crucial for studying interventions that target heterogeneous bias. The empirical challenge is that we do not observe overoptimism directly. Instead, we observe raw prediction errors, which are a combination of underlying overoptimism and noise. Noise arises from the intrinsic randomness of gambling outcomes and elicitation noise in the prediction question (Kahneman, 1965; Gillen et al., 2019). We apply a shrinkage procedure to recover estimates of the underlying overoptimism parameter for each participant (Chen, 2024). The shrinkage procedure requires us to measure the variance of both sources of noise, which we accomplish using the underlying bet-level microdata and the panel structure of our data for gambling outcome noise and elicitation noise respectively. To measure outcome noise, we assume that the outcomes of bets are independent. To measure elicitation noise, we assume that elicitation errors are mean-zero and i.i.d. across surveys. We provide details on our procedure and discuss robustness to alternative assumptions in Appendix B.3.

Panel C of Figure 6 shows that individual-specific estimates are quite dispersed. The estimates imply that 9% of participants are *underoptimistic* about their expected returns, and 10% are overoptimistic by more than 20¢/dollar. This result illustrates the importance of targeting in our context: optimal policy would treat these two kinds of bettors differently.

To support our interpretation of prediction errors as evidence of overoptimism, we provide evidence against four alternative explanations. First, people might have misunderstood the nature of our survey question in a way that caused them to report higher predictions. Since we used consistent language across all expected returns questions, this explanation would imply that people consistently overestimate all net returns. Instead, we find that overestimation is local to predictions about the future: when we ask people to recall past net returns, they *underestimate* their winnings (Appendix Figure B.1).[26] Second, people might skew their reported predictions upwards out of a desire to protect ego-relevant beliefs. If this were true, we would expect raising the stakes of predictions to decrease predicted winnings (Prior et al., 2015). In contrast, we find that incentivizing predictions did not affect responses (Appendix Figure B.2). Third, we show in Appendix B.1 that rounding in survey responses does not drive our overoptimism estimate. Fourth, we note that our pre-registered truncation of predicted net returns to the range $[-0.4, 0.25]$ does not reduce our average predicted returns measure.[27]

While our experiment was not designed to study the behavioral mechanisms behind overoptimism, our descriptive patterns provide suggestive evidence supporting some mechanisms from the literature. First, in this sample of highly experienced bettors, overoptimism about future returns persists even though people did not overestimate past returns. This pattern is inconsistent with models of motivated reasoning where agents forget unfavorable signals (Bénabou and Tirole, 2002; Huffman et al., 2022; Sial et al., 2023). Instead, our results are consistent with models where agents selectively interpret past signals (Thaler, 2024), for example by attributing bad news to luck and good news to skill (Ross, 1977; Stipek and Gralinski, 1991). Second, parlay betting is associ-

[26] The finding that people do not overestimate past net returns is consistent with previous work by gambling studies scholars (Braverman et al., 2014).

[27] Of the 790 participant-months in periods 1 and 2 where participants placed at least one wager, we have 36 predicted returns that were above 0.25¢/\$ (4.6% "truncated down") and 37 predicted returns that were below -0.4¢/\$ (4.7% "truncated up"). Truncation increases our measure of average predicted returns and therefore our measure of average overoptimism: the average of the untruncated belief variable is 2.7¢/\$.

ated with overoptimism. The main differences between parlays and other kinds of bets are that parlays have multiple legs and longer odds. We find that when we measure the association conditioning on bet odds, the parlay effect remains large and significant while bet odds have no independent effect (Appendix Table B.1). This pattern suggests that overoptimism about parlay wagers is not caused only by a generic over-weighting of small probabilities (Kahneman and Tversky, 1979). One plausible alternative is parlays in particular are more complex than standard wagers because of their multiple components (Enke and Shubatt, 2023).

6.2 Self-control Problems

As described in Section 4.3, we measure both people's perceptions of their self-control problems and their naivete about their self-control problems following the procedure outlined in Carrera et al. (2022). We combine these measures to get an estimate of total self-control problems.

Our measure of perceived self-control problems is the difference between observed valuations of the Bet Less Bonus and the counterfactual valuations that agents would have reported in the absence of self-control problems. Intuitively, a participant who wants his future self to bet less will value the Bet Less Bonus more than a participant with no self-control problems. To measure the counterfactual valuation of the Bonus for a participant with no self-control problems, we compute the unconditional transfer value of the Bonus plus the approximate consumer surplus loss from raising the price of wagering by 2¢ per dollar for each participant. We provide further details in Appendix B.2.

Panel A of Figure 7 shows that, on average, participants valued the bonus at \$17.58, which exceeds the no self-control problems benchmark by \$1.86 (12%). We refer to this difference as the *behavior change premium* (BCP) because it represents the amount of money a participant is willing to pay for the behavior change induced by the Bonus. We next convert the BCP into price-metric units by dividing it by the average predicted Bonus-induced consumption reduction. The average participant was willing to pay \$0.007 for their future self to reduce consumption by a marginal dollar. Following Carrera et al. (2022) and Allcott et al. (2022a), we use this statistic as our estimate of average perceived self-control problems ($\tilde{\gamma}^{SC}$ in our model). In Panel B of Figure 7, we show that the BCP per dollar wagered is larger for agents who self-report a desire to reduce future betting consumption. The fact that the BCP per dollar wagered coincides with this qualitative measure of perceived self-control problems across participants validates our use of the BCP per dollar as an estimate of $\tilde{\gamma}^{SC}$.

We also find that participants are not naive about their self-control problems. Intu-

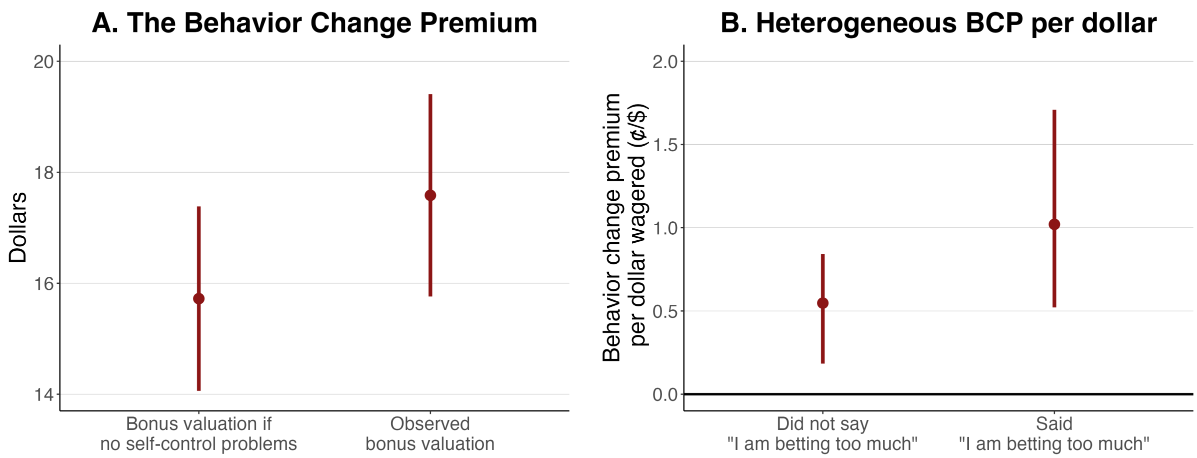

Figure 7: Evidence on self-control problems from Bonus valuations

Notes: Panel A plots average observed bonus valuations against counterfactual average bonus valuations for participants without self-control problems. The difference between observed valuations and the counterfactual valuations is the behavior change premium. Error bars represent 95% confidence intervals. A paired t-test of the difference gives a p-value of 4e-05 (the measures are highly correlated across subjects). Panel B illustrates heterogeneity in the behavior change premium per dollar with respect to a qualitative measure of self-control problems. Error bars represent bootstrapped 95% confidence intervals. The randomization p-value for the difference in estimates is 0.055. For both plots, the sample is the subset of participants whose predicted consumption under both the bonus control and bonus treatment conditions were below the benchmark and above the value at which the price increase from the Bet Less Bonus would no longer bind (N=352, 79% of the analysis sample).

itively, if people underestimate their future self-control problems and are otherwise unbiased, they will also underestimate their future dollars wagered (Augenblick and Rabin, 2019). Appendix Figure B.3 shows that participants do not underestimate their future consumption. Therefore, we use perceived self-control problems ($\tilde{\gamma}^{SC}$) as a measure of overall self-control problems (γ^{SC}).

Summary of bias estimates Taken together, our results imply that $E_i[\gamma_i^{SC}] = E_i[\tilde{\gamma}_i^{SC}] =$ 0.7¢ per dollar wagered. The first equality, showing that perceived self-control problems and true self-control problems coincide, comes from participants not underestimating future consumption. The second equality, quantifying the magnitude of self-control problems, comes from the BCP per dollar.

We emphasize that our estimate of average price-metric self-control problems is an order of magnitude smaller than our estimate of average price-metric overoptimism. This result implies that in our sample, the primary motive for corrective policy is overoptimism correction rather than self-control problem correction.

6.3 Price-sensitivity

Price-based corrective policies, such as taxes, can only improve welfare to the extent that consumption responds to prices. We provide two kinds of evidence on own-price demand responses: the randomized effects of the Bet Less Bonus and predicted responses to other price changes. While the Bet Less Bonus represents actual price changes, it may overstate the actual price-sensitivity because it made the expected impact on consumption salient[28]. On the other hand, hypothetical price changes are not actual price changes, but they may represent more realistic saliency of prices on consumption.

Our first set of results is the treatment effects of the Bet Less Bonus, which was an experimentally induced 2¢ increase in the price of wagering $1. The first point in Panel A of Figure 8 illustrates the main result: the Bet Less Bonus reduced consumption by 34% for the average participant. We provide details on our regression specification in Appendix B.4. Despite the claims of gambling industry groups, we also find that treated participants did not substitute to illegal alternatives nor other forms of gambling in economically significant amounts.[29]

Our second set of results concern predicted responses to various hypothetical price changes. The main pattern is that predicted responses to hypothetical price changes were about half as large as the Bonus effect. Panel A of Figure 8 illustrates this result by plotting all of our price-sensitivity estimates on the same scale. This pattern holds across three kinds of hypothetical price changes: changes in the house cut, changes in skill, and versions of the Bonus treatment where differences in the payment rate were salient (as described in Section 4.5). To validate these predictions, we show that predicted changes in consumption correlate with observed changes in consumption across both the Bonus control and treatment conditions (Appendix Table B.2)[30].

We illustrate how price-sensitivity is associated with important participant characteristics in Panel B of Figure 8. For brevity, we show the results using price-sensitivity implied by hypothetical price changes. We estimate heterogeneous price-sensitivities by four subgroups of the analysis sample, splitting at the median of pre-study wager vol-

[28] We added multiple-choice questions to survey 3 specifically to check for these non-price effects. Responses suggest that non-price effects were relevant for a substantial minority of participants (Appendix Figure B.4).

[29] Industry groups often emphasize substitution to illegal markets as a limitation of policies that restrict legal gambling (see for example (Association, 2024)). Appendix Figure B.5 illustrates our finding that the Bet Less Bonus did not cause increases in untracked forms of gambling, including illegal gambling. Two important caveats are that our measurements of untracked gambling are from unincentivized self-reports and that they measure short-run rather than long-run elasticities.

[30] We also note that the literature on hypothetical choice generally finds that people *overstate* sensitivity to salient changes in conditions, so standard hypothetical response bias cannot explain the direction of our effect (Bernheim et al., 2021).

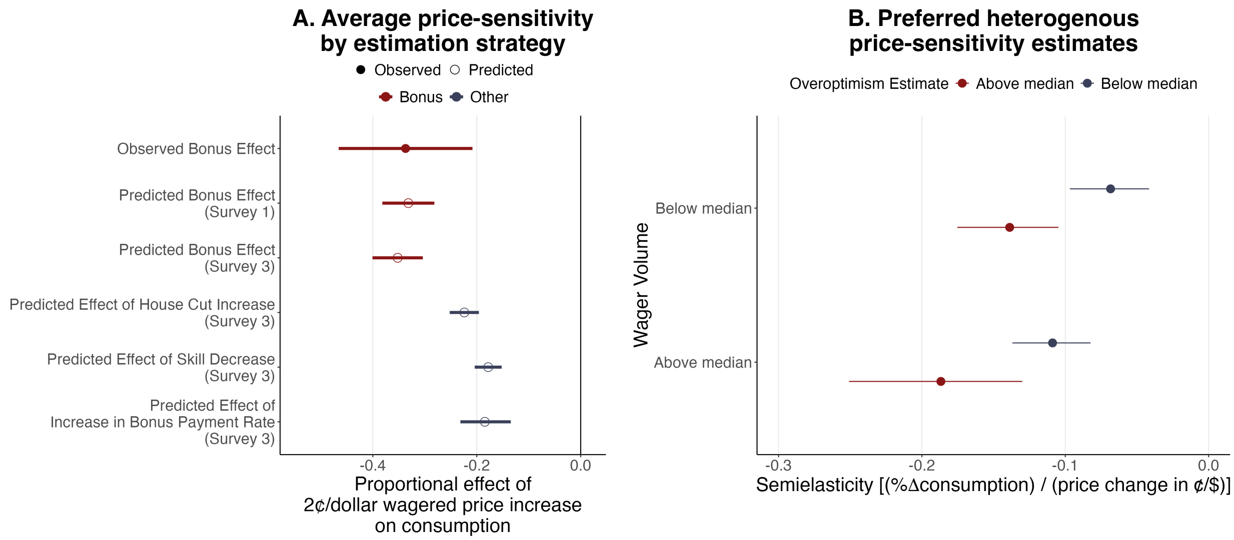

Figure 8: Evidence on price-sensitivity

Notes: Panel A plots several estimates of how a 2¢/$ price increases would affect consumption. Estimates are reported as proportional changes. The first row is the estimated Bonus treatment effect. The second and third row are predicted Bonus treatment effects, elicited on surveys 1 and 3 respectively. The fourth and fifth row are predicted effects of hypothetical naturalistic price changes. The fourth row is the effect of a hypothetical 2¢/dollar increase in the house cut, and the fifth row is the effect of the participant hypothetically learning that their own expected returns were 2¢/dollar worse than they had previously thought. The estimate in the sixth row comes from predicted responses to hypothetical Bonus treatments of various sizes. Panel B plots our preferred price-sensitivity estimates, which are estimated from the predicted effect of a hypothetical change in the house cut. Points represent group-specific semielasticity estimates. Units are proportional effects of a 1¢ price change on consumption. Our subgroups are defined by splitting the sample above and below median for dollars wagered and estimated overoptimism. Error bars represent 95% confidence intervals.

ume and of the individual-specific overoptimism estimate. All estimates are reported as semielasticities (proportional effects of a 1¢/dollar price increase). We find that high-volume bettors have larger proportional consumption responses to price changes, and that high overoptimism is associated with larger responses, conditional on the wager volume subgroup. This second result is good news for the targeting efficiency of price-based policies: all else equal, more biased bettors respond more to price changes. [31]

6.4 Experimental Evidence on Bias-correction Interventions

Heterogeneity in overoptimism and self-control problems provide a strong theoretical justification for interventions that correct those biases. In principle, such interventions could asymmetrically reduce the consumption of the most biased bettors while preserving entertainment value for unbiased consumers (Camerer et al., 2003). In practice, policy-

[31] Given our modest sample size, the error bars in Figure 8 are reasonably tight. This precision highlights an important benefit of using hypothetical price responses: they improve our power to detect heterogeneous price-sensitivity by letting us observe choices under different price levels within each subject.

makers and sportsbooks have advocated for and implemented numerous bias-correcting interventions. We analyze two of these interventions, historical winnings transparency and voluntary wager limits. These interventions do not achieve their stated goals, in the sense that their effects are small in magnitude and imperfectly targeted at heterogeneity in bias.

Addressing Overoptimism: History Transparency People respond to information about their own past returns by updating their beliefs in the direction of the signal. Panel A of Figure 9 shows that people who received good (bad) news about past returns updated positively (negatively) about future returns. The slope of the best fit line implies that every 1¢ of good news about past returns increased beliefs about future returns by 0.089¢. This result demonstrates that beliefs are at least somewhat malleable, and that it could be possible to correct overoptimism by sending people the right signal.

Unfortunately, information about past returns is not the right signal: while the history transparency treatment affected beliefs, it did not do so *in a way that reduced bias.* Two results allow us to draw this conclusion.

Our first result is that the average treatment effect did not counteract average bias. Instead, the treatment induced *positive* (though statistically indistinguishable from zero) updates on average. Even though the treatment had no average effect, it could still have had well-targeted treatment effects – that is, it could have reduced predicted future returns more for the overoptimistic bettors. Our second result is that this was not true: treatment effects were not well-targeted at bias. Panels B and C of Figure 9 show that bettors with high overoptimism according to two measures (raw prediction errors and shrunk overoptimism estimates, respectively) do not update more negatively. To understand these result, we examine the backward-looking errors about past returns. For the treatment effect to be well-targeted, more overoptimistic bettors should receive worse news. Panel D of Figure 9 shows that this is not the case: over-estimation of future returns is not correlated with over-estimation of past returns.

These findings suggest that interventions aiming to address forward-looking overoptimism by correcting backward-looking biased beliefs in the sports betting context face fundamental challenges. We view the study of alternative approaches to overoptimism correction as an important area for future research on sports betting.

Addressing Self-control Problems: Voluntary Wager Limits We find that voluntary binding wager limit tools only partially address self-control problems, even in our experimental setting where all participants make active choices about limits. In the left panel of Figure 10, we show how the chosen limits reduced consumption relative to participants' stated ideals. While the average participant stated that they would like to re-

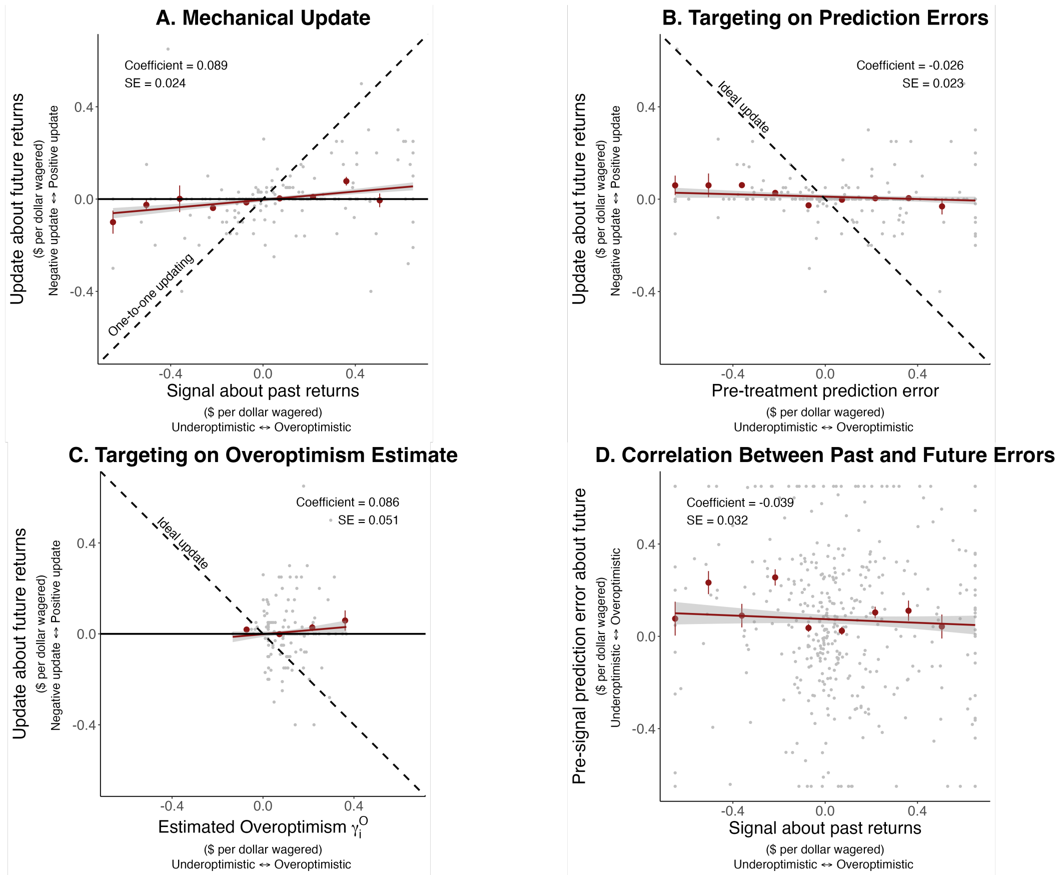

Figure 9: The effects of history transparency on misperceptions

Notes: The figure summarizes the effects of our History Transparency Treatment on beliefs. Panel A shows how updates correlate with the observed signal about past returns. The update variable is equal to the future prediction minus the initial prediction, or zero if the agent chose not to change the prediction. The Signal variable is equal to reported past returns minus the true past returns, truncated to lie within $[-0.65, 0.65]$ (since this is also the domain of the update variable). Panels B and C show how updates depend on two measures of misperceptions, the raw prediction errors in panel B and our shrunk misperception estimates in panel C. If the intervention were well-targeted, participants would update along the dotted line (so that overoptimistic bettors become less overoptimistic). We do not observe this pattern. Panel D shows that the signal about past returns was not correlated with pre-signal prediction errors. Points correspond to treated participant in panels A through C, while panel D includes a point for both treated and control participants to improve precision. Lines, coefficients, and standard errors all come from univariate regressions.

duce consumption by about 20%, they only set limits that would reduce consumption by 3%.[32] For robustness, we also compare the chosen limits to the counterfactual ideal of setting self-control problems γ^{SC} to zero. This exercise implies that limit tools eliminate less than half of the overconsumption from self-control problems in our sample.

Unlike history transparency, voluntary limits are well-targeted at bias. The right panel of Figure 10 shows that people without self-control problems almost never set binding

[32]Such low take-up is consistent with qualitative evidence that people do not express interest in using limit tools (Appendix Figure B.6).

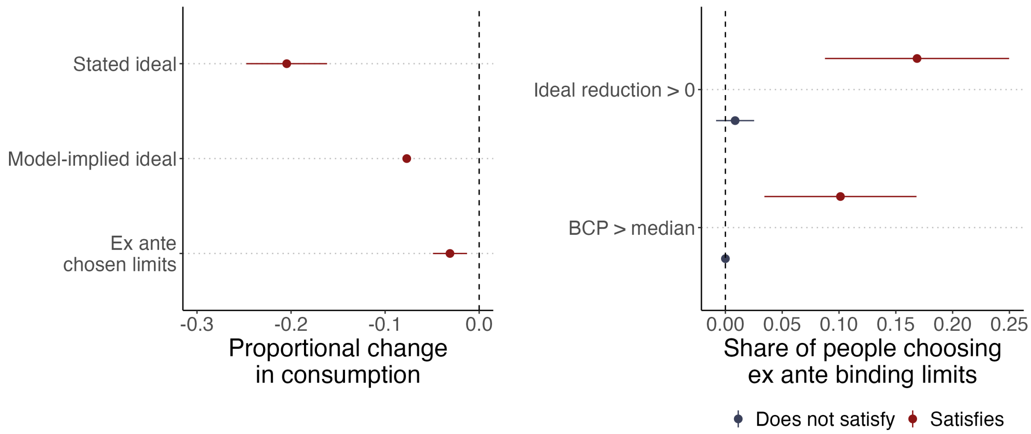

Figure 10: Limit choices

Notes: This figure summarizes limit take-up among the subgroup of bettors in the limits treatment who did not have difficulty finding limit screens on their tracked apps (N=200). The left panel shows three average proportional consumption reductions. The stated ideal reduction and ex ante chosen limits reduction are computed by dividing stated ideal consumption and the sum of all wager limits by predicted dollars wagered, respectively. The model-implied ideal is the model-implied reduction in consumption from eliminating self-control problems for the average partipant. It is equal to $E[\tilde{\gamma}^{SC}] \cdot E[\eta_i] = 0.007 \cdot 11\% = 7.7\%$. The right panel shows how the share of participants setting binding limits varies according to two measures of self-control problems. Ideal reduction > 0 is the subset of people whose stated ideal consumption was less than predicted consumption, and BCP $>$ median is the subset with an above-median behavior change premium.

limits, so limit tools only reduce consumption for biased bettors. Because of this targeting, limit tools could improve welfare even when taxes are set optimally (Allcott et al., 2025).

Why do some agents who perceive themselves to have self-control problems not set binding limits? One answer, consistent with the theoretical discussion in Laibson (2015), is that uncertainty about the future demand for betting causes people to avoid making binding commitments even when they want to reduce future consumption for any given demand realization. Qualitative evidence shows that many bettors are uncertain about the ideal amount to bet in the future (Appendix Figure B.7) and that this uncertainty played a role in limit decisions (Appendix Figure B.8). Ultimately, our small sample size and the paucity of self-control problems in the sample make it difficult for us to draw firm conclusions about the determinants of limit take-up. Our results point to the demand for flexibility as a potentially important and underappreciated barrier.

7 Model Estimation

We now combine our empirical evidence with our model to illustrate the implications of our results for policy analysis. We use our data to estimate a specialized version of the model and use the estimated model to simulate sports betting consumption under various counterfactual policies. This section estimates the structural parameters of our model; Section 8 reports on the simulations.

The main specialization is a constant semielasticity of demand functional form.[33] To impose constant semielasticity, we assume that the functional form of nonfinancial utility is $z_i(x; \tilde{F}) = z_{1i} x \log(x) + z_{2i} x + g_i(\tilde{F})x + h_i(\tilde{F})$. The demand curve is then defined as follows. Normalize the status quo average price of betting to τ_0. Let $x_i^*(\tau_0)$ denote the normative consumption at this price. We consider consumption as a function of bias and other unit prices τ_i. We show in Appendix C.1 that, under our functional form assumption, demand is characterized by the following equation:

$$\log(x(\tau_i)) = \log\left(\underbrace{x_i^*(\tau_0)}_{\text{Normative consumption at status quo}}\right) - \underbrace{\eta_i}_{\text{Semielasticity}} \cdot \left(\underbrace{(\tau_i - \gamma_i^O - \gamma_i^{SC}) - \tau^0}_{\text{As-if price change}}\right) \tag{3}$$

For estimation with our panel data, we modify equation (3) by imposing structure on log normative consumption: $\log(x^*(\tau_0)_{it}) = \xi_i + \delta_t + \varepsilon_{it}$. The parameters ξ_i and δ_t represent the individual-specific taste for betting and seasonal trends in normative demand respectively, and ε_{it} are i.i.d. normative demand shocks.

Our estimates of bias parameters and price-sensitivity come directly from the reduced-form experimental results already presented. Specifically, the statistics illustrated in Panel B of Figure 6, Panel B of Figure 7, and Panel B of Figure 8 are our estimates of overoptimism γ_i^O, self-control problems γ_i^{SC}, and semielasticity η_i, respectively.

Given these estimates, we can now estimate the individual and time-specific components of normative consumption ξ_i and δ_t. We use the following estimating equation:[34]

$$E\left[\frac{x_{it}}{\exp\left(\eta_i \cdot (\tau^0 + \gamma_i^O + \gamma_i^{SC} - \tau_i)\right)}\right] = \exp(\xi_i + \delta_t) \tag{4}$$

The left side of equation (4) represents residualized consumption – observed demand adjusted by the effects of bias and prices. We construct a panel of residualized consumption over the three pre-study periods $t = -2, -1, 0$, using observed consumption and our estimates of γ_i^O, γ_i^{SC}, and η_i. Then, we use a poisson fixed effects regression to estimate ξ_i and δ_t for this panel, normalizing $\delta_{-1} = 0$.

This residualization procedure clarifies our key assumption: all consumption that is not driven by estimated overoptimism and self-control problems must be driven by preferences. An obvious but important implication is that if other unmeasured biases drive consumption, our model will treat those biases as preferences and overstate the

[33] This functional form is motivated by the shape of the predicted consumption demand curve (Appendix Figure B.9).

[34] To derive equation (4), we exponentiate both sides of (3), divide by the effect of the as-if price change, and take the expectation of both sides over the normative demand shocks.

normative utility from betting. More broadly, misspecification would generally cause us to attribute consumption incorrectly to bias versus preferences.[35]

8 Quantitative Policy Analysis

8.1 A Planner's Problem With Heterogeneous Consumer Bias

We introduce a simple social planner's problem to structure our analysis of corrective interventions. The planner can set sports betting taxes and recycle revenues into lump-sum transfers that are distributed equally to consumers. Sportsbooks produce wagers at constant marginal cost in a competitive market. Therefore, unit taxes pass through fully to consumer prices, and we assume that consumers perceive any tax-induced price changes. The planner places equal welfare weights on all consumers, so under these assumptions her problem is to maximize average money-metric normative utility subject to a balanced budget constraint:

$$\max_{\tau_i} \sum_i E_i[u_i^{normative}(x_i(\tau_i))] + R(\tau) \qquad \text{s.t.} \quad R(\tau) = E_i[\tau_i x_i(\tau_i)] \tag{5}$$

where τ_i denotes the unit tax on consumer i's betting.

We begin with a first-best benchmark. The optimal personalized tax for agent i is the rate that internalizes both biases: $\tau_i^* = \gamma_i^O + \gamma_i^{SC}$. Under such a regime, every agent consumes the normatively optimal level of gambling, and total surplus is maximized.

We next turn to the optimal uniform tax, which is the solution to equation (5) under the constraint that $\tau_i = \tau$ for all i. The optimal uniform tax is a weighted average of individual biases, with weights proportional to the slope of the demand curve (Diamond, 1973; Allcott and Taubinsky, 2015). Specifically, the optimal tax must satisfy $\tau^* = E[w_i \tau_i^*]$, where the demand weights are specified as $w_i = \frac{\eta_i x_i(\tau^*)}{E_i[\eta_i x_i(\tau^*)]}$. Intuitively, the planner trades off the desire to levy high taxes on biased bettors while allowing unbiased bettors to continue consuming. The optimal uniform tax falls short of the first-best benchmark to the extent that this tradeoff bites.

While a uniform tax cannot differentially target biased consumers, bias-correcting nudges could circumvent this tradeoff by asymmetrically reducing the consumption of biased bettors (Camerer et al., 2003; Allcott et al., 2025). To quantify these potential targeting benefits, we simulate interventions that reduce a bias by some fraction. For example, to simulate a nudge that reduces overoptimism by 10%, we set each individual's overoptimism parameter to $\gamma_i^O \cdot 0.9$. This stylized intervention represents a perfectly targeted

[35] For example, we restrict heterogeneity in self-control problems to be across two subgroups.If unobserved heterogeneity in bias drives variance in consumption, our model will incorrectly attribute this variance to heterogeneous preferences.

nudge. As we illustrated in Section 6.4, many real-life nudges fall short of this ideal.

Discussion of assumptions We make restrictive assumptions for tractability. One assumption that is likely violated is perfect competition.[36] A second is that taxes may be imperfectly salient in the sense of Chetty et al. (2009): since the price of betting is framed as betting odds, which consumers may misunderstand, taxes may not pass through fully to perceived prices. Extending the model along these lines would require additional evidence on firm conduct or consumer perceptions, but our analysis of the targeting efficiency of taxes and nudges would still be relevant for understanding these richer models.

We also emphasize that our planner has a particular normative objective: maximize total surplus without distributional weights. While our empirical bias estimates would be key inputs for policy analysis under a variety of social welfare functions, planners with different objectives would of course select different optimal taxes.[37] When analyzing nudges, we also make the additional normative assumption that the intervention does not change the normative utility function. For example, an intervention that reduces overoptimism does not reduce the *nonfinancial* utility of gambling.[38]

8.2 Counterfactual Results

Table 3 summarizes the results of our counterfactual simulations.

We begin by describing the impacts of first-best personalized taxes. Column (1) shows that aggregate betting falls by 39% relative to a benchmark without taxes. One interpretation is that in our sample, 39% of consumption is driven by bias. Total surplus increases by \$25.1 per consumer-week (Column (5)). We think of this effect as an upper bound on the potential benefits of corrective policy. Columns (2), (3), and (4) then report the impacts of this policy on the components of total surplus. Taxes necessarily reduce *decision consumer surplus*, which is consumer surplus (CS) evaluated according to the decision utility function (1), because taxes make consumption more expensive. But *normative consumer surplus*, which is CS evaluated according to the normative utility function (2), decreases by much less than decision consumer surplus, because personalized taxes reduce the harms from overconsumption of sports betting.

We use a back-of-the-envelope calculation to extrapolate our local estimate of total

[36] The U.S. sports betting market is dominated by two firms, DraftKings and Fanduel. See O'Connell and Smith (2024) for an analysis of sin taxes under market power.

[37] A previous version of this paper formalized how a planner would set optimal uniform taxes under generalized versions of the model with less restrictive normative and positive assumptions (Brown et al., 2025).

[38] In our microfoundation, this assumption amounts to the restriction that the $g_i(\tilde{F})$ and $h_i(\tilde{F})$ terms in nonfinancial utility are both constant. See Glaeser (2006) for an early discussion of how interventions might impact normative utility and Butera et al. (2022) for empirical evidence in a different context.

	Policy	(1) Δ% Cons.	(2) Δ Decision CS ($/cons-wk)	(3) Δ Norm. CS ($/cons-wk)	(4) Δ Revenue ($/cons-wk)	(5) Δ Tot. Surplus ($/cons-wk)	(6) % of First-Best
	Zero tax (baseline)	0%	0.0	0.0	0.0	0.0	—
Taxes							
(1)	First-best ($\tau_i = \gamma_i$)	-39%	-43.0	-3.2	28.2	25.1	100%
(2)	Status quo ($\tau = 2.02\%$)	-21%	-25.6	-9.0	22.7	13.6	54%
(3)	Optimal uniform ($\tau = 4.97\%$)	-45%	-53.6	-20.1	39.4	19.3	77%
Nudges							
(4)	Full bias correction	-39%	—	25.1	0.0	25.1	100%
(5)	Full overoptimism correction	-34%	—	24.9	0.0	24.9	99%
(6)	10% overoptimism correction	-5%	—	7.1	0.0	7.1	28%
(7)	Full self-control correction	-7%	—	5.6	0.0	5.6	22%

Table 3: Counterfactual Results

Notes: This table reports on simulations of policies defined in Section 8.1. We simulate taxes (rows 1 through 3) and hypothetical bias-correcting nudges (rows 4 through 7). All policies are evaluated relative to a no regulation benchmark. We report six statistics. Column (1) is the percentage change in sports betting consumption. Columns (2) and (3) are weekly changes in decision consumer surplus and normative consumer surplus, as defined in section 8.2. We do not report a change in decision consumer surplus for nudges, because it is not obvious whether to use the original bias parameters or the nudge-affected bias parameters to compute decision consumer surplus. Column (4) is the weekly revenue generated from gambling taxes. Column (5) is the change in total surplus, defined as the sum of (3) and (4). Column (6) is the change in total surplus as a percentage of the total surplus improvement from a first-best personalized policy. All simulations fix the time-specific demand for betting at the level from period -1 (February 8 to March 9), so weekly surplus magnitudes apply to weeks in that time period.

surplus benefits for the experimental sample to the broader market. The first-best policy delivers 1.77¢ of surplus benefits per dollar wagered in the experimental sample. If we assume that a first-best policy would also deliver 1.77¢ of surplus benefits per dollar wagered in the U.S. sports betting market, that would imply an aggregate surplus gain of $2.63 billion per year.[39] We view this statistic as a rough upper bound for the benefits of corrective policy compared to a no-tax benchmark.[40] As a point of comparison, Allcott et al. (2019) estimate that an optimal sugar-sweetened beverages tax delivers welfare gains between $2.4 billion and $6.8 billion (depending on the specification) relative to a status quo with no taxes.[41]

Personalized taxes are not implementable, so we turn next to the feasible alternative: uniform taxes. Table 3 reports on the effects of status quo sports betting taxes and optimal uniform sports betting taxes in rows (2) and (3) respectively. There are two key takeaways. The first is that surplus-maximizing tax rates are higher than current U.S. tax rates. Our

[39] In the simulation period, the average wager volume under status quo taxes is $1421 per person-week, so the first-best policy delivers 25.1/1421 = 1.77 of total surplus gains for every dollar wagered. Multiplying this number by the $149 billion wagered by Americans in 2024 yields a total surplus gain of $2.63 billion.

[40] Recall that our sample is intentionally non-representative; we oversampled high-volume sports bettors. Whether this extrapolation overestimates or underestimates potential surplus gains depends in part on whether casual bettors are more or less biased than high-volume bettors.

[41] The Allcott et al. (2019) calculation includes both externality and internality benefits, while ours includes only internality benefits.

estimates imply an optimal uniform tax of 4.97%, which is more than twice as large as the 2023 status quo average rate of 2.02%.[42] Other jurisdictions, including Germany, have taxes that are closer to our estimated surplus-maximizing rate.[43] The second is that uniform policies, like all policies that do not target heterogeneity in bias, leave surplus gains on the table. Moving from the status quo tax (row (2)) to the optimal uniform tax (row (3)) delivers $5.7 per consumer-week of total surplus benefits, but this is less than half the distance from the status quo uniform tax to the first-best benchmark (row (1)).

One way to target heterogeneous bias is to use bias-correcting nudges. Rows (4) through (7) simulate stylized nudges. A nudge that fully eliminates bias delivers the same outcomes as first-best personalized taxes, albeit with different distributional consequences.[44] Reducing overoptimism is more valuable than reducing self-control problems, since overoptimism is quantitatively larger. Row (6) illustrates that even a modest reduction in overoptimism can deliver significant benefits. Overall, these results illustrate the social value of designing improved bias-correction interventions that overcome the challenges described in Section 6.4.

Robustness We summarize the implications of alternative approaches to estimation in Table 4. Estimation strategies that impose homogeneous bias greatly increase our optimal tax estimate. With homogeneous bias, we would (obviously) have concluded that a uniform tax was first best. These results demonstrate the empirical importance of carefully attending to heterogeneity in bias. By contrast, our results are stable when we use different price-sensitivity estimates. The optimal tax rate remains in the relatively narrow range of $[4.64\%, 4.97\%]$. Specifications that use steeper demand slopes (such as row (7), which uses the observed effect of the Bet Less Bonus) do have larger consumption responses to taxes and, therefore, larger welfare effects of taxes (Harberger, 1964).

8.3 Extensions

We now use our machinery to study the regulation of parlays and bans on sports betting. These extensions rely on additional assumptions and evidence.

Parlay-specific taxes One way for a planner to target heterogeneity in bias is to con-

[42] The optimal uniform rate is lower than the sum of the unweighted average bias estimates from Section 6 (7.8% + 0.7% = 8.5%). The reason is that high-volume bettors are less overoptimistic on average (Panel A of Figure 6), and high-volume bettors are weighted more heavily in the demand-slope weighted average of biases. This shades the optimal rate down.

[43] Germany imposes a 5% tax on all dollars wagered. Some high-tax states within the U.S. with higher taxes, like New York State, also come close to our estimate.

[44] One interesting takeaway from our simulations is that the surplus-maximizing tax would also increase revenues. Much of the public debate around sports betting focuses on the impacts of tax revenues for state governments. In many settings, policies that raise revenues have efficiency costs. Table 3 illustrates that raising revenues through increased sports betting taxes can have efficiency *benefits*.

	Specification			Uniform tax		First-best
	(1)	(2)	(3)	(4)	(5)	(6)
	Identification of η	Heterogeneity	Optimal rate (%)	Consumption reduction (%)	Total surplus gains ($/consumer-week)	Total surplus gains ($/consumer-week)
(1)	Change in house cut	Bias and Semielasticity	4.97%	45%	$19.3	$25.1
(2)	Change in house cut	None (Homogeneous)	9.12%	69%	$55.1	$55.1
(3)	Change in house cut	Bias only	4.59%	44%	$15.9	$21.8
(4)	Change in Bonus rate	Bias and Semielasticity	4.96%	40%	$16.4	$21.9
(5)	Change in Bonus rate	None (Homogeneous)	9.12%	60%	$44.9	$44.9
(6)	Change in Bonus rate	Bias only	4.59%	37%	$12.7	$18.0
(7)	Effect of Bet Less Bonus	Bias and Semielasticity	4.80%	53%	$23.2	$29.8
(8)	Effect of Bet Less Bonus	None (Homogeneous)	9.12%	85%	$87.0	$87.0
(9)	Effect of Bet Less Bonus	Bias only	4.59%	61%	$25.8	$33.3

Table 4: Sensitivity of corrective tax analysis to estimation strategy

Notes: This table illustrates how our analysis depends on whether we allow for heterogeneous parameters and on our estimation strategy for the semielasticity term η. Columns (1) and (2) describe the specification of interest. Columns (3), (4), and (5) report the implied optimal uniform corrective tax on dollars wagered, the effect of that tax on dollars wagered in percent terms, and the total surplus gains in dollars per consumer-week. Column (6) shows the total surplus gains from a first-best personalized tax. Column (1) varies the identification strategy for the semielasticity η, using the three approaches outlined in Section 6.3 and visualized in Panel A of Figure 8. For the "Effect of Bet Less Bonus" estimation strategy, we use the predicted bonus effect in survey 1. Column (2) varies our approach to heterogeneity. Row (1) is the specification that we use in the rest of the paper.

dition regulations on observable characteristics of wagers. Given our empirical evidence that parlay bettors are much more overoptimistic than others, we consider restrictions on parlay betting. To illustrate the impacts of parlay-specific regulation, we study a planner who can levy two uniform taxes, one on parlay wagers and one on non-parlay wagers. We view this exercise as informative about the benefits of both literal parlay taxes and other targeted regulations such as advertising restrictions or bans on parlays.[45]

To evaluate parlay-specific taxes, we must specify two empirical objects that are not required for the rest of the paper. The first is the substitutability between parlay and non-parlay betting. The second is the nature of the association between parlay betting and overoptimism (Panel B of Figure 6) – that is, whether it results from generally overoptimistic individuals sorting into parlay betting or from true effects of parlay betting on overoptimism. Since our experiment was not designed to estimate these objects, we consider multiple plausible empirical scenarios. We detail our assumptions and analytical approach, which extends Jacobsen et al. (2020)'s method for quantifying the benefits of conditioning taxes on product characteristics to our setting, in Appendix C.2.

We find that it is generically optimal to levy higher taxes on parlays than on other

[45]These kinds of interventions have been considered in the United States. For advertising restrictions, the state of Massachusetts prohibits sports betting advertisements that "would reasonably be expected to confuse and mislead patrons." For banning specific kinds of wagers, the state of Ohio banned player-specific prop-bets on NCAA contests in 2024. The SAFE Bet Act, which was introduced in Congress in 2024, would impose federal advertising restrictions and make the ban on player-specific prop-bets national.

wagers (Appendix Table C.1). Even if the association between parlay betting and bias is driven by sorting rather than a causal effect of parlays on bias, high taxes on parlays are still a way for the planner to target biased consumers. If parlays do cause overoptimism, then the optimal additional tax on parlays is even higher and the targeting benefits from parlay-specific taxes are large. Specifically, when we attribute the full association between parlays and overoptimism to a causal effect, the optimal tax on parlays is more than three times higher than the optimal tax on non-parlays. Under this assumption, taxing parlays and non-parlays separately captures about half of the potential surplus gains that are "left on the table" by uniform taxes.

Bans While an outright ban on sports betting is a blunt instrument, it is evidently politically feasible: sports betting is still illegal in many states, including California and Texas. To evaluate the welfare effects of bans, we incorporate additional evidence on the willingness to accept (WTA) to completely stop betting on tracked apps for 30 days (Brynjolfsson et al., 2019). We treat this WTA as a measure of the total perceived net benefits from legal sports betting. On average, this measure of perceived net benefits exceeds the internality from overoptimism (Appendix Figure B.10), which suggests that outright bans would reduce welfare. We provide details on our incentivized elicitation in Appendix A.3 and present detailed results and caveats in Appendix B.5.

9 Conclusion

Policymakers around the world are engaged in debates about whether and how to regulate online sports betting. This paper documents that sports betting is partially driven by behavioral biases rather than a normative preference to gamble. By connecting observational data from sportsbook accounts to a survey panel, we show that high volume sports bettors do not internalize the average financial costs of betting and are willing to pay to reduce their own future betting. A Pigouvian tax to correct these biases would be larger than status quo sports betting taxes in the U.S.

This study has several limitations; here are five. First, our positive analysis is narrowly focused on overoptimism and self-control problems in our sample of high-volume bettors. To the extent that other externalities or internalities distort sports betting consumption, in either direction, these distortions would need to be incorporated into policy analysis alongside our bias estimates. Second, our study was not designed to measure the harms of severe problem gambling, which are concentrated in a relatively small group of consumers. Third, while we find little evidence of substitution between legal and illegal sports betting during our experiment, we do not rule out substitution to illegal markets as an unintended consequence of regulation over longer time horizons. Fourth, we do not claim that our experimental sample is representative. Within our sample, lower-volume bettors are more overoptimistic than their counterparts, and survey measures indicate that our experimental participants are less biased than a representative sample of bettors. Both of these results suggest that bias estimates in our sample may be conservative for the population. Finally, in our normative analysis, we emphasize that optimal corrective policy in the sports betting market depends on parameters that we do not estimate, including market power in sports betting and the salience of sports betting taxes. Though the price-metric biases that we measure are key statistics for optimal policy under a variety of social welfare functions, a richer welfare analysis would incorporate evidence on these additional parameters.

Despite these limitations, our results have implications for the design of sports betting policy. Our results on the heterogeneity of overoptimism highlight the benefits of targeting: instruments that reduce the consumption of biased consumers while allowing unbiased consumers to enjoy gambling are more efficient than blunt instruments like uniform taxes. When evaluating policies like responsible gaming nudges, advertising regulations, or restrictions on particular kinds of bets, regulators should consider the extent to which policy-induced consumption reductions are likely to be concentrated among biased consumers. One useful next step for future research would be to empirically analyze the targeting of these and other realistic interventions.

References

Acland, D. and Levy, M. R. (2015). Naiveté, Projection Bias, and Habit Formation in Gym Attendance. *Management Science*, 61(1):146–160.

Akee, R. K. Q., Spilde, K. A., and Taylor, J. B. (2015). The Indian Gaming Regulatory Act and Its Effects on American Indian Economic Development. *Journal of Economic Perspectives*, 29(3):185–208.

Allcott, H., Braghieri, L., Eichmeyer, S., and Gentzkow, M. (2020). The Welfare Effects of Social Media. *American Economic Review*, 110(3):629–676.

Allcott, H., Cohen, D., Morrison, W., and Taubinsky, D. (2025). When Do "Nudges" Increase Welfare? *American Economic Review*, 115(5):1555–1596.

Allcott, H., Gentzkow, M., and Song, L. (2022a). Digital Addiction. *American Economic Review*, 112(7):2424–2463.

Allcott, H., Kim, J., Taubinsky, D., and Zinman, J. (2022b). Are High-Interest Loans Predatory? Theory and Evidence from Payday Lending. *The Review of Economic Studies*, 89(3):1041–1084.

Allcott, H., Lockwood, B. B., and Taubinsky, D. (2019). Regressive Sin Taxes, with an Application to the Optimal Soda Tax. *The Quarterly Journal of Economics*, 134(3):1557–1626.

Allcott, H. and Taubinsky, D. (2015). Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market. *American Economic Review*, 105(8):2501–2538.

Ambuehl, S., Bernheim, B. D., and Lusardi, A. (2022). Evaluating Deliberative Competence: A Simple Method with an Application to Financial Choice. *American Economic Review*, 112(11):3584–3626.

American Gaming Association (2024). State of the states 2024: The aga commercial casino industry report. Technical report, American Gaming Association, Washington, DC.

American Gaming Association (2025a). Commercial gaming revenue tracker. https://www.americangaming.org/resources/commercial-gaming-revenue-tracker/. Archived 6 March 2025 at https://web.archive.org/web/20250306024649/https://www.americangaming.org/resources/commercial-gaming-revenue-tracker/. Accessed 11 April 2026.

American Gaming Association (2025b). State of the states 2025: The aga commercial casino industry report. Technical report, American Gaming Association, Washington, DC.

American Gaming Association, T. (2024). CY 2023 AGA Commercial Gaming Revenue Tracker.

American Psychiatric Association (2013). *Diagnostic and Statistical Manual of Mental Disorders*. Fifth edition edition.

Association, A. G. (2024). Supporting the Legal Sports Betting Market - American Gaming Association. https://www.americangaming.org/advocacy/supporting-the-legal-sports-betting-market/.

Auer, M. and Griffiths, M. D. (2022). Predicting Limit-Setting Behavior of Gamblers Using Machine Learning Algorithms: A Real-World Study of Norwegian Gamblers Using Account Data. *International Journal of Mental Health and Addiction*, 20(2):771–788.

Augenblick, N., Niederle, M., and Sprenger, C. (2015). Working over Time: Dynamic Inconsistency in Real Effort Tasks. *The Quarterly Journal of Economics*, 130(3):1067–1115.

Augenblick, N. and Rabin, M. (2019). An Experiment on Time Preference and Misprediction in Unpleasant Tasks. *The Review of Economic Studies*, 86(3):941–975.

Baker, S. R., Balthrop, J., Johnson, M. J., Kotter, J. D., and Pisciotta, K. (2024). Gambling Away Stability: Sports Betting's Impact on Vulnerable Households.

Banerjee, A. and Mullainathan, S. (2010). The Shape of Temptation: Implications for the Economic Lives of the Poor.

Becker, G. M., Degroot, M. H., and Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3):226–232.

Bénabou, R. and Tirole, J. (2002). Self-Confidence and Personal Motivation*. *The Quarterly Journal of Economics*, 117(3):871–915.

Bernheim, B. D. (2016). The Good, the Bad, and the Ugly: A Unified Approach to Behavioral Welfare Economics. *Journal of Benefit-Cost Analysis*, 7(1):12–68.

Bernheim, B. D., Björkegren, D., Naecker, J., and Pollmann, M. (2021). Causal Inference from Hypothetical Evaluations.

Bernheim, B. D. and Taubinsky, D. (2018). Chapter 5 - Behavioral Public Economics. In Bernheim, B. D., DellaVigna, S., and Laibson, D., editors, *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 1 of *Handbook of Behavioral Economics -*

Foundations and Applications, pages 381–516. North-Holland.
Braverman, J., Tom, M. A., and Shaffer, H. J. (2014). Accuracy of self-reported versus actual online gambling wins and losses. *Psychological Assessment*, 26(3):865–877.
Brown, M., Grasley, N., and Guido, M. (2025). Do Sports Bettors Need Consumer Protection? Evidence from a Field Experiment Link to latest version.
Brynjolfsson, E., Collis, A., and Eggers, F. (2019). Using massive online choice experiments to measure changes in well-being. *Proceedings of the National Academy of Sciences*, 116(15):7250–7255.
Butera, L., Metcalfe, R., Morrison, W., and Taubinsky, D. (2022). Measuring the Welfare Effects of Shame and Pride. *American Economic Review*, 112(1):122–168.
Calvosa, P. (2017). Responsible Gambling Strategies for Internet Gambling: An Empirical Investigation into the Italian Gambling Market. *International Journal of Business and Management*, 12(7):p17.
Camerer, C., Issacharoff, S., Loewenstein, G., O'Donoghue, T., and Rabin, M. (2003). Regulation for Conservatives: Behavioral Economics and the Case for "Asymmetric Paternalism". *University of Pennsylvania Law Review*, 151(3):1211–1254.
Carrera, M., Royer, H., Stehr, M., Sydnor, J., and Taubinsky, D. (2022). Who Chooses Commitment? Evidence and Welfare Implications. *The Review of Economic Studies*, 89(3):1205–1244.
Chaloupka, F., Levy, M., and White, J. (2019). Estimating Biases in Smoking Cessation: Evidence from a Field Experiment. Technical Report w26522, National Bureau of Economic Research, Cambridge, MA.
Chater, N. and Loewenstein, G. (2023). The i-Frame and the s-Frame: How Focusing on Individual-Level Solutions Has Led Behavioral Public Policy Astray. *Behavioral and Brain Sciences*, 46(4046264).
Chegere, M., Falco, P., Nieddiu, M., Pandolfi, L., and Stein, M. (2022). It's a Sure Win! Experimental evidence on overconfidence in betting behavior. *CSEF Working Papers*, (655).
Chen, J. and Roth, J. (2024). Logs with Zeros? Some Problems and Solutions*. *The Quarterly Journal of Economics*, 139(2):891–936.
Chen, K. (2024). Empirical Bayes When Estimation Precision Predicts Parameters.
Chetty, R. (2015). Behavioral Economics and Public Policy: A Pragmatic Perspective. *American Economic Review*, 105(5):1–33.
Chetty, R., Looney, A., and Kroft, K. (2009). Salience and Taxation: Theory and Evidence. *American Economic Review*, 99(4):1145–1177.
Commonwealth of Australia (2022). National Consumer Protection Framework for Online Wagering in Australia – National Policy Statement. Technical report, Commonwealth of Australia and State/Territory Governments.
Connolly, A. J., Stall, L. M., Floyd, C. G., Kraus, S. W., and Grubbs, J. B. (2024). Exploring the Differences in Positive Play among Various Sports Wagering Behaviors. *Journal of Gambling Studies*.
Couture, C., Cross, J., and Wu, S. (2024). Impact of Sports Gambling on Mental Health. *Economics Letters*, 243(4865678):111922.
DellaVigna, S. and Malmendier, U. (2006). Paying Not to Go to the Gym. *American Economic Review*, 96(3):694–719.

Diamond, P. A. (1973). Consumption Externalities and Imperfect Corrective Pricing. *The Bell Journal of Economics and Management Science*, 4(2):526–538.

Donkor, K., Goette, L., Müller, M. W., Dimant, E., and Kurschilgen, M. J. (2025). The Price of Identity: Overoptimism and Congruence Concerns.

Enke, B. and Shubatt, C. (2023). Quantifying Lottery Choice Complexity.

Evans, W. N. and Topoleski, J. H. (2002). The Social and Economic Impact of Native American Casinos.

Feeney, D. (2023). National Survey on Gambling Attitudes and Gambling Experiences. Technical report, National Council on Problem Gambling.

Forrest, D. and McHale, I. G. (2024). The Dependence of Online Gambling Businesses on High-Spending Customers: Quantification and Implications. *Journal of Gambling Studies*.

Gambling Commission (2025). Industry statistics: Annual report, financial year April 2024 to March 2025.

Gerstein, D., Murphy, S., Toce, M., Volberg, R., Harwood, H., Tucker, A., Christiansen, E., Cummings, W., and Sinclair, S. (1999). Gambling impact and behavior study: Report to the National Gambling Impact Study Commission.

Gillen, B., Snowberg, E., and Yariv, L. (2019). Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study. *Journal of Political Economy*, 127(4):1826–1863.

Glaeser, E. (2006). Paternalism and Psychology. *University of Chicago Law Review*, 73(1).

Goudriaan, A. E., van den Brink, W., and van Holst, R. J. (2019). Gambling Disorder and Substance-Related Disorders: Similarities and Differences. In Heinz, A., Romanczuk-Seiferth, N., and Potenza, M. N., editors, *Gambling Disorder*, pages 247–269. Springer International Publishing, Cham.

Greenberg, N. (2022). Analysis | Parlays are big business for sportsbooks — and big trouble for bettors. https://www.washingtonpost.com/sports/2022/08/31/sports-betting-parlay-odds/.

Grubbs, J. B. and Kraus, S. W. (2022). Symptoms of Problem Gambling Among US Adults Who Wager on Sports. *JAMA Network Open*, 5(10):e2239670.

Grubbs, J. B. and Kraus, S. W. (2023a). Cryptocurrency and addictive behaviors in a census-matched u.S. Sample. *International Gambling Studies*, pages No Pagination Specified–No Pagination Specified.

Grubbs, J. B. and Kraus, S. W. (2023b). The relative risks of different forms of sports betting in a U.S. sample: A brief report. *Comprehensive Psychiatry*, 127:152420.

Grubbs, J. B. and Kraus, S. W. (2023c). Sports wagering in the context of addictive disorders: Results from a census-matched U.S. sample. *Cogent Mental Health*, 2(1):2231497.

Guryan, J. and Kearney, M. S. (2008). Gambling at Lucky Stores: Empirical Evidence from State Lottery Sales. *American Economic Review*, 98(1):458–473.

Guryan, J. and Kearney, M. S. (2010). Is Lottery Gambling Addictive? *American Economic Journal: Economic Policy*, 2(3):90–110.

Haaland, I., Roth, C., and Wohlfart, J. (2023). Designing Information Provision Experiments. *Journal of Economic Literature*, 61(1):3–40.

Harberger, A. C. (1964). The Measurement of Waste. *The American Economic Review*, 54(3):58–76.

Heirene, R. M., Wang, A., and Gainsbury, S. M. (2022). Accuracy of self-reported gambling frequency and outcomes: Comparisons with account data. *Psychology of Addictive Behaviors*, 36(4):333–346.

Hollenbeck, B., Larsen, P., and Proserpio, D. (2024). The Financial Consequences of Legalized Sports Gambling.

Holtgraves, T. (2009). Evaluating the Problem Gambling Severity Index. *Journal of Gambling Studies*, 25(1):105–120.

Huffman, D., Raymond, C., and Shvets, J. (2022). Persistent Overconfidence and Biased Memory: Evidence from Managers. *American Economic Review*, 112(10):3141–3175.

Humphreys, B. and Ruseski, J. (2024). Legalized Sports Betting and Mental Health. *Economics Faculty Working Papers Series*.

Ivanova, E., Magnusson, K., and Carlbring, P. (2019). Deposit Limit Prompt in Online Gambling for Reducing Gambling Intensity: A Randomized Controlled Trial. *Frontiers in Psychology*, 10:639.

Jacobsen, M. R., Knittel, C. R., Sallee, J. M., and van Benthem, A. A. (2020). The Use of Regression Statistics to Analyze Imperfect Pricing Policies. *Journal of Political Economy*, 128(5):1826–1876.

Kahneman, D. (1965). Control of spurious association and the reliability of the controlled variable. *Psychological Bulletin*, 64(5):326–329.

Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2):263–291.

Kaur, S., Kremer, M., and Mullainathan, S. (2015). Self-Control at Work. *Journal of Political Economy*, 123(6):1227–1277.

Kearney, M. S. (2005). The Economic Winners and Losers of Legalized Gambling. *National Tax Journal*, 58(2):281–302.

Laibson, D. (2015). Why Don't Present-Biased Agents Make Commitments? *American Economic Review*, 105(5):267–272.

Levitt, S. D. (2004). Why are Gambling Markets Organised so Differently from Financial Markets? *The Economic Journal*, 114(495):223–246.

List, J. A., Rodemeier, M., Roy, S., and Sun, G. K. (2023). Judging Nudging: Understanding the Welfare Effects of Nudges Versus Taxes.

Lockwood, B., Allcott, H., Taubinsky, D., and Sial, A. (2025). What Drives Demand for State-Run Lotteries? Evidence and Welfare Implications. *Review of Economic Studies*, 92(w28975):2578–2623.

Massachusetts Gaming Commission (2023). MGC Open Meeting– November 27, 2023.

Matsuzawa, K. and Arnesen, E. (2024). Sports Betting Legalization Amplifies Emotional Cues & Intimate Partner Violence.

Möbius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2022). Managing self-confidence: Theory and experimental evidence. *Management Science*, 68(11):7793–7817.

Muggleton, N., Parpart, P., Newall, P., Leake, D., Gathergood, J., and Stewart, N. (2021). The association between gambling and financial, social and health outcomes in big financial data. *Nature Human Behaviour*, 5(3):319–326.

O'Connell, M. and Smith, K. (2024). Optimal Sin Taxation and Market Power. *American Economic Journal: Applied Economics*, 16(4):34–70.

Potenza, M. N., Balodis, I. M., Derevensky, J., Grant, J. E., Petry, N. M., Verdejo-Garcia, A., and Yip, S. W. (2019). Gambling disorder. *Nature Reviews Disease Primers*, 5(1):1–21.

Prior, M., Sood, G., and Khanna, K. (2015). You Cannot be Serious: The Impact of Accuracy Incentives on Partisan Bias in Reports of Economic Perceptions. *Quarterly Journal of Political Science*, 10(4):489–518.

Read, D. and van Leeuwen, B. (1998). Predicting Hunger: The Effects of Appetite and Delay on Choice. *Organizational Behavior and Human Decision Processes*, 76(2):189–205.

Rose-Berman, I. (2024). Essay | 'Responsible Gambling' Campaigns Always Fail. That's the Point. https://www.wsj.com/business/hospitality/responsible-gambling-campaigns-always-fail-thats-the-point-13a972de.

Ross, L. (1977). The Intuitive Psychologist And His Shortcomings: Distortions in the Attribution Process1. In Berkowitz, L., editor, *Advances in Experimental Social Psychology*, volume 10, pages 173–220. Academic Press.

Sadoff, S., Samek, A., and Sprenger, C. (2020). Dynamic Inconsistency in Food Choice: Experimental Evidence from Two Food Deserts. *The Review of Economic Studies*, 87(4):1954–1988.

Sial, A. Y., Sydnor, J. R., and Taubinsky, D. (2023). Biased Memory and Perceptions of Self-Control.

Simonetti, K. S. a. I. (2025). America Has Fallen in Love With Long-Shot Sports Bets. https://www.wsj.com/business/what-are-parlays-sports-betting-gambling-881bdaee.

Snowberg, E. and Wolfers, J. (2010). Explaining the Favorite–Long Shot Bias: Is it Risk-Love or Misperceptions? *Journal of Political Economy*, 118(4):723–746.

Stipek, D. J. and Gralinski, J. H. (1991). Gender differences in children's achievement-related beliefs and emotional responses to success and failure in mathematics. *Journal of Educational Psychology*, 83(3):361–371.

Strack, P. and Taubinsky, D. (2021). Dynamic Preference "Reversals" and Time Inconsistency.

Tamny, J. (2023). Billy Walters' Excellent 'Gambler' - Advice From the World's Greatest. https://www.forbes.com/sites/johntamny/2023/10/12/billy-walters-excellent-gambleradvice-from-the-worlds-greatest/.

Taylor, W., McCarthy, D., and Wilbur, K. C. (2024). Online Gambling Policy Effects on Tax Revenue and Irresponsible Gambling.

Thaler, M. (2024). The Fake News Effect: Experimentally Identifying Motivated Reasoning Using Trust in News. *American Economic Journal: Microeconomics*, 16(2):1–38.

Thaler, R. H. and Sunstein, C. R. (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Nudge: Improving Decisions about Health, Wealth, and Happiness. Yale University Press, New Haven, CT, US.

Vice (2022). The Rise of Mobile Gambling Is Leaving People Ruined and Unable to Quit. https://www.vice.com/en/article/the-rise-of-mobile-gambling-is-leaving-people-ruined-and-unable-to-quit/.

Wood, R. T. A., Wohl, M. J. A., Tabri, N., and Philander, K. (2017). Measuring Responsible Gambling amongst Players: Development of the Positive Play Scale. *Frontiers in Psychology*, 8.

Online Appendices

A Experimental Design Appendix

Survey instruments are available at mattbrownecon.github.io.

A.1 Timeline and recruitment details

While the vast majority of participants found the study through our advertisements, a few were recruited via a snowball design, in which current participants were asked to referrals friends who were regular mobile sports bettors. Referred participants were always assigned to the same treatment condition as the person who recruited them. We refer to the original participant and the people they recruited as a "referral group," and we always cluster standard errors by referral group. In practice only 8 members of the analysis sample (2%) came from the snowball recruitment procedure.

We undertook several procedures to minimize attrition among the 533 participants who completed survey 1. Participants received text and email reminders to take the surveys, and we delayed most payments until after survey 3 to encourage participants to return. Since we wanted to observe full bet histories, we also emailed survey 3 to all participants who completed survey 1, regardless of whether they completed survey 2. Some participants had technical problems refreshing their account data. We provided extensive support to help participants resolve such issues over email. When a participant was unable to refresh one of their accounts, we still allowed them to complete the survey.

A.2 Index definitions

Problem Gambling Severity Index The Problem Gambling Severity Index (Problem Gambling Severity Index) is a survey instrument from the psychiatric literature that is designed to provide a summary measure of problem gambling risk in non-clinical contexts (Holtgraves, 2009). In survey 1, participants report how often they experienced each of nine consequences in the past year: {*Have you bet more than you could really afford to lose?*, *Have you needed to gamble with larger amounts of money to get the same feeling of excitement?*, *When you gambled, did you go back another day to try to win back the money you had lost?*, *Have you borrowed money or sold anything to get money to gamble?*, *Have you felt that you might have a problem with gambling?*, *Has gambling ever caused you any health problems, including stress or anxiety?*, *Have people criticized your betting or told you that you had a gambling problem, regardless of whether or not you thought it was true?*, *Has your gambling caused any financial problems* *for you or your household?*, *Have you felt guilty about the way you gamble or what happens when you gamble?*}. As pre-registered, we construct the Problem Gambling Severity Index score by assigning scores from 0 to 3 to responses, and summing across the nine questions to compute the Problem Gambling Severity Index score. Among gambling policymakers, Problem Gambling Severity Index scores in the ranges of $\{0, [1, 2], [3, 7], [7, 21]\}$ are considered indicators of "no risk," "low risk," "moderate risk," and "high risk" for problem gambling, respectively.

Self-control scale We adapt a positive play behavior scale from Wood et al. (2017). The scale is designed to measure the extent to which a participant engages with sports betting in a healthy manner. We elicit agreement/disagreement with five statements: {*In the last 30 days, I felt in control of my sports betting behavior.*, *In the last 30 days, I was honest with my family and/or friends about the amount of money I spent sports betting.*, *In the last 30 days, I was honest with my family and/or friends about the amount of time I spent sports betting.*, *In the last 30 days, I only bet on sports with money that I could afford to lose.*, *In the last 30 days, I only spent time sports betting that I could afford to lose.*} The options were a five-point Likert scale ((Strongly disagree (-2), disagree (-1), neither agree nor disagree (0), agree (1), strongly agree (2))). We sum the question scores to create the self-control scale, which ranges from -12 to 12.

Gambling literacy scale We adapt a gambling literacy scale from (Wood et al., 2017). We elicit agreement/disagreement with three statements in survey 2: {*Gambling is not a good way to make money for most people.*, *My chances of winning get better after I have lost (reverse coded).*, *If I gamble more often, it will help me to win more than I lose (reverse coded).*} We use the same five-point Likert scale as above. We sum the question scores to create a gambling literacy score in the range [-6, 6].

Sports betting makes life better Following Allcott et al. (2022a), we asked in all surveys: *To what extent do you think sports betting makes your life better or worse?* Participants could respond on a scale from -5 ("makes life much worse") to 5 ("makes life much better"). The question is designed to qualitatively capture the extent to which people derive value from sports betting.

A.3 Overall consumer surplus elicitation

We elicited the WTA to stop betting entirely on tracked apps in the 30 days after survey 3 via an incentivized Becker-DeGroot-Marschak (BDM) (Becker et al., 1964) mechanism that was implemented for one participant. We interpret this WTA as the perceived valuation of the ability to place bets on tracked apps for 30 days. This valuation is useful for analysis of policies like total bans on sports betting or uniform caps on sports betting consumption.

Since all of our policy analysis focuses on the period between survey 1 and survey 3, a more appropriate object of interest is “the total surplus from betting on all apps in the 30 days after survey 1.” We say “all apps” because to the extent that participants stated a low value utilized untracked apps, low valuations may reflect easy cross-app substitution rather than a genuine low WTA. To extrapolate from our incentivized WTAs to the target, we ask a set of hypothetical WTAs: 1) the WTA to stop betting on all apps for 30 days, 2) the WTA to stop betting on tracked apps for 30 days (unincentivized), and 3) the WTA to stop betting on tracked apps for 30 days if the sports schedule had been “similar to the schedule between survey 1 and survey 2 (April 9th to May 9th).”

A.4 Tables and Figures

Audience Group	Example Targeted Interests	Number of clicks
General Sports Betting (GSB)	*Any of:* {Sportsbooks, Daily fantasy sports, DraftKings, FanDuel, Barstool Sports}	3,752
Sports Fans	*Any of:* {Soccer, Premier League, Baseball, Basketball, ESPN, DraftKings, FanDuel, NBA}	1,634
GSB ∩ Sports Fans	≈ *At least one interest in GSB AND a separate sports interest*	3,743
GSB ∪ Casinos	*Any of GSB or:* {Casino Games, Lotteries, Online gambling}	1,514
GSB ∪ Beer	*Any of GSB or:* {Pabst Blue Ribbon, Budweiser, Coors Light, Miller Light, Natural Light}	1,478
Bill Simmons	Bill Simmons	968
GSB ∪ Day Trading	*Any of GSB or:* {Day Trading}	396
Total		13,485

Table A.1: Targeted audience groups for Facebook and Instagram Ads

Notes: The reported interests are summaries of multiple interest categories as defined by Facebook (e.g., “Sportsbooks” includes “Sportsbook (game)” and “Sports betting (gambling)”). The number of unique link clicks are the number of unique users who clicked on an ad that was targeted. The total number of unique link clicks in this table is slightly higher than that reported in Table 1 because a user who was targeted with and clicked on ads from two categories will be double-counted in this table.

Figure A.1: Example advertisement

Notes: We fielded this advertisement and qualitatively similar advertisements from March 13 through April 3 on Facebook and Instagram. Clicking on the ad directed users to our intake survey.

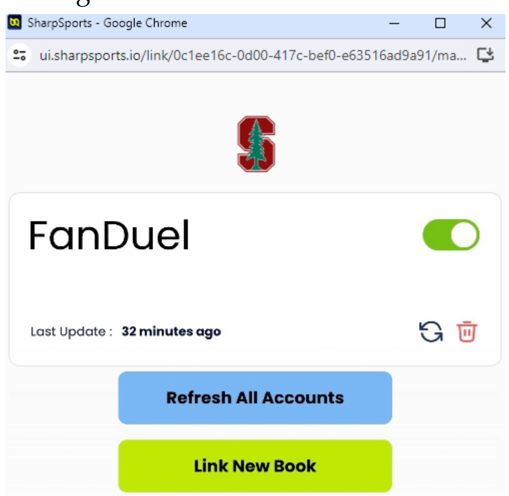

Figure A.2: Example syncing portal

Notes: The image is a screenshot of the SharpSports syncing portal for a participant who only used FanDuel.

B Reduced-form Results Appendix

B.1 The impact of rounding on prediction data

Our predictions data show excess mass at zero and at multiples of 5 cents. This fact suggests that self-reported data may be only a rounded representation of the true underlying belief. Following (Allcott et al., 2022b), we conduct a test to see how, if at all, this rounding impacts our average misperception estimates and find that rounding causes our predictions data to somewhat understate average misperceptions.

We illustrate CDFs of rounding-adjusted belief distributions for various focal beliefs and bandwidths in Figure B.11. The rounding-adjusted beliefs are on average to the right of the unadjusted distribution. This is because zero is by far the largest focal belief, and many more people report predictions slightly above zero than slightly below zero. We interpret this as evidence that more people round down than round up.

B.2 Estimating self-control problems

We use the strategy from Carrera et al. (2022) to estimate perceived self-control problems. Several papers provide formal proofs clarifying when and why this approach identifies perceived time-inconsistency. Here, we provide details on our implementation and briefly discuss some of its advantages.

Implementation As we note in the main text, the first step is to construct the *behavior change premium* (BCP), which is the difference between observed bonus valuations and counterfactual valuations of the bonus that we would observe for participants who do not have self-control problems. To measure these average counterfactual bonus valuations, we add the unconditional transfer value of the Bonus to the consumer surplus loss from an increase in the price of wagering a dollar by 2¢. The average value of the unconditional transfer was \$27.75. We measure the consumer surplus loss for an agent without perceived self-control problems with the trapezoid under a linear demand curve from a 2¢ price change, using participants' predictions of consumption in the treatment and control conditions to trace out demand.[46] This consumer surplus loss was \$12.03 on average. Therefore, an average participant without self-control problems would value the bonus at $\$27.75 - \$12.03 = \$15.72$. In fact, the average participant valued the bonus at \$17.58. The difference between the average observed valuation and the average counterfactual valuation is the BCP, as illustrated in Figure 7.

The behavior change premium captures the willingness to pay to reduce next month's betting by $\$K$, where K is the participant's prediction of how much the bonus reduces their betting. In our case, the average participant predicted that the bonus would reduce next month's betting by \$274. Allcott et al. (2022a) show that the BCP per dollar is a direct estimate of price-metric perceived self-control problems for our model.[47] This result does rely on our quasilinear utility assumption, which we validated in a previous version of

[46]Concretely, given predicted consumption x_i^C in the control condition and treatment consumption x_i^B, we approximate the consumer surplus loss with the trapezoid $\frac{x_i^C + x_i^B}{2} \cdot 0.02$.

[47]Carrera et al. (2022) study a quasihyperbolic discounting model; Allcott et al. (2022a) studies a model with linear temptation utility that maps more directly to our model.

the paper.[48] Our estimate of the average BCP per dollar is \$0.007. To compute subgroup specific estimates, we simply apply this procedure within each subgroup.

We make one modification to the standard Carrera et al. (2022) estimation. The estimation strategy relies on an assumption the incentive is a linear price change applying to the region of the demand curve where the participant will be consuming in the next month. In our experiment, the price change only applies to dollars wagered below the benchmark value. We also set a maximum bonus payment of \$90, so for some participants, the price change only applied to the first \$4500 in reduced betting below the benchmark. We do two things to deal with these subtleties. First, we restrict to participants whose predicted consumption in both treatment and control lie within the region where the price change is binding, $[B - 4500, B]$. This is a subsample of 352 participants (79% of the analysis sample). Second, for participants with $B > 4500$, we use $B - 4500$ as the left edge of the trapezoid defining CS loss rather than 0. For these participants, the trapezoid is approximated with $\left(\frac{x_i^C + x_i^B}{2} - (B - 4500)\right) \cdot 0.02$.

Advantages of the approach We chose this measurement strategy in part because it is robust when bettors are uncertain about future normative demand. A related but different strategy studies the take-up of binding commitment devices.[49] This strategy faces challenges when future demand is uncertain. Under such uncertainty, agents who want to correct self-control problems might still avoid taking up commitment (Laibson, 2015). Intuitively, if an agent does not know their future optimum, they will not take up binding commitments that might cause consumption to be suboptimally low. In that case, the commitment device take-up strategy can create false negatives (Strack and Taubinsky, 2021).[50] Such uncertainty is relevant in our context. A majority of bettors agree that "It's hard for me to predict in advance how much I'll want to bet in a given week." (Appendix Figure B.7). Our strategy avoids these issues because it measures the extent to which

[48] For this estimation, the quasilinear utility assumption amounts to an assumption that participants are risk-neutral over experimental payments. If participants are risk-averse (risk-loving) over experimental payments, the BCP per dollar understates (overstates) perceived self-control problems. Risk preferences matter only to the extent that participants are ex ante uncertain about their own future betting consumption. We view risk-neutrality as a sensible assumption given the stakes of the bonus payments and given that we intentionally did not emphasize uncertainty in the framing of the elicitation. In Brown et al. (2025) we show that the average participant is somewhat risk-averse over experimental payments, and that this measure of risk aversion did not correlate with bonus valuations.

[49] See, for example DellaVigna and Malmendier 2006; Kaur et al. 2015; Augenblick et al. 2015.

[50] Another prominent strategy for measuring self-control problems is *choice-reversal* designs (Read and van Leeuwen, 1998; Sadoff et al., 2020; Augenblick et al., 2015; Augenblick and Rabin, 2019). These designs show that choices made in-the-moment differ from choices made in advance. While these designs are attractive in some contexts, Strack and Taubinsky (2021) show that choice reversal designs are generally not robust to uncertainty about future demand. They also show that the approach we implement requires weaker assumptions for identification.

people would trade future cash payments for future changes in betting consumption, regardless of the level of future normative demand.

An additional and separate advantage of our approach is that it is robust to mean-zero elicitation noise in reported willingness to pay. This is not the case for other popular measurement strategies (Carrera et al., 2022).

B.3 Details on shrinkage procedure and results

The goal of our shrinkage procedure is to recover individual-specific estimates of latent overoptimism from noisy data on prediction errors. In our setting, there is quantitatively important and non-random heterogeneity across individuals in the *noisiness of raw prediction errors*, because the returns of low-volume bettors are noisier than the returns of high-volume bettors. These facts imply that traditional shrinkage methods can produce biased estimates.[51] We implement a procedure from Chen (2024) that allows for this kind of heterogeneity in precision across units.

To implement the Chen (2024) estimator, we define our overoptimism estimand γ_i^O as the difference between perceived expected net returns $E_{\tilde{F}}[a] = r_i^{Perc}$ and true expected net returns $E_F[a] = r_i$, both of which are time-invariant. We observe noisy signals of true expected returns and of perceptions in periods $t = 1, 2$. Noise in realized net returns arises because of the randomness of sports outcome. Noise in predicted net returns is survey elicitation noise (Kahneman, 1965; Gillen et al., 2019; Haaland et al., 2023). Formally, we write observed prediction errors as $\hat{\gamma}_{it} = \hat{r}_{it}^{Perc} - \hat{r}_{it}$, where $\hat{r}_{it}^{Perc}$ and $\hat{r}_{it}$ are defined as follows:

$$\hat{r}_{it} = r_i + \nu_{it}^{Realization} \qquad E[\nu_{it}^{Realization}] = 0; \quad V[\nu_{it}^{Realization}] = \sigma^2_{Realization,it} \tag{B.1}$$

$$\hat{r}_{it}^{Perc} = r_i^{Perc} + \nu_{it}^{Response} \qquad E[\nu_{it}^{Response}] = 0; \quad V[\nu_{it}^{Response}] = \sigma^2_{Response} \tag{B.2}$$

We assume prediction and realization errors are independent: $E[\nu^{Realization}\nu_{it}^{Response}] = 0$.

We measure variances from the panel structure of our data. The variance of responses $\sigma^2_{Response}$ is identified from the variance of changes in responses across surveys.[52] Observation-specific realization variances $\sigma^2_{Realization,it}$ are computed directly from observed wager histories. For a bettor who places K bets in a period, we model realized returns as a weighted sum of K independent Bernoulli trials where the weights are proportional to dollars wagered in the bet and the trial success probability is given by the bet

[51] We found that low-volume bettors and parlay bettors are more overoptimistic. These are exactly the bettors who have less precise realizations, and therefore their average prediction errors are more noisy. Traditional shrinkage procedures inappropriately shrink these bettors' prediction errors towards the (lower) mean estimate among the high-volume, non-parlay bettors who have more precise signals.

[52] We estimate that response errors have a standard deviation of $\hat{\sigma}_{Response} = 12.3$¢.

odds. Specifically, we approximate the probability of winning bet k with $\frac{1}{q_{itk}}$ where q_{itk} is the bet's payout per dollar for winners. Then, the variance of net returns per dollar is

$$\sigma^2_{Realization,it} = \sum_k s^2_{itk}(q_{itk} - 1) \tag{B.3}$$

where s_{ikt} is the share of dollars allocated to bet k.

As the input to our Chen (2024) "CLOSE" shrinkage procedure, we use the average individual-specific prediction error across both periods of our data. We then use these variance estimates described above to compute individual-specific standard errors for these estimates. Our implementation of CLOSE assumes that the underlying overoptimism γ_i^O is Gaussian conditional on these standard errors. The output is the individual-level overoptimism estimates illustrated in Panel C of Figure 6

To validate our use of these CLOSE estimates, we compare them to the distribution of overoptimism estimated from a random effects model. We estimate the following model on our panel of prediction errors.

$$\hat{\gamma_{it}^O} = \gamma_i^O + \nu_{it} \tag{B.4}$$

$$\text{Where} \quad \gamma_i^O \sim N(\mu_{\gamma^O}, \sigma^2_{\gamma^O}) \quad \text{and} \quad \nu_{it} \sim N(0, \sigma^2_\nu) \tag{B.5}$$

Unlike CLOSE, this model does not exploit information on differences across people in the precision of raw prediction errors, and it also imposes a functional form assumption. One benefit of this approach is that it does not rely on any assumptions about how betting returns arise (recall that in the CLOSE procedure we model returns as the result of independent Bernoulli trials). We illustrate the random effects estimates alongside the CLOSE estimates of overoptimism with Figure B.12. The distributions have similar means and standard deviations. We conclude that our qualitative takeaways are robust to these alternative assumptions.

B.4 Bonus treatment effects

To estimate the treatment effects of the Bet Less Bonus, we use the following regression specification for participants i and periods $t \in \{1, 2\}$:

$$Y_{it} = \tau_t^B B_i + \beta_t X_i + \delta_t + \varepsilon_{it} \tag{B.6}$$

where Y_{it} is an outcome variable, B_i is a bonus treatment indicator, X_i is a vector of prespecified individual-level covariates, and δ_t is a period fixed effect.[53] Coefficients $\hat{\tau}_t^B$ are

[53] The covariates are: a limits treatment indicator, log winsorized baseline wagers, elicited beliefs about period-1 winnings, a measure of information conveyed in the information treatment, white, income terciles, and randomization stratum indicators.

estimates of the contemporaneous and long-run impacts of the Bet Less Bonus for $t = 1$ and $t = 2$ respectively.

When implementing this regression, we make two choices to ensure that our estimates can be interpreted as proportional average treatment effects of the price change. First, we define a normalized dependent variable Y_{it}^{Norm} to be the period-t outcome as a proportion of the period-0 outcome, and then scaled such that the control group always has an average of 1. Treatment effects on Y_{it}^{Norm} are interpreted as average proportional treatment effects relative to the control group (Chen and Roth, 2024). Second, we restrict the sample to the set of participants who could have been exposed to the price change. Recall that the Bet Less Bonus was a payment for reducing consumption below a personalized benchmark; the 6% of treated participants who wagered more than the benchmark were not exposed to the price change and so were dropped from the sample. By construction, these are the treated participants with the highest normalized consumption outcomes, so to ensure that the treated and control groups are comparable we also drop the 6% of control group participants with the highest normalized consumption outcomes.

Appendix Figure B.5 displays our results on substitution to other kinds of gambling. The top row represents the 34% treatment effect of the Bet Less Bonus on tracked sports betting. The rows below plot treatment effects for various other categories of gambling other than sports bets on tracked apps. We elicited spending on these other categories via an unincentivized question at the end of each survey, which participants knew would not affect their payments. The second row represents the treatment effect on the sum of spending on the individual categories, while the rows below represent treatment effects on individual categories.

B.5 Bans

To evaluate if a ban on sports betting would be welfare-enhancing, we must assess if overall consumer surplus is positive after accounting for bias. Extrapolating this from our main demand estimates is unreliable. Instead, we elicit participants' perceived consumer surplus from sports betting directly, and then compare it to the uninternalized costs.

Our measure of perceived consumer surplus is the willingness to accept (WTA) to stop sports betting entirely for a 30-day period. We elicited this WTA in survey 3 as described in Section A.3. We interpret this measure as perceived by the long run self. Since our participants are sophisticated about self-control problems, the only uninternalized cost in this perceived measure comes from overoptimism. Therefore, by comparing the sizes of perceived consumer surplus and internalities from overoptimism, we can evaluate the welfare effects of bans.

Our results imply that a ban would reduce welfare. Figure B.10 shows that perceived consumer surplus exceeds uninternalized costs for 93% of participants. The median WTA to stop betting for 30 days was $100. Such high valuations imply that participants place high value on inframarginal bets.

This result should be interpreted with particular caution, since it relies on the WTA to completely stop sports betting on tracked accounts. It is well-known that results from this kind of experimental elicitation of WTAs are sensitive to framing choices (Allcott et al., 2020).[54] Also, the monthly cost of giving up gambling for a single month may differ from the monthly cost of giving up gambling forever. Overall, we view our results as suggestive evidence that the biases that we measure are not large enough to make a ban welfare-enhancing in our sample of participants who have already become high-volume bettors. Our results have little to say about the welfare effects of a ban on other populations, such as non-bettors.

B.6 Tables and Figures

[54]We do provide some evidence that our high valuations were not driven by confusion about the BDM incentivization mechanism: we elicited hypothetical WTAs before the incentivized WTAs, and these were even larger.

	Misprediction			Prediction	Realization
	(1)	(2)	(3)	(4)	(5)
Constant	0.0177	0.0304	0.0215	0.0319*	0.0011
	(0.0638)	(0.0662)	(0.0643)	(0.0189)	(0.0622)
Above median parlay share	0.1800***	0.1945***	0.1729***	-0.0293**	-0.2133***
	(0.0466)	(0.0507)	(0.0487)	(0.0138)	(0.0454)
No college degree	0.1357**	0.1362**	0.1311*	0.0205	-0.1057
	(0.0677)	(0.0677)	(0.0684)	(0.0200)	(0.0660)
Graduate degree	-0.0075	-0.0109	-0.0087	0.0137	0.0260
	(0.0495)	(0.0498)	(0.0496)	(0.0146)	(0.0483)
Above median live bet share	-0.0745	-0.0714	-0.0761*	-0.0098	0.0708
	(0.0452)	(0.0454)	(0.0454)	(0.0134)	(0.0440)
Above median wagers	-0.0220	-0.0253	-0.0198	0.0094	0.0344
	(0.0465)	(0.0467)	(0.0467)	(0.0137)	(0.0453)
Above median age	-0.0311	-0.0328	-0.0285	-0.0100	0.0267
	(0.0457)	(0.0458)	(0.0460)	(0.0135)	(0.0445)
Above median income	0.0231	0.0242	0.0223	-0.0287**	-0.0467
	(0.0482)	(0.0483)	(0.0483)	(0.0143)	(0.0470)
Above median bet riskiness		-0.0362			
		(0.0502)			
Standardized bet riskiness			0.0124		
			(0.0246)		
R^2	0.06614	0.06749	0.06680	0.03505	0.08368
Observations	371	371	371	371	371

Table B.1: Individual-level predictors of overoptimistic predictions

Notes: Coefficients are in units of dollars per dollar wagered. All covariates are binary, except for "standardized bet riskiness," which is the standardized average bet riskiness for wagers placed in $t = 1$ and $t = 2$. An observation is a member of the analysis sample who placed wagers both in the pre-study period ($t = -1, -2$) and in the study period ($t = 1, 2$).

	True consumption (as share of t=0 consumption)	
	(1)	(2)
Constant	0.2263***	0.0506
	(0.0780)	(0.1298)
Control Prediction × Control	0.3173***	0.3480***
	(0.0436)	(0.0781)
Bonus Prediction × Bonus	0.2093***	0.1516***
	(0.0478)	(0.0575)
Control Prediction × Bonus		0.1462*
		(0.0848)
Bonus Prediction × Control		0.0758*
		(0.0444)
R^2	0.11373	0.12591
Observations	442	442

Table B.2: Predicted consumption predicts observed consumption

Notes: The table reports coefficient estimates and standard errors from multivariate regressions. Coefficients are in units of dollars per dollar wagered. An observation is a member of the analysis sample who placed wagers both in the pre-study period ($t = -1, -2$) and in the study period ($t = 1, 2$).

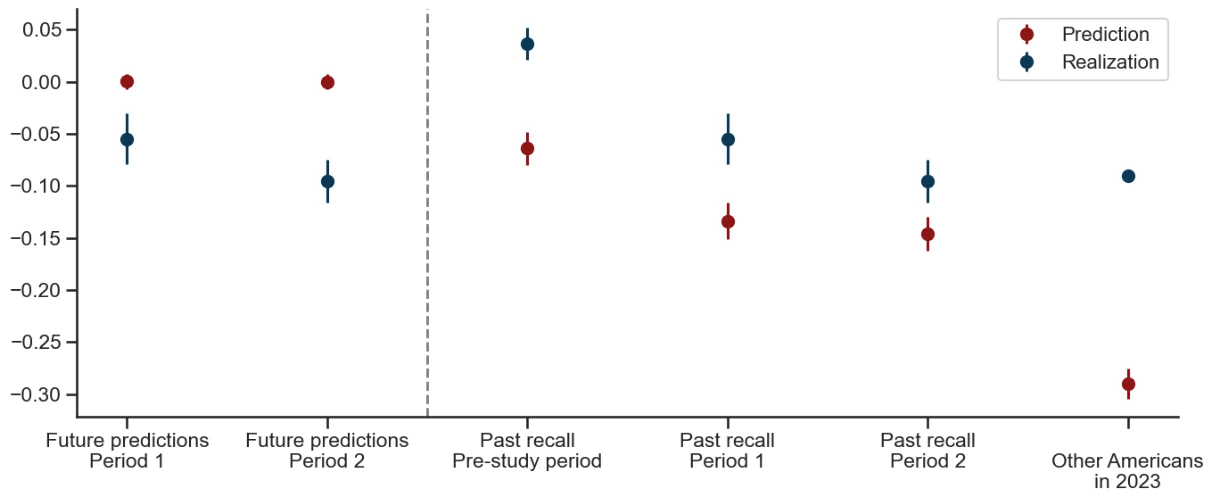

Figure B.1: Mispredicted financial returns in the future and the past

Notes: The figure illustrates mean predicted net returns and realized net returns across participants for various periods. The points to the left of the dotted line represent forward-looking prediction, and the points to the right of the line represent backward-looking recollections. The "Pre-study period" is the period from January 1, 2024 until April 8, 2024 (the day the study began). Future predictions about periods 1 and 2 took place in surveys 1 and 2 respectively. Recollections about the pre-study period, period 1, and period 2 took place during surveys 1, 2, and 3 respectively. The question about other Americans took place in survey 1. The ground truth (-0.09) for that question comes from American Gaming Association (2024). Future predictions and realized returns are truncated to lie in $[-0.25, 0.4]$; returns for questions are not truncated. Error bars represent 95% confidence intervals.

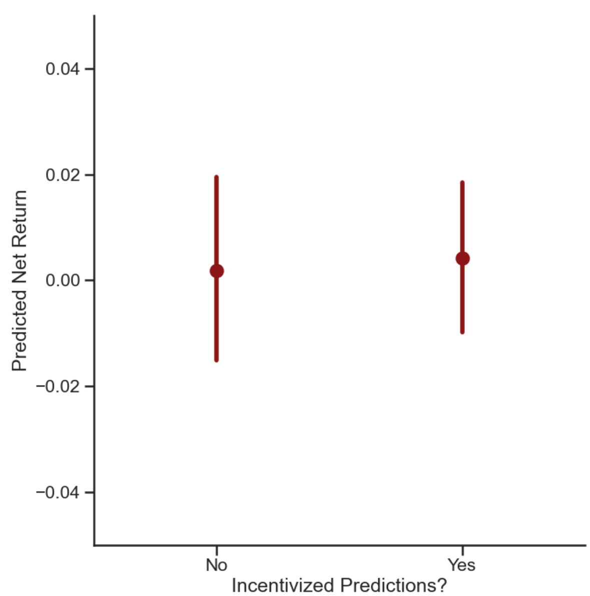

Figure B.2: Impact of incentivization on predictions of financial returns

Notes: The figure plots the means of predicted net returns by whether predictions were incentivized. Error bars represent 95% confidence intervals. Predictions were incentivized for 65% of participants.

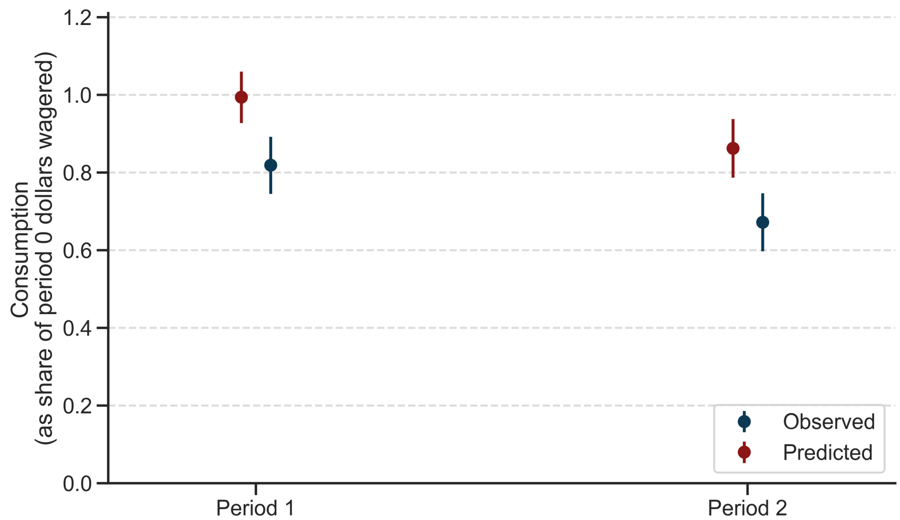

Figure B.3: Evidence against naivete from predictions of future dollars wagered

Notes: The figure compares predicted dollars wagered to the truth. We restrict to observations in the bonus control condition. Points represent mean wager volume, as a share of period 0 dollars wagered. Error bars represent 95% confidence intervals. The red series represents predictions, and the blue series represents observations. Both the observed and predicted variables are truncated to lie in $[0, 2]$

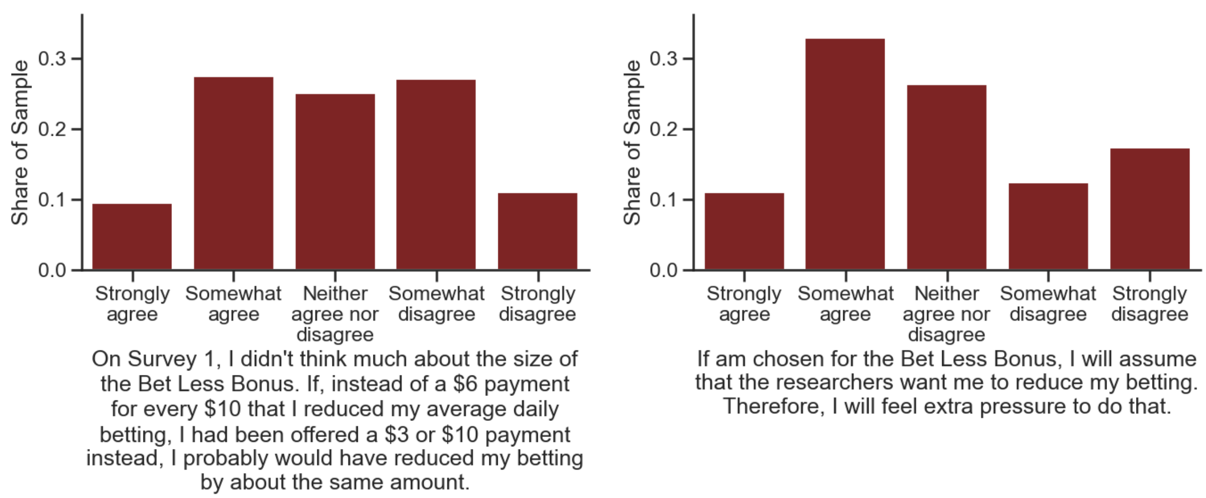

Figure B.4: Qualitative evidence on how participants interpreted the bonus

Notes: The figure shows the share of people giving each response to the likert scale questions about the Bet Less Bonus. Both questions were asked on survey 3.

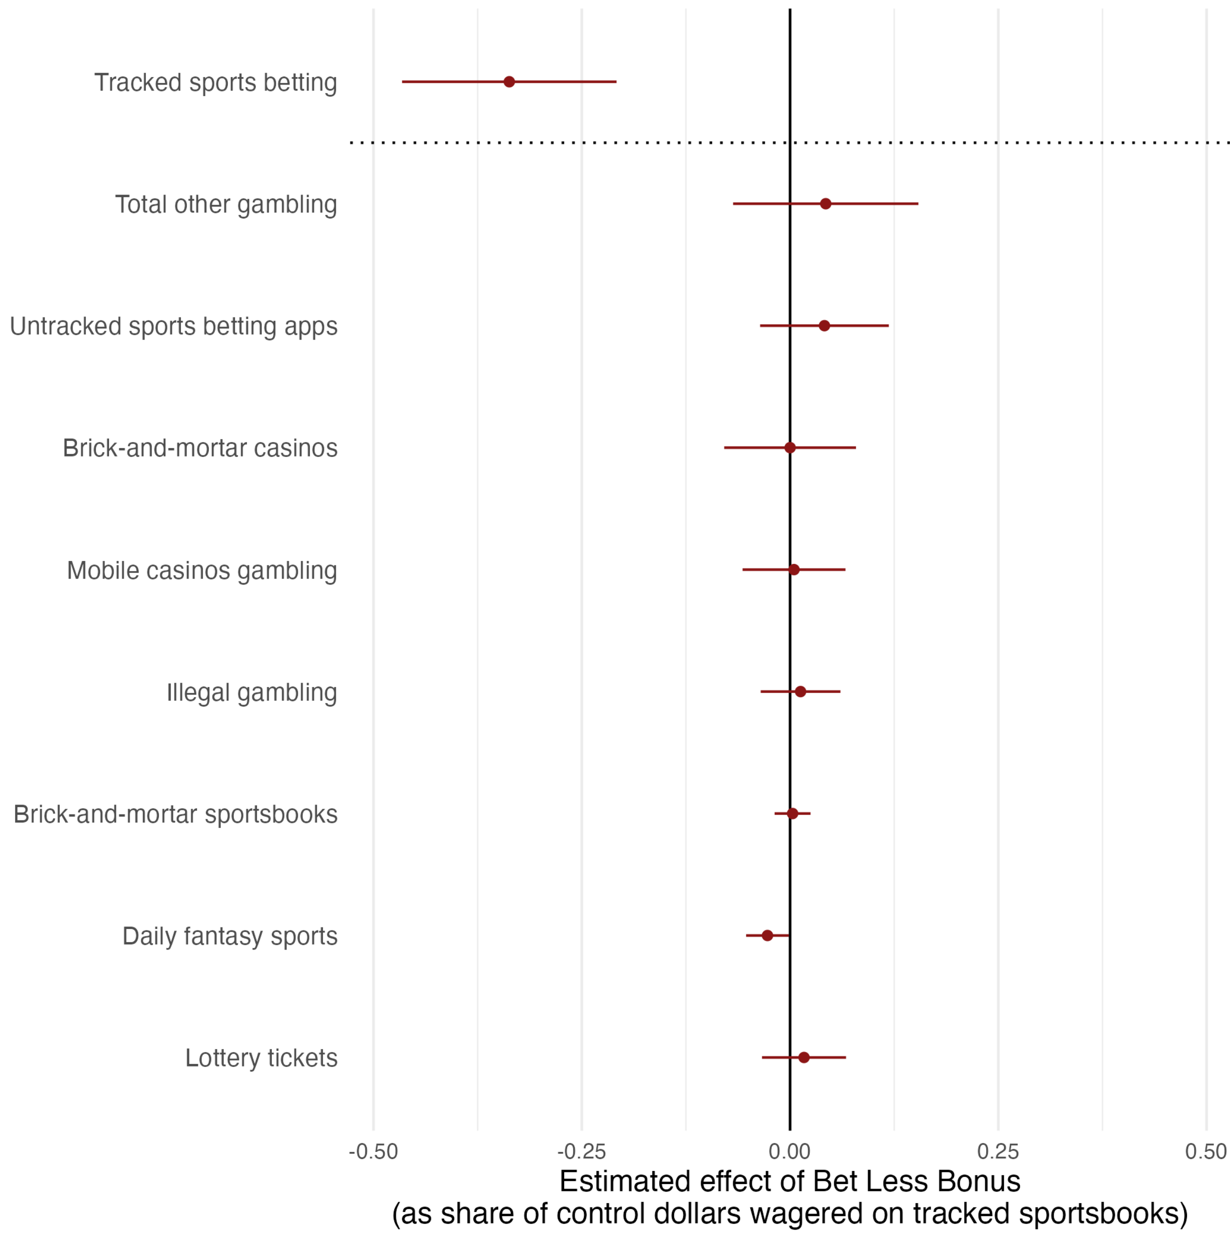

Figure B.5: Effect of Bet Less Bonus on self-reported untracked gambling

Notes: The figure presents coefficients from bonus treatment effect regressions as specified in equation B.6. For tracked sports betting, the dependent variable is normalized as described in that section. For untracked gambling, the dependent variable is normalized by the same factor that the tracked sports betting variable. Therefore, all coefficients are interpreted on a common scale as changes proportional to the magnitude of tracked sports betting consumption in the control condition. The normalized dependent variable ratios are truncated above at 2 to improve precision. Error bars represent 95% confidence intervals.

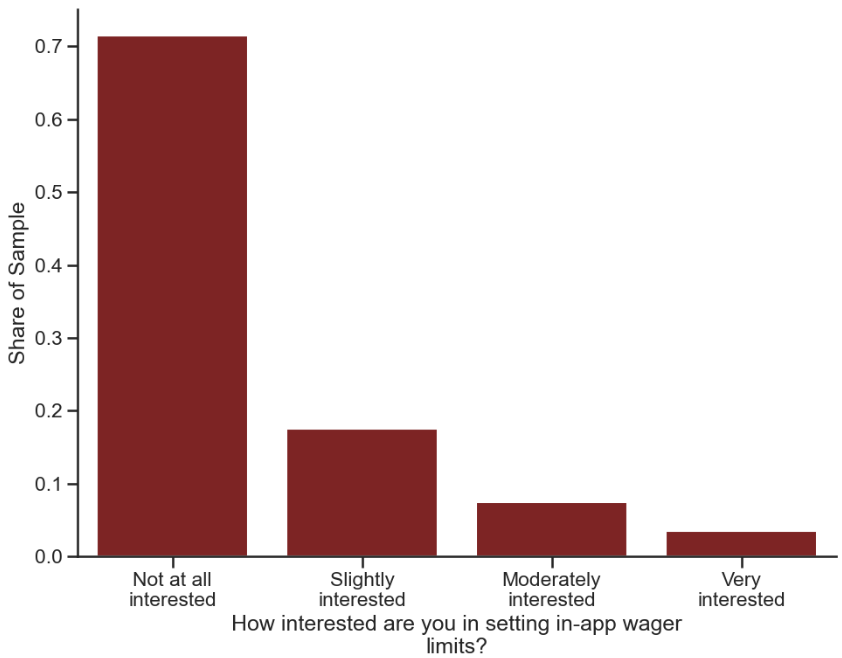

Figure B.6: Qualitative interest in limits

Notes: The figure reports the frequency of responses to the question, "How interested are you in setting in-app wager limits?" We asked this question before participants set their limits. We restrict to the subsample of bettors in the limits treatment who did not have difficulty finding limit screens for all apps (N=200).

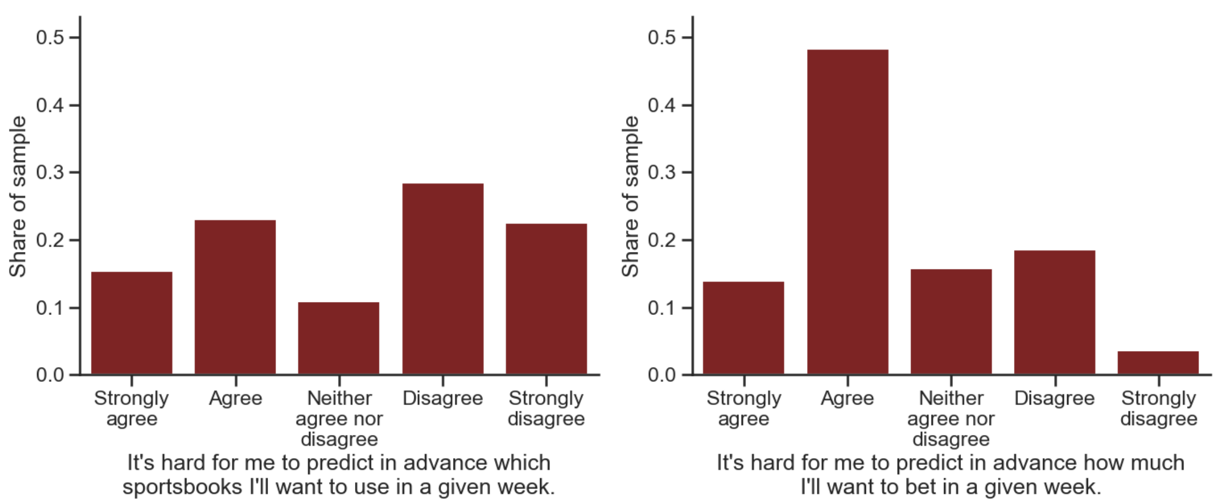

Figure B.7: Qualitative evidence on uncertainty about future demand

Notes: The figure shows the share of people giving each response to the likert scale questions about uncertainty about the future. Both questions were asked on survey 3.

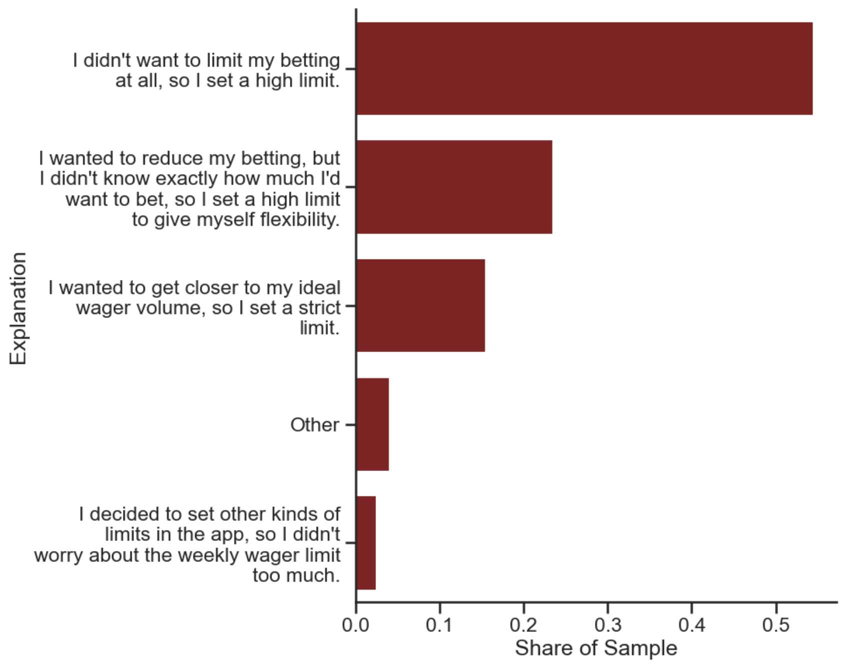

Figure B.8: Qualitative explanations of limit choices

Notes: The figure reports the frequency of responses to the question, "Please select the statement that best describes your thinking when choosing the weekly wager limit." We asked this question after participants set their limits. We restrict to the subsample of bettors in the limits treatment who did not have difficulty finding limit screens for all apps (N=200).

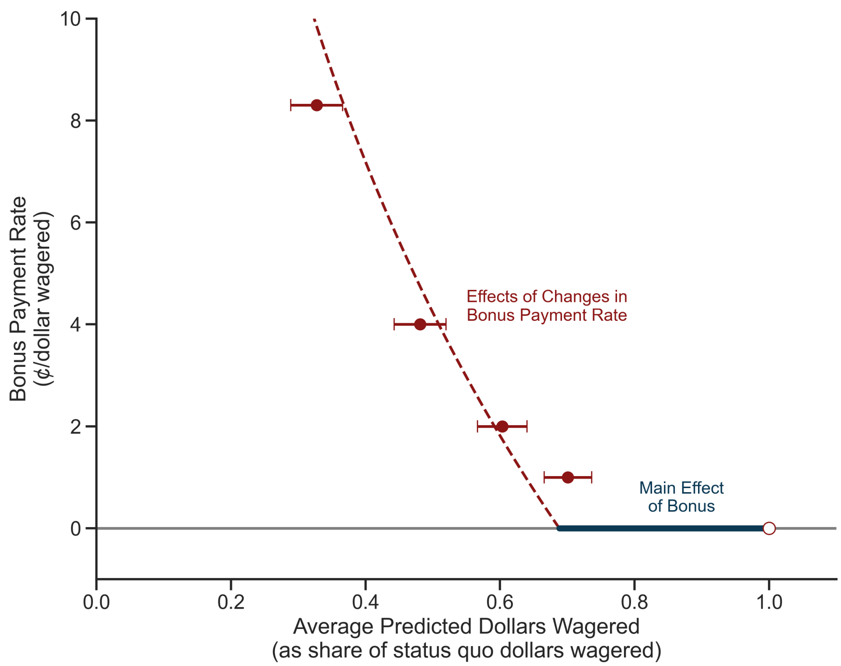

Figure B.9: Effects of varying the Bonus payment rate on predicted consumption

Notes: The figure shows the average predicted consumption as a share of control consumption for Bet Less Bonuses with varying payment rates. Error bars represent 95% confidence intervals. We asked this question to all participants in survey 3. The dotted line is a constant semielasticity fit through the aggregated shares. We drop participants who predict that they will wager $0 in the control condition, since the proportional effect of a price change on consumption is undefined for them. We truncate predictions from above at 1, to ensure that there are no upward sloping demand curves.

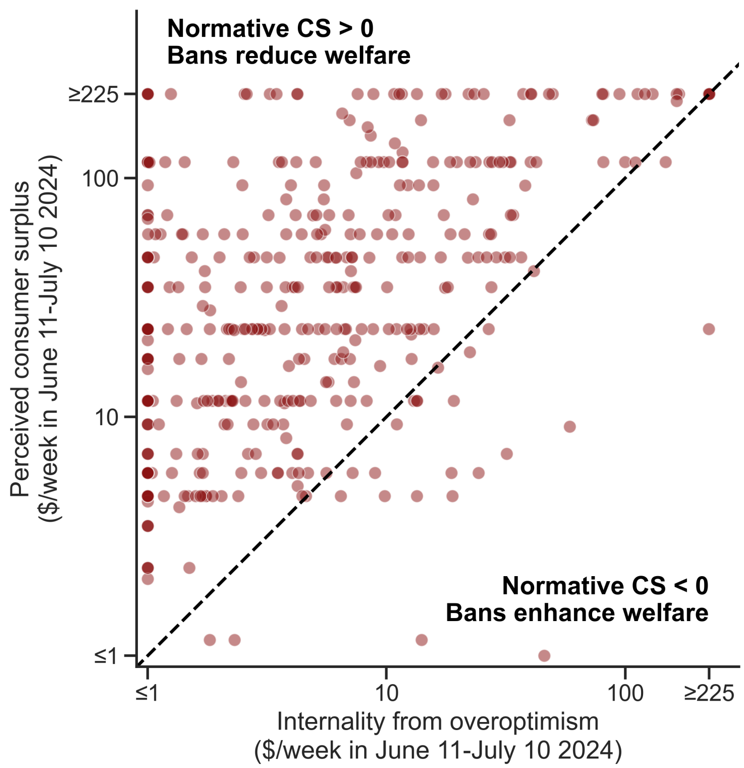

Figure B.10: Evaluating bans

Notes: The figure compares the perceived value of betting to uninternalized costs from overoptimism. Approximating self-control problems to zero, normative consumer surplus is the sum of these two components. Bans are welfare enhancing if normative consumer surplus is negative. The perceived consumer surplus measure is the WTA to stop betting on tracked apps, as elicited via an incentivized BDM mechanism. The internality from overoptimism is the expected total uninternalized financial costs. We compute it by multiplying our estimate of overoptimism γ_i^M by predicted consumption. For both measures, values are reported in units of dollars per week in the 30 days following survey 3.

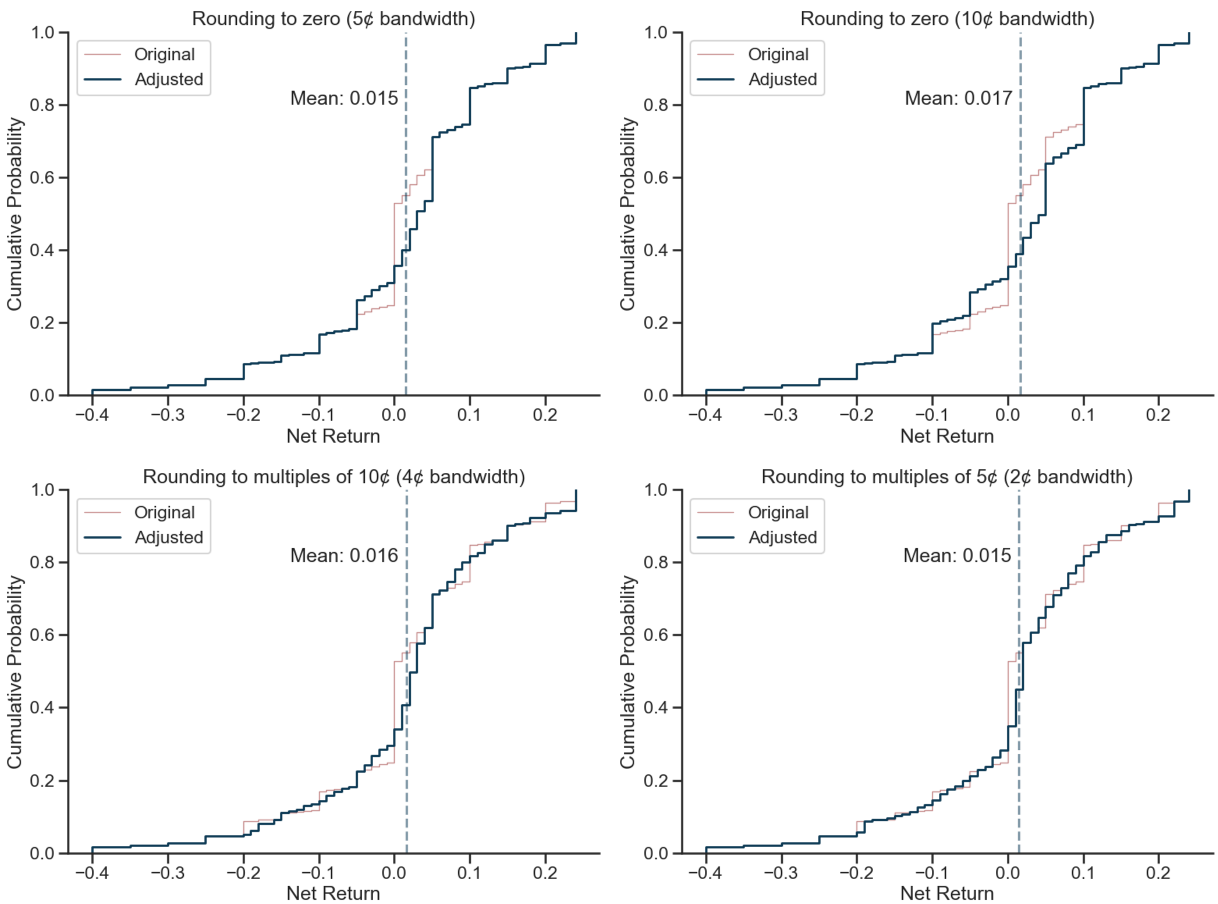

Figure B.11: Rounding-adjusted belief distributions

Notes: Let j index "focal" predictions which agents round to. We assume that an agent whose underlying belief is within bandwidth d of a focal prediction reports their true belief with probability p_j and the focal prediction with probability $(1 - p_j)$. To estimate the p_js, we measure the excess mass at each focal prediction. Then, we construct a counterfactual unrounded distribution that re-allocates excess mass to non-focal predictions. This procedure follows Allcott et al. (2022b).

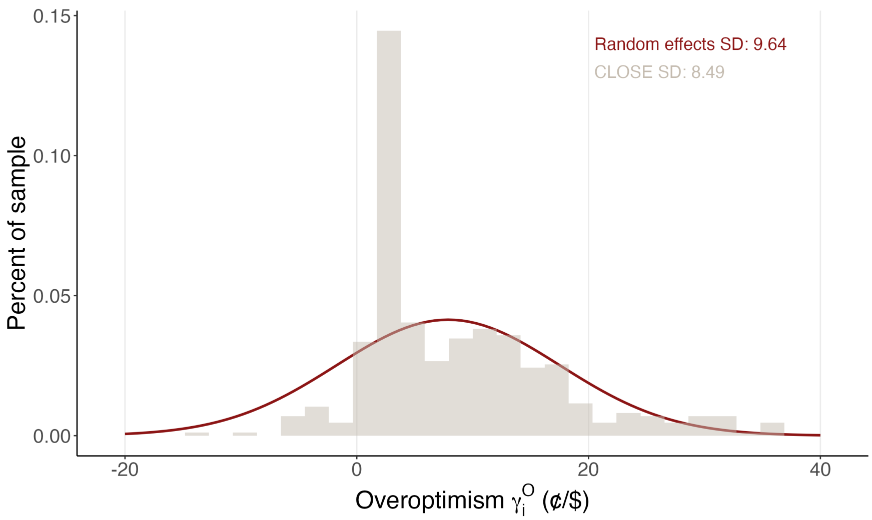

Figure B.12: Distribution of overoptimism from a random effects specification

Notes: The figure illustrates the distribution of overoptimism across individuals estimated from the random effects specification in equations B.4 and B.5. The distribution from the random effects model is illustrated with the red density function, and the distribution of individual-specific overoptimism as estimated with our illustrated with the gray histogram. We report the estimated standard deviation of overoptimism as well as the standard deviation of the CLOSE estimates.

C Model-based Analysis Appendix

C.1 Deriving the constant semielasticity demand curve

We show here how we derive the demand function (3). We rely on the specialization of nonfinancial utility to $z_i(x;\tilde{F}) = z_{1i}x\log(x) + z_{2i}x + g_i(\tilde{F})x + h_i(\tilde{F})$ where $z_{1i} < 0$.

For notational convenience, let $\gamma_i = \gamma_i^O + \gamma_i^{SC}$ denote the sum of the biases. We consider two first order conditions for consumption: the decision FOC and the normative status quo FOC (given $\tau_i = \tau_0$, $\gamma_i^O = \gamma_i^{SC} = 0$). These are the derivatives of (1) and (2). We define $x_i^*(0)$ as the consumption that satisfies the normative status quo FOC and $x_i^{choice}(\tau)$ as the consumption that satisfies the choice FOC.

$$E_F[a] - \tau_0 + z_i'(x_i^*(0)) = 0 \tag{C.1}$$

$$(E_F[a] + \gamma_i^O - \tau_i) + z_i'(x_i^{choice}(\tau_i)) + \gamma_i^{SC} = 0 \tag{C.2}$$

Taking the difference of these FOCs, using our functional form assumption, and rearranging yields the following expression:

$$x^{choice}(\tau_i) = \exp\left(\log(x^*(0)) + \frac{1}{z_{1i}} \cdot (\tau_i - \gamma_i - \tau_0)\right) \tag{C.3}$$

which is exactly the form of equation (3). This result shows that parameter estimates from log-linear demand specifications admit a structural interpretation. The slope of log demand in prices is a semielasticity term which arises from the curvature of the nonfinancial utility of betting, as usual. The intercept is the log of normative demand at the reference tax.

C.2 Parlay-specific taxes

To study the welfare effects of parlay-specific taxes, we adapt an approach from Jacobsen et al. (2020) to our context. We compute welfare effects under various assumptions about the substitutability between bet types and the relationship between parlays and bias.

This extension allows bias to vary across individuals and two product types: single bets s and parlays p. Let γ_{ij} denote the price-metric internality from individual i's consumption of product j, and let τ_{ij} denote an arbitrary set of individual and product-specific taxes. The first-best tax schedule is $\tau_{ij}^* = \gamma_{ij}$. The key result of Jacobsen et al. (2020) is that for linearized demand curves, the deadweight loss from a second-best corrective tax regime is equal to the sum of squared residuals from a GLS regression of unit-specific externalities on the set of unit-level policy variables on which the planner can condition the taxes. In our context, a unit is a product-individual and the second-best regime of interest is one where the planner sets optimal uniform singles taxes and parlays taxes.

To implement this approach empirically, we need to specify two empirical objects that were not necessary for the analysis in the body of the paper: product-specific biases γ_{is}, γ_{ip} and demand matrices that may allow for substitution between products. The experiment was not designed to measure these objects; in all other analyses we aggregate bias γ_i and price-sensitivity η_i across products. Therefore, we present results for multiple scenarios.

Product-specific bias assumptions Panel B of Figure 6 shows that parlay bettors are more overoptimistic than singles bettors. This association could be driven by selection of biased individuals into parlay betting, or it could be driven by the causal effect of parlay betting on bias. We consider two extreme cases: one in which the association is "all selection," and one in which it is "all causal."

Specifically, we model i's bias for good j as follows. We allow for general heterogeneity in bias for singles wagers γ_{is}. The effect of parlays on bias is an additive treatment effect δ: $\gamma_{ip} = \gamma_{is} + \delta$. The main paper delivers estimates of overall bias for each individual, which is a weighted sum of bias across bet types: $\overline{\gamma}_i = \sigma_i \gamma_{is} + (1 - \sigma_i)\gamma_{ip}$. Rearranging yields the following equation:

$$\overline{\gamma}_i = \underbrace{E_i[\gamma_{is}]}_{\text{Avg. singles overoptimism}} + \underbrace{(1-\sigma_i)}_{\text{i's parlay share}} \cdot \underbrace{\delta}_{\text{Parlay effect}} + \underbrace{(\gamma_{is} - E_i[\gamma_{is}])}_{\text{Idiosyncratic overoptimism}} \tag{C.4}$$

Our scenarios make assumptions on the within-individual effect of parlays δ. An individual-level regression of bias on the parlay share yields a highly significant association of $\hat{\beta} =$ 11.0. In the "all causal" scenario, we attribute the full association to the causal effect of parlays, and assume $\delta = \hat{\beta}$. In the "all selection" scenario, we assume $\delta = 0$. Given a specification of δ, we simulate the product-specific biases given our estimate of overall biases and observed consumption shares.

Substitutability assumptions The key question here is the extent to which parlay bets and singles bets are substitutes. For each individual, we specify a 2×2 demand matrix D_i where D_i^{ss} is the own-price derivative for singles, D_i^{pp} is the own-price derivative for parlays, and $D_i^{sp} = D_i^{ps}$ are the cross-price derivatives. Specifically, we parameterize the system as follows:

$$\text{Own-price singles slope:} \qquad D_i^{ss} = \eta_i x_i \theta_i - \delta_i \tag{C.5}$$

$$\text{Own-price parlays slope:} \qquad D_i^{pp} = \eta_i x_i (1-\theta_i) - \delta_i \tag{C.6}$$

$$\text{Cross-price slope:} \qquad D_i^{sp} = D_i^{ps} = \delta_i = -\lambda \eta_i x_i \sigma_i (1-\sigma_i) \tag{C.7}$$

where η_i are the own-price semielasticities of overall sports betting demand that were estimated for the main paper and θ_i denotes the baseline share spent on single bets for consumer i. These demand matrices are specified to make substitution between parlays and singles possible while still being consistent with aggregate price-sensitivities from the main paper. The λ parameter scales the substitution intensity between parlays and singles.

We simulate results for two scenarios: $\lambda = 0$ (no substitution) and $\lambda = 1.85$. The latter case implies substantial substitution. To see this, we computed aggregate diversion ratios in our sample. The $\lambda = 1.85$ parameterization implies that for every \$1 reduction in parlay wagers induced by a parlay tax, 50¢ is diverted to non-parlays while 50¢ is diverted to the numeraire.

C.2.1 Results

Table C.1 presents the optimal tax rates and total surplus benefits under our empirical scenarios. While quantitative estimates of surplus benefits vary across the scenarios, we highlight qualitative takeaways and economic intuition.

Even in the scenarios where parlays do not cause bias (Panels A and C), the planner chooses to set higher taxes on parlays. The reason is that the planner can use parlay-specific taxes to target individual-level heterogeneity in bias. In Panel A, with no substitution, this targeting motive alone causes the planner to set a tax on parlays that is 50% higher than the tax on singles. In Panel C, though, bettors can respond to the tax by substituting from parlays to singles. Since internalities are equal across bet types, such substitution is distortionary. The planner mitigates this distortion by decreasing the difference in tax rates. Overall, the welfare benefits from targeting when parlays do not cause bias are small – targeting on parlays captures less than 5% of the first-best surplus benefits that are left on the table by uniform taxes.

When parlays do cause bias, parlay taxes are much more beneficial. In the scenario with no substitution (Panel B), parlay-specific taxes capture 45% of the remaining potential surplus gains from targeting. The planner achieves these gains by levying a parlay tax that is more than three times higher than the tax on non-parlays. When people do substitute between wagers (Panel D), the benefits are even larger. Intuitively, the substitution from parlays to non-parlays is now a *benefit* to the planner, since non-parlays have much smaller internalities. The simulations hold the aggregate price-sensitivity of gambling demand fixed (since this is pinned down by our empirical estimates), so accounting for this kind of substitution increases the attractiveness of parlay-specific taxes.

C.3 Tables and Figures

Panel A: No substitution, Parlay result is from selection into parlays

Policy	Tax on non-parlays	Tax on parlays	Surplus gain (as share of first-best surplus gain)
No tax	0.0%	0.0%	0.00
Status quo tax	2.0%	2.0%	0.36
Optimal uniform tax	5.3%	5.3%	0.58
Optimal bet-specific tax	4.8%	7.5%	0.60

Panel B: No substitution, Parlay result is from causal effects of parlays on bias

Policy	Tax on non-parlays	Tax on parlays	Surplus gain (as share of first-best surplus gain)
No tax	0.0%	0.0%	0.00
Status quo tax	2.0%	2.0%	0.29
Optimal uniform tax	5.3%	5.3%	0.47
Optimal bet-specific tax	3.6%	13.5%	0.71

Panel C: Yes substitution, Parlay result is from selection into parlays

Policy	Tax on non-parlays	Tax on parlays	Surplus gain (as share of first-best surplus gain)
No tax	0.0%	0.0%	0.00
Status quo tax	2.0%	2.0%	0.36
Optimal uniform tax	5.3%	5.3%	0.58
Optimal bet-specific tax	5.1%	6.3%	0.59

Panel D: Yes substitution, Parlay result is from causal effects of parlays on bias

Policy	Tax on non-parlays	Tax on parlays	Surplus gain (as share of first-best surplus gain)
No tax	0.0%	0.0%	0.00
Status quo tax	2.0%	2.0%	0.21
Optimal uniform tax	5.3%	5.3%	0.35
Optimal bet-specific tax	3.5%	14.0%	0.78

Table C.1: Gains from bet-specific tax schemes

Notes: The "No substitution" and "Yes substituiton" scenarios correspond to parameterizations with $\lambda^{\text{Demand}} = \{0, 1.85\}$ respectively. The "Parlay result is from selection into parlays" and "Parlay result is from causal effects of parlays on bias" scenarios correspond to parameterizations with $\lambda^{\text{Parlay}} = \{0, \hat{\beta}\}$ respectively, where $\hat{\beta}$ is the coefficient from an unweighted cross-participant regression of total bias γ_i on parlay share. We compute the optimal tax rates and surplus gains using the approach described in section C.2. Surplus gains are reported as the share of the gains from a first-best policy consisting of optimal individual-product specific taxes. The optimal uniform rates reported in this section differ slightly from the optimal rates reported in the body of the paper because the linear functional form results in slightly different demand weights w_i.

D Pre-registration Appendix

D.1 Sample restrictions

Some participants began the study with multiple synced accounts but dropped some accounts as the experiment proceeded. We pre-registered that we would treat participants with dropped accounts differently depending on the stated reason for dropping. On each survey, we asked participants to explain the reason for their dropped account. If the participant indicated that they had closed their account, then we retain them in the analysis sample and assume that they placed no wagers on that account. In all other cases, including cases where no reason for deactivation was given, we drop the participant from the analysis sample. In practice, only four participants were retained in the sample with incomplete data; all other cases were dropped.

D.2 Pre-registered analyses

The analysis presented in the body of the paper deviated at times from the pre-registered analyses. In the interest of transparency, we justify our choices and present the initial pre-registered analyses in this subsection. There are three important deviations.

First, we pre-registered randomized treatment effect regressions with winsorized log dependent variable for dollars wagered. Unfortunately, our data contains many zeros, which meant that regressions with log dependent variables are sensitive to arbitrary scaling and truncation choices. We illustrate this fact in Table D.1. The first column is our pre-registered specification for the bonus treatment effect, winsorizing at the 5th and 95th percentiles of the $t = 0$ consumption distribution of monthly wagers. Other columns report results for different winsorization choices, and coefficient estimates vary wildly. Tables D.2 and D.3 contain all pre-registered treatment effect regressions for the Bonus and limit treatments respectively.[55]

Second, the pre-analysis plan described a second study, the *screenshot study*. In the screenshot study, we collected information on betting activity with self-reports and screenshots. We also collected predicted net returns and qualitative survey responses as in the main study. Participants in the screenshot study made predictions as in the main study, and half of them were randomized into the *limits treatment*, as in the main study. None of them were randomized to receive the Bet Less Bonus or the history transparency treatment, because data limitations made these impossible to implement.

[55]Due to a coding error in the Qualtrics survey, the *self-control scale* and *gambling literacy scales* did not appear in survey 1. Therefore, we omit them as covariates from the pre-registered specification. We elicited them in survey 2.

We did not include the screenshot study data in the main paper for ease of exposition, improved data quality, and because the sample size was smaller than anticipated: 54 participants completed the final screenshot study. Since we pre-registered that we would conduct relevant analyses pooling the main sample and the screenshot sample, we replicate our main overoptimism result in the pooled sample in Figure D.1. We cannot replicate our self-control problems or Bonus treatment effect results with this sample, since the Bet Less Bonus was not included in the screenshot study.

Third, we pre-registered an average misperceptions estimate that weighted by the inverse variance of net winnings. In the main analysis sample, this estimate is 4.22¢/$. Directionally, this estimate is lower than our raw average difference presented in Figure 5, because low-variance participants overestimate net winnings by less. However, for reasons related to the discussion of shrinkage in Section B.3, this kind of weighted average is not an unbiased estimator of average overoptimism in a heterogeneous population when true overoptimism is correlated with the precision of winnings. The Chen (2024) shrinkage procedure handles the correlation between precision and parameters explicitly and also allows us to estimate heterogeneous overoptimism, so we use it for our main estimates.

D.3 Tables and Figures

	Log weekly wagers (winsorized)					
	(1)	(2)	(3)	(4)	(5)	(6)
t=1 × Bonus	-0.556***	-0.823***	-0.781***	-0.629***	-0.416***	-0.306***
	(0.114)	(0.165)	(0.156)	(0.126)	(0.094)	(0.072)
Bonus × t=2	-0.132	-0.153	-0.149	-0.135	-0.122	-0.112
	(0.116)	(0.168)	(0.159)	(0.129)	(0.097)	(0.075)
Winsorization	[p5, p95] of t=0	[p0.5, p99.5] of t=0	[p1, p99] of t=0	[p3, p97] of t=0	[p10, p90] of t=0	[p20, p80] of t=0
Upper bound (dollars/wk)	3,993.4	23,573.7	13,684.4	6,591.3	1,847.2	701.7
Lower bound (dollars/wk)	9.36	1.17	1.69	5.83	18.8	36.0
Treatment group mean (t=1)	4.18	3.70	3.77	4.05	4.40	4.57
Control group mean (t=1)	4.77	4.59	4.62	4.73	4.84	4.87
Treatment group mean (t=0)	5.09	5.09	5.09	5.09	5.09	5.09
Control group mean (t=0)	5.14	5.14	5.14	5.14	5.14	5.14
R^2	0.61534	0.54066	0.55531	0.60314	0.61997	0.59613
Observations	884	884	884	884	884	884

Table D.1: The sensitivity of regressions with log dependent variables to winsorization choices

Notes: This table reports results from regressions of the following form: $\tilde{Y}_{it} = \alpha \tilde{Y}_{i0} + \tau_t^B B_i + \beta_t X_i + \varepsilon_{it}$ where Y_{it} is dollars wagered per week in period t, $\tilde{Y}_{it}$ is a winsorized version of $\log(Y_{it})$, B_i is a bonus treatment indicator, and X_i is a vector of controls as defined in the main text. We handle zero values by truncating the observed $t = 1, 2$ wagers at quantiles of the $t = 0$ wager distribution (this distribution does not have zeros by construction, since it was an eligibility criterion that participants must have placed wagers in Period 0). Columns differ in the restrictiveness of the winsorization bounds. Column 3 reports our pre-registered specification that truncates at the 5th and 95th percentile.

	Log weekly wagers (winsorized)			Positive play index				Makes life better survey question			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
t=1 × Bonus	-0.549***	-0.604***	-0.524***	0.910***				0.208			
	(0.115)	(0.178)	(0.151)	(0.319)				(0.128)			
t=2 × Bonus	-0.128	-0.039	0.013	0.869***				0.170			
	(0.117)	(0.152)	(0.164)	(0.313)				(0.136)			
Above median t=0 wagers × t=1 × Bonus		0.068									
		(0.231)									
Above median t=0 wagers × t=2 × Bonus		-0.206									
		(0.240)									
Above median t=0 returns × t=1 × Bonus			-0.040								
			(0.224)								
Above median t=0 returns × t=2 × Bonus			-0.275								
			(0.232)								
Bonus					0.889***	0.477	0.953**		0.189	0.087	0.220
					(0.289)	(0.397)	(0.424)		(0.119)	(0.170)	(0.166)
Above median t=0 wagers × Bonus						0.813				0.198	
						(0.584)				(0.238)	
Above median t=0 returns × Bonus							-0.147				-0.067
							(0.585)				(0.233)
Dep. var. Range	-	-	-	[-10,10]	[-10,10]	[-10,10]	[-10,10]	[-5,5]	[-5,5]	[-5,5]	[-5,5]
Dep. var. SD	1.83	1.83	1.83	3.50	3.50	3.50	3.50	1.74	1.74	1.74	1.74
Dep. var. Mean	4.44	4.44	4.44	7.10	7.10	7.10	7.10	1.42	1.42	1.42	1.42
R^2	0.62262	0.62604	0.62414	0.09127	0.09126	0.09436	0.09346	0.40031	0.40028	0.40104	0.40122
Observations	884	884	884	884	884	884	884	884	884	884	884

Table D.2: Pre-registered Bonus treatment effect regressions

Notes: The table contains all pre-registered Bonus treatment effects, as specified in section 3.4 of the pre-analysis plan. An observation is an individual $\times\, t = \{1, 2\}$. We drop two participants who did not place any wagers in the 30 days before the study began. The dependent variables are log weekly wagers (winsorized at the 5th and 95th percentile), the positive play index (defined in Section A.2), and responses to the "sports betting makes life better" survey question (defined in Section A.2), for columns (1-3), (4-7), and (8-11) respectively. All columns include the following covariates: a limits treatment indicator, log winsorized pre-study wagers, elicited beliefs about period-1 winnings, a measure of information conveyed in the information treatment, white, income terciles, and randomization stratum indicators. The "makes life better" columns also include the baseline outcome measured on survey 1 as an outcome. Columns (1), (4), and (8) implement regressions as specified in equation B.6 in Section 6.3, and coefficients are interpreted as in that section. Columns (5) and (9) pool the bonus treatment effect across periods. Columns (2), (6), and (9) add heterogeneity by pre-study wagers. Columns (3), (7), and (11) add heterogeneity by pre-study returns. Heterogeneity is pooled across periods for the survey outcomes.

	Log weekly wagers (winsorized)			Positive play index				Makes life better survey question			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
t=1 × Limits	0.044	0.019	-0.047	0.093				0.223*			
	(0.107)	(0.153)	(0.152)	(0.314)				(0.127)			
t=2 × Limits	0.140	0.086	-0.055	-0.275				-0.105			
	(0.117)	(0.146)	(0.169)	(0.330)				(0.133)			
Above median t=0 wagers × t=1 × Limits		0.019									
		(0.211)									
Above median t=0 wagers × t=2 × Limits		0.086									
		(0.234)									
Above median t=0 returns × t=1 × Limits			0.188								
			(0.220)								
Above median t=0 returns × t=2 × Limits			0.398*								
			(0.229)								
Limits					-0.091	0.202	-0.410		0.059	0.255*	0.114
					(0.294)	(0.389)	(0.426)		(0.119)	(0.149)	(0.170)
Above median t=0 wagers × Limits						-0.560				-0.428*	
						(0.586)				(0.240)	
Above median t=0 returns × Limits							0.612				-0.117
							(0.581)				(0.237)
Dep. var. Range	-	-	-	[-10,10]	[-10,10]	[-10,10]	[-10,10]	[-5,5]	[-5,5]	[-5,5]	[-5,5]
Dep. var. SD	1.83	1.83	1.83	3.50	3.50	3.50	3.50	1.74	1.74	1.74	1.74
Dep. var. Mean	4.44	4.44	4.44	7.10	7.10	7.10	7.10	1.42	1.42	1.42	1.42
R^2	0.61156	0.61379	0.61435	0.09734	0.09665	0.09904	0.10038	0.39653	0.39429	0.40050	0.39548
Observations	884	884	884	884	884	884	884	884	884	884	884

Table D.3: Pre-registered Limits treatment effect regressions

Notes: The table contains all pre-registered Limits treatment effects, as specified in section 3.4 of the pre-analysis plan. An observation is an individual $\times$ $t = \{1, 2\}$. We drop two participants who did not place any wagers in the 30 days before the study began. The dependent variables are log weekly wagers (winsorized at the 5th and 95th percentile), the positive play index (defined in Section A.2), and responses to the "sports betting makes life better" survey question (defined in Section A.2), for columns (1-3), (4-7), and (8-11) respectively. All columns include the following covariates: a limits treatment indicator, log winsorized pre-study wagers, elicited beliefs about period-1 winnings, a measure of information conveyed in the information treatment, white, income terciles, and randomization stratum indicators. The "makes life better" columns also include the baseline outcome measured on survey 1 as an outcome. Columns (1), (4), and (8) implement regressions as specified in equation B.6 in Section 6.3, and coefficients are interpreted as in that section. Columns (5) and (9) pool the bonus treatment effect across periods. Columns (2), (6), and (9) add heterogeneity by pre-study wagers. Columns (3), (7), and (11) add heterogeneity by pre-study returns. Heterogeneity is pooled across periods for the survey outcomes.

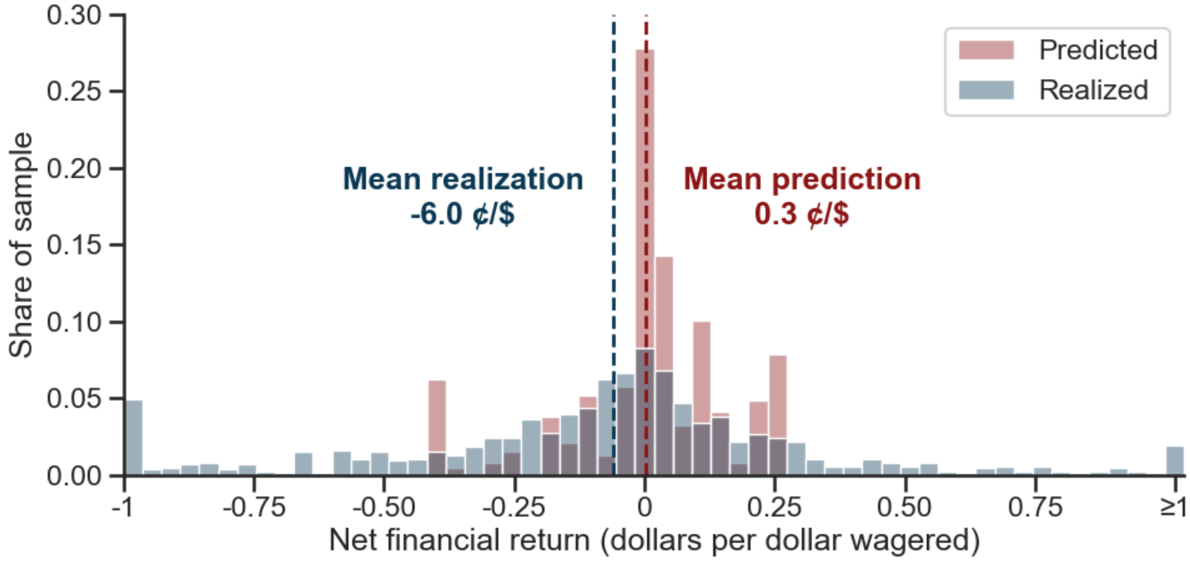

Figure D.1: Comparing predicted financial returns to realized financial returns (pooled sample)

Notes: This figure is the pooled-sample analog of Figure 5. It contains prediction and realization data for the participants in our main analysis sample as well as 52 additional participants from the screenshot study. For participants in the screenshot study, realized returns come from self-reported monthly dollars wagered and winnings on DraftKings and Fanduel. We use screenshots of account history screens on those apps to validate these self-reports. Predictions are censored to lie within $[-0.4, 0.25]$. Dotted lines and annotations represent averages.