Pre-analysis plan for: Price Misperceptions and Self-control Problems in Mobile Sports Betting

Matt Brown, Nick Grasley, and Mariana Guido

April 25, 2024

1 Introduction

Mobile sports betting is expanding rapidly in the United States. A 2018 Supreme Court ruling allowed states to legalize and regulate mobile sports betting.[1] In 2023, mobile sports betting was fully legal in 28 states, and Americans wagered over 100 billion dollars at mobile sportsbooks (American Gaming Association, 2024).

Proponents of deregulation argue that widespread and growing consumption reflects the entertainment value of mobile sports betting. However, critics of deregulation argue that untransparent prices and compulsive behavior cause sports bettors to harmfully overconsume gambling. The optimal gambling regulation depends on the extent to which consumption is driven by bias. However, we lack clear evidence about the implications of bias for welfare in the mobile sports betting market.

How can society design regulations for mobile sports betting that reduce welfare losses from biased consumption while preserving the consumer surplus from sports betting's entertainment value? In this paper, we formulate a simple model of mobile sports betting consumption that allows self-control problems and price misperceptions to cause overconsumption. First, sports bettors might misperceive the costs of betting. Misperceptions could arise because people think they are better at picking winners than they actually are (Chegere et al., 2022; Donkor et al., 2023), or because they neglect that the sportsbook takes a cut from every bet. Second, self-control problems

[1] The ruling was *Murphy vs. NCAA*. It struck down the Professional and Amateur Sports Protection Act, which banned states (except Nevada) from authorizing sports betting.

could drive consumption. Some sports bettors report experiencing irresistible gambling compulsions – they cannot stop themselves from gambling, even though they would like to. To the extent that cost misperceptions and self-control problems cause overconsumption, corrective interventions have the potential to improve welfare. We then provide experimental evidence on the extent to which price misperceptions and self-control problems drive consumption of mobile sports betting in a sample of real bettors. Finally, we use our empirical estimates of model parameters to evaluate the welfare effects of various mobile sports betting regulations.

To study cost misperceptions, we conceptualize the price of wagering on sports as the negative expected net winnings from placing a $1 sports wager. The average net winnings in the population are negative since betting operators ("the house") take a cut from every bet.[2] However, more skilled bettors will face lower prices than lower-skilled bettors. Since neither bettor skill nor the average house cut is easily legible from the information posted on betting platforms, bettors may misperceive the prices that they face. In the experiment, we ask participants to predict their average net winnings from sports betting in the future. We identify cost misperceptions by comparing these predictions to participants' realized net winnings.

We next turn to self-control problems. The psychiatric literature regards compulsive gambling behaviors as a symptom of gambling disorders (Potenza et al., 2019). We model compulsive gambling as arising from temptation (Allcott et al., 2022; Banerjee and Mullainathan, 2010). Temptation causes the short-run self to choose more betting in the moment than they would have planned in the long run.[3] In our experiment, we randomly assign some participants to the *Wager Limits Treatment*, which is designed to allow the long-run self to choose consumption rather than the short-run self. In the Wager Limits Treatment, participants set monthly limits using built-in wager limit tools in their sports betting apps. We require treated participants to make a choice about what wager limit to set over the course of the study. Following the behavioral economics literature on commitment devices and self-control problems, we identify the effect of temptation on consumption by comparing consumption in the Wager Limits Treatment to consumption in a control condition.

Our experiment is also designed to provide evidence about other key parameters. Follow-

[2]The average price of wagering $1 in 2023 was $0.09 (American Gaming Association, 2024).

[3]By modeling self-control problems as arising from temptation instead of from the quasi-hyperbolic discounting model of Laibson (1997), we can derive similar results but with greater analytical tractability in our empirical context.

ing Allcott et al. (2022), we randomly assign participants to a *Bonus Treatment*. Participants in the bonus treatment are paid to reduce their mobile sports wagers. The contemporaneous effect of the bonus treatment on consumption allows us to measure the elasticity of demand. We also track consumption for a 30-day period after the incentive is active. The effect of the bonus treatment in this post-incentive period provides evidence about the extent to which mobile sports betting is habit forming. We use predictions of future wager volume to understand whether consumers are sophisticated about their self-control problems. Finally, we ask survey questions about problematic gambling behaviors and the subjective value of sports betting, which we use to validate our main estimates.

One straightforward intervention to address cost misperceptions is to require sports betting platforms to make past net winnings transparent and salient for users. If participants were aware of their past outcomes, their beliefs about future net winnings would be closer to the truth. We include the *Winnings Information Treatment* to measure the effects of an intervention in this class. Participants in the Winnings Information Treatment are informed of their past net winnings. We measure the effect of this information on beliefs about future net winnings both in the short run and in the long run. We further measure the effect of the Winnings Information Treatment on consumption.

There are several factors beyond the scope of our study that influence the optimal gambling policy. First, we focus only on measuring internalities from sports betting consumption – the uninternalized costs to the sports bettor. To the extent that sports betting causes negative externalities, for example, through spillovers to other negative behaviors[4] or indirect effects on the sporting leagues themselves,[5] the optimal policy is more restrictive. Second, we focus on two important biases: self-control problems and cost misperceptions. To the extent that other biases cause more (less) overconsumption, optimal policies will be more (less) restrictive.[6] Third, we are narrowly focused on the harms associated with overconsumption. We provide no evidence about rare but extremely costly harms from problem gambling, such as suicide. Fourth, our analysis obviously

[4]There is evidence about the effects of sports and gambling problems on intimate partner violence Card and Dahl (2011); Korman et al. (2008). Problem gambling is also linked to alcoholism, though evidence about causal effects is limited.

[5]There have been several scandals involving outcome manipulation in 2024 alone. One NBA coach, J.B. Bickerstaff, reported threats on his life from sports bettors (source).

[6]For example, if sports betting is addictive and biases affect the dynamics of sports betting consumption as in Loewenstein et al. (2003), bettors may overconsume even more.

has nothing to say about non-welfarist reasons for restricting sports betting. Fifth, our estimates apply to a selected sample of sports bettors who were willing to participate in an online study, and both cost misperceptions and self-control problems may differ in the population. Despite these limitations, our study is the first empirical investigation that measures the impact of two key biases on mobile sports betting consumption with an economic model suitable for policy analysis. We hope that future work will continue to paint an even clearer picture for regulators.

In section 2, we outline the details of our experimental design. In section 3, we explain how we will analyze our experimental results.

2 Experimental design

2.1 Sports betting context and terminology

The most prominent firms offering legal mobile gambling in the United States specialize in three types of gambling: *sportsbooks, daily fantasy sports (DFS)*, and *casino games*.[7] The legal status of each form of gambling can differ across jurisdictions, so a single firm will often use different mobile apps and websites to offer different gambling products. This study focuses on betting at mobile sportsbooks.

There are more than 40 legal mobile sportsbooks operating in the United States, but most betting occurs on a few popular platforms. Our study focuses on the following six platforms: DraftKings, FanDuel, BetMGM, Caesars, ESPNBET, and Hard Rock Bet. We focus on these platforms because they are large (they combine for around 90% market share[8]) and because the platforms' designs allow for straightforward implementation of our experimental treatments. Throughout, we refer to these six platforms collectively as the *supported platforms*.

Our primary measure of sports betting activity is in units of *dollars wagered*. The amount wagered is the amount of money a bettor places at risk in a given bet. The amount wagered is distinct from whether the bettor wins or loses, and is distinct from the bet's odds.

[7] Casino games include baccarat, blackjack, slots, roulette, and more. Online poker is covered by different laws than these casino games, and the main platforms in the online poker market (e.g, PokerStars) are different from the largest platforms in the rest of the mobile betting market. Sportsbook betting involves placing bets on an outcome or outcomes of a sports contest or set of contests. DFS is more complex. DFS players pay money to enter tournaments. Then, they select a group of sports players to form a "team." The DFS player earns points when players on their team perform well, and DFS players who do well in tournaments can earn cash prizes. In the U.S., DFS is exempted from some online gambling regulations, and the legal status of DFS can be different from that of sportsbook betting.

[8] https://igamingbusiness.com/sports-betting/sportsbook/h2-us-sports-betting-3/

2.2 Overview

Participants in our experiment take three main surveys, each 31 days apart. We refer to the main surveys as Surveys 1, 2, and 3, and we refer to period t as the 30-day period following survey t.

2.2.1 Recruitment and intake procedures

Advertisements. We recruited participants by running Facebook and Instagram ads over a period from March 1, 2024, to April 3, 2024. We targeted our advertisements at users who were likely to be eligible for the study. Specifically, we targeted:

- Men over the age of 18 (the vast majority of sports bettors are male).
- Users who lived in states where mobile sports betting was legal, but mobile casino gambling was not legal.
- Users who had engaged with social media content related to mobile sports betting.

We also designed our advertisements to appeal specifically to sports bettors. Figure 1 presents a representative ad. We emphasize that the study is about sports betting. We also include our institutional affiliation to build our credibility.

Referral program. We recruit additional participants using a snowball sample design. All participants who were recruited via Facebook ads were invited to refer eligible friends to the study. We refer to a group of participants recruited in this manner by a single person as well as that person as a *referral group*. People in the same referral group were assigned into the same treatment condition for the *Wager Limits Treatment*, *Bonus Treatment*, and *Winnings Information Treatment*.

Screening survey. Upon clicking the ad, participants were taken to a brief screening survey. A participant was eligible to participate in the study if they satisfied the following inclusion criteria:

- The participant said that they wagered at least $100 on sports contests in the last 30 days.
- The participant said that they were a "regular sports bettor" (we explain that "regular" means betting at least once a week on average).
- The participant said that they used at least one of the supported platforms.

In our screening survey, we also elicited the set of accounts at supported platforms that the participant had used in the past 30 days. We refer to these as the *baseline supported accounts*. We then explained our study procedures, obtained consent, and recorded contact information.

Syncing Survey. All participants are required to sync their sports betting accounts to our research study so that we can collect data on their betting activity. We provide details on the data that we collect in Section 2.3. Upon completion of the screening survey, we emailed participants a survey with instructions about how to sync their betting accounts to the research study. We requested that participants sync all baseline supported accounts that they regularly used for sportsbook betting before taking Survey 1. Some participants voluntarily synced additional supported accounts beyond the ones they told us about in the screening survey. We add these to the set of baseline supported accounts for that participant. Some participants also synced unsupported accounts, but we do not include the betting from these unsupported accounts in the main analysis.

2.2.2 Survey 1 (baseline)

Participants take Survey 1 on April 9. All participants are required to refresh the link between their sports betting accounts and the research study before beginning the survey. Participants also have the opportunity to add new accounts at this time. We drop participants from the study if none of their baseline supported accounts have been active in the last 30 days.

Perceived net winnings and winnings information treatment. We begin by measuring participants' beliefs about their net winnings per \$100 wagered. We elicit beliefs about three versions of this number: net winnings per \$100 wagered for all American bettors in 2023,[9] net winnings per \$100 wagered for themselves in 2024 so far, and finally participants' predictions of average net winnings per \$100 wagered over the next 30 days going forward. We use a graphical representation of net winnings and attention checks to ensure that participants understand the object of interest. For the predictions about own future winnings only, we prompt participants to revise their answer if it falls outside of $[-\$40, \$25]$, since we view these values as implausibly far from breaking even. All belief elicitations about past events in the study are incentivized. For 65% of participants, belief elicitations about future events are also incentivized. For the remaining 35%,

[9] The correct number for this question is that Americans lost \$9 for every \$100 they bet in 2023: Source

the belief elicitations about future events are not incentivized.

We now introduce the first main experimental treatment in the study. We show participants in the *winnings information treatment* information about their net winnings in the past. These participants observe their true net winnings per $100 wagered over the past year, as well as the number of bets they placed. They then get the chance to revise their prediction about their own net winnings. Participants in the control condition do not view this information, nor do they revise their predictions.

Predicted future wagers. We turn to predictions of future wager volume. We begin by informing participants of their average weekly wagers over the past 30 days. We also elicit self-reported average weekly wagers for the following other types of betting: non-supported sportsbooks, daily fantasy sports, mobile casinos, in-person casinos, lottery tickets, and all other types of betting. Then, we ask them to predict their average weekly wagers over the next 30 days.[10]

Demand for commitment. After we elicit predictions of future wagers, we introduce participants to the *Bet Less Bonus*. Following Allcott et al. (2022), the Bet Less Bonus is described as follows.

In this part of the survey, we'll explain the Bet Less Bonus. 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.

For example:

- *If you bet $B or more per day over the next 30 days, you'd receive $0 (since you didn't reduce your betting).*
- *If you bet $(B-10) per day over the next 30 days, you'd receive $6.*
- *If you didn't wager at all over the next month, you'd receive $M.*

[10] Sports betting is seasonal. In particular, the 30-day period before April 9 coincides with the NCAA basketball tournaments, which are popular betting events. To ensure that misunderstandings of seasonal patterns do not drive mistaken predictions, we ask participants to think about and write down examples of sporting events that they have bet on in the past 30 days and examples of sporting events they will bet on in the next 30 days before making predictions.

We specify the variables as follows. The benchmark B is the average daily wagers on baseline supported accounts in 2024, rounded up to the nearest \$10. We define the maximum payment M to be $\max\{B/10 \cdot 6, 90\}$. This says that we cap bonus payments at \$90.

Next, again following Allcott et al. (2022), we elicit the predicted consumption reduction from the Bet Less Bonus. We ask participants whether they'd prefer to receive the Bet Less Bonus or a certain payment equal to their expected earnings from the Bet Less Bonus under their own predictions. We finally elicit a fixed payment that would make participants indifferent between the Bet Less Bonus and the fixed payment. All of these elicitations use incentive-compatible procedures, such as a multiple price list. We randomly assign 1 percent of participants to have their Bet Less Bonus condition determined by these questions.

We then inform the remaining 99 percent of participants about whether they'll receive the Bet Less Bonus. Participants in the *Bonus Treatment* receive the bonus, while participants in the Bonus Control do not. Participants must pass an attention check to make sure they understand whether they'll receive the bonus. We also send text message reminders about the Bet Less Bonus to participants in the Bonus Treatment group between surveys 1 and 2.

Wager Limits Treatment. We begin the next module by asking all participants to report their ideal weekly wager volume on baseline supported accounts. Then, we guide participants in the *Wager Limits Treatment* through a procedure that is designed to reduce the impact of self-control problems on betting from baseline supported accounts.

In the Wager Limits Treatment, we ask treated participants to plan a weekly wager limit over all supported accounts. We inform participants that sports betting apps have limit tools to control their gambling behavior. We explain how these limits work: If you choose to set a weekly wager limit of \$K , then you can't wager more than \$K in a given week, and you can't reverse the limit until the week is over. After a week, the wager limit automatically renews unless they choose to change the wager limit. We then explain that it's a study requirement that participants place some weekly wager limit for all of their accounts. We repeatedly highlight that they do not have to choose a *binding* limit – if they want, they can choose to limit themselves to a very high number, like \$999,999. We then ask participants to report their planned weekly wager limits for each of their accounts. We use these self-reported weekly wager limits as evidence of about the extent to which participants desire to restrict their own weekly betting.[11]

After planning a weekly wager limit, we then have participants implement weekly wager limits in all of their apps. Participants view video guides to the limit-setting procedure for each of the baseline supported accounts. In some cases, in-app wager limit tools may differ from a simple weekly wager limit[12]. In these cases, we require them to implement a closely related style of limit in the app. We ask participants to confirm that they set a limit on each possible account, and we also ask them to report the numerical value of the limits they chose. We use these stated limit values as our primary measure of limit choice.

Survey Questions. We conclude by recording responses to a set of questions described in section 2.4 and a set of demographics including: household income, race, age, zip code, and education.

2.2.3 Surveys 2 and 3

The next two surveys arrive on May 10 and June 10, respectively. As in Survey 1, participants are required to refresh the data for their synced accounts before starting. In both surveys, we record predicted net winnings per $100 wagered, stated past consumption from other types of betting, predicted future wagers on baseline supported apps, and the outcomes described in section 2.4.

In survey 2, for participants in the Limits Treatment, we include a set of survey questions about whether the past limits were too high, too low, or just right. We ask them whether they changed their weekly limits during the month, and if they did, what those changes were. We also ask them if they want to revise their limits for the next period, If they say they do want to revise their limits, we provide video instructions about how to do so.

In survey 3, we ask an additional set of questions to measure overall value of sports betting, to validate set limits, and to understand how participants qualitatively perceived changes in their behavior.

WTP for blocking software. In order to estimate the total surplus from sports gambling,

[11] The interpretation is similar to that of the *limit tightness* variable from Allcott et al. (2022).

[12] For example, some apps have a wager limit that bind across both sports betting and casino games. In addition, some apps do not implement wager limits in every state. In practice, the vast majority of participants will have access to wager limits, and all participants can can still opt to set some kind of limit for nonstandard wager limit apps. Since apps can implement different wager limit tools by state, it was not feasible to account for all possible differences in wager limit tools across app-state pairs.

we elicit participants' willingness to pay to use blocking software for the next 30 days. Blocking software, once downloaded, prevents users from accessing any sports betting websites or apps and cannot be removed for a prespecified period, in our case, 30 days.[13] We first describe the blocking technology and then implement an incentivized BDM mechanism where the choice will be implemented for one participant, with an incentive of up to $199.

Limit history screenshot. We require individuals in the Wager Limits Treatment to submit a screenshot of their limit history for one of their main apps. We provide video instructions for participants to find this page and inform them that their final payment is contingent upon uploading a valid screenshot. We use these screenshots as a manipulation check to validate that our limits treatment actually caused participants to set limits.

Open-ended questions: We finally ask a few of open-ended questions for participants to self-report how their behavior has changed by being in the study and how they perceive the consequences of these changes in their everyday lives. We further ask them about what characteristics of sports betting apps they would want to change.

2.3 Synced betting data

We contracted with SharpSports, a company that specializes in allowing individual sports bettors to sync their accounts to third-party platforms. By submitting account credentials in a SharpSports portal, participants shared the full betting activity history with our research team. We maintained static records of these betting histories, which we updated whenever a user triggered a "refresh request." We require that users refresh their betting histories at the beginning of surveys 1, 2, and 3. At the end of survey 3, we unsync all betting accounts. Therefore, we obtain betting activity from before the study began, as well as in periods $t = 1$ and 2. We do not obtain betting activity in period $t = 3$, since this period comes after we unsynced the accounts.

The betting history data contains a detailed record of every bet that a user places. The betting history data does not contain information about deposits, withdrawals, time use, mobile casino activity, daily fantasy sports activity, or limit-setting activity.

[13] We will use BetBlocker as the blocking technology.

2.4 Survey-based outcome measures

Problem Gambling Severity Index (Survey 1 only). The Problem Gambling Severity Index (PGSI) 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). Participants report how often they experienced the following nine consequences in the past year (options: *Never, Sometimes, most of the time, almost always*):

- 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?

Following standard procedures in the psychiatric literature on problem gambling, we construct the *PGSI score* by assigning *Never, Sometimes, Most of the time, and Almost always* scores of 0,1,2 and 3, respectively, and summing across the nine questions.

Gambling literacy index (Survey 1 only). We adapt a gambling literacy scale from **?**. We elicit agreement with the following three items.

- 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)

The options are a five-point Likert scale (*Strongly disagree (-2), disagree (-1), neither agree nor disagree (0), agree (1), strongly agree (2)*). The numeric scores written here correspond to the first question where "strongly agree" is correct; we reverse the order of the scores for the other questions. We sum the question scores to create a *gambling literacy score* which ranges from 6 to -6.

Sports betting makes life better (All surveys). Following Allcott et al. (2022), we asked: *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.

Positive play (all surveys). 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 the following items.

- 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 are 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 *positive play behavior score*, which ranges from -12 to 12.

Baseline risk index All four of these survey scores capture a different dimension of potentially problematic mobile sports betting behavior. We define the *baseline risk index* as the first principal component of the combination of these four scores.

2.5 Screenshot study

We are targeting a sample size of $N = 525$ for the main study.[14] We plan to also recruit a supplemental sample of participants for the *Screenshot Study*. Many participants consent to participation, but then at the last minute decide that they do not actually want to sync their sportsbook accounts. On April 7, we invited these participants to join a future Screenshot Study. Similar to the main study, this consists of three surveys spaced 30 days apart. It begins in the last week of April.

The main difference between the screenshot study and the main study is that participants in the screenshot study do not sync their sportsbook accounts to the study. Instead, we ask them to access screens on their sportsbook accounts that contain summaries of their past wagering activity. We then ask them to report the numbers they see on these screens in survey fields. Finally, we ask them to upload screenshots of their account activity to verify that they were reporting numbers honestly.

We use the self-reported past wagering activity data to construct the amount wagered in period $t = 0, 1, 2$. The exact procedure for constructing amount wagered differs by sportsbook, since different sportsbooks make different information publicly asccessible on summary screens.

The procedures in the Screenshot Study follow those of the main study with two exceptions. First, participants of the Screenshot Study can never view the Winnings Information Treatment, since that treatment relies on information about past winnings. Second, participants in the Screenshot Study do not complete the module about the Bet Less Bonus. We do not elicit WTP for the bonus, nor do we assign participants to the Bonus Treatment.

2.6 Randomization

We are targeting 525 total participants for the main study. Before survey 1, we cross-randomize participants into: *Wager Limits Treatment* (50%) vs. *Wager Limits Control* (50%), *Bonus Treatment* (35%) vs. *Bonus Control* (65%), and *Winnings Information Treatment* (50%) vs. *Winnings Information Control* (50%).

We stratify on two baseline gambling measures. First, we stratify on whether wagers in February 2024 were above or below the median (February is the most recent month for which we have

[14] There is some uncertainty about the sample size, because we do not know how many people will want to sync their accounts at the last minute.

complete data). Second, we stratify on whether realized net winnings in February were above or below the median.

2.7 Constructing the analysis sample

Our main analysis sample consists of participants who complete all three surveys. Whenever possible, we will pool participants from the screenhsot study and the main study into a single analysis. However, for some analyses, we do not collect all relevant variables from screenshot study participants. In this case, we restrict to participants in the main study.

By default, all participants' accounts will remain synced unless the participant deactivates their account or explicitly takes action to unsync their accounts. We require that participants do not unsync their accounts during the study, but we nevertheless anticipate that we may lose data from some accounts during the study. In some cases, we can determine whether we lost data because the account was completely deactivated or the participant specifically severed the link with SharpSports. In the first case, we will infer that the participant placed zero bets on the account going forward. In the second case, if the participant refuses to re-sync their accounts, we will drop the participant's observation from the final sample, since we view this behavior as evidence of attempted data manipulation (for example, a participant in the Bet Less Bonus treatment may have tried to un-sync one account to earn a larger bonus). If we are unable to determine the reason for deactivation, we will drop the participant's observation.

3 Empirical strategy

3.1 Descriptive evidence

We first show baseline summary statistics and balance across treatment conditions. The summary statistics variables include baseline log wagers, baseline realized net winnings per dollar, income, education, age, and the components of the baseline risk index. We also show summary statistics on the gambling characteristics of participants, such as the number of platforms used. We present balance across treatment conditions for these covariates.

Our preferred measure of sports betting consumption is $\log(w_{it})$ where w_{it} is amount wagered in period t on baseline supported accounts. We winsorize this variable at its 5th and 95th percentiles in $t = 0$.

3.2 Price misperceptions

We define price misperception as the difference between perceived expected net winnings per dollar wagered $\tilde{\pi}_{it}$ and the expected net winnings per dollar wagered $\hat{\pi}_{it}$ for person i in period t. In the experiment, we record an individual's predicted net winnings per dollar wagered $\tilde{p}_{i1}$ and their realized net winnings per dollar wagered $\hat{p}_{i1}$. We anticipate that some respondents will misunderstand the question or give implausible predictions of net winnings. Therefore, we restrict attention to observations between -\$0.40 and \$0.25 predicted net winnings per dollar wagered.

We first we present a histogram of predictions $\tilde{p}_i$. In the histogram, we compare predictions to both the average net winnings in the sample $\mathrm{E}[\hat{p}_i]$ and the net winnings of Americans in 2023, which is \$0.09. This plot transparently represents how perceived net winnings compare to average winnings.

We next compute two estimates of population average overconfidence. First, we present the simple average difference between predicted and realized net winnings across participants: $\frac{1}{N}\sum_i \tilde{p}_i - \hat{p}_i$. Second, we present the same average, but weighting participants by the inverse variance of the expected net winnings. This weighting scheme reduces the impact of participants who make few gambles and therefore have noisy net winnings on the final estimate.

We approximate the variance of participant expected net winnings using the implied house odds. We assume that bet $b_j \in B$ is distributed Bernoulli with implied house probabilities q_j. Upon winning, the house pays $\frac{1}{q_j}$ multiplied by the wager amount w_j. Assuming the bets that a participant places are independent, the variance of realized net winnings per dollar wagered is

$$Var(\hat{p}_i) \approx \left(\frac{1}{\sum_j w_j}\right)^2 \cdot \sum_j \left(\frac{w_j}{q_j}\right)^2 q_j(1-q_j).$$

This expression is an approximation of the true variance for two reasons. First, the true probability of a participant winning depends on participant skill. Second, the payoff from winning is not $\frac{w_j}{q_j}$, it depends on the cut taken by the house. However, we expect participant skill to not fluctuate too far from the mean and we expect house cuts to be small, so this expression will be a reasonable approximation.

3.3 Limit treatment effects

We estimate the effects of the limit treatment using the following regression in a pooled sample of limit and screenshot study participants:

$$Y_{it} = \tau_t^L L_i + \beta_t X_i + \alpha_t + \varepsilon_{it}. \tag{1}$$

where i is an individual and t is a 30-day time period ($t \in \{1,2\}$); Y_{it} is an outcome variable (winsorized log wagers, sports betting makes life better question, and the positive play index); L_i is an indicator for the limit treatment respectively; X_i is a vector of baseline covariates; and α_t is a time fixed effect. The covariates we use are: an indicator for the screenshot study; an indicator for the bonus treatment; a set of randomization stratum indicators; winsorized log wagers in period 0; perceived period-0 past winnings; a measure of the shock from the winnings information treatment[15]; income bins; and baseline survey outcomes (if and only if the outcome is a survey variable).

For this and all treatment effect regressions, we cluster standard errors at the referral group level.

For the survey outcomes, we additionally implement a specification where we force τ_1^L to equal τ_2^L. We do this to increase statistical power and because we expect the limit treatment effect on survey outcomes to be nearly constant over time.

3.4 Bonus treatment effect

We estimate the effects of the bonus treatment using the following regression, restricting to participants in the main study who did not have their bonus treatment condition assigned by the multiple price list question:

$$Y_{it} = \tau_t^L B_i + \beta_t X_i + \varepsilon_{it}. \tag{3}$$

[15]We define a variable

$$InfoShock_i = \begin{cases} 0 & \text{if } WinningsInformationTreat_i = 0 \\ TruePastNetWinnings - ReportedPastNetWinnings & \text{if } WinningsInformationTreat_i = 1 \end{cases} \tag{2}$$

and include this variable in the covariate vector.

where the notation is as in the previous section except that B_i is a bonus treatment indicator, and we now include the limit treatment L_i in the set of covariates.

For both the bonus and limit treatments, we report heterogeneous treatment effects along two dimensions: log wagers and past net winnings.

4 References

References

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

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

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

Card, D. and Dahl, G. B. (2011). Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior*. *The Quarterly Journal of Economics*, 126(1):103–143.

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*. Number: 655 Publisher: Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy.

Donkor, K., Goette, L., Müller, M. W., Dimant, E., and Kurschilgen, M. J. (2023). Identity and Economic Incentives.

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

Korman, L. M., Collins, J., Dutton, D., Dhayananthan, B., Littman-Sharp, N., and Skinner, W. (2008). Problem Gambling and Intimate Partner Violence. *Journal of Gambling Studies*, 24(1):13–23.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics*, 112(2):443–477. Publisher: Oxford University Press.

Loewenstein, G., O'Donoghue, T., and Rabin, M. (2003). Projection Bias in Predicting Future Utility. *The Quarterly Journal of Economics*, 118(4):1209–1248.

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. Number: 1 Publisher: Nature Publishing Group.

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. Publisher: Frontiers.

5 Tables and Figures

Figure 1: Sample Facebook Advertisement