

UC Berkeley

CEGA Working Papers

Title

Gamblers Learn from Experience

Permalink

<https://escholarship.org/uc/item/70n6t72s>

Authors

Blumenstock, Joshua E.
Olckers, Matthew

Publication Date

2020-12-02

DOI

10.26085/C3V302

Series Name: WPS
Paper No.: 144
Issue Date: 2 Dec 2020

Gamblers Learn from Experience

Joshua E. Blumenstock and Matthew Olckers



CEGA

Center for Effective Global Action

Working Paper Series

Center for Effective Global Action
University of California



This paper is posted at the eScholarship Repository, University of California. http://escholarship.org/uc/cega_wps Copyright © 2020 by the author(s).

The CEGA Working Paper Series showcases ongoing and completed research by faculty affiliates of the Center. CEGA Working Papers employ rigorous evaluation techniques to measure the impact of large-scale social and economic development programs, and are intended to encourage discussion and feedback from the global development community.

Recommended Citation:

Blumenstock, Joshua E., Olckers, Matthew (2020): Gamblers Learn from Experience. CEGA Working Paper Series No. WPS-144. Center for Effective Global Action. University of California, Berkeley. Text. <https://doi.org/10.26085/C3V302>

Gamblers Learn from Experience

Joshua E. Blumenstock

Matthew Olckers

UC Berkeley

Monash University

November 3, 2020

Abstract

Mobile phone-based gambling has grown wildly popular in Africa. Commentators worry that low ability gamblers will not learn from experience, and may rely on debt to gamble. Using data on financial transactions for over 50 000 Kenyan smartphone users, we find that gamblers do learn from experience. Gamblers are less likely to bet following poor results and more likely to bet following good results. The reaction to positive and negative feedback is of equal magnitude, and is consistent with a model of Bayesian updating. Using an instrumental variables strategy, we find no evidence that increased gambling leads to increased debt.

We thank Milo Bianchi, Jonathan Guryan, Kai Barron and Xiaojian Zhao for helpful and detailed feedback. This project started while Matthew was a PhD student at the Paris School of Economics and a visiting student researcher at UC Berkeley. We are grateful to researchers from both institutions for helpful comments. In particular, we thank Francis Bloch, Margherita Comola, Fabrice Etilé, Simon Gleyze, Sylvie Lambert, Liam-Wren Lewis, Karen Macours and David Margolis.

Blumenstock: jblumenstock@berkeley.edu; Olckers: matthew.olckers@monash.edu

1. Introduction

Enabled by mobile phone technology, sports betting has grown wildly popular in Kenya in recent years. For instance, the term “SportPesa”, the name of Kenya’s largest betting platform, was the most popular search term on Google in 2018.¹ A 2017 GeoPoll survey of 3 879 youth in sub-Saharan Africa found that 76% of Kenyan youth had used their mobile phone for gambling (GeoPoll, 2017). These betting apps are designed to work on even the most basic mobile phones, thus creating new opportunities to gamble for millions of low-income individuals.

Mobile phone-based sports betting has spread from Kenya, an early adopter of mobile money, throughout Africa and other low and middle-income countries. Sports betting is also popular in wealthy nations, with the global sports betting market expected to grow by 140 billion dollars by 2024 (Technavio, 2020). In the United States alone, where sports betting was only recently legalized, an estimated 8 billion dollar market is projected by 2024 (Hancock, 2019).

The rapid growth of sports betting has stoked concerns that sports betting could drive people into financial ruin—especially when credit is easily available. Commenting on the recent legalisation of sports betting in the United States, Victor Matheson, an expert on the economics of sports, raises concern of whether gamblers can learn from experience.²

“Think of all those sports fans who say, ‘You know, I would never buy a lottery ticket. That’s just luck. But I know everything about sports. I should be able to win this.’ And guess what? You can’t beat the casino. These amateurs who think they’re experts don’t stand a chance, but do stand a chance of really getting sucked in. And the question is how quickly can they extricate themselves and realize that ‘Yeah, I’m actually not any good at this.’”

This paper shows that gamblers learn from experience. Gamblers are less likely

¹Of the top five search queries, three were related to sports betting (“sportpesa”, “livescore”, “bet-pawa”). The other two were “kenya” and “news.” Facebook was sixth, with only a quarter of the search volume as SportPesa. See trends.google.com

²See Freakonomics podcast, episode 388.

to bet following negative feedback on their gambling ability and more likely to bet following positive feedback. The magnitude of the reaction is similar for positive and negative feedback.

We demonstrate this result using a rich dataset that captures all mobile money transactions, including detailed information on sports bets and consumer loans, for thousands of Kenyan mobile phone owners. We study how gamblers respond to experience by focusing our analysis on one extremely popular type of bet, the SportPesa ‘jackpot bet’, where the gambler simultaneously makes predictions on 17 soccer matches. The gambler only wins money if he or she correctly predicts 12 or more outcomes correctly. By focusing on how the gambler reacts to performance below the 12-match threshold, we separate the signal of a correct predictions from the income effect of winning a bet.

Our first set of results show that gamblers increase betting outlays in weeks after making successful predictions and reduce outlays after making incorrect predictions. We find gamblers respond with equal magnitude to positive and negative feedback. Gamblers are 2.8 percentage points more likely to place a bet when they predict a relatively high share of matches correctly and 2.9 percentage points less likely to place a bet when they predict a relatively low share of matches correctly. The symmetric response persists for at least three weeks.

The symmetry in response to positive and negative feedback stands in contrast to conventional wisdom. To provide an indication of how people expect gamblers to behave, we conducted a survey of 32 academics and business professionals before writing this paper. The majority (18 respondents) predicted gamblers would be more responsive to positive than to negative feedback. Only 4 respondents predicted the correct result—that gamblers respond equally to positive and negative feedback.

Our second set of results investigate the effects of increased betting on other financial activities of the gambler. We adopt an instrumental variables strategy where we use random variation in gambling outcomes from prior weeks to instrument for the gambler’s current betting outlays. Our identifying assumption is that, conditional on the number of bets placed and individual fixed effects (which control for

gambling ability and other unobserved individual characteristics), the number of correct predictions in a given week is random.

Analyzing financial transactions in the week of a small positive shock to gambling expenditures, we find that gamblers more actively make deposits and withdrawals from their savings account, but observe no significant net effect on savings accumulation (or decumulation). Analyzing loan applications and repayments, we find that the shocks to gambling expenditures do not cause gamblers to apply for consumer credit. These results stand in contrast to a popular narrative that borrowers are gambling on credit (The Economist, 2018). Instead, they are more consistent with recent evidence from Uganda, which suggests that gambling may be a means to generate liquidity for lump sum purchases (Herskowitz, 2020). In addition to entertainment, sports betting may be a substitute to traditional financial investment.

These results contribute to an active debate about the role and regulation of sports betting in developing economies. The negative expected returns from gambling have motivated policymakers to discourage gambling.³ Recent field experiments, which test whether gambling demand can be reduced by better informing gamblers of the likelihood of winning, have shown mixed results (Zenker et al., 2018; Abel et al., 2020).⁴

These findings also relate to studies of investor behavior and stock market speculation. Sports betting shares a similar structure to stock market speculation (Sauer, 1998; Levitt, 2004). And like sports betting, the large number of small investors who

³In Kenya, the government has introduced taxes on revenue and winnings and the gambling regulator has introduced bans on outdoor advertising and celebrity endorsements. In Tanzania, religious leaders have pushed for sports betting to be banned altogether. The Ugandan government has suspended all new gambling licenses and pledged not to renew existing licenses.

⁴In South Africa, researchers asked participants to roll sixes with dice to demonstrate the tiny probability of winning the national lottery, and found that those who took longer to get a six decreased demand for lottery tickets (Abel et al., 2020). The experiment in Thailand demonstrated the probability of winning the national lottery with a poster containing one million dots representing the number of tickets sold and several pins representing number of winning tickets per prize category (Zenker et al., 2018). The demonstration lead to improved knowledge but no change in the willingness to pay for lottery tickets. One way to reconcile these results would be if an individual's belief in their own probability of winning differed from their belief about the population average. Abel et al.'s (2020) experiment, which found changes in behavior, focused on the individual whereas Zenker et al.'s (2018) experiment, which found no change in behavior, focused on the population average.

make low or even negative returns has ignited much debate on whether these investors learn rationally from experience or with a bias (Barber et al., 2020). Our setting allows us to separate feedback on ability from income effects. We find support for a rational model of learning (Mahani and Bernhardt, 2007; Linnainmaa, 2011).

More broadly, our results provide insight into how people respond to signals of their ability. Lab experiments find that gamblers tend to accept correct predictions but explain away incorrect predictions (Gilovich, 1983).⁵ Recent literature has found that people react more to positive than to negative feedback (Eil and Rao, 2011; Sharot et al., 2011; Mobius et al., 2014; Zimmermann, 2020) and tend to forget negative feedback (Chew et al., Forthcoming; Zimmermann, 2020; Gödker et al., 2019; Huffman et al., 2020). Some studies do find people react more to negative feedback (Ertac, 2011; Coutts, 2019), especially concerning an investment decision (Kuhnen, 2015). We find that gamblers react symmetrically to good and bad results in the previous week, which supports a model of Bayesian updating.⁶ A recent experiment focusing only on financial decisions also finds support for Bayesian updating (Barron, 2020). Perhaps ability in financial contexts, such as sports betting, may not be important for self-image (Gotthard-Real, 2017). Therefore, gamblers have no need to discount negative feedback about their gambling ability to retain a positive self-image.⁷

In summary, our paper makes three contributions. First, we use a novel empirical setting to provide evidence that gamblers do learn from experience. The structure of the jackpot bets allows us to separate feedback on ability from income effects. Second, in contrast to a growing literature on biased updating, we find that gamblers display responses of equal magnitude to positive and negative feedback. Third, we use the relationship between past betting results and current betting expenditure to provide suggestive evidence of how increases in betting expenditure affect financial health.

⁵People attribute success to their own skill and explain away failure, even when the task has no element of skill—such as flipping a coin (Langer and Roth, 1975).

⁶Symmetric updating does not rule out biased priors (Grossman and Owens, 2012; Buser et al., 2018). In our data, we cannot observe the gambler's prior on his or her betting ability.

⁷Benjamin (2019) emphasizes that the evidence on biased updating is “confusing” and there is no unified explanation for the range of findings.

2. Model of Gambling Behavior

We provide a model of gambling behavior to study how gamblers learn from experience.

A gambler can predict the outcome of a match with probability $\theta \in [0, 1]$, which represents the gambler's ability. Sports betting differs from other gambles such as lotteries in that θ can differ between gamblers. In sports betting, some gamblers may have more skill in determining the outcome of a match, whereas in a lottery, the probability of winning is identical for all gamblers.

The gambler does not know his ability. He has a prior belief $\Pr[\theta] \sim \text{Beta}(s_w, u_w)$. Using the Beta distribution to specify the gambler's prior allows for an intuitive interpretation. Suppose, before starting to gamble, the gambler watches m_w matches to test how many he can predict correctly. His prior belief of θ has the distribution $\text{Beta}(s_w, u_w)$ if the gambler predicts the outcome of s_w matches successfully and u_w matches unsuccessfully.⁸

Each week, a gambling operator offers a bet on the outcome of a sports match. The gambler chooses whether to bet as a function of his expected ability, $E[\theta] = \frac{s_w}{m_w}$ (where $m_w = s_w + u_w$). The gambler only bets if

$$E[\theta] \geq c_t = \mathcal{N}(\mu_c, \sigma_c),$$

where c is a cutoff in week t . The cutoff c_t captures factors other than the gambler's belief of his ability which impact his decision to bet.⁹

Once the gambler starts betting, he uses Bayes Rule to update his belief of his betting ability θ . Let s_b be the number of successful bets in m_b matches, $s_b \sim \text{Binomial}(m_b, \theta)$. Let u_b be the number of unsuccessful bets. Since the Beta distribution is a conjugate

⁸We assume a Haldane prior for gamblers who have never watched any soccer matches before.

⁹We could also include the gambler's uncertainty about his ability in the model, as measured by the variance of his beliefs, $\text{Var}[\theta] = \frac{(s_w + s_b)(u_w + u + b)}{(m_w + m_b)^2(m_w + m_b + 1)}$. This variance approaches zero as the number of matches the gambler watches, m_w , and bets on, m_b , increases.

of the Binomial distribution, the posterior belief also has a Beta distribution.

$$\Pr[\theta|s_b] \sim \text{Beta}(s_w + s_b, u_w + u_b)$$

$$E[\theta|s_b] = \frac{s_w + s_b}{m_w + m_b}$$

This relationship between the posterior and prior beliefs has an intuitive interpretation. The prior simply adds s_w successful bets and u_w unsuccessful bets to the gambler's record of s_b successful bets and u_b unsuccessful bets.

After watching or betting on many matches, the gambler will learn his ability. A streak of lucky outcomes may cause him to overestimate his ability for a period, but eventually his belief $E[\theta|s_b]$ will converge to θ .

How many matches does the gambler need to observe to be confident of his betting ability? The standard deviation of the posterior provides an indication. Let $m = m_w + m_b$ be the number of matches observed. The standard deviation σ_θ of the posterior with ability θ is

$$\sigma_\theta = \sqrt{\frac{\theta(1-\theta)}{m+1}}.$$

Suppose the gambler wants σ_θ to be smaller than ϵ . We have

$$\begin{aligned} \sigma_\theta &\leq \epsilon \\ \sqrt{\frac{\theta(1-\theta)}{m+1}} &\leq \epsilon \\ \frac{\theta(1-\theta)}{\epsilon^2} - 1 &\leq m. \end{aligned}$$

As an example, if we let $\theta = 0.4$, the gambler would need to observe approximately 600 matches to be 95 percent confident (within two standard deviations) that his ability lies between 0.36 and 0.44. The gambler needs to watch or bet on a large number of games to be confident in his ability.

If the gambler fails to learn from experience, by ignoring past results or only remembering successful bets, he may never learn his ability and become trapped

by his belief that he will eventually win big. Our unique data allows us to test how gamblers react to feedback. We find that gamblers do learn from experience and react in a similar way to the Bayesian gambler in our model.

3. Context and Data: Mobile Money and Sports Betting

In 2007, the Kenyan telecom company Safaricom launched M-Pesa, a mobile phone-based financial platform that allows users to conduct basic financial transactions over the mobile phone network. Since 2007, mobile money has proliferated in the developing world and today there are more than a billion registered mobile money accounts across 290 mobile money services in at least 95 countries (GSMA, 2019). Currently, mobile money accounts are more common than bank accounts in most African nations.

The introduction of mobile money has been associated with important welfare effects. In Kenya, mobile money has been linked to improved risk-coping (Jack and Suri, 2014) and is estimated to have lifted as many as 194 000 Kenyans out of poverty (Suri and Jack, 2016). Businesses also benefit from mobile phone adoption as payments for goods and services can be collected with mobile money. Consequently, many in the policy and aid communities view mobile financial services as key to improving financial inclusion among the poor (GSMA, 2019; Lauer and Lyman, 2015).

The widespread adoption of mobile money has been accompanied by the rapid spread of mobile phone-based gambling. Gambling operators have leveraged the M-Pesa mobile money network to collect bets on soccer matches and other sports. Since the transaction cost of collecting bets and disbursing winnings has been lowered by mobile phone technology, the companies offer bets for as little as 1 KSH (0.01 USD).

Many Kenyans have strong interest in sport, especially soccer. Gambling operators have tapped into this interest and sports betting has been widely adopted in a short period of time.¹⁰ In just five years of operation, SportPesa, the leading sports

¹⁰A survey of 1 130 Kenyan youth between the ages of 17 and 35 conducted by GeoPoll in 2017 found that 60 percent of respondents had placed bets on football matches in the past and a further 16 per-

betting operator, has gathered enough revenue to sponsor two English Premiership soccer teams, deals estimated at 4 and 12 million dollars per year.¹¹

We use a dataset that contains basic metadata on all mobile money transactions conducted by an anonymized sample of Kenyan smartphone owners. These data were collected by a smartphone app, installed by 93 565 Kenyans between April and June 2016 and in July 2017. No personally identifiable information was used in this study.

We use a subsample of 58 215 individuals for whom we observe a payment to a gambling company. The large number of gamblers in our sample matches survey evidence on gambling incidence in Kenyan youth (GeoPoll, 2017). Among the subsample of gamblers, 77 percent are men and 75 percent are 35 years of age or younger. Appendix A contains additional descriptive statistics.

Our analysis focuses on transactions between individuals and private companies. For each transaction, we observe the value of the transaction and the name of the company involved. Since sports betting in Kenya operates almost exclusively via mobile money, we can observe the deposits and withdrawals the gambler makes. We also observe the details of transactions with the betting company, SportPesa.

4. Measuring Gambling Ability with Jackpot Performance

Our model of gambling assumes that there is heterogeneity in gambling ability, such that high ability gamblers have a greater probability of predicting the outcome of a sports match correctly than low ability gamblers. To support this assumption, we show that certain gamblers consistently predict more matches correctly than others.

cent had tried other forms of betting (GeoPoll, 2017).

¹¹A leaked spreadsheet from Kenya's betting regulator revealed that wagers totalled 300 million USD in May 2019 and SportPesa is estimated to hold two thirds of this market.

4.1 Jackpots as a Standardized Measure of Gambling Ability

One empirical challenge is to define a standardized measure of betting ability. Sports betting companies offer a range of different bets. For example, gamblers can bet on single matches or they can chain matches together to increase the payout (and decrease the probability of winning). Also, individual matches differ in odds. It is easier to predict the outcome of a match with a clear favorite than to predict the outcome of a match with similar strength teams. We address this challenge by studying the number of correct predictions on the weekly jackpots offered by SportPesa, Kenya's leading sports betting operator.

Each week, SportPesa selects 13 matches for a midweek jackpot and 17 matches for a weekend jackpot (called the MegaJackpot). The jackpot is awarded if the outcome of all the pre-selected matches are predicted correctly. Bonuses are awarded for 10 or more correct predictions on the midweek jackpot and 12 or more correct predictions on the weekend jackpot. The amount of the bonus increases exponentially with the number of correct predictions. For example, the bonuses for the weekend jackpot on 31 March 2019 were approximately USD \$460, \$2 190, \$8 140 and \$64 300 for 12, 13, 14 and 15 correct predictions.

The jackpots provide a standardized measure of betting ability, both between gamblers and across time. Since SportPesa selects the matches each week we need not worry about how the gambler picks the matches. The types of matches SportPesa selects each week are also very similar. SportPesa has an incentive to select matches with no clear favorite, where the odds are balanced across the three possible outcomes of home team wins, away team wins, or a draw. The more difficult it is to predict the outcome of a single match in the jackpot, the less likely a gambler will predict all the matches correctly.

The jackpots are extremely popular. In our data, we observe midweek or weekend jackpot bets placed by over 17 000 individuals, which represents 30 percent of all individuals for whom we observe some type of betting transaction.

4.2 Estimates of Gambling Ability

We directly estimate the gambler’s ability θ and find significant differences in gambling ability between gamblers. Each jackpot match provides a standardized type of bet to compare gamblers and estimate their ability. As described in Section 2, the expectation of ability, θ , for s correct predictions in m matches is $E[\theta|s, m] = \frac{s}{m}$. Provided s is generated from a binomial distribution with a θ probability of success, $E[\theta|s, m]$ follows a Beta distribution. We can derive Clopper-Pearson confidence intervals around $E[\theta|s, m]$ from the Beta distribution.


The range of the confidence interval for a particular gambler is a function of the number of matches we observe for that gambler. We focus initially on a smaller group of gamblers for which we observe a large number of matches.¹² In this sample, the smallest confidence interval is 0.06, a gambler for whom we observe 1 129 match predictions.

To sharpen intuition, Figure 1 shows the expected ability for the subsample of gamblers with at least 500 predictions—for whom our ability estimates are most precise. The figure orders individuals from left to right by the number of matches. As the number of matches increases, the confidence interval around the estimates becomes smaller. We show a histogram of the estimated ability for the larger group of 2 016 gamblers for whom we observe predictions on at least 100 matches. We plot the expected ability from random guessing with an orange horizontal line. The gamblers who have a confidence interval above this line, have a better expected betting ability than random guessing. In this sample, 47 percent do better than random guessing.

The figure highlights the heterogeneity in gambling ability. We can observe several gamblers whose confidence intervals do not intersect, which shows that some gamblers have higher ability than others. Unlike many other types of gambling, sports betting has a role for skill.

¹²In Appendix D, we use the full sample of gamblers to document correlations in each gambler’s betting results over time. We find that gamblers who predict a high number of matches correctly in the current week are more likely to predict a high number of matches correctly in the following week.

Figure 1: Estimates of betting ability for frequent gamblers



Notes: The histogram drawn on the vertical axis shows the distribution of betting ability (defined as the long-run proportion of correct predictions) for the 2016 gamblers for whom we observe bets on at least 100 matches. To the right of the histogram we show the subsample of gamblers for whom we observe bets on at least 500 matches. For each of these gamblers, we plot their estimated ability with a black dot and the Clopper-Pearson 95 percent confidence interval as a line with whiskers. The fact that many of the confidence intervals do not intersect illustrates heterogeneity in ability. The orange line shows the fraction of correct predictions that would result from random guessing.

5. Empirical Evidence of Learning

To study how a gambler responds to feedback on his ability, we must separate the positive signal of a correct prediction from the income effect of a winning bet. Sport-Pesa's midweek and weekend jackpots provide an ideal setting to separate these effects. The midweek jackpot awards prizes for 10 or more correct predictions out of 13 matches and the weekend jackpot awards prizes for 12 or more correct predictions out of 17 matches. Any variation in the number of correct predictions below 10 for the midweek jackpot and below 12 for the weekend jackpot has no income effect. We can isolate the effect of feedback by focusing on this range of the jackpot results. In so doing, we note that there is likely an even stronger signal of quality in weeks when the gambler actually wins money—but we cannot use this variation because it is confounded by the direct effect of the money won.

The probability of winning a prize for the jackpots is very low, so focusing on the

range of results below the cutoff for prizes still provides us with ample variation in signals of ability. If a gambler selected the favorite (the team most likely to win) for every match in the jackpot, he would have a less than one percent chance of winning a prize. In our sample, 0.68 percent of midweek jackpot bets and 0.94 percent of weekend jackpot bets won a prize.

5.1 Gamblers Respond to Past Results

We study gamblers' betting behavior in response to the share of correct predictions on the jackpot in the previous week. We use the following specification:

$$Betting_{it} = \beta \text{Jackpot}_{i(t-1)} + \gamma \text{Jackpot Tickets}_{i(t-1)} + v_i + \epsilon_{it} \quad (1)$$

where v_i are individual fixed effects.¹³

We use two measures of betting behavior as dependent variables. First, we use an indicator for placing either a midweek or weekend jackpot bet. Second, we use the inverse hyperbolic sine (which is similar to a log transformation) of mobile money transfers to betting accounts. $\text{Jackpot}_{i(t-1)}$ measures the share of correct predictions on jackpot bets in the previous week. We are interested in the magnitude of β , the impact in the share of correct predictions in the previous week on the propensity to bet in the current week. We control for $\text{Jackpot Tickets}_{i(t-1)}$, which indicates the number of jackpot tickets purchased in the previous week.

Results, shown in Table 1, indicate that gamblers react significantly to the previous week's betting results. A one standard deviation increase in the share of correct predictions in the previous week increases the probability that the gambler plays the jackpot in the current week by 1.78 percentage points and increases mobile money transfers to the betting account by 5.01 percent on average.

¹³Our main specification includes individual fixed effects to control for individual differences in betting ability and the propensity to gamble. Following recent work by Imai and Kim (Forthcoming) and others, we do not include time fixed effects as the use of individual and time fixed effects can obscure interpretation, except under conditions of linearly additive effects (Kropko and Kubinec, 2020; de Chaisemartin and D'Haultfœuille, 2020). For reference, we show results with both fixed effects in Appendix Table A5.

Table 1: Response to previous week's betting results

	<i>Dependent variable:</i>	
	Placed jackpot bet (1)	Betting expenditure (2)
Share correct in week $t - 1$	0.148 (0.017) [0.116, 0.180]	0.424 (0.095) [0.237, 0.610]
Individual fixed effects	Yes	Yes
Week fixed effects	No	No
Individuals	15 715	15 715
Weeks	119	119
Observations	70 390	70 390

Notes: All specifications control for the number of midweek jackpot and weekend jackpot tickets. The sample excludes individual-week observations where the individual won a prize on the jackpot in the previous week. Robust standard errors are shown in round brackets and the 95 % confidence interval is shown in square brackets.

5.2 Symmetric Reaction to Positive and Negative Feedback

Many experiments find biased learning from feedback. People react to positive feedback and ignore or forget negative feedback—especially when feedback involves measures of intelligence or performance.¹⁴ In contrast, we find that gamblers react symmetrically to positive and negative feedback.

To study the difference between positive and negative feedback, we construct a categorical variable from the share of correct predictions on jackpot matches. We start by calculating the mean share of correct predictions for each individual. We then define three categories: (i) positive feedback as more than 10 percent above the mean, (ii) negative feedback as 10 or more percent below the mean and (iii)

¹⁴See, for example, Eil and Rao (2011); Sharot et al. (2011); Mobius et al. (2014); Godker et al. (2019); Zimmermann (2020); Huffman et al. (2020) and Chew et al. (Forthcoming).

between 10 percent above or below the mean as the base category. In Appendix E we show that results are not sensitive to the choice of a 10 percent threshold.¹⁵

We test for biased learning using the following specification:

$$\begin{aligned} \text{Bet on jackpot}_{i(t+\tau)} = & \beta_p \text{Positive Feedback}_{it} + \beta_n \text{Negative Feedback}_{it} + \\ & \gamma \text{Jackpot Tickets}_{it} + v_i + \epsilon_{i(t+\tau)} \end{aligned} \quad (2)$$

where v_i are individual fixed effects and we control for the number of jackpot tickets purchased in a given week. We are interested in comparing the coefficients β_p and β_n to compare how gamblers respond to positive versus negative feedback of their betting ability. We investigate how betting results in week t change the likelihood the gamblers places a jackpot bet in week $t + 1, t + 2, t + 3$ and $t + 4$.

Table 2 shows our estimates of the positive and negative feedback effects. Gamblers are more likely to bet following positive feedback and less likely to bet following negative feedback of their gambling ability. Remarkably, the magnitude of the coefficients are very similar in all specifications. The effect remains statistically significant for three weeks.

As this symmetry was unexpected, we conducted a survey of 32 people in which we asked the question, “Each week a gambler bets on a number sports matches. How will the gambler respond to the past week’s betting results?” The possible responses are shown below, along with the number of respondents who selected that answer in parentheses.

- *The gambler will be more responsive to a positive result in the previous week than a negative result.* (18)
- *The gambler will respond equally to positive and negative feedback.* (4)
- *The gambler will be more responsive to a negative result in the previous week than a positive result.* (8)
- *The gambler will not respond to the past week’s betting results.* (2)

¹⁵We restrict analysis to gamblers for whom we observe at least one week in all three categories. Otherwise, the coefficients on the positive and negative feedback indicators would be estimated on different sets of individuals. We explain this choice in more detail in Appendix C.

Table 2: Response to positive and negative feedback

	Place jackpot bet in week:			
	$t + 1$	$t + 2$	$t + 3$	$t + 4$
Positive feedback (β_p)	0.028 (0.005)	0.011 (0.006)	0.014 (0.006)	0.008 (0.006)
Negative feedback (β_n)	-0.029 (0.005)	-0.011 (0.006)	-0.016 (0.005)	-0.012 (0.006)
P-value $ \beta_p = \beta_n $	0.433	0.487	0.349	0.281
Individual Fixed Effects	Yes	Yes	Yes	Yes
Week Fixed Effects	No	No	No	No
Individuals	4 399	4 111	3 900	3 704
Weeks	119	118	117	116
Observations	48 948	45 626	43 216	40 959

Notes: Dependent variable is an indicator for whether the individual places a jackpot bet in the weeks following a week of positive feedback or negative feedback on their gambling ability. Positive (or negative) feedback is defined as when the fraction of correct predictions made by the individual is more than 10% higher (or lower) than their average rate of correct predictions. All specifications control for the number of midweek jackpot and weekend jackpot tickets. The sample excludes individual-week observations where the individual won a prize on the jackpot in week t . We report robust standard errors in parenthesis below each estimate.

Only 4 respondents predicted our result while the majority predicted that gamblers would be more responsive to positive than to negative feedback.

6. How Do Gamblers Fund Gambling Expenditure?

Our final set of results explore the causal impact of gambling on the financial decisions of gamblers. Policymakers worry that gamblers may use credit to fund betting—a recipe for financial ruin. Anecdotes suggest that betting increases the demand for loans and causes bankruptcy, but we are not aware of causal evidence of the impact of increased betting expenditure on gamblers’ other financial behaviors.¹⁶

6.1 Identification and Estimation

We are interested in estimating the causal effect of increased betting expenditure on the use of savings and credit that we observe in our data. Since betting expenditure, in general, is not random, we use an instrumental variables strategy. For an instrument to be valid in this context, it must be relevant (i.e., correlated with betting), and it must satisfy the exclusion restriction that it should be related to loan default only through the endogenous measure of betting.

We use the share of correct jackpot predictions in week $t - 1$ as an instrument for betting expenditure in week t . In Section 5.1, we demonstrated the relevance of this instrument by showing that an increase in the number of correct predictions on SportPesa’s jackpot increased the propensity to bet in the following week. The second specification in Table 1 shows the strong relationship between our instrument, the share of correct predictions in the previous week, and the endogenous variable, betting expenditure. The partial F-statistic for this specification is 19.87, well above the standard benchmark of 10 (Stock and Yogo, 2005).

Our exclusion restriction requires that the number of correct predictions on the jackpot is random conditional on the gambler’s ability (approximated with an individual-

¹⁶For instance, Dahir (2017) notes that “gambling addiction is on the rise in Kenya and leaving young people bankrupt and suicidal.” An article by The Economist (2018) warns that sports betting may be linked to high default rates on consumer loans: “Anecdotal evidence is mounting of abuses—most notoriously of young Kenyans borrowing to splurge on online betting sites.”

specific fixed effect) and the number of jackpot tickets purchased (the endogenous regressor). In our context, it is difficult to imagine how outcomes may be impacted by the previous week’s jackpot results other than through the current week’s betting behavior. One concern with the exclusion restriction is that winnings from prior weeks could directly impact outcomes in future weeks. However, as noted in Section 4.1, we only consider observations when the gambler’s predictions fell below the threshold where the gambler wins money (10 in the midweek jackpot and 12 in the weekend jackpot) . We exclude these observations from the sample to ensure our instrument is not driven by an income effect—though Appendix E shows that results are qualitatively unchanged when these observations are not excluded. Since the probability of winning money from a jackpot bet is small, the sample size reduces by less than one percent.

Empirically, we estimate the impact of increased betting expenditure on measures of gambler i ’s use of savings and credit at time t , denoted by Y_{it} , as:

$$Y_{it} = \beta \operatorname{arsinh}(\widehat{\text{Betting expenditure}}_{it}) + v_i + \gamma \text{jackpot tickets}_{i(t-1)} + \epsilon_{it} \quad (3)$$

where v_i are individual fixed effects and betting expenditure is instrumented by the share of correct jackpot predictions in the previous week. For all non-negative continuous outcomes, we use the inverse hyperbolic sine transformation, $\operatorname{arsinh} = \ln(x + \sqrt{x^2 + 1})$, and interpret β as the elasticity between betting expenditure and the outcome (Bellemare and Wichman, 2020).

We focus our analysis on a few specific margins of financial account use that we observe in our data:

- *Savings withdrawals*: The value of withdrawals from an individual’s M-Shwari savings account during the week, in Kenyan Shillings (KSH), scaled with the inverse hyperbolic sine transformation. M-Shwari is the digital banking service offered by M-Pesa, Kenya’s dominant mobile money money service (FinAccess, 2019).
- *Savings deposits*: The value of deposits into the individual’s M-Shwari account

in a given week (KSH), scaled with the inverse hyperbolic sine transformation.

- *Net savings deposits*: The value, in KSH, of all M-Shwari deposits minus the value of all withdrawals.
- *Applied for a loan*: An indicator for whether the individual applied for a loan from one of several popular lending companies.¹⁷
- *Loans received*: The value of loans received in a given week, defined as a payment from a loan company to the individual's mobile money account, scaled by the inverse hyperbolic sine transformation.
- *Loan repayments*: The value of loans repaid in a given week, defined as a payment from the individual to a loan company, scaled by the inverse hyperbolic sine transformation.

This is not the full set of savings and credit options available to gamblers—they could be transacting on accounts that are not mediated by their phone—but mobile money is the primary formal financial ecosystem used by most Kenyans, and one of the key drivers of financial inclusion in Kenya.¹⁸

6.2 Results

Instrumental variables estimates of the impact of betting expenditures on use of savings and credit are presented in Table 3. In columns 1 and 2, we find that increases in gambling expenditures cause gamblers to more actively use their savings accounts—both increasing the value of withdrawals and deposits to their accounts. Specifically, a one percent increase in gambling expenditure increases withdrawals from a savings account by 0.547 percent and increases top-ups into the savings account by 0.388 percent. The increase in withdrawals is of similar magnitude to the increase in deposits, and we cannot reject the null hypothesis that there is no effect on net savings accumulation (column 3).

¹⁷The list of companies includes: M-Shwari, Tala, Branch, KCB, Equity Bank and Co-op Bank.

¹⁸For instance, nationally representative survey evidence by FinAccess (2019) indicates that while 79% of Kenyans have mobile money accounts, only 30% have bank accounts.

Increases in gambling expenditures do not have a statistically significant impact on borrowing behavior. The estimates in columns 4-6 are imprecise, but we can reject with 95 confidence that a one percent increase in gambling expenditure would increase loan applications by more than 0.10 percentage points.

Taken together, the results in Table 3 suggest that gamblers are not relying primarily on debt to fund their gambling activities. If anything, gamblers appear to be paying for gambling with the balance in their savings account—but without a clear negative effect on their net deposits. These results are consistent with the interpretation that sports betting is a type of risky investment activity. Gamblers may move money between their risk-free savings accounts and buy risky sports bets to generate large lump sums (Herskowitz, 2020).

It should be noted that our instrumental variables strategy identifies a local average treatment effect, and should thus be interpreted as the causal effect of relatively small increases in betting expenditures. Large shocks that dramatically alter a gambler's betting expenditures may have qualitatively different effects on financial behavior.

Table 3: Instrumental variables estimates of the impact of increased betting expenditure

	<i>Dependent variable:</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	Savings withdrawal	Savings deposit	Net savings deposit	Applied for loan	Loans received	Loan repayments
<i>Units</i>	<i>Elasticity</i>	<i>Elasticity</i>	<i>KSH</i>	<i>Indicator</i>	<i>Elasticity</i>	<i>Elasticity</i>
Betting expenditure	0.547 (0.202) [0.151, 0.944]	0.388 (0.187) [0.021, 0.754]	124.358 (353.13) [-567.77, 816.49]	0.025 (0.038) [-0.049, 0.099]	-0.204 (0.316) [-0.823, 0.416]	-0.442 (0.325) [-1.080, 0.196]
Individual fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Week fixed effects	No	No	No	No	No	No
Individuals	15 715	15 715	15 715	15 715	15 715	15 715
Weeks	119	119	119	119	119	119
Observations	70 390	70 390	70 390	70 390	70 390	70 390

Notes: Independent variable is the inverse hyperbolic sine of betting expenditures, instrumented with the past week's jackpot performance. The dependent variable in columns (1) and (2) are the inverse hyperbolic sine of withdrawals from and deposits to the M-Shwari savings account. The dependent variable in column (3) is the total deposits minus withdrawals, in Kenyan Shillings (KSH), with an approximate exchange rate of 100 KSH to \$1 USD. Dependent variables in columns (4)-(6) capture loan behavior on all loan providers who transact with mobile money, including M-Shwari, Branch and Tala. All specifications control for the number of midweek jackpot and weekend jackpot tickets purchased in the previous week. The sample excludes individual-week observations where the individual won a prize on the jackpot in the previous week. We report robust standard errors in round brackets and the 95 percent confidence intervals in square brackets.

7. Conclusion

Our analysis indicates that sports bettors differ in ability, react to past results and react symmetrically to positive and negative feedback. We also provide carefully identified, although imprecise, estimates of the impact of increased betting outlays on other types of financial activity. We do not find strong support for the hypothesis that sports betting systematically drives people into financial ruin.

Our results contradict a common intuition that gamblers continue betting without any regard for past performance, or that they react asymmetrically to wins and losses. Gamblers do learn from experience. However, as our model shows, this learning may require a large number of bets before the gambler has an accurate understanding of his or her own ability.

References

- Abel, Martin, Shawn Cole, and Bilal Zia, "Changing Gambling Behavior through Experiential Learning," *The World Bank Economic Review*, 06 2020. lhaa016.
- Barber, Brad M, Yi-Tsung Lee, Yu-Jane Liu, Terrance Odean, and Ke Zhang, "Learning, Fast or Slow," *The Review of Asset Pricing Studies*, 2020, 10 (1), 61–93.
- Barron, Kai, "Belief Updating: Does the 'Good-News, Bad-News' Asymmetry Extend to Purely Financial Domains?," *Experimental Economics*, 2020, pp. 1–28.
- Bellemare, Marc F and Casey J Wichman, "Elasticities and the Inverse Hyperbolic Sine Transformation," *Oxford Bulletin of Economics and Statistics*, 2020, 82 (1), 50–61.
- Benjamin, Daniel J, "Errors in Probabilistic Reasoning and Judgment Biases," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 2, Elsevier, 2019, pp. 69–186.
- Buser, Thomas, Leonie Gerhards, and Joël van der Weele, "Responsiveness to Feedback as a Personal Trait," *Journal of Risk and Uncertainty*, 2018, 56 (2), 165–192.
- Chew, Soo Hong, Wei Huang, and Xiaojian Zhao, "Motivated False Memory," *Journal of Political Economy*, Forthcoming.
- Coutts, Alexander, "Good News and Bad News are Still News: Experimental Evidence on Belief Updating," *Experimental Economics*, 2019, 22 (2), 369–395.
- Dahir, Abdi Latif, "Gambling Addiction is on the Rise in Kenya and Leaving Young People Bankrupt and Suicidal," *Quartz*, 2017.
- de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, September 2020, 110 (9), 2964–96.
- Eil, David and Justin M Rao, "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself," *American Economic Journal: Microeconomics*, 2011, 3 (2), 114–38.
- Ertac, Seda, "Does Self-Relevance Affect Information Processing? Experimental Evidence on the Response to Performance and Non-Performance Feedback," *Journal of Economic Behavior & Organization*, 2011, 80 (3), 532–545.

- FinAccess, “2019 FinAccess Household Survey,” Technical Report, National Bureau of Statistics, Kenya 2019.
- GeoPoll, “Mobile Gambling among Youth in Sub-Saharan Africa,” <https://www.geopoll.com/blog/mobile-gambling-among-youth-in-sub-saharan-africa/> 2017. Accessed: 2019-01-08.
- Gilovich, Thomas, “Biased Evaluation and Persistence in Gambling,” *Journal of Personality and Social Psychology*, 1983, 44 (6), 1110–1126.
- Gödker, Katrin, Peiran Jiao, and Paul Smeets, “Investor Memory,” *Available at SSRN 3348315*, 2019.
- Gotthard-Real, Alexander, “Desirability and Information Processing: An Experimental Study,” *Economics Letters*, 2017, 152, 96–99.
- Grossman, Zachary and David Owens, “An Unlucky Feeling: Overconfidence and Noisy Feedback,” *Journal of Economic Behavior & Organization*, 2012, 84 (2), 510–524.
- GSMA, “State of the Industry Report on Mobile Money,” Technical Report, GSM Association 2019.
- Hancock, Alice, “A Bet on America: The Sports Gambling Gold Rush,” *Financial Times*, October 2019.
- Herskowitz, Sylvan, “Gambling, Saving, and Lumpy Liquidity Needs,” *American Economic Journal: Applied Economics*, 2020.
- Huffman, David, Collin Raymond, and Julia Shvets, “Persistent Overconfidence and Biased Memory: Evidence from Managers,” *Working Paper*, 2020.
- Imai, Kosuke and In Song Kim, “On the Use of Two-way Fixed Effects Regression Models for Causal Inference with Panel Data,” *Political Analysis*, Forthcoming.
- Jack, William and Tavneet Suri, “Risk Sharing and Transactions Costs: Evidence from Kenya’s Mobile Money Revolution,” *American Economic Review*, 2014, 104 (1), 183–223.
- Kropko, Jonathan and Robert Kubinec, “Interpretation and Identification of Within-Unit and Cross-Sectional Variation in Panel Data Models,” *PLoS One*, 2020, 15 (4), e0231349.
- Kuhnen, Camelia M, “Asymmetric Learning from Financial Information,” *The Journal of Finance*, 2015, 70 (5), 2029–2062.

- Langer, Ellen J and Jane Roth, "Heads I Win, Tails It's Chance: The Illusion of Control as a Function of the Sequence of Outcomes in a Purely Chance Task," *Journal of Personality and Social Psychology*, 1975, 32 (6), 951–955.
- Lauer, Kate and Timothy Lyman, "Digital Financial Inclusion: Implications for Customers, Regulators, Supervisors, and Standard-Setting Bodies," Technical Report, Consultative Group for Alleviating Poverty (CGAP) 2015.
- Levitt, Steven D, "Why are Gambling Markets Organised So Differently from Financial Markets?," *The Economic Journal*, 2004, 114 (495), 223–246.
- Linnainmaa, Juhani T, "Why Do (Some) Households Trade So Much?," *The Review of Financial Studies*, 2011, 24 (5), 1630–1666.
- Mahani, Reza and Dan Bernhardt, "Financial Speculators' Underperformance: Learning, Self-Selection, and Endogenous Liquidity," *The Journal of Finance*, 2007, 62 (3), 1313–1340.
- Mobius, Markus, Muriel Niederle, Paul Niehaus, and Tanya Rosenblat, "Managing Self-Confidence: Theory and Experimental Evidence," Technical Report, UC San Diego July 2014.
- Sauer, Raymond D, "The Economics of Wagering Markets," *Journal of Economic Literature*, 1998, 36 (4), 2021–2064.
- Sharot, Tali, Christoph W Korn, and Raymond J Dolan, "How Unrealistic Optimism is Maintained in the Face of Reality," *Nature Neuroscience*, 2011, 14 (11), 1475–1479.
- Stock, James H and Motohiro Yogo, "Testing for Weak Instruments in Linear IV Regression," in "Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg," Cambridge University Press, 2005, pp. 80–108.
- Suri, Tavneet and William Jack, "The Long-Run Poverty and Gender Impacts of Mobile Money," *Science*, 2016, 354 (6317), 1288–1292.
- Technavio, "Global Sports Betting Market 2020-2024," Technical Report IRTNTR40575 April 2020.
- The Economist, "Borrowing by Mobile Phone Gets Some Poor People into Trouble," 2018.

Zenker, Juliane, Sebastian Vollmer, and Andreas Wagener, "Better Knowledge Need Not Affect Behavior: A Randomized Evaluation of the Demand for Lottery Tickets in Rural Thailand," *The World Bank Economic Review*, 11 2018, 32 (3), 570–583.

Zimmermann, Florian, "The Dynamics of Motivated Beliefs," *American Economic Review*, February 2020, 110 (2), 337–61.

Online Appendix

A. Descriptive Statistics

Table A1 provides descriptive statistics of our dataset. Our sample consists of mostly young men. Among the subsample of gamblers, 77 percent are men and 75 percent are 35 years of age or younger. Gamblers spend 481.67 KSH (approximately 4.81 USD) and receive 364.40 KSH (approximately 3.64 USD) from gambling on average per week.

The distribution of gambling activity is fat-tailed, with the top 20 largest spenders betting an average of KSH 84 349.41 (approximately 843 USD) per week. Betting income has fat tails because the largest prizes are awarded for bets with extreme odds. Also, betting income only measures withdrawals from the betting account. Small wins may fund subsequent bets rather than being withdrawn from the betting account.

B. Popularity of Sports Betting in Kenya

To emphasize the scale of sports betting, we use data on on internet search queries on Google. In Figure A1 we plot the relative popularity of SportPesa against Facebook.¹⁹ Facebook serves as a good benchmark as it is the most popular search query worldwide. Users of online services typically search for the name of service rather than type out the web address so the index of search queries serves as a good proxy for the relative number of users. For example, a user may type in “facebook” in the search bar rather than typing out “www.facebook.com”. Figure A1 shows a clear pattern. Since 2014, sports betting has grown rapidly in popularity. In 2018, SportPesa was the most popular search query in Kenya.

¹⁹Visit trends.google.com to access a current version of the graph.



Figure A1: Popularity of Google search queries in Kenya

C. More Details on our Test for Biased Learning

In Section 5.2 of the main paper, we test for biased learning using the following specification:

$$\begin{aligned}
 \text{Bet on } \text{jackpot}_{i(t+\tau)} = & \beta_p \text{ Positive feedback}_{it} + \beta_n \text{ Negative feedback}_{it} + \\
 & \gamma \text{ jackpot tickets}_{it} \\
 & v_i + \epsilon_{i(t+\tau)}
 \end{aligned}$$

The indicators for positive and negative feedback have three buckets. The share of correct predictions in week $t - 1$ can be

1. positive (10 percent above the gambler's mean),
2. negative (10 percent below the gambler's mean),
3. or base (between 10 percent above and below the gambler's mean).

We drop all gamblers who do not have at least one observation in each of the positive, negative and base buckets. In this appendix, we explain why we need to restrict the sample in this way.

Suppose there are two types of gamblers: stubborn low skill S and Bayesian high skill B . The S gamblers get L percent of predictions correct whereas as B gamblers get H correct and $H > L$.

We observe some bets and results for each type of gambler. Since $H > L$, the B gambler is more likely to “fill” the higher buckets of the categorical variable measuring the number of correct predictions in Week $t - 1$. In contrast, the S gambler is more likely to “fill” the low buckets.

If for a given gambler, one of the indicators is zero for all weeks we observe, this gambler does not contribute to the estimate of this coefficient. Therefore S gamblers will contribute more to the estimation of the coefficients of the low result indicator and the B gamblers will contribute more to the estimation of the high result indicators.

Assume S gamblers do not consider past performance when betting and B gamblers use rational Bayesian updating. This means that the higher bins will reflect Bayesian updating whereas the lower bins will reflect the stubborn betting. This will generate a biased updating result even though no single gambler is biased towards positive feedback.

D. Correlation in Correct Predictions Across Time

The approach in Section 4 uses variation for the subsample of 73 gamblers for which we observe at least 500 match predictions. In this appendix, we use variation for the 6 953 gamblers for which we observe jackpot bets in at least two consecutive weeks.

If all gamblers are identical and predict the outcome of matches with some fixed probability, there will be no correlation in the number of correct predictions across time. If a gambler who predicts a high number of matches correctly this week is more likely to predict a high number of matches correctly next week, this suggests heterogeneity in gambling ability.

To test for correlation in the number of correct predictions across time, we use

the following specification:

$$\begin{aligned} \text{Correct prediction}_{s_{it}} = & \sum_{l=1}^L \beta_l \text{Correct prediction}_{s_{i(t-l)}} \\ & + \theta_t + \sum_{l=1}^L \gamma_l \text{jackpot tickets}_{s_{i(t-l)}} + \epsilon_{it} \end{aligned}$$

The model checks if the number of correct prediction in week t by gambler i is positively correlated with the number of correct predictions in previous weeks up to a lag of L weeks. We include week fixed effects, θ_t , so the estimates use only within week variation to compare gamblers. We also control for the number of midweek and the number of weekend jackpot tickets purchased.

We report the results in Table A2. In all specifications, the number of correct predictions on jackpot matches is positively correlated with the number of correct predictions in previous weeks. If all gamblers had identical chances of success, this correlation would not be present. The positive correlation can only be explained by heterogeneity in the chance of success between gamblers, which we interpret in our model as heterogeneity in gambling ability.

E. Robustness Checks

Cutoffs for the low and high categories

In Section 5.2, we use 10 percent above and below the mean number of correct predictions as the cutoff to form the positive feedback, negative feedback and base categories. Here, we test the robustness of our results to cutoffs of 5 percent and 15 percent. Note that, as described in Appendix C, each gambler must have at least one observation in each of the three categories. By adjusting the cutoff for the three categories, the number of gamblers used in the estimation changes slightly with the cutoff threshold.

Results, presented in Tables A3 and A4 indicate that gamblers react to both positive and negative feedback. While the absolute value of the coefficient is often larger for negative feedback than for positive feedback, in most cases the difference in ab-

solute magnitude is not statistically significant.

Including time fixed effects

We use individual fixed effects to control for individual differences in betting ability and other time invariant characteristics. We choose not to use time fixed effects in addition to the individual effects. For reference, Table A5 shows the results with the inclusion of time fixed effects.

Sample selection

In our instrumental variables estimates in Section 6, we excluded from our regressions any observations in which the gambler won money from the jackpot in the previous week. As discussed, this restriction helps limit the possibility of an income effect from winning, which would create scope for violations of the exclusion restriction. However, this sample restriction may also affect our results. For this reason, Table A6 re-estimates equation (3) with the full sample, without excluding any observations. Results are qualitatively unchanged from those presented in Table 3.

Table A1: Descriptive statistics

	Gamblers	Non-Gamblers	Full sample
Number of individuals	58 215	35 350	93 565
Male	77.0 %	52.0 %	67.6 %
Age	30.66 [25, 34]	31.74 [25, 36]	31.06 [25, 35]
Betting expenditure (weekly)	481.67 [18, 306]	0.00 [0, 0]	299.69 [0, 132]
Betting income (weekly)	364.40 [0, 177]	0.00 [0, 0]	226.73 [0, 46]

Notes: Mean values reported, with 25th and 75th percentiles in brackets. We define gamblers as individuals with at least one mobile money transfer to or from a betting company. We exclude users with less than three weeks of mobile money transactions from the sample. Monetary amounts are measured in Kenyan Shillings (KSH).

Table A2: Testing for heterogeneity in betting ability

	<i>Dependent variable:</i>			
	Share of correct predictions in week t			
	(1)	(2)	(3)	(4)
Share of correct predictions in week:				
$t - 1$	0.077 (0.005)	0.075 (0.007)	0.075 (0.008)	0.068 (0.009)
$t - 2$		0.064 (0.007)	0.062 (0.008)	0.056 (0.009)
$t - 3$			0.067 (0.008)	0.074 (0.009)
$t - 4$				0.059 (0.009)
Individual fixed effects	No	No	No	No
Week fixed effects	Yes	Yes	Yes	Yes
Individuals	6 953	4 008	2 589	1 825
Weeks	119	118	117	116
Observations	35 642	22 370	15 749	11 854

Notes: All specifications control for the number of midweek jackpot and weekend jackpot tickets. The sample excludes individual-week observations where the individual won a prize on the jackpot in week t . Robust standard errors in parenthesis.

Table A3: Response to positive and negative feedback with a 5% cutoff

	Place jackpot bet in week:			
	$t + 1$	$t + 2$	$t + 3$	$t + 4$
Positive feedback (β_p)	0.015	0.013	0.011	0.004
	(0.005)	(0.005)	(0.005)	(0.005)
Negative feedback (β_n)	-0.026	-0.011	-0.010	-0.009
	(0.005)	(0.005)	(0.005)	(0.005)
P-value $ \beta_p = \beta_n $	0.0169	0.3959	0.413	0.200
Individual Fixed Effects	Yes	Yes	Yes	Yes
Week Fixed Effects	No	No	No	No
Individuals	5 309	4 953	4 678	4 454
Weeks	119	118	117	116
Observations	53 601	49 904	47 196	44 875

Notes: Dependent variable is an indicator for whether the individual places a jackpot bet in the weeks following a week of positive feedback or negative feedback on their gambling ability. Positive (or negative) feedback is defined as when the fraction of correct predictions made by the individual is more than 10% higher (or lower) than their average rate of correct predictions. All specifications control for the number of midweek jackpot and weekend jackpot tickets. The sample excludes individual-week observations where the individual won a prize on the jackpot in week t . We report robust standard errors in parenthesis below each estimate.

Table A4: Response to positive and negative feedback with a 15% cutoff

	Place jackpot bet in week:			
	$t + 1$	$t + 2$	$t + 3$	$t + 4$
Positive feedback (β_p)	0.020	0.007	0.010	0.003
	(0.008)	(0.009)	(0.008)	(0.009)
Negative feedback (β_n)	-0.033	-0.015	-0.013	-0.017
	(0.008)	(0.008)	(0.008)	(0.008)
P-value $ \beta_p = \beta_n $	0.117	0.239	0.380	0.101
Individual Fixed Effects	Yes	Yes	Yes	Yes
Week Fixed Effects	No	No	No	No
Individuals	2 409	2 233	2 118	2 024
Weeks	119	118	117	116
Observations	34 027	31 388	29 646	28 265

Notes: Dependent variable is an indicator for whether the individual places a jackpot bet in the weeks following a week of positive feedback or negative feedback on their gambling ability. Positive (or negative) feedback is defined as when the fraction of correct predictions made by the individual is more than 10% higher (or lower) than their average rate of correct predictions. All specifications control for the number of midweek jackpot and weekend jackpot tickets. The sample excludes individual-week observations where the individual won a prize on the jackpot in the previous week. We report robust standard errors in parenthesis below each estimate.

Table A5: Response to positive and negative feedback with time fixed effects

	Place jackpot bet in week:			
	$t + 1$	$t + 2$	$t + 3$	$t + 4$
Positive feedback (β_p)	0.035 (0.006)	0.016 (0.006)	0.015 (0.006)	0.009 (0.006)
Negative feedback (β_n)	-0.030 (0.006)	-0.014 (0.006)	-0.017 (0.006)	-0.011 (0.006)
P-value $ \beta_p = \beta_n $	0.277	0.402	0.381	0.391
Individual Fixed Effects	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes
Individuals	4 399	4 111	3 900	3 704
Weeks	119	118	117	116
Observations	48 948	45 626	43 216	40 959

Notes: Dependent variable is an indicator for whether the individual places a jackpot bet in the weeks following a week of positive feedback or negative feedback on their gambling ability. Positive (or negative) feedback is defined as when the fraction of correct predictions made by the individual is more than 10% higher (or lower) than their average rate of correct predictions. All specifications control for the number of midweek jackpot and weekend jackpot tickets. The sample excludes individual-week observations where the individual won a prize on the jackpot in the previous week. We report robust standard errors in parenthesis below each estimate.

Table A6: Instrumental variables estimates of the impact of increased betting expenditure (including winners)

	<i>Dependent variable:</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	Savings withdrawal	Savings deposit	Net savings deposit	Applied for loan	Loans received	Loan repayments
<i>Units</i>	<i>Elasticity</i>	<i>Elasticity</i>	<i>KSH</i>	<i>Indicator</i>	<i>Elasticity</i>	<i>Elasticity</i>
Betting expenditure	0.699 (0.234) [0.240, 1.158]	0.578 (0.218) [0.150, 1.005]	183.99 (374.106) [-549.24, 917.23]	0.029 (0.040) [-0.049, 0.107]	-0.191 (0.332) [-0.842, 0.460]	-0.418 (0.340) [-1.085, 0.249]
Individual fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Week fixed effects	No	No	No	No	No	No
Individuals	15 748	15 748	15 748	15 748	15 748	15 748
Weeks	119	119	119	119	119	119
Observations	71 099	71 099	71 099	71 099	71 099	71 099

Notes: Independent variable is the inverse hyperbolic sine of betting expenditures, instrumented with the past week's jackpot performance. The dependent variable in columns (1) and (2) are the inverse hyperbolic sine of withdrawals from and deposits to the M-Shwari savings account. The dependent variable in column (3) is the total deposits minus withdrawals, in Kenyan Shillings (KSH), with an approximate exchange rate of 100 KSH to \$1 USD. Dependent variables in columns (4)-(6) capture loan behavior on all loan providers who transact with mobile money, including M-Shwari, Branch and Tala. All specifications control for the number of midweek jackpot and weekend jackpot tickets purchased in the previous week. The sample includes individual-week observations where the individual won a prize on the jackpot in the previous week. We report robust standard errors in round brackets and the 95 percent confidence intervals in square brackets.