Gambling, Saving, and Lumpy Expenditures: Sports Betting in Uganda

Sylvan Herskowitz*

November 23, 2016

JOB MARKET PAPER

Link to the most recent version

Abstract

Demand for large and indivisible, or "lumpy", expenditures creates need for liquidity. For people in developing countries, acquiring this liquidity often requires choosing among high-cost strategies. I conduct a study with 1,715 bettors in Kampala, Uganda, to show that sports betting is being used as an alternative to conventional liquidity generation strategies such as saving or credit. First, I document that, despite expected losses of 35-50%, participants view betting as a likely source of liquidity for desired lumpy expenditures and use a natural experiment to show that this is not just cheap talk: winnings increase both the size and likelihood of making such expenditures. Second, I use a randomized field experiment to show that provision of a simple commitment-savings technology causes a 26% reduction in a revealed preference measure of betting demand. I then conduct two lab-in-the-field experiments to isolate the role of betting as a mode of liquidity generation. Increasing the salience of a desired lumpy expenditure causes an increase in betting demand by 17.2%, and a budgeting exercise decreases betting demand by 34.9% for people who learn that they could save more than previously believed. Back-of-the-envelope calculations suggest that betting to create liquidity may be a rational response for people with low ability to save.

*Department of Agricultural and Resource Economics, University of California, Berkeley. I am deeply grateful to my adviser, Jeremy Magruder, and give special appreciation to Elisabeth Sadoulet. I have benefited from invaluable feedback and support from Michael Anderson, Pierre Bachas, Patrick Baylis, Peter Berck, Lauren Bergquist, Josh Blonz, Fiona Burlig, Alain de Janvry, Aluma Dembo, Seth Garz, Tarek Ghani, Marieke Kleemans, Ken Lee, David Levine, Ethan Ligon, Aprajit Mahajan, Craig McIntosh, Ted Miguel, Jon Robinson, Leo Simon, Chris Udry, Liam Wren-Lewis, and Bruno Yawe. I would also like to thank Antoine Guilhin, Jackline Namubiru, Amon Natukwatsa, and Innovations for Poverty Action for excellent work implementing the study. This project would not have been possible without generous funding support from the NSF, CEGA's EASST Collaborative, the Rocca Center, the AAEA, and the Institute for Money, Technology, and Financial Inclusion. All errors are my own. The most recent version and online appendix can be found at www.sylvanherskowitz.com/jmp.html. Email: sherskowitz@berkeley.edu. Cell: (207) 512-6649.

1. Introduction

When people want to make large and indivisible, or "lumpy" expenditures, they must first determine how to raise the required liquidity. Many products and approaches have been developed in order to help people meet these liquidity demands. Credit and saving strategies are the most common and take many forms, with both formal and informal variations.[1] Although less frequently considered, gambling presents an additional alternative where participants can risk money for a chance to win a payout. In this paper, I present evidence from Uganda that the recent global growth of sports betting is driven, in part, by liquidity demands for lumpy expenditures and as a consequence of costly alternative options for liquidity generation.

Sports betting has boomed in popularity over the past decade and is now a global industry valued at over one hundred billion dollars.[2][3] Betting is a bundled product. It is a source of enjoyment for many participants, but it is also centered on a financial gamble, offering the possibility of sizable payouts. In recent years, new technologies have enabled international betting companies to enter previously untouched markets. This growth has been fastest across Africa and in developing countries, where financial institutions are often weak.[4][5][6] In these settings, limited liability increases the prevailing cost of available credit, while the lack of positive interest accounts, inflation, and high transaction costs are among the factors contributing to a negative effective return on saving.[7] This results in an unappealing menu of options for liquidity generation and may make people more willing to tolerate high expected losses from betting.

In Kampala, Uganda, a recent policy report found that 36% of men had participated in sports betting during the previous year, spending an average of 12% of their income on betting (Ahaibwe et al. 2016). This is a substantial expenditure, especially for a population sitting at or near the poverty line. Despite expected losses of 35-50% per dollar spent, three out of four respondents reported "making money" as their primary motivation for betting.

[1] A recent paper by Casaburi and Macchiavello (2016) shows that Kenyan dairy farmers are willing to sacrifice a portion of their total income in return for *less* frequent payments from buyers as a form of commitment saving. Excellent surveys of the saving and microcredit literature are provided by Karlan et al. (2014) and Banerjee (2013) respectively. Also see Besley et al. (1993), Brune et al. (2015), Dupas and Robinson (2013a,b), and Kast et al. (2014) for additional examples related to savings.

[2] http://www.statista.com/topics/1740/sports-betting/

[3] http://www.bbc.com/sport/0/football/24354124

[4] African Development Bank (2011) and Beck and Cull (2014)

[5] PricewaterhouseCoopers (2014)

[6] MORSS Global Finance: http://www.morssglobalfinance.com/the-global-economics-of-gambling/

[7] Collins et al. (2009) and Banerjee and Duflo (2007)

Using a similar sample, participants in this study listed betting as the second most likely source of money for a desired expenditure, following saving. In response to these stated motivations, I build on existing theory to create a model linking lumpy expenditures, saving ability, and demand for betting. In particular, I show that demand for lumpy expenditures increases demand for gambles while improving saving ability results in a decrease.

I conduct a set of complementary field experiments to test these implications of the model while collecting data on reported consumption, income, and betting behavior and a revealed preference measurement of betting demand. 1,003 men were included in a full study with five bi-weekly visits, creating a unique, high-frequency panel of betting behavior, business performance, and expenditure data. I supplemented this sample with a separate group of 712 participants in a condensed single-visit study. Ultimately, I generate four pieces of empirical evidence in support of the hypotheses derived from the model, using a randomized field experiment, a natural experiment, and two lab-in-the-field experiments.

I begin by presenting descriptive evidence that bettors both perceive and use sports betting as a method of liquidity generation. To add credibility to these survey responses, I use a natural experiment to show that expenditure behavior in response to winnings is consistent with participants' stated motivation. Next, I use a randomized field experiment to show that introducing a commitment saving technology lowers betting demand. I then conduct a pair of lab-in-the-field experiments in order to directly test whether people demand bets as a mode of liquidity generation and in response to low ability to save. In the first experiment, betting demand increases after experimentally inducing greater salience of a lumpy expenditure. In the second, betting demand decreases following a budgeting exercise for participants who learn that their capacity to save was better than previously believed. These findings could not be explained by an alternative hypothesis that betting demand is purely driven by its value as a consumption good. Finally, I conclude with a set of back-of-the-envelope calculations and demonstrate that the expected returns for betting and saving are similar for many people within a reasonable range patience levels and returns to saving.

To give credence to the hypothesis that sports betting is both viewed and used as a means of liquidity generation, I begin by presenting descriptive evidence from the bettors in the sample on their stated motivation and reported betting behaviors. Their responses suggest that we should see an impact of winnings on lumpy expenditures. The size and likelihood of winnings are not random, but are linked to peoples' betting choices. However, winnings should be effectively random after conditioning on the number and types of bets an individual makes in a given time period. Implementing this selection on observables design, I

test for the impact of winnings on lumpy expenditures and find that winnings increase both the likelihood and size of lumpy expenditures. Results are strongest among respondents categorized as having a low ability to save, consistent with the theory that using betting as a way to generate liquidity is most appealing for people with limited alternatives.

In the second result, I test the main predictions from the model. Improved ability to save should reduce betting demand. This should result from two channels: crowding out of all present-day expenditures and a drop in the relative appeal of betting as a mode of liquidity generation. Randomly selected participants were offered a wooden saving box, similar to a piggy bank, to assist them in their ability to save. This basic technology contains features common to many saving products: a component of ex-ante commitment to saving and a reduction in exposure to spending pressure and temptation. At the end of the study one month later, participants were offered a choice between cash and betting tickets in a revealed preference measure of betting demand. Recipients of the saving box were 26% less likely to demand the full amount of tickets offered.

I then use two lab-in-the-field experiments in order to isolate the role of betting as a method of liquidity generation that makes it distinct from other normal goods. The third result uses a randomized priming dialog in conjunction with the revealed preference measure of betting demand. Interviewers asked respondents a set of questions related to a previously identified and desired expense, in order to increase its salience. Respondents who were randomly selected to receive the prime before the betting ticket offer were 17.2% more likely to demand the maximum number of betting tickets. This large and significant increase confirms that many study participants view betting as a means of acquiring liquidity for their lumpy expense. If betting demand were purely driven by consumption, increased salience of a lumpy expenditure should not have caused a large increase in betting demand, and may have reduced it if people anticipated using the offered cash for saving. In addition, respondents categorized with low saving ability drove this effect, increasing their likelihood of demanding the maximum number of tickets by 29.1%, while the effect was below 5% for people with high saving ability. This parallel heterogeneous response with the analysis of winning usage provides further evidence that both betting demand and winning usage are linked with liquidity needs for lumpy expenditures among this group of bettors.

The final empirical result shows that a positive update on perceived saving ability decreases demand for bets. Before the betting ticket offer, randomly selected respondents were guided through a brief budgeting exercise that assisted them in making realistic assessments of their weekly saving potential. The results show that respondents who learned that they

had more capacity to save than previously believed were 34.9% less likely to demand the maximum number of betting tickets. If betting were purely a consumption good (and not in competition with saving as a mode of liquidity generation), new information revealing the availability of additional disposable income should have led to an increase in demand for all normal goods, including betting, and not the decrease that we observe.

Together, these results tell a consistent story: sports betting is in competition with saving as a mode of liquidity generation in pursuit of lumpy expenditures. However, negative returns suggest that sports betting results in substantial losses of expected income. Whether betting is rational depends on the available alternatives. Limited access to affordable credit and demand for non-creditable expenditures make borrowing infeasible or unappealing to most people in the sample. Compared to saving, back-of-the-envelope calculations suggest that, after accounting for future discounting as well as the impact of inflation, exposure to temptation, social pressures, and risk of loss or theft on peoples' return to saving, many people may rationally prefer betting as a way to generate liquidity.

Local media and political figures in Uganda have expressed increasing concern about the social effects of sports betting, including crowding out of scarce household resources, dissaving, domestic violence, and even suicide.[8][9][10] If policy-makers want to reduce betting, this paper suggests that improving financial services and alternative strategies of liquidity generation for this population may be an effective strategy. Even the simple saving interventions in this study, focused on budgeting and commitment saving, lowered betting demand. More ambitious initiatives like low-cost secure banking and mobile saving services or broadened access to affordable credit could have larger impacts.

This paper contributes to at least two broad areas of literature in economics. First, it makes a number of contributions to the development literature and the sub-fields looking at the financial lives of the poor, saving constraints, and temptation goods. By showing that betting is being used as a second-best strategy of liquidity generation, this paper contributes to a growing literature on financial management strategies of the poor under saving and credit constraints. This study extends the work of Collins et al. (2009) and Banerjee and Duflo (2007) by providing a new example of how poor families or individuals often use unconventional strategies to meet their financial needs.

Second, this paper is among the first to show that, by inhibiting income aggregation,

[8] http://allafrica.com/stories/201603150296.html

[9] http://www.monitor.co.ug/Business/Prosper/The-price-of-betting-on-Ugandans/-/688616/2107602/-/k7i4bh/-/index.html

[10] http://www.monitor.co.ug/News/National/Soccer-fan-kills-self-over-Arsenal-s-loss-to-Monaco/-/688334/2639990/-/dn6tkoz/-/index.html

saving constraints push people toward other, low-return liquidity generation strategies. In this regard, Karlan et al. (2014) provide an excellent overview of the saving literature, while Casaburi and Macchiavello (2016) show another example of saving constraints leading to the adoption of second-best saving strategies among Kenyan dairy farmers.

Third, this paper builds on Banerjee and Mullainathan (2010), who synthesize a growing and related literature on temptation goods. Their work shows that these goods have a disproportionate impact on the poor and their ability to save. This paper distinguishes betting from other consumption and temptation goods and shows that its financial properties put it in direct competition with saving as a mode of liquidity generation.

Fourth, this paper also contributes to a separate literature in economics on gambling, providing one of the first tests of a debate over the importance of financial constraints on demand for gambles. In 1948, Milton Friedman and Leonard Savage first presented a model of rational demand for gambles among people facing non-concavities in their indirect utility function, an insight that was later extended to include demand for lumpy expenditures (Kwang 1965). Others argued against the plausibility of this motivation for betting and claimed that well-functioning credit markets and ability to save should render this source of betting demand inconsequential (Bailey et al. 1980). While cross-sectional studies have found that financial circumstances and services are important determinants of gambling participation, there is limited empirical work testing these causal linkages.[11] Although some work has shown consumption behavior of lottery winners consistent with these causal hypotheses, this paper reproduces and extends that finding, adding three empirical results directly testing the causal mechanisms of betting demand within the same sample (Crossley et al. 2016; Imbens et al. 2001).

Finally, the existing literature on gambling is almost exclusively set in developed countries. This paper makes a further contribution by looking at gambling in a developing country and is the first to study sports betting in Africa, the region where it has grown fastest.

This paper proceeds as follows. Section 2 provides further background and details on sports betting and the specific context of Kampala, Uganda. Section 3 motivates the research question with an illustrative model of rational betting behavior. Section 4 provides details on the research design and data collection. Section 5 presents descriptive evidence of demand for betting and betting behavior in the sample. Section 6 discusses the main empirical results of the project. Section 7 presents back-of-the-envelope calculations comparing betting and saving in Uganda. Section 8 concludes.

[11] See Ariyabuddhiphongs (2011) and Grote and Matheson (2013) for recent surveys of the literature.

2. Background and Motivation

2.1 Global Expansion of Sports Betting

The global sports betting industry is already valued at over one hundred billion dollars. The last ten years have seen its most rapid expansion to date.[12] A recent report from the European Gaming and Betting Association estimates that sports betting grew at a rate of 5.4% per year across Europe from 2001-2013.[13] In the US, where most sports betting is still illegal, monetized fantasy sports is now itself a multi-billion dollar industry marked by the emergence of companies like Fan Duel and Draft Kings.[14] But growth has been fastest in many developing countries within Africa.

Adaptation of online betting technology in the form of internet-linked, vendor-operated betting consoles and betting shops has broadened access to new betting products with higher payoffs and a wider range of betting options than have previously been available. These platforms allow investors to offer internationally calibrated odds on sporting matches and have facilitated the entry of these firms into new markets. A 2009 consultant's report by MORSS Global Finance estimates that between 1999 and 2007 Africa experienced a 114% increase in betting revenues, a faster rate of growth than any other region.[15] A 2014 Pricewaterhouse-Coopers report estimates that sports betting in South Africa quintupled between 2009 and 2013, from 15.8 to nearly 80 million USD in gross revenues.[16] However, scarcity of reliable data makes it difficult to precisely estimate the size of the sports betting industry across Africa.

What is known is that international companies are rapidly entering and expanding in African markets.[17] Regulation varies widely by country, but the appeal of new tax revenue streams is a strong incentive for local governments to permit its entry and growth. The expansion of sports betting across Africa is likely to continue.

[12] www.statista.com/topics/1740/sports-betting/ and www.bbc.com/sport/0/football/24354124

[13] European Gaming and Betting Association (n.d.)

[14] http://fortune.com/2015/04/06/draftkings-and-fanduel-close-in-on-massive-new-investments/

[15] www.morssglobalfinance.com/the-global-economics-of-gambling/

[16] PricewaterhouseCoopers (2014)

[17] Recent media articles from Ghana, Nigeria, Senegal, Malawi, Sierra Leone, Tanzania, Liberia, Zimbabwe, and Kenya all observe a sharp rise in sports betting in their respective countries. Click on the country name for a linked article.

2.2 Sports Betting in Uganda

Sports betting is a legal, large, and rapidly expanding industry in Uganda. Throughout Kampala, Uganda's capital, nearly every commercial center features one or more betting shops. Although shops can typically accommodate more than fifty people at a time, in peak hours they overflow with customers.[18] Gambling in different forms has long been a part of Ugandan culture, but this format and extent of popularity are new. The arrival and expansion of international betting companies throughout the country began less than ten years ago. As of June 2015, there were 23 licensed betting companies operating in Uganda with over 1,000 betting outlets spread across the country (Ahaibwe et al. 2016).

A local policy report recently analyzed a representative sample of Kampala residents and found high participation rates among young men (18-40) across all income levels (Ahaibwe et al. 2016). But it is the lower income quintiles who devote the largest share of their earnings to betting. According to the report, 36% of men in Kampala had gambled at some point in the last year, devoting 12% of their income to betting on average. The impacts of sports betting have not been rigorously identified, but survey respondents suggest that betting is most likely to displace household expenditures and investments. Meanwhile, local media coverage has reported cases of bankruptcy, loss of school fees, and suicide as a result of accrued debt and shame attributed to betting.[19] [20] [21]

The format of betting in Uganda is the same as that spreading across the rest of the continent and available on most online sports betting sites. First, a bettor chooses which matches to include on his betting ticket from a list of options, typically featuring over 100 games. He then predicts a result or outcome for each of these matches such as "Team A defeats Team B". Predicting less-likely outcomes or adding additional games to a ticket is rewarded with a higher possible payout should the ticket win. If every predicted outcome on the ticket occurs, it can be redeemed for its payout value. If any single outcome is incorrect, the ticket is worth nothing.[22] Even by local standards, the minimum cost of placing a bet is relatively low, at just 0.18 USD per ticket. While bettors can target extremely large payouts if they choose, companies often cap the maximum payout at around 2000 USD, and most

[18] http://www.monitor.co.ug/Business/Prosper/The-price-of-betting-on-Ugandans/-/688616/2107602/-/k7i4bh/-/index.html

[19] http://allafrica.com/stories/201603150296.html

[20] http://www.monitor.co.ug/Business/Prosper/The-price-of-betting-on-Ugandans/-/688616/2107602/-/k7i4bh/-/index.html

[21] http://www.monitor.co.ug/News/National/Soccer-fan-kills-self-over-Arsenal-s-loss-to-Monaco/-/688334/2639990/-/dn6tkoz/-/index.html

[22] Additional details on the structure and format of betting are presented in Online Appendix B.

bettors target amounts much lower, around 50 USD, making the magnitude of sports betting payouts resemble scratch tickets more closely than national lotteries or Megabucks. Section 5 provides further details.

2.3 Literature Review

This paper contributes to multiple sub-fields within the development economics literature, including financial strategies of the poor, savings, and temptation goods. First, by demonstrating that liquidity needs are linked to both betting demand and winning usage, this paper contributes directly to a growing literature on financial management strategies of the poor. Work by Collins et al. (2009) and Banerjee and Duflo (2007) has shown that poor families must often use multiple and sometimes unconventional strategies to meet their liquidity needs. This paper provides a new, and previously undocumented, example.

This paper also makes a direct contribution to the literature on saving and saving constraints. Karlan et al. (2014) summarize the literature on saving that investigates barriers to saving and the adoption of saving technologies. Other work has shown the impact of saving constraints on investments and on resiliency to negative shocks (Brune et al. 2015; Dupas and Robinson 2013a,b). This paper makes a unique contribution by showing that limited saving ability impedes one's ability to accumulate available liquidity and pushes people toward betting, a low-return alternative. This finding is similar to a recent paper by Casaburi and Macchiavello (2016), showing that Kenyan dairy farmers are willing to sacrifice a portion of their income in return for *less* frequent payments as a contractual form of commitment saving with upstream buyers. While not directly linked to saving, a recent paper by Brune (2016) also presents evidence of high demand for lumpy income, where a lottery-based bonus scheme with large, low probability payouts increases labor supply of Malawian tea plantation workers more than a flat bonus of equivalent expected value.

Part of betting demand may also be similar to "temptation goods", as characterized in work by Banerjee and Mullainathan (2010). Temptation is likely to have a particularly strong effect on savings patterns among the poor because these goods typically have low absolute levels of satiation for consumers. Similar to other temptation goods, betting may not be valued prior to consumption, or may be regretted after purchase. A recent paper by Schilbach (2015) examines the relationship between saving and another temptation good with unconventional features: alcohol. However, whereas alcohol chemically alters peoples' time preferences away from saving, the distinguishing feature of betting is the financial gamble that puts it in direct competition with saving as a way to generate liquidity. This

gamble generates tension between sports betting and saving that goes beyond conventional temptation goods, that affect saving principally through over-consumption and crowd-out.

This paper also contributes to a separate literature on gambling and is one of the first to provide direct tests for two predictions in the literature. Existing work has proposed a wide range of reasons why people could engage in gambling despite negative expected returns. These factors include direct utility from gambling, misperceptions about the games themselves, and addiction (Becker and Murphy 1988; Bordalo et al. 2012; Conlisk 1993; Heath and Tversky 1991; Raylu and Oei 2002). However, some of the earliest theoretical work on gambling from Friedman and Savage (1948) suggests that demand for gambles can also be generated by non-concavities in an individual's utility function. In their example, illustrated by Figure 1, an individual with income c increases expected utility by purchasing a lottery ticket that costs $c-\underline{c}$ with a 50% chance of winning $\bar{c}-\underline{c}$. The authors suggest that the separate concave regions of this utility function could result from impeded social mobility between economic classes. Demand for indivisible expenditures can create another type of non-concavity and similarly generate demand for gambles (Kwang 1965).

More recently, others have argued that non-concave utility is not a credible explanation for real-world gambling with high expected losses, and that ability to save and access to credit should render this source of demand insubstantial. Bailey et al. (1980) conclude that "risk preference due to a rising marginal utility of income could occur, if at all, only in remarkable conditions" (p. 378).[23] Whether these "remarkable conditions" exist in the real world and create demand for gambles, as well as the causal link between financial constraints and gambling, are ultimately empirical questions. This paper adds to the limited existing evidence.

Three recent review articles from Grote and Matheson (2013), Bruce (2013), and Ariyabuddhiphongs (2011) all refer to the robust negative relationship that has been documented between income levels and betting intensity. They also all highlight a general lack of credible identification in non-lab settings in the literature. While no papers to my knowledge have aimed at directly testing the motive of liquidity generation as their primary hypothesis, the few related papers that do exist present mixed evidence. Snowberg and Wolfers (2010) examine American horse betting and find that misperceptions of odds drive the well-established long-shot bias more than demand for high payouts. However, their evidence does not rule out demand for liquidity as a contributing factor. The other two papers focus on

[23]Twenty years later, Hartley and Farrell (2002) pushed back on this finding by showing that demand for gambles can persist even with complete savings and credit markets in certain ranges of a non-concave utility function when rates of interest and time preference differ.

usage of winnings. Imbens et al. (2001) show that lottery winners purchase large durable goods following wins, a finding consistent with Friedman and Savage. Crossley et al. (2016) present similar evidence in the United Kingdom, showing that credit-constrained people who buy lottery tickets use large inheritances to make lumpy expenditures, suggestive that these bettors face binding liquidity constraints. While suggestive, neither of these studies is able to show that this ex-post behavior is a driver of betting demand. In this paper, I link ex-post usage of winnings with ex-ante betting demand driven by liquidity needs, to directly test the causal link between saving and gambling.

Finally, this paper fills a gap in the existing empirical research on gambling through its choice of setting. In a review by Ariyabuddhiphongs (2011) of more than 100 gambling studies, three are based in developing countries, where the stakes and context of gambling are likely distinct from those of relatively wealthier gamblers in developed countries. To my knowledge, only one is set in Africa (Abel et al. 2015). Looking at how gambling behavior relates to financial constraints in Africa, the region where sports betting is growing fastest, is an important contribution of this paper.

3. Model

Although betting is a bundled good that includes both direct enjoyment and a financial gamble, this model focuses on the gambling component that makes sports betting distinct from other normal goods. This section generalizes and extends a model by Crossley et al. (2016) of demand for gambles resulting from demand for lumpy expenditures. This extension allows for the flexible form of gambles offered in Uganda and illustrates a central tension between gambling and saving. I use the model to illustrate four main predictions. First, demand for lumpy expenditures generates demand for gambles. Second, increased valuation of a lumpy expenditure increases demand for gambles among people who can not afford to make the lumpy expenditure. Third, demand for lumpy expenditures also creates demand for saving in an overlapping range of income levels as for betting. Finally, improvement in saving ability decreases demand for gambles.

3.1 Demand for Gambles

An agent wants to maximize expected utility subject to a budget constraint. His weekly income is Y and utility is derived from the consumption of one divisible good, D, and the possible purchase of a single unit of a lumpy good: $L \in \{0, 1\}$. Consumption of the divisible

good yields utility $u(D)$, where $u'(D) > 0$ and $u''(D) < 0$. Purchase and consumption of the lumpy good yields a discrete utility payoff, η, and costs a price, P. An agent's utility is therefore: $v() = u(D) + \eta L$.

Figure 2a shows that utility without purchase of L is conventional, concave utility: $u(D)$. However, if the individual has enough income, $Y > P$, then he must decide whether the extra utility from consuming L is worth the loss in utility from reducing consumption of D. Purchase of L is represented by a jump onto the upper curve in Figure 2a. However, having spent P on the lumpy good, he gets the discrete utility payoff of η but can only spend the remainder of his income, $Y - P$, on the divisible good. Given his income level, the agent optimizes his utility by selecting the higher curve. The crossing point of the two curves, Y^*, is therefore the threshold at which individuals switch from not making to making the lumpy expenditure. The envelope of these two pieces is the utility maximizing value function for non-gamblers such that optimal utility is:

$$\begin{aligned} U^1() &= u(Y) && if && Y < Y^* \\ U^2() &= u(Y - P) + \eta && if && Y \geq Y^* \end{aligned}$$

Next, I allow for the possibility of making bets (or gambles). There are two stages of this single time period. In the first stage, an individual assesses his income, Y, and has the option of purchasing a betting ticket of any value, B. The ticket has a likelihood, σ, of resulting in net winnings of W. If purchased, the outcome of the lottery is immediately realized. Those who win purchase the lumpy good, while those who lose do not.[24] Therefore, the utility following a loss is $U^1(Y - B) = u(Y - B)$ while utility following a win is $U^2(Y + W) = u(Y + W - P) + \eta$. A betting choice of $[B, W]$ will result in expected utility somewhere on the segment between $U^1(Y - B)$ and $U^2(Y + W)$ determined by the likelihood of winning the bet, σ, such that expected utility for the bettor is:

$$E[v(Y)] = \underbrace{\sigma[u(Y - P + W) + \eta]}_{If\ \ Win} + \underbrace{(1 - \sigma)u(Y - B)}_{If\ \ Lose}$$

Because bets in this setting are fully flexible, an agent can choose his "optimal" bet constituting the amount he risks, B, and the net amount he tries to win, W. This means that the

[24]The decision of whether or not to buy L is deterministic once the result of the bet has been realized. Additionally, only people who will buy the lumpy good after a win have an incentive to make a gamble. This is because the concavity of $u(D)$ makes it so that using expected winnings on more of the divisible good gives less expected additional utility than the expected loss of utility when the gamble does not win.

best possible bet he could make, $[B^*, W^*]$, will be on the segment that is tangent to U^1 and U^2. These points of tangency will define the optimal bet for everyone with income levels between these endpoints, although the amounts wagered, B, and the targeted net winnings, W, as well as the likelihood of winning, will depend on an individual's income level.

If betting companies offered actuarially fair bets, then expected net winning or losses would be the same, such that $\sigma W = (1-\sigma)B$. Figure 2b illustrates this optimal bet with fair odds for an individual with income $\tilde{Y}$. Utility after a loss is indicated at point A, while utility following a win is at point C. The likelihood of this fair bet winning is such that $\frac{EA}{EC} = \frac{\sigma}{1-\sigma}$. A fair bet of $[B^*, W^*]$ results in expected utility at point E, which is an increase in expected utility for this bettor from F to E. People with income level $Y < \tilde{Y} - B^*$ will be too poor to bet; no available fair bet will increase expected utility. Similarly, no one with income level $Y > \tilde{Y} + W^*$ will bet because no fair gamble improves on his direct consumption of L and D. I define the lower and upper endpoints of the range of income levels that demand fair bets as Y_m^B and Y_M^B, respectively.

Of course, betting shops do not offer fair odds. Instead, they decrease expected payouts in order to make profits by reducing the likelihood of winning. Figure 3a shows that there is also demand for unfair gambles where the amount bet, B, and won, W are held constant but the likelihood of winning, σ, has been reduced below that offered by a fair bet. This can be seen by tracing horizontally from the starting utility at point D toward the vertical axis until it reaches the convex segment at point F. Win likelihoods as low as $\sigma_{min} = \frac{B-H}{B+W}$, indicated on the figure, will still improve expected utility for an individual with income $\tilde{Y}$.

3.2 Increased Valuation of a Lumpy Expenditure

Next, increasing the valuation of a lumpy expenditure increases demand for bets. Anticipating the empirical strategy for my third result, I claim that increasing the salience of an expense is equivalent to an increase in its anticipated value.[25]

As before, Figure 3a shows the range of income levels within which individuals demand fair bets, with the envelope the "convexified" expected utility of the agent with fair bets. As before, the endpoints of this range are Y_m^B and Y_M^B for the minimum and maximum, respectively. An increase in the valuation of L shifts η upward, in turn increasing anticipated

[25]This is consistent with experimental evidence from diverse settings whereby random variation in the salience of an item amplifies the valuation of that item. Barber and Odean (2007) show this phenomenon in the stock market when companies have unusually large or small single-day performances. Further, Ho and Imai (2008) show how salience of a third party political candidate resulting from random ordering on ballots leads to an increase in the candidate's resulting vote share.

utility for all income levels at which the lumpy good is purchased. Figure 3b shows that this increase in η also shifts the location of the tangent line defining the range of bettors and their optimal bets. Both the top and bottom endpoints of this range shift downward such that $\frac{\partial Y_m^B}{\partial \eta} < 0$ and $\frac{\partial Y_M^B}{\partial \eta} < 0$. The downward shift in the upper bound shows that some people who could already afford the good are no longer willing to risk the possibility of losing their bet and no longer being able to make the lumpy expenditure. For the empirical tests in this study, the relevant shift will be on the expansion of the lower bound of people who now demand gambles. This is because the lumpy expenditures used in the study were identified as being expenses that respondents could not afford at the time of the interview. For those who demand bets both before and after the shift, expected utility from betting and the amount spent on betting have both increased.

3.3 Demand for Saving

Saving is an alternative liquidity generation strategy. To allow for saving, I switch to a two time-period model. Keeping the model as simple as possible, gambling and saving decisions take place in the first period only and income, Y, is the same in both periods. Under these assumptions, the previous result defining an income range of betting demand is unaffected. However, there may be a range of incomes, also around Y^*, where saving to purchase L in the second period is preferred to spending all income on the divisible good.

Utility over two time periods is structured similarly to the single period, except that the second period is discounted by a factor $\delta \leq 1$. When saving, the agent chooses how much income to set aside for use in the next period, S, such that $S \leq Y$. However, all of S may not make it to the second period. γ represents the loss of savings between time periods. $\gamma > 1$ would suggest positive interest on savings. However, given the population and setting of the study, γ is likely below one as the result of possible loss, theft, inflation, or expenditure on non-valued temptation goods in future time periods.

Without positive interest, the agent would never save if saving did not result in purchase of L.[26] Therefore, two-time-period utility for a saver (purchasing L in the second period) is maximized with the choice of S^*:

$$\max_S V_s(Y) = u(Y - S) + \delta[u(Y + \gamma S - P) + \eta]$$

[26]This is the result of the concavity of $u()$ such that, even before considering time discounting or savings losses, additional marginal utility from consumption of D in period two would be less than the utility from spending that money on consumption of D in the first period.

For graphical clarity, I have set $\delta = 1$ and $\gamma = 1$ in the figures. Figure 4a shows that the same individual with income $\tilde{Y}$ would be willing to sacrifice S^* of consumption in the first period for additional consumption in the second. The horizontal axis is still the income level, as it was for betting, but the vertical axis is now average utility over two periods. Period 1 utility will be at N and period 2 utility at Q, leading to average two-period utility M, and a gain of utility over not saving equal to $M - R$. Figure 4b shows the locus of optimized saving utilities for each income level. Again, the envelope of the non-saving utility function and utility from optimal saving will constitute the new, maximized indirect utility function of potential savers. The region $Y \in [Y_m^S, Y_M^S]$ defines the range of income levels for which saving is welfare improving. Similar to betting, we observe that, if this region is non-empty, then $Y^* \in [Y_m^S, Y_M^S]$ and betting and saving will both be welfare improving in some area around Y^*.

When both betting and saving are welfare improving, the agent's choice will be determined by parameter assumptions in the model. In particular, higher levels of patience will make saving relatively more attractive, whereas less fair bets (lower σ given a choice of B and W) will lower the value of betting relative to saving.

3.4 Changes in Saving Ability

The ability to transfer income from the first to the second period is captured by the parameter γ. A rise in γ will lead to an increase in utility from a saving strategy at all income levels. This is simply because there is now more potential income to be spread across the two periods. An increase in saving ability also pushes the locus of optimal saving utilities upward, as shown in Figure 4b. Figure 4c illustrates this shift, showing that the end points of this range for saving have moved outward such that $\frac{\partial Y_m^S}{\partial \gamma} < 0$ and $\frac{\partial Y_M^S}{\partial \gamma} > 0$.

When $[Y_m^B, Y_M^B] \cap [Y_m^S, Y_M^S]$ is non-empty, and both strategies of liquidity generation are preferred to direct consumption, parameter assumptions will determine which strategy is preferred. An increase of γ will expand this range of potential overlap while also resulting in more utility from saving. This will lead to a weak decrease in demand for bets as they become a relatively less appealing method of liquidity generation.

As mentioned earlier, betting is a bundled good. The other component of its appeal is direct enjoyment, which should behave like other normal goods captured in the model by D. As saving ability improves, an individual increases the total amount set aside for saving such that $\frac{\partial S^*}{\partial \gamma} > 0$. Because consumption of divisible goods in period one is equal to $Y - S$, the increase in saving ability decreases today's consumption. Therefore, a positive

change in saving ability affects betting both by reducing the relative appeal of gambles and also by shifting consumption of normal goods toward the future. These results are derived in Appendix E. If betting is also a temptation good, then the effect of improved saving ability could be even stronger than for other normal goods if it lowers on-hand liquidity and therefore reduces exposure to temptation.

4. Experimental Design and Data

Given the absence of existing data on betting in this context, I designed a study to test the predictions resulting from the model. Between September 2015 and July 2016, I conducted a set of field experiments with 1,715 bettors in Kampala and created a unique data set able to provide evidence on these hypotheses.

4.1 Overview

Field work for the project was conducted over eleven months between September 2015 and July 2016, involving three phases of data collection. First, between October and December of 2015, a set of 483 participants were identified and included in "Wave 1" of the study. These respondents were visited and interviewed in person five times, once every two weeks. A second group of 520 participants were identified and included in Wave 2, between April and June 2016, following similar protocols. I refer to these 1,003 participants as being part of the "full study". To further explore the link between saving ability and demand for lumpy expenditures, a complementary "condensed" study was conducted with 712 additional respondents over three weeks in July 2016, with activities contained in a single visit.

4.2 Listing/Targeting

The study targeted young men between the age of 18-40, self-employed in small micro-enterprises or services, with weekly incomes below $50. Earlier piloting, as well as previous assessments of betting in Uganda, suggested that this group was likely to have a high incidence and intensity of betting along with unmet liquidity needs (Ahaibwe et al. 2016; Ssengooba and Yawe 2014). This is also a demographic of inherent interest, as they constitute a significant portion of Uganda's informal economy and serve important roles as key contributors of household income.

Each survey round began with a listing exercise in selected parishes around Kampala in

order to identify suitable respondents and invite them to participate in the study.[27] Participants were identified at their place of work and asked a short set of initial questions to determine whether they met the targeting criteria of gender, age, employment, and income.

Overall, listing from both waves of the full study included 5,522 people. Their characteristics are consistent with piloting, policy papers, and review of media coverage. Sports betting is extremely popular in this demographic, as 32% reported betting in most weeks.[28] After completing the listing, a randomized selection of respondents was chosen among those who bet regularly. The full study was launched at the beginning of October. The condensed study was conducted in July 2016 using a new sample of 712 respondents and followed the same criteria for inclusion. Suitable respondents were interviewed immediately upon identification instead of returning to them later. Additional technical details on field protocols are included in Appendix D.

4.3 Data Collection

Full study participants were interviewed in person five times, in alternating weeks. In addition, brief phone check-ins were conducted on the weeks between visits. The surveys captured a wide range of background characteristics and information, including household composition, education, numeracy, literacy, savings background, credit background, and risk and time preferences. For topics whose answers were not likely to change over the study period, the questions were asked only once. In addition, certain recurrent survey modules were conducted at each in-person interview, including consumption, household shocks, business investments, earnings, transfers, anticipated expenses, anticipated earnings, betting expenditures, and winnings. Phone check-ins were restricted to the noisiest and most important recurrent variables: weekly earnings, major expenditures, and betting participation.

During the third visit, members of the research team gave wooden saving boxes to randomly selected respondents. These boxes are a simple soft commitment savings device similar to piggy banks. During the final visit of the full study, as well as the first visit for participants in Wave 2, field team members conducted a revealed preference measure of betting demand. Additional details are provided below. Randomized primes were conducted in conjunction with these betting ticket offers. Additional details of these activities are included in Section 5. Figure 5 depicts the data collection timeline with "V1" signifying "visit 1,"

[27] In Wave 1, parishes were randomly chosen from the full set of parishes in Kampala that had viable commercial centers where the target population was likely to be found at their workplaces. In Wave 2, parishes closer to the city center were targeted due to logistical challenges and budget constraints.

[28] Appendix Table A.2 summarizes the listing data.

"PC2" signifying "phone check-in 2," etc.

Additional randomized treatments unrelated to the hypotheses in this paper were also conducted during the second and fourth visits of the project.[29] All treatments were randomized and included as controls in final estimating regressions.

The condensed study was designed to test a number of hypotheses that could not be included in the full study. In particular, it expanded on the priming experiment with a brief budgeting exercise designed to test the effect of perceived savings ability on demand for betting. It was conducted over three weeks following the conclusion of the full study. Details on this budgeting activity will be provided in the discussion section.

4.4 Measuring Betting Demand

Field team members collected a revealed preference measure of betting demand in the fifth and final visit for participants in the full study, as well as in the first visit for participants in Wave 2 of the full study. It was also included at the end of the condensed study and is an important outcome variable for three of the four empirical results.

Respondents were offered the choice between pre-filled betting tickets and a designated amount of cash. They were told the amount spent on the ticket as well as the approximate size of the payout should the ticket win, but they were not permitted to see the actual outcomes predicted on the ticket.[30] The amount of cash offered was less than the price of the ticket, preventing respondents from taking the money and purchasing a new ticket themselves, but it was similar to the expected value of the ticket. The cost of the ticket (or stake) was 1,000 Ugandan Shillings (UGX, approximately 0.35 USD), which is the most common value for bets in Uganda. The bets were placed with well-known and trusted betting companies, familiar to all respondents.

Respondents were then asked how many units of cash or betting tickets they would like to choose. Participants in the full study could select up to four, whereas participants in the

[29]The second round contained a randomized offer of a wallet with which respondents were encouraged to set aside money and budget for betting. The fourth round contained a randomized information treatment whereby selected respondents were given a detailed accounting and aggregating of their betting expenses and winnings up to that point in the study.

[30]The decision not to show them the tickets was made because of participants' strong beliefs about the outcomes of matches, such that they might value a betting ticket at zero if it chose an outcome at odds with their strongly held priors.

condensed study were limited to two.[31] [32] The analyses in Section 6.2, 6.3, and 6.4 use the binary outcome of "maximum tickets demanded" out of the total number offered This choice was made because maximum ticket demand is the highest powered outcome with just over 40% of respondents having demanded the full number. Results using alternative continuous measures of the outcome variable are provided in the appendix.

5. Descriptive Evidence

5.1 Background Characteristics

Descriptive statistics from the survey provide context on the financial situation and constraints shaping peoples' betting, saving, and expenditure decisions. Panel (a) in Table 1 shows considerable heterogeneity of income levels, betting intensity, household situation, age, and education. It also shows that the full and condensed study samples are broadly similar along most of these dimensions.[33] For the full study, weekly income and betting expenditures were calculated as weekly averages of reported betting and income over the course of the study. In the condensed study, respondents were asked how much they spent in a "normal" week. Although the condensed study sample appears wealthier on average than the full study sample, they otherwise look similar.

Overall, individual and household incomes are low. Adjusting for children in the household, median per capita income is beneath the two dollar per day poverty line. Meanwhile, betting intensity is high, as the median bettor spent 8.6% of his weekly income on betting during the course of the full study and those in the condensed study estimated their expenditures at 8.3%. The mean for both is around 12.5%, indicating that some people in the sample bet at very high levels of intensity while many others participate more moderately.

Survey responses also identify a number of obstacles people face in their financial lives.

[31] The basic setup of the betting ticket offer was the same across all groups; however, there were two additional differences between the full and condensed study. First, during the full study, participants were given the additional choice of whether they wanted tickets that targeted low, medium, or high payouts, whereas in the condensed study the payout size was always medium. Second, the amount of money they could choose in place of a betting ticket was held fixed during the full study but was experimentally varied during the condensed study. All of these varying factors are controlled for by using time and price fixed effects.

[32] Appendix Table A.1 shows that there is a positive and highly significant relationship between this measure of betting demand and respondents' reported levels of betting.

[33] Differences between the full and condensed study samples are not a point of primary concern. The two samples were drawn from different communities and therefore are likely to differ along certain dimensions. In addition, randomization was conducted by survey round and, therefore, these differences do not threaten the identification strategies implemented.

Risk of theft (49/59%), pressure to spend (27/34%), and existing debt (43/23%) are all cited with high frequency by participants in both the full and condensed samples. Although roughly 90% have mobile money, only 41% have bank accounts and less than 50% felt they could get a loan from a bank if they wanted one.[34]

For betting and saving to be relevant strategies for meeting liquidity needs, pre-existing access to liquidity should be low. Table 1 Panel (b) confirms this. In particular, the variable "Available Liquidity" is respondents' answer to the question, "What is the biggest expense you could make without needing to borrow?" The majority of participants could not afford an expense greater than 1.5 times their normal weekly income without borrowing money. The table also shows that the distribution of targeted betting payouts is above the level of respondents' available liquidity. Associated correlations between the log of the targeted win amount and the log of available liquidity or mean income are 0.082 and 0.213, respectively. Both are significant at the 99% confidence level, providing further suggestive evidence that financial capabilities relate to betting behavior.

Finally, respondents were asked directly, "what is the biggest reason that you bet?" Over 79% of respondents said that their primary motivation for betting was "a way to get money". The second most common answer was simply "fun", cited as the top reason for betting by just 15% of respondents. This overwhelming response suggests that bettors themselves consider the possibility of a financial payout as the most important feature of sports betting in determining their participation. While this response, and survey responses in general, should always be treated with skepticism, the dominance of this response suggests that it should be taken seriously.

5.2 Lumpy Expenditures and Liquidity Generation

A primary assumption of this paper is that people want to make lumpy expenditures that they are presently unable to afford. Without this demand for a lumpy good, utility should not have a non-concavity of the form outlined in the model, making this source of demand for gambles irrelevant. In the full study, interviewers asked respondents about three categories of potential desired lumpy expenditures: business investments, household expenditures, and personal expenditures. Enumerators explained that these should be indivisible expenses and that they should be realistically attainable in order to avoid purely aspirational reported targets. Table 2, Panel (a) shows their responses. Most notably, the majority of respondents could readily list an expense for each of the three categories. Only 5.8% were unable to

[34] Appendix Table A.3 summarizes these factors.

identify a desired expenditure from any of the three categories.

During the condensed study, after being asked to identify a desired large expenditure, interviewers asked respondents how thought they they were most likely to get funds for this purchase. Table 2, Panel (b) shows these responses. Although saving is viewed as the most likely source (by a considerable margin), betting is second. 25.3% of respondents considered betting to be a likely source of funding, almost as much as family and friends, bank loans, or money lenders combined. Considerable work focuses on credit access as a key method of helping people to overcome liquidity constraints. However, even after adding together all three, very different, categories of credit, this population considers betting to be just as likely a source of liquidity for a desired expenditure.

5.3 Availability of Credit

The data provide evidence of why credit is viewed so unfavorably in the sample. First and foremost, the cost of bank credit is high. Respondents report that they expect to pay 20-25% interest on a six-month bank loan. In terms of income losses, this would be the equivalent of a betting ticket that returned 80-83 cents per dollar spent. Access is also an issue. Only 41% of respondents have a bank account, a pre-requisite for most loans. Of those who do not have an account, one out of three said they were deterred by high usage fees, while 15% said that they did not have the required documents or had been refused in the past. When respondents were asked directly about their prospects of getting a loan, just 48% thought that bank credit was a possibility.

The previous section also showed that people demand many different lumpy expenditures. While business expenses may be eligible for bank loans, banks do not typically give credit for furniture, repairs, phones, or clothes, all frequently cited in this population. 85% of respondents had a non-business expenditure that they were eager to make in the coming months. Although bank credit would be a more efficient way of generating liquidity than betting, it is only available to a minority of the population and only applicable to a limited subset of the expenditures they hope to make.

Another possible source of credit is from money lenders. In contrast to banks, money lenders can be found throughout Kampala and do not restrict how borrowers use their loans. However, despite their availability and low barriers to borrowing, they are not viewed favorably; just 2.1% of respondents cited money lenders as the likely source of money for their desired expense, and they were the only source cited less frequently than bank loans. The standard money lender rate in Kampala is 50% interest on a six-month loan. This is

equivalent to 33% expected losses. Although this is still slightly better than betting, after factoring in the possibility of default, penalties, and risk of losing collateral, the expected losses from money lender credit are comparable with betting.[35] Credit from money lenders is likely so unpopular precisely because it is not an improvement over betting.

5.4 Saving Ability

Saving ability is the key dimension of heterogeneity used in this paper. In particular, it is used in the analysis of the impact of winnings on lumpy expenditures and in the response of betting demand to an increase in the salience of a desired lumpy expenditure. Table 1 Panel (b) shows "saving ability" as the response to the question, "How much money do you think you could save in a normal week without straining your regular household finances?" Median saving ability relative to income in the sample is between 20 and 30% of weekly income.

Dividing the sample between people who are able to devote a relatively high or low portion of their income to saving is appealing but problematic. The correlation of this measure with income level is 0.282. In order to remove this correlation from the measure, I only compare people to others in their survey round within the same income ventile (five percentile group). This process removes 98% of the correlation between this indicator and income level to just 0.005. Although this measure may still correlate with other characteristics, it serves as a reasonable first approximation of people who are relatively more or less able to use saving as a way to generate liquidity.

This section has provided some descriptive credence to the hypotheses and assumptions of the theory presented in this paper. People have limited liquidity available, face a high cost of credit, have considerable constraints on their ability to save, view betting as a means of acquiring liquidity, and consider it a likely source of liquidity for a set of desired expenditures. Taking these responses seriously suggests that we should revisit and test theories of rational demand for gambles as a financial asset, setting aside the component of betting resulting

[35]Missing a scheduled repayment to a money lender typically results in penalties raising the overall expected cost of the loan repayment. At a minimum, money lenders double the interest fees for missed monthly repayments. Someone who wanted to borrow 60 USD would be expected to pay back 15 USD per month (10 USD toward the principal and 5 USD toward the interest) for six months, a total repayment of 90 USD. A single missed repayment would raise the repayment total to 95 USD, dropping the rate of return from 67% to 63%, a rate that is now almost the same as from betting. The greater risk is that the money lender decides not to be lenient, and keeps whatever collateral is being held as leverage for repayment. For collateral to work, it must be more valuable to the borrower than the full value of repayment. Even a low risk of missing a repayment and losing one's collateral will again push the expected value of credit below 60% and into the range estimated for betting.

from fun. The first empirical result will be to see if these survey responses are just cheap talk or if winnings do in fact increase lumpy expenditures, while the remaining three will be direct tests of responses in betting demand to different randomized treatments.

6. Results

This paper contains four main empirical results. The first result looks at expenditure behavior and shows that winnings increase the size and frequency of lumpy purchases, consistent with the motivations for betting articulated in the previous section. Second, I test the effect of a change in saving ability on betting demand, using the randomized distribution of a simple commitment saving technology, and find a significant reduction in response to the treatment. In order to isolate the role of betting as a means of liquidity generation, I then conduct two lab-in-the-field experiments. I use a randomized priming treatment to show that increased salience of a desired lumpy expenditure increases betting demand. Finally, I use a randomized budgeting exercise to update participants' perceived saving capacity and find that those who receive a positive update reduce their demand for betting.

Analysis samples in this section differ by result. This is because different identification strategies have different data requirements that are not available for all groups of participants. I use the largest sample possible for each analysis.[36]

6.1 Use of Winnings

The previous section showed that respondents view betting as a likely source of liquidity for desired lumpy expenditures; that they claim to bet because it is a way to get money; and that the amounts that they target correlate with their income levels and the amount of liquidity to which they have access. However, if winnings do not increase the likelihood of making lumpy expenditures and the size of such expenditures, then this stated motivation could just be cheap talk. These results confirm that winnings do increase lumpy expenditures, and that this response is driven primarily by people with low ability to generate savings, consistent with the hypothesis that betting is a strategy for liquidity generation most appealing to those with limited alternatives.

Figure 6 shows the biggest win observed in the data for each participant in the full study, scaled by his mean income. Over 60% of the population won some amount over the course of

[36] Appendix Table A.4 summarizes which groups were used for which analyses, along with an explanation of why this choice was made.

the study, with many having won substantial sums relative to weekly earnings. If winnings were randomly assigned, we could look directly at the impact of these wins on expenditure behavior. However, the amount that one can win is conditional on the number and types of bets that he places. To identify a valid causal effect of winnings on expenditures, I implement a selection on observables approach. This is done by characterizing every individual's bets in each time period of the study. Knowing the amount spent on betting, the number of tickets purchased, the payoff they were targeting, and the number of independent matches included on the betting tickets, I can also infer the bookmakers' assessment of the likelihood that a bet will win. I then characterize the full distribution of betting realizations for each bettor in each time period by their moments. Controlling for these moments and their higher order terms, I consider the realized amount of winnings to be effectively random.[37] Additional details about the structure of betting in Uganda are contained in Appendix B1, while Appendix B3 provides details on how I converted reported betting into the moments of a betting portfolio.

Using this identification strategy, I look at the effect of winnings on whether an individual made a lumpy purchase above a given threshold in that time period. The estimating equation used for the analysis is:

$$Y_{i,t} = \beta_0 + \beta_1 W_{i,t} + \sum_{b=1}^{3} BetMoments_{i,t}^{b} + \lambda X_{i,t} + \gamma_i + \delta_t + \psi_s + \epsilon_{i,t}$$

$Y_{i,t}$ is whether a lumpy purchase costing more than a given threshold relative to the individual's mean income was made by individual i in time period t. $W_{i,t}$ is the amount of reported winning in a given two-week period. $BetMoments_{i,t}$ is the calculated moments (mean, variance, skewness, and kurtosis) and higher-order terms (quadratic and cubic) of the betting portfolio for individual i in period t, and $X_{i,t}$ are time-varying covariates. Individual fixed effects, survey round fixed effects, and time fixed effects are also included. Standard errors are clustered at the individual level.

Ex-ante, it is not clear what the threshold for a "large" lumpy expenditure should be. Instead of taking an arbitrary stance, I try multiple thresholds. In Table 3, I show results for thresholds of 0.5, 1, 2, and 4 times weekly income with the win amount scaled relative to income. The win amount is similarly scaled relative to mean weekly income.

At all thresholds, the effect of winnings on the likelihood of making large purchases

[37] This approach is similar to that utilized by Anderson (2016) in his analysis of the impact of college sports success on fund-raising ability, where he argues that, conditional on bookmaker spreads, winning is uncorrelated with potential outcomes.

is positive. Column (1) shows that additional winnings equal to mean income raise the likelihood of making an expenditure equal to half of mean income by 2 percentage points, significant at the 95% confidence level. Column (2) suggests that this effect on expenditures equal to or greater than mean income is 1.8 percentage points, or 6%, significant at the 90% confidence level. Despite meaningful magnitudes relative to their mean incidence, the estimates in Columns (3) and (4) for thresholds of double and quadruple mean income are no longer statistically significant. This is largely because there are not enough wins of sufficient size in the data to plausibly affect these infrequent purchases, lowering statistical power for these outcomes. Figure 7a shows these regression results graphically, focusing on thresholds up to two times mean income. In the figure, the x-axis represents the threshold for the biggest expenditure in that time period, while the y-axis is the estimated coefficient on the win amount, reflecting the regression results from Table 3. Figure 7c rescales Figure 7a by the mean of the outcome variable so that it can be interpreted as the proportional increase in likelihood of a purchase of a given size as a result of additional winnings. For many of these thresholds, winnings significantly increase the likelihood of making large expenditures.

Gambling for liquidity generation should be more likely among people with low saving ability. Figure 7b recreates the results in Figure 7a after splitting the sample between low and high ability savers. The effect for high ability savers is drawn in red, while low ability savers are shown in blue, with the 95% confidence interval represented with blue dashed lines. For nearly all of these thresholds, additional winnings have a positive and significant effect on low ability savers, always larger than for those with high saving ability for whom the effect is not distinguishable from zero at any threshold. We now see that additional winnings equal to an individual's mean income are associated with a 4.1 percentage point increase in the likelihood of making a purchase equal to or greater than his mean income among people with low saving ability; this is an increase of approximately 14.2%, significant at the 95% confidence level. The effect on those with high saving ability is estimated at slightly below zero and cannot be distinguished from no effect. The rest of these regression results are presented in Appendix Table A.5. Figure 7d uses the same regression coefficients, but scales them by the mean of the outcome variable, showing an increasing magnitude as the size of the expense gets larger.

These results used a binary outcome variable indicating whether or not a lumpy purchase was made above a given threshold. I also conduct similar analysis using a continuous measure of the size of the biggest expense made in that period. The findings are similar to those above and contained in Appendix F. Additional winnings increase the size of an individual's

largest expense by 0.052-0.33 cents per dollar on average, with effect sizes between 0.11 and 0.53 for people with low saving ability.

Although ex-post usage of winnings does not definitively confirm ex-ante motivation for betting, this evidence suggests that lumpy expenditures and betting are tightly linked in this context. Additionally, individuals likely anticipate their own future consumption decisions in deciding whether or not to bet.

6.2 Commitment-Savings Treatment

The model predicts that an improvement in saving ability should decrease demand for betting. To test this causal relationship, I use a randomized field experiment to introduce a new saving technology and create exogenous variation in saving ability. Randomly selected participants were chosen to receive a soft commitment-savings device in the form of a wooden savings box. These boxes are nailed closed and have a small slit in the top so that, like a piggy bank, money can be easily deposited but cannot be retrieved without breaking it open. These boxes are commonly found in Ugandan markets and are therefore familiar to the study participants. At the end of the third visit, field team members gave randomly selected respondents a saving box and assisted them in writing down their saving target on the outside.

In the first wave, 25% of respondents were selected to receive the boxes, whereas 50% of participants in the second wave were selected. Panels (a) and (b) of Table 4 show baseline balances by wave. Despite random assignment, the endline lumpy prime weakly correlates with the saving box treatment in the second wave, significant at the 10% level. All analyses control for the effect of the lumpy prime and additional robustness checks are conducted to ensure that the observed effect is not being driven either by this correlation or by an interaction between treatments.[38]

At the endline, one month after the savings boxes were distributed, interviewers asked participants if they had used a savings box at any time in the preceding month. People in the treatment group were 53 percentage points more likely to have used a saving box compared a control group mean of 16%, a difference significant at the 99% confidence level. Although low and high ability savers have slightly different propensity to use saving boxes in the absence of the intervention, both respond similarly to the treatment. When estimating

[38]The effect of these treatments is in opposite directions. Therefore, the positive correlation works against finding a measurable effect for either result. Both results survive robustness checks to ensure that the interaction is not driving either result.

the treatment on the treated effect of the saving boxes, random assignment to the treatment group will be used as an instrument for saving box use.[39]

I estimate the effect of the saving box treatment using a difference in differences strategy with participants in Wave 2 for whom I have both a baseline and endline measure of betting demand. I use the following equation:

$$B_{i,t} = \beta_0 + \beta_1 SaveBox_{i,t} + \lambda X_{i,t} + \gamma_i + \delta_t + \epsilon_{i,t}$$

$B_{i,t}$ is an indicator of whether the respondent chose the maximum number of betting tickets from the offer. $SaveBox_{i,t}$ is an indicator of whether or not individual, i, had been offered a saving box by time t. $X_{i,t}$ are time-varying covariates for individual i. Individual fixed effects and time fixed effects are all also included. Standard errors are clustered at the individual level.

Table 5 shows these estimation results. Using within variation and a full set of covariates, as shown in in Column (2), we see that receiving a savings box reduced the likelihood of demanding the maximum number of tickets by 13.15 percentage points from a control group average of 49.88. This constitutes a reduction of demand by 26%, significant at the 95% confidence level. The results show strong heterogeneity. As shown in Column (3), low saving ability respondents reduce demand by 39%, significant at the 99% confidence level. For people with high saving ability, this falls to just 1.12 percentage points and is no longer statistically significant, while the difference between these groups shows significance at the 90% confidence level in Column (5). In Columns (6) - (8), we see that the treatment on the treated estimate is nearly double the size of the average treatment effect. In the pooled estimate, the likelihood of demanding the maximum number of tickets falls by 25 percentage points.[40]

These are large magnitudes, particularly for people with low saving ability, shown in Columns (3) and (7). However, despite randomization, the baseline measure of betting demand for people with low saving ability in the treatment group was significantly higher than for those in the control group. The difference-in-differences specification used above

[39]Table A.9 shows that the treatment had a very high takeup rate.

[40]Switching to a continuous measure of tickets demanded reveals consistent results although significance is lost in some regressions from using a lower powered outcome variable. These results are shown in Table A.10. The average treatment effect (as in Column (3)) shows a reduction of 0.3 tickets demanded, down from an average of 2.5, a reduction of 12%. The treatment on the treated effect (as in Column (6)) is again almost twice as big as the average treatment effect, causing a reduction of 0.57 tickets or roughly 23% among those who were induced to use a savings box by the treatment. The effect on the low saving ability participants is significant at the 95% level while the others are slightly below 90% significance.

should appropriately adjust for these baseline differences and still be a valid causal effect. However, to ensure that the estimated effect is not purely the result of a baseline irregularity, I re-estimate the effects of the saving box using a cross-sectional analysis of all participants in the full study, ignoring the baseline measure.

This alternative estimation strategy results in an average treatment effect of a 7.27 percentage point (16%) reduction in betting demand, significant at the 95% confidence level. Instrumental variables estimation gives a treatment on the treated effect of a 33% decrease in betting demand. However, we no longer see clear heterogeneity by saving ability. These results are shown in Appendix Table A.8.[41]

Regardless of specification, the analysis shows a large and significant reduction in betting demand, between 16 and 26% in the cross-sectional and difference-in-differences estimations, respectively. The unstable response by saving ability across the two specifications suggests that this dimension of heterogeneity should be treated with caution. A treatment designed to reduce exposure to temptation and social pressures on money designated for saving is not necessarily well-suited for a group of people characterized by limitations on the portion of income they have available for saving. Other saving products may be better tailored to the needs of this group of "low ability" savers.

The model suggested that improved saving ability should affect demand for betting through two main mechanisms: crowding out normal goods and undermining the financial appeal of betting. There does appear to be a reduction in on-hand liquidity as well as an increase in available liquidity (not necessarily on-hand) for people with low saving ability, but neither of these effects is significantly different from zero.[42] Regardless of mechanism, the reduction in betting demand of between 16-26% already demonstrates an important link

[41] Appendix Table A.12 confirms that these results are not driven by correlation with the lumpy expenditure prime. Interacting the saving box treatment with the lumpy prime *increases* the magnitude and significance of the overall average treatment effect by approximately 33%. Interestingly, the effect now seems stronger among people with high saving ability, which might be because they have more discretionary income to put toward saving. Table A.12 also suggests that the interaction of the saving box and the lumpy prime is positive and significant for people with high saving ability. This could result from people feeling that responsible saving decisions in one part of their weekly budget frees them to take riskier choices with the remainder of their discretionary spending. Appendix Table A.11 shows that switching the outcome to the count of tickets demanded greatly reduces statistical power and loses significance, although the sign of the effect remains the same.

[42] There is some evidence that the treatment group had more liquidity available to them at the endline, which could diminish the appeal of betting as a method of liquidity generation (see Appendix Table A.13). There is also a statistically insignificant but potentially economically meaningful reduction in on-hand liquidity among people who received the saving box (see Appendix Table A.14, showing a small (insignificant) reduction in having 3 USD on hand, although the effect is positive (insignificant) on a threshold of 15 USD. A reduction in on-hand liquidity could raise the opportunity cost of betting and constitute evidence that saving crowds out all normal goods.

between saving and betting. The final two results use lab-in-the-field experiments in order to isolate the importance of the financial motive for betting.

6.3 Prime on Lumpy Good

The third result uses a lab-in-the field experiment to show that increasing the salience of a desired lumpy expenditure increases demand for betting. During the baseline of the full study, interviewers asked respondents to identify a large business, household, or personal expense they wanted to make in the next few months. During the condensed study, these questions were asked at the beginning of the survey. For randomly selected respondents, interviewers went through a dialog referring to these desired expenditures just before the betting ticket offer. They stated, "Earlier, you mentioned that you wanted to buy ____. How much would it cost? How much more money do you think you would need in order to be able to make that expense? Do you know where you would go to purchase it?"[43] These questions were designed to make the respondent reflect on this expenditure, increasing its salience, before measuring betting demand. Right after these primes, respondents were offered the choice between betting tickets and cash. Respondents in the control group were asked these same questions immediately *after* the betting ticket offer. Figure 8 illustrates the structure of this randomization.

To the extent that the appeal of betting results from its potential as a strategy of liquidity generation, the model predicts that an increase in salience should increase betting demand. However, the portion of betting demand that is purely based on consumption may fall with increased salience of a desired expenditure because the money offered could be set aside and potentially saved in order to make that expenditure later. These predictions go in different directions and reflect different dimensions of the underlying demand for betting.[44]

The regression equation for this analysis is:

$$B_i = \beta_0 + \beta_1 LumpyPrime_i + \lambda X_i + \epsilon_i$$

B_i is again whether the maximum number of tickets were demanded. β_1 is the coefficient of interest and measures the effect of having received the lumpy prime before the betting

[43]For respondents in the full study, the enumerators first checked to see whether the large expenditure had already been made and, if so, whether they needed to make that expense again (as in the case of rent or school fees). These answers were controlled for in the analysis.

[44]It is also possible that salience shifts time preferences. If this were the case, then the increase in salience would make saving less appealing. However, we would only see a big increase in betting demand if betting were, in fact, seen as a credible alternative strategy of getting that good.

ticket offer, where $LumpyPrime_i = 1$ for those in the treatment group and equal to zero for those in the control group. X_i is a set of covariates for individual i including an indicator for the time period in which that person was offered the betting tickets. The full specification also includes the amount of cash offered and treatment status for the other randomized treatments in the study. Robust standard errors are used to adjust for heteroskedasticity in the error term. Randomization balance is shown in Table 4 Panel (c). None of the 22 variables checked are significantly imbalanced across treatment and control groups.

Because $LumpyPrime_i$ is randomly assigned, the estimate of β_1 can be interpreted as the causal effect of the prime on betting demand. Table 6 shows the results. The first four columns show that, regardless of which covariates are included, the lumpy good prime has a highly significant, stable, and positive effect on the likelihood that a respondent chose the maximum number of tickets. The preferred full specification in Column (4) shows an effect size of 7.2 percentage points relative to a control group mean of 41.8, a treatment effect of 17.2% that is significant at the 99% confidence level. Again, I expect that betting will be particularly appealing as a mode of liquidity generation for people with limited ability to save. Therefore, I split the sample by the same dimension of heterogeneity as used in the analysis of winning usage. Column (5) shows that the effect of the prime on people with low saving ability climbs to 12.2 percentage points, an effect size of 25.6% relative to the mean of the relevant control group. For people with relatively high saving ability in Column (6), this effect falls to 2.1 percentage points and I cannot reject the null hypothesis of no effect. The interaction term in Column (7) between saving ability and the lumpy prime confirms that we can reject the null that these groups have the same response to the treatment.[45]

These results show that an increase in the salience of a lumpy expenditure increased betting demand. Participants themselves view betting as a mechanism by which they can generate liquidity needed for their expenditures. Additionally, the results exhibited an identical heterogeneous response to that observed in an analysis of the effect of winnings on lumpy expenditures. In the group of people for whom saving is least likely to be a feasible strategy of liquidity generation, winning usage and betting demand both appear strongly linked to lumpy expenditures.

[45]This also can be checked using a triple interaction with raw saving ability, mean income, and the lumpy good prime. The triple interaction term remains significant at the 95% confidence level and climbs to an estimated difference of 14.6 percentage points. These results remain significant if I switch the outcome variable to the proportion of total tickets demanded; these results are included in Appendix Table A.15.

6.4 Budgeting for Savings

The final result uses another lab-in-the-field research design to identify the causal effect of a change in perceived ability to save on betting demand. Changing perceived ability to save should also affect the relative appeal of saving to betting so long as the update is credible. To do this, I built a budgeting exercise into the condensed study. Interviewers guided randomly selected respondents through a brief budgeting exercise nested inside the lumpy prime treatment, aimed at assessing ability to save.[46] Figure 9 shows the structure of this randomization. Table 4, Panel (d) shows the balance by covariates.[47]

Early in the survey, interviewers asked respondents about their typical weekly earnings and essential expenditures on food, transportation, and rent. They were asked about a large lumpy expenditure that they hoped to make within the next few months and how much money they thought they could save per week without putting excessive strain on their personal or household finances. At the end of the survey, after the lumpy good prime, interviewers guided randomly selected respondents in the treatment group through a brief budgeting activity before the betting ticket offer. Those in the control group went through these same questions after the offer. In the budgeting dialog, respondents were told, "Earlier in this interview you said that you earn _____ UGX in a typical week. You also said you spent _____ on food, _____ on transportation, and _____ on rent in a normal week. This leaves you with _____ UGX per week. How much money do you think you could realistically save per week?" Tablets used for data collection automatically calculated and filled in the blanks based on earlier responses. Respondents were unconstrained in their answers to this final question and were free to ignore this information.[48]

The anticipated effect of this treatment depends on whether a respondent's updated estimate is above or below his original naively estimated saving ability. In other words, the sign of the treatment will depend on whether the respondent is learning encouraging or discour-

[46] This group was omitted during the lumpy good prime analysis above.

[47] Baseline proportionate saving ability is significantly different across treatment despite randomization. This is to be expected, having checked for balance across 20 variables. What matters most for this result is that saving updates are balanced across treatment status and that we do not see any statistical differences by raw saving ability, raw saving update amount, and proportionate update amount. Baseline saving ability levels are included in all regressions to account for this imbalance.

[48] After the respondent gave an answer, the enumerator said, "At that rate of saving, it would take you _____ weeks/months to have enough money to make that expense." Analysis similar to that detailed below was also conducted with respect to the time update. Learning that saving would take twice as long as previously anticipated does increase demand for the maximum number of tickets by 11.6%, although it is only significant at the 90% confidence level. This is consistent with the theory that betting and saving are competing strategies of liquidity generation. I do not see any clear heterogeneity by measures of Beta and Delta discounting. These results are in Appendix Table A.18.

aging information. In the data, 48% of respondents decreased their estimated saving ability, 27% did not update their estimate after the discussion, and 25% increased their estimate. The median raw positive update was 15,000 Ugandan shillings (approximately 4.25 USD) and the median proportionate update was 10% of income. The median raw negative update was 17,000 Ugandan shillings (approximately 4.85 USD). The median negative proportionate update was 17% of income. Figure 10a shows the raw update size (in thousands of Ugandan shillings) and Figure 10b shows the update scaled relative to mean income. By having both a naive and assisted estimate of saving ability for each person in the sample, I can assess the impact of receiving this update on betting demand while controlling for the appropriate counterfactual of people who would have gotten an update of the same size.

Because the content of the saving ability update determines whether saving has become more or less appealing, I anticipate heterogeneous treatment effects, estimated using the following regression equation:

$$B_i = \beta_0 + \beta_1 LumpyPrime_i + \beta_2 Budget_i + \beta_3 (Budget \times Update)_i + \beta_4 Update_i + \lambda X_i + \epsilon_i$$

B_i is the outcome measure of betting demand from the betting ticket offer. $LumpyPrime_i$ indicates whether the individual received the lumpy prime. $Budget_i$ is an indicator for being assigned to the budgeting treatment group and doing the budgeting activity before the ticket offer. β_2 is the effect of the budgeting activity independent of the update. β_3 is the effect of the content of the update. $Update_i$ is calculated as the difference between a respondent's new, assisted estimate of saving ability and his original, naive estimate. It is positive if the new estimate is larger than his original, naive estimate. In certain specifications, the update is scaled by mean income. β_4 captures differences in people with different update sizes independent of whether they did the budgeting exercise before or after the ticket offer. X_i is a set of covariates for individual i. The full specification also included the amount of cash offered, as well as all other treatments included in the study and a full set of covariates. Robust standard errors are used to adjust for heteroskedasticity in the error term.

Table 7 shows the results. Column (1) shows that the budgeting exercise had a negative but insignificant effect on demand for betting tickets. The estimates in Columns (2) and (3) show that, regardless of whether the updating is estimated in raw local currency or converted to proportion of income, improving perceived saving ability lowers betting ticket demand. Column (3) suggests that, when an individual learns that he can save 10% more of his income than previously thought, the likelihood that he will demand the maximum

number of betting tickets decreases by 5.4 percentage points or approximately 13%.[49]

As discussed, responses are likely different depending on the content of the update. Column (4) codes the saving updates as positive or negative (zero is omitted). This analysis shows strong heterogeneity, with a large and negative effect for people who receive positive information about their saving ability and an insignificant effect for those learning negative information. A positive update causes a 34.9% reduction in the likelihood of demanding the maximum number of betting tickets relative to the relevant control mean. The effect for the negative update is indistinguishable from zero. Columns (5) and (6) test for a linear relationship between the update amount and betting demand, again splitting the treatment at zero. The raw update in Column (4) estimates a similar treatment effect on both sides of zero, though the proportionate measure suggests that this effect is driven by positive updates. The median positive update of 15,000 shillings would cause a six percentage point decrease in the likelihood of demanding the maximum number of betting tickets, as shown in Column (5), or approximately 14% of the mean. A median positive proportional update of 0.10 is estimated to cause nearly a ten percentage point decrease using the estimates in Column (6) or just under 25% of the mean.

Imposing linearity on the estimates is a rigid assumption and so I conduct a non-parametric estimation of the treatment effect of the budgeting exercise with the proportionate update measure. Figure 11 shows the non-parametric lowess regression of the saving update, scaled to weekly income, on demand for the maximum number of tickets. In addition, a linear model with a spline at zero is included as a reference point. These non-parametric estimates suggest that there is no clear effect of the saving prime on people learning negative information, whereas positive information decreases demand for betting tickets with greater magnitude effects for larger update sizes.

Attributing the response of the saving box treatment to a change in the relative appeal of saving and betting as competing methods of liquidity generation was confounded by other factors, including on-hand liquidity and the possibility that saving crowds out all current expenditures. However, an update revealing that a person has *more* disposable income available for saving does not face the same challenges and would predict an *increase* in demand for betting if betting were exclusively a consumption good. We see the opposite.

[49]This approach will be valid so long as the betting ticket offer did not affect peoples' responses in the budgeting exercise. If it had, then those who went through the exercise after the betting ticket offer could have systematically different potential updates, which would result in invalid controls for people who did the exercise prior to the offer. Robustness checks show that the timing of the betting ticket offer did not have a significant effect on the raw update size, size of the update relative to income, or likelihood of having a positive, negative, or zero update. These robustness checks are shown in Appendix Table A.17.

The overall reduction in betting demand for people who received a positive saving update suggests that the perceived increase in feasibility of saving as a liquidity generation strategy undercut that source of appeal for betting.

7. Comparing Betting and Saving

The model in this paper framed saving as the primary alternative to betting, consistent with survey responses citing them as the two most likely sources of liquidity for a desired expense (see Table 2, Panel (b)). The results in Section 6 showed shifts in betting demand consistent with participants viewing betting as a mode of liquidity generation and in competition with saving. This section looks at whether betting could be a purely rational, utility maximizing, liquidity generation strategy without allowing for any direct utility from participation. Even stashing money under the mattress should be better than accepting 35-50% expected losses from betting. However, participants reported a number of challenges impeding their ability to save, including risk of theft, pressure from family or friends, or personal temptation. These all factor into expected "losses" when money is set aside with the intention of saving. Inflation will lower this effective interest rate further. Ultimately, given his level of patience, an individual has to consider this expected return on saving as well as the return on betting in order to determine his preferred strategy.

People save in many ways and each technology or saving strategy has its own benefits and drawbacks. For example, formal bank accounts are a common vehicle for saving and should substantially reduce expected losses and theft while also reducing exposure to temptation and social pressure. However, transaction fees in this setting are high and counteract most of these benefits, consistent with evidence by Dupas et al. (2016) in nearby Kenya. Based on focus group discussions, deposit fees in Uganda are equal to approximately 3% of the median study participant's weekly income. Given that Ugandan banks often have long lines and are only open during regular working hours, making deposits requires time away from productive activity, conservatively estimated at one hour per transaction, equivalent to an additional 2% of weekly income based on a 50 hour working week. For someone saving 10% of his weekly income, weekly deposits, which would best reduce exposure to theft, temptation, or social pressure, would impose expected losses from transaction costs of 50% of his deposit value, even before accounting for losses from inflation or withdrawal fees. A less frequent deposit strategy would reduce transaction costs but at the expense of limiting the other benefits of this strategy.

Rotating saving groups (roscas) are also very common in Uganda. 50% of respondents reported using them. Their primary advantage is a 50% reduction in expected wait time before a full payout relative to saving independently (Anderson and Baland 2002). In addition, social pressure acts as a commitment device for people who may not always follow through with their saving goals. However, roscas also depend on the efficacy of social sanctions among their members for the group to survive (Anderson et al. 2009). The rigidity of the weekly payment structure and the threat of social sanctions impose their own risks. Someone who has a bad week at work could risk a sanction for missing a payment and might be forced to either sacrifice needed short-term consumption or to borrow elsewhere at high cost to cover the payment. Additionally, the amount of the contributions and payout are the result of a bargaining process among group members and might not correspond with the optimal targets and contribution levels for an individual's unique liquidity needs.

Acknowledging that different saving strategies have different cost and benefit structures, I use the simplest form of saving, cash savings, as a benchmark to compare saving and betting strategies. From what is known about the structure of betting, I calculate the expected payoff to a person using betting to pursue a desired expenditure. I similarly calculate the expected payoff to that same person if he instead pursues a (cash) saving strategy. Setting these two approaches equal, I can then identify the minimum return on saving required for a person to be willing to save instead of bet, over a range of reasonable patience levels.

I first expand my model to allow for more than two time periods. Additionally, I need to make some assumptions about a number of parameters. I set the amount to be either saved or bet at 10% of weekly income (between the mean and median betting expenditures in the population) and the payout of a win equal to the price of the lumpy expenditure at twice the size of the individual's income (approximately 200,000 UGX or 60 USD). I also must make assumptions about the size of forgone utility from the divisible good while pursuing the lumpy expenditure and the payoff to the large expenditure, η, from the original model. Forgone utility is treated essentially as a numeraire, and I try a wide range of utility payoffs to the lumpy expenditure; results are not sensitive to this parameter choice.[50] The return on betting is captured in the likelihood of winning the bet, p, and estimated based on knowledge of the structure of betting in Uganda for a payout of this size.

Research in neighboring Kenya by Mbiti and Weil (2013) estimates a yearly discount factor of 0.64, suggesting a weekly discount factor of approximately 0.9915. With these

[50]The ratio between forgone utility and the utility payoff to the lumpy good essentially captures the individual's risk aversion in that it defines how steep or flat an individual's utility function is at that level of income. I show the results again with two different choices for η in Appendix Table A.2.

values, I can examine a range of reasonable weekly discount factors, δ, and solve for the "tipping point" rate of return on saving, γ^*, that leaves people indifferent between betting and saving. People with $\gamma \geq \gamma^*$ prefer to save and those with $\gamma < \gamma^*$ prefer to bet. Additional details and derivation are contained in Appendix C1.

In Figure 12, I trace the locus of weekly discount factors and threshold return to saving. As expected, less patient people are willing to tolerate less expected losses from saving than people with greater patience. For Mbiti and Weil's estimated discount factor, threshold gamma is 0.716, suggesting that these people would prefer betting if they expect more than 28% losses from money set aside for saving. Of course, people who entirely discount the future, on the left side of the diagram, will always prefer betting over saving. For people with $\gamma < 0.63$ who expect to lose 37% of each dollar saved, then even perfect patience ($\delta = 1$) will not be sufficient to sustain saving. This shows a convergence toward the expected return on betting.

What is a reasonable estimate for the rate of return on saving, γ? In my review of the literature, I have not seen this parameter directly estimated. Ideally, it should incorporate at least five main components: (1) Inflation, (2) Expenditures on Temptation Goods, (3) Social Pressures, (4) Theft/Loss, and (5) Transaction Costs. For each of these components, I estimate a reasonable range of discount sub-factors γ_{1-5} and take their product in order to construct a range of return to saving such that $\gamma = \Pi_1^5 \gamma_i$.[51] Table 8 summarizes the range of sub-discount factors resulting from my estimates. This exercise suggests that the rate of return to saving likely falls between 0.6415 and 0.9215 for people in this population. Details of the assumptions and sources of these estimates are provided in Appendix C2. Comparing this range with the "tipping point" values of return on saving calculated earlier suggests that, for a substantial portion of people, betting will be the preferred strategy.[52] [53] Even if betting is not strictly preferable to saving, these differences may not be perceptible to the bettors themselves or may be close enough that direct enjoyment of betting pushes them over the edge.

[51] Taking the product of these components is a conservative assumption. If, instead, these taxes are removed from saving simultaneously, then the range of estimated γs could be considerably lower.

[52] Without a clear strategy for how to map these discount factors to my sample, it is difficult to more precisely estimate the size of this proportion, but these calculations broadly suggest that there are likely to be many people for whom betting is preferable.

[53] In calculating the return to betting, I assumed that bettors are following an optimal betting strategy and thus getting expected losses on the low end of my estimates (around 35%). Making less favorable assumptions about betting strategies would put the expected losses from betting closer to 50%. This would reduce the tipping point γs downward by roughly 15-20% and could tip the balance back toward saving for some people. This adjustment would not affect the broad observation that betting and saving strategies may result in very similar expected payoffs for many people.

8. Conclusion

Over the past decade, sports betting has seen considerable growth, but has expanded most rapidly in African markets. This paper has looked closely at the behavior of 1,715 sports bettors in Kampala, Uganda, using a range of different empirical methods. The findings contribute to a number of strands of literature in economics and provide clear policy implications.

Sports betting is distinct from other consumption or temptation goods as a result of the gamble contained at the center of the bet. This financial appeal, along with unmet liquidity needs increases demand for betting. This was demonstrated in an analysis of the effect of winnings on lumpy expenditures as well as an analysis of the response to a prime on the salience of a desired large expenditure. These responses were particularly strong among people with low ability to save and who therefore had limited alternatives to generate this liquidity. In addition, a treatment that increased ability to save and the appeal of saving as as an alternative strategy of liquidity generation, reduced betting demand. This has direct policy implications even without disentangling the mechanisms of this effect. The results using the experimental prime on desired lumpy expenditures and a budgeting activity suggest that the financial mechanism is a significant factor driving betting demand and that the effect of improving saving ability on betting is more than just a mechanical reduction of all current expenditures, but reflects the diminished appeal of betting as a way to get liquidity.

Betting is an enjoyable activity for many of its participants, but financial constraints and demand for liquidity are also significant drivers of betting demand that should be taken seriously. This should not be surprising; this is why bettors themselves say that they bet. However, expected losses between 35 and 50% make betting an inefficient way to generate liquidity. This is particularly concerning because the population in this study sits near or below the poverty line and provides critical income for both themselves and their families. With average participation levels between 8.5-12.5% of weekly income, these losses have serious implications for bettors and their families.

If policy-makers are interested in deterring gambling, this population needs better alternatives to undermine the financial motive for betting. Saving is a struggle for many people across the globe, but in developing countries these challenges may be particularly severe. And yet, even the simple interventions tested in this study reduced betting demand. More ambitious interventions, such as lowering the cost of secure saving or expanding access to affordable credit, may have stronger effects. Broadly, making sure that financial services are designed for and reach this vulnerable population could substantially lower this source of

demand for betting in a vulnerable population.

1 References

Abel, Martin, Shawn Cole, and Bilal Zia (2015). "Debiasing on a Roll: Changing Gambling Behavior through Experiential Learning." In: *Policy Research Working Paper* 7195, pp. 187–202.

African Development Bank (2011). "Institutions for Credit Markets in Africa." In: *Africa Capacity Development Brief* 2 (5).

Ahaibwe, Gemma et al. (2016). "Socio Economic Effects of Gambling: Evidence from Kampala City, Uganda." In: *Economic Policy Research Centre* (126).

Anderson, Michael (2016). "The Benefits of College Athletic Success: An Application of the Propensity Score Design with Instrumental Variables." In: *Review of Economics and Statistics* (Accepted).

Anderson, Siwan and Jean-Marie Baland (2002). "The Economics of Roscas and Intrahousehold Resource Allocation." In: *The Quarterly Journal of Economics*, pp. 963–995.

Anderson, Siwan, Jean-Marie Baland, and Karl Ove Moene (2009). "Enforcement in Informal Saving Groups." In: *Journal of Development Economics* 90, pp. 14–23.

Ariyabuddhiphongs, V. (2011). "Lottery Gambling: A Review." In: *Journal of Gambling Studies* 27, pp. 15–33.

Bailey, Martin J., Mancur Olson, and Pau. Wonnacott (1980). "The Marginal Utility of Wealth Does Not Increase: Borrowing, Lending, and Friedman-Savage Gambles." In: *The American Economic Review* (703), pp. 372–9.

Banerjee, Abhijit (2013). "Microcredit Under the Microscope: What Have We Learned in the Past Two Decades, and What Do We Need to Know?" In: *Annual Review of Economics* 5, pp. 487–519.

Banerjee, Abhijit and Esther Duflo (2007). "The Economic Lives of the Poor." In: *Journal of Economic Perspectives* (21(1)), pp. 141–168.

Banerjee, Abhijit and Sendhil Mullainathan (2010). "The Shape of Temptation: Implications for the Economic Lives of the Poor." In: *NBER Working Paper Series* 15973.

Barber, Brad M. and Terrance Odean (2007). "All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors." In: *The Review of Financial Studies* 21 (2), pp. 785–818.

Beck, Thorsten and Robert Cull (2014). "Small- and Medium-Sized Enterprise Finance in Africa." In: *Africa Growth Initative at Brookings - Working Paper* 16.

Becker, Gary S. and Kevin M. Murphy (1988). "A Theory of Rational Addiction." In: *Journal of Political Economy* 96 (4), pp. 675–700.

Besley, Timothey, Stephen Coate, and Glenn Loury (1993). "The Economics of Rotating Savings and Credit Associations." In: *The American Economic Review* 83 (4), pp. 792–810.

Bordalo, Pedro, Nicola Gennaioli, and Andre Shleifer (2012). "Salience Theory of Choice Under Risk." In: *The Quarterly Journal of Economics*, pp. 1243–1285.

Bruce, Alistair (2013). "Betting Motivation and Behavior." In: ed. by Leighton Vaughan-Williams and Donald S. Siegel. Chap. The Oxford Handbook of the Economics of Gambling, pp. 451–473.

Brune, Lasse (2016). "The Effect of Lottery-Incentives on Labor Supply: Evidence from a Firm Experiment in Malawi." In: *Working Paper*.

Brune, Lasse. et al. (2015). "Commitments to Save: A Field Experiment in Rural Malawi." In: *Economic Development and Cultural Change*, pp. 187–220.

Casaburi, Lorenzo and Rocco Macchiavello (2016). "Firm and Market Resonse to Saving Constraints: Evidence from the Kenyan Dairy Industry." In: *Working Paper*.

Collins, Daryl et al. (2009). *Portfolios of the Poor: How the World's Poor Live on Two Dollars a Day*. Princeton University Press.

Conlisk, John (1993). "The Utility of Gambling." In: *Journal of Risk and Uncertainty* (6), pp. 255–275.

Crossley, Thomas F., Hamish Low, and Sarah Smith (2016). "Do Consumers Gamble to Convexify?" In: *Journal of Economic Behavior and Organization* (forthcoming).

Dupas, Pascaline and Jonathan Robinson (2013a). "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya." In: *American Economic Journal: Applied Economics* 5.1, pp. 163–92.

— (2013b). "Why Don't the Poor Save More? Evidence from Health Savings Experiments." In: *American Economic Review* 103 (4), pp. 1138–1171.

Dupas, Pascaline et al. (2016). "Challenges in Banking the Rural Poor: Evidence from Kenya's Western Province." In: ed. by Sebastian Edwards, Simon Johnson, and David N. Weil. Chap. African Successes, Volume III: Modernization and Development, pp. 63–101.

European Gaming and Betting Association. *Sports Betting: Commercial and Integrity Issues*. http://www.egba.eu/media/Sports-Betting-Report-FINAL.pdf. Accessed: 2016-05-05.

Friedman, Milton and L. J. Savage (1948). "The Utility Analysis of Choices Involving Risk." In: *The Journal of Political Economy* 56.4, pp. 279–304.

Grote, Kent and Victor A. Matheson (2013). "The Economics of Lotteries: A Survey of the Literature." In: ed. by Leighton Vaughan-Williams and Donald S. Siegel. Chap. The Oxford Handbook of the Economics of Gambling, pp. 670–691.

Hartley, Roger and Lisa Farrell (2002). "Can Expected Utility Theory Explain Gambling." In: *The American Economic Review* 92.3, pp. 613–624.

Heath, Chip and Amos Tversky (1991). "Preference and Belief: Ambiguity and Competence in Choice under Uncertainty." In: *Journal of Risk and Uncertainty* (4:4-28), pp. 5–28.

Ho, Daniel E. and Kosuke Imai (2008). "Estimating Causal Effects of Ballot Order from a Randomized Natural Experiment." In: *Public Opinion Quarterly* 27 (2), pp. 216–240.

Imbens, G., D.B. Rubin, and B.I. Sacerdote (2001). "Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players." In: *The American Economic Review* (91(4)), pp. 778–794.

Jakiela, Pamela and Owen Ozier (2015). "Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies." In: *Review of Economic Studies* 83(1), pp. 231–268.

Karlan, Dean, Aishwarya Lakshmi Ratan, and Jonathan Zinman (2014). "Savings By and For the Poor: A Research Review and Agenda." In: *The Review of Income and Wealth* 60(1), pp. 611–616.

Kast, Felipe, Stephan Meier, and Dina Pomeranz (2014). "Under-Savers Anonymous: Evidence on Self-Help Groups and Peer Pressure as a Savings Commitment Device." In: *Harvard Business School - Working Paper* (12-060).

Kwang, Ng Yew (1965). "Why do People Buy Lottery Tickets? Choices Involving Risk and the Indivisibility of Expenditure." In: *Journal of Political Economy* 73.5, pp. 530–535.

Mbiti, Isaac and David N. Weil (2013). "The Home Economics of E-Money: Velocity, Cash Management, and Discount Rates of M-Pesa Users." In: *American Economic Review: Papers and Proceedings* (103(3)), pp. 369–374.

PricewaterhouseCoopers (2014). *Raising the Stakes in Africa, Gambling Outlook: 2014-2018*. http://www.pwc.co.za/en_ZA/za/assets/pdf/gambling-outlook-2014.pdf.

Raylu, Namrata and Tian P.S. Oei (2002). "Pathological Gambling: A Comprehensive Review." In: *Clinical Psychology Review* (22), pp. 1009–1061.

Schilbach, Frank (2015). "Alcohol and Self-Control: A Field Experiment in India." In: *Working Paper*.

Snowberg, Erik and Justin Wolfers (2010). "Explaining the Favorite-Longshot Bias: Is It Risk-Love or Misperceptions?" In: *Journal of Political Economy* 118.4.

Ssengooba, Kizito and Bruno L. Yawe (2014). "Gambling and Mobile Money Payments: A Case Study of Sports Betting in Uganda." In: *IMTFI - Working Paper*.

Figures

Figure 1: Friedman-Savage Utility Function

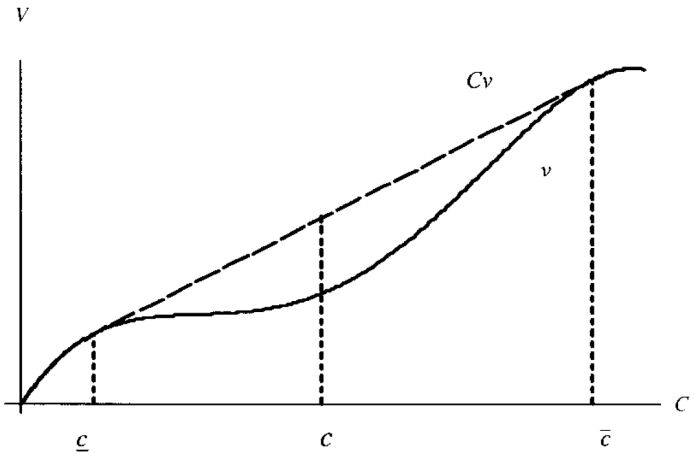

Notes: This figure is from Friedman and Savage's original 1948 paper showing a non-concave indirect utility function where an individual with income level c would be willing to accept a gamble with a 50% chance of income level $\underline{c}$ and a 50% chance of $\bar{c}$.

Figure 2: Indirect Utility with Lumpy Good and Demand for Gambles

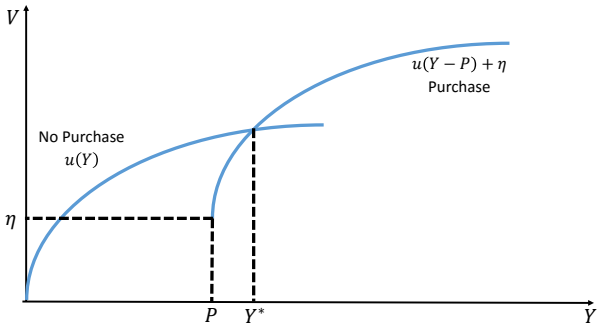

(a) Indirect utility with a lumpy good

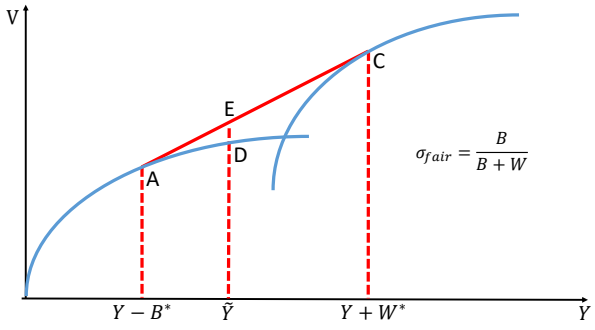

(b) Fair Gambles

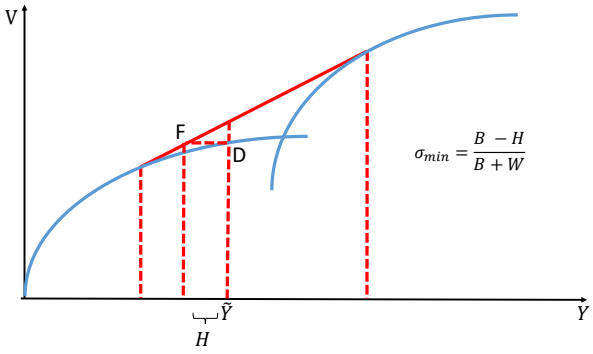

(c) Unfair Gambles

Notes: Panel (a) shows indirect utility from income with the possibility of consuming a lumpy good L. For income levels $Y > Y^*$, people will pay a price P to consume L with a utility payoff of η. Maximized utility is determined by income endowment and will be the envelope of the two pieces of the utility function. Panel (b) shows that someone with income level $\tilde{Y}$ will demand a fair gamble that risks reducing his income by B^* for a chance to go up to $\tilde{Y} + W$ with a likelihood of winning at σ. Expected utility from the gamble is at point E on the convex combination of winning and losing utilities. Panel (c) shows that there is also demand for *un*fair gambles with the same loss and win amounts but win likelihood as low as σ_{min}.

Figure 3: Effect of Increased Salience of Lumpy Good on Demand for Gambles

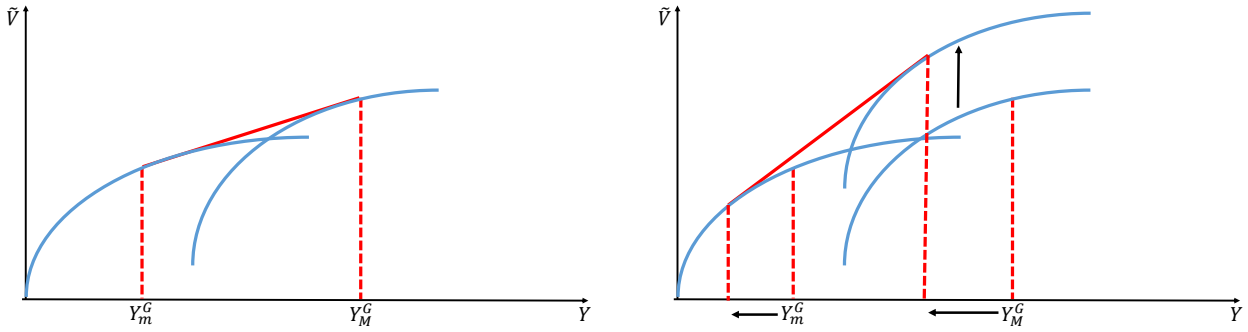

(a) Regular Salience

(b) Increased Salience

Notes: Panel (a) shows demand for gambles with normal salience of the lumpy good. Increasing the lumpy good's salience is modeled as an increase in its valuation represented by the upward shift in the payoff to the lumpy good in Panel (b). This shift in valuation of the lumpy good increases demand for betting among people who could not afford the good.

Figure 4: Effect of Reduced Saving Ability on Demand for Saving

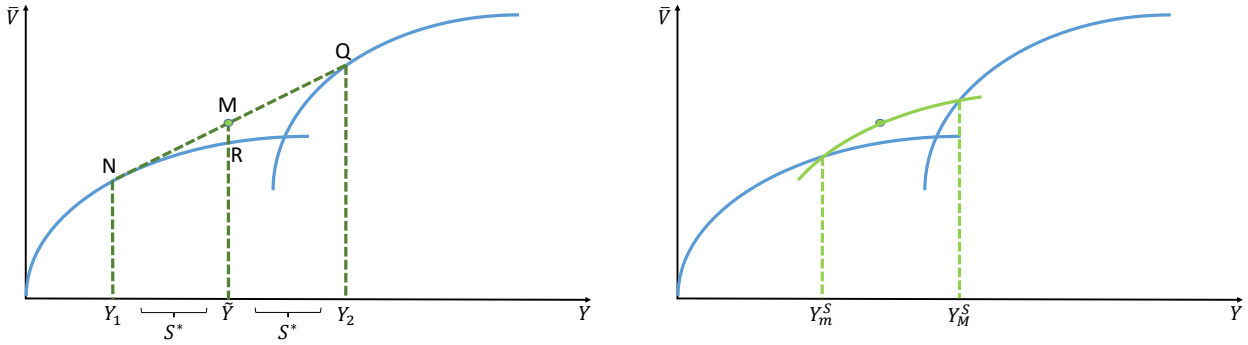

(a) Saving Decision for Y

(b) Optimal Saving Utilities

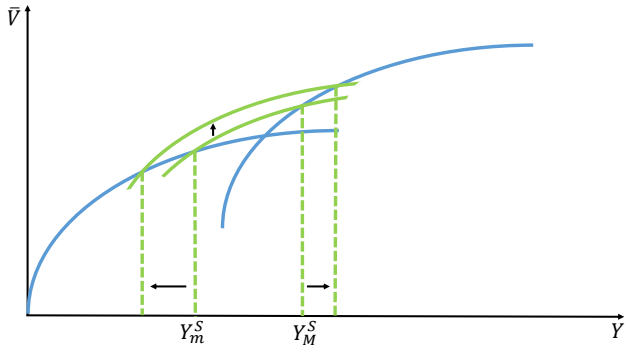

(c) Utility with improved saving ability

Notes: Panel (a) shows the optimal saving decision for an individual with income level $\tilde{Y}$, shifting S^* from T1 to T2 and increasing average expected utility from point R to M. Panel (b) shows optimal saving utilities for all income levels and defines the income range where saving is a welfare improving strategy. Panel (c) shows how this range of incomes and utility from saving increases as saving ability improves.

Figure 5: Full Study Timeline

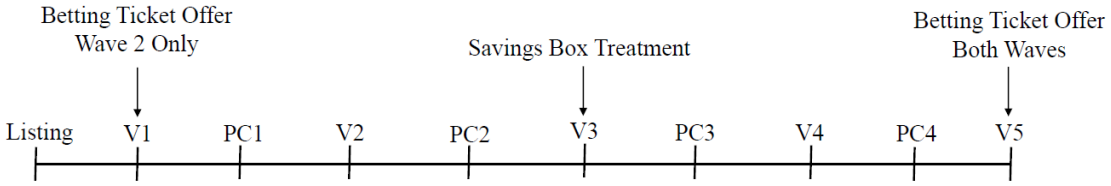

Notes: The figure above illustrates the study timeline for the 1,003 participants in the full study.

Figure 6: Biggest Win During Study Period

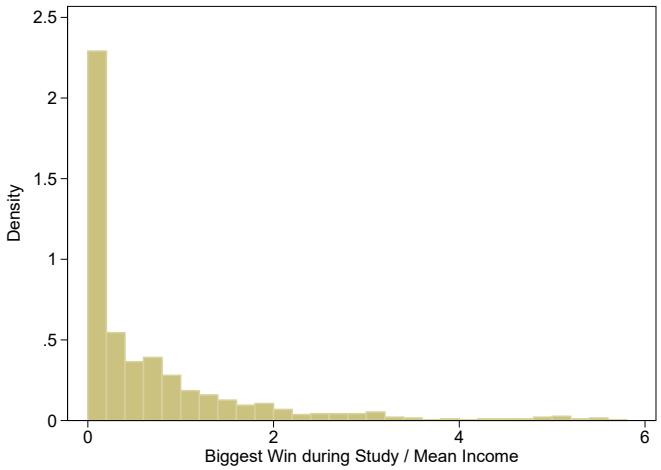

Notes: This figure shows the biggest recorded win for each respondent in the full study over the course of the nine weeks of participation. 10 people had wins bigger than six times their weekly income.

Figure 7: Effect of Winnings on Expenditure Thresholds

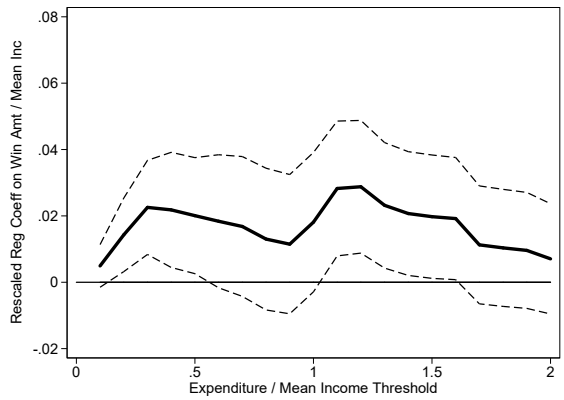

(a) All - Regression Raw Coefficients

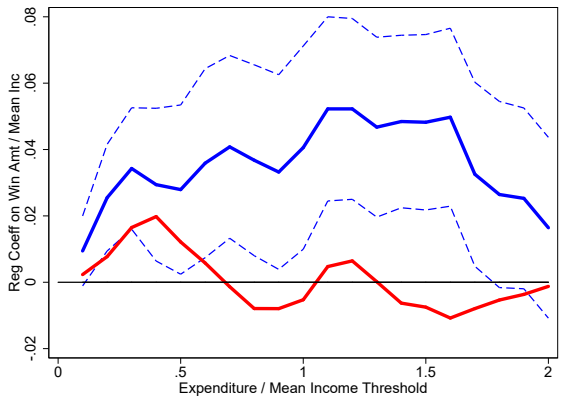

(b) Split - Regression Raw Coefficients

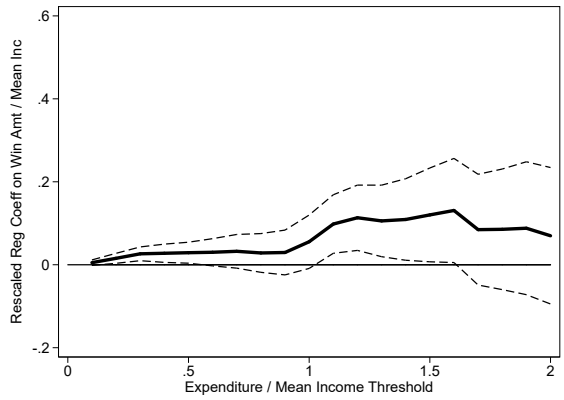

(c) All - Rescaled Regression Raw Coefficients

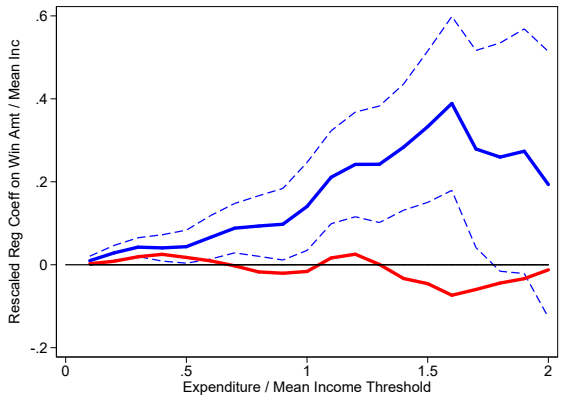

(d) Split - Rescaled Regression Raw Coefficients

Notes: Each graph shows the coefficient estimates for β_1 from a regression of $Y_{i,t} = \beta_0 + \beta_1 W_{i,t} + \sum_{b=1}^{3} BetMoments_{i,t}^{b} + \lambda X_{i,t} + \gamma_i + \delta_t + \psi_s + \epsilon_{i,t}$ where the outcome variable is making a purchase above a threshold in that time period (indicated on the x-axis). The magnitude of the estimate for β_1 is captured on the y-axis. $BetMoments_{i,t}$ are the moments and higher order terms of an individual's bets for that time period. $X_{i,t}$ are time-varying individual covariates. I include time, survey round, and individual fixed effects in all regressions. Standard errors are clustered at the individual level. Panels (a) and (c) are the estimates for all respondents together with the 95% confidence interval dotted around the estimates. Panels (b) and (d) split the sample by people with relatively low and high capacity to save. Low saving ability is the top solid line in blue in both sub-figures. Confidence intervals are only included for people with low saving capacity. Panels (a) and (b) show the raw regression coefficients, whereas Panels (c) and (d) rescale the coefficient by the mean of the outcome variable (incidence rate of an expenditure of that size).

Figure 8: Lumpy Expenditure Prime (LP) Setup

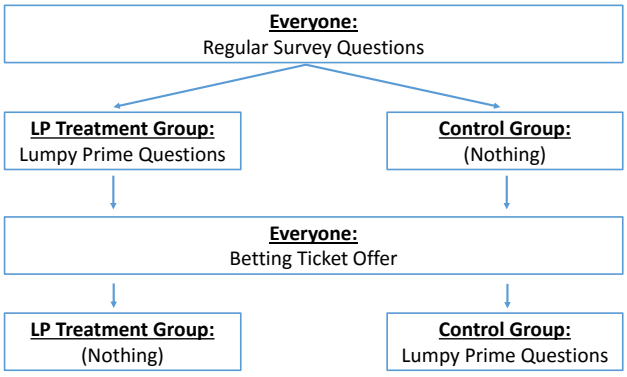

Notes: This figure shows the setup of the lumpy expenditure priming experiment. At the end of the survey, randomly selected respondents from all phases of the project (full and condensed study participants) were led through a priming dialog designed to increase the salience of a desired lumpy expenditure. People in the control group did not receive the dialog at that time. Then, a revealed preference measure of betting demand was elicited from all participants in the form of an offer of betting tickets or cash. After the offer, the participants in the control group were guided through the same dialog.

Figure 9: Saving Budgeting Exercise (BE) Setup

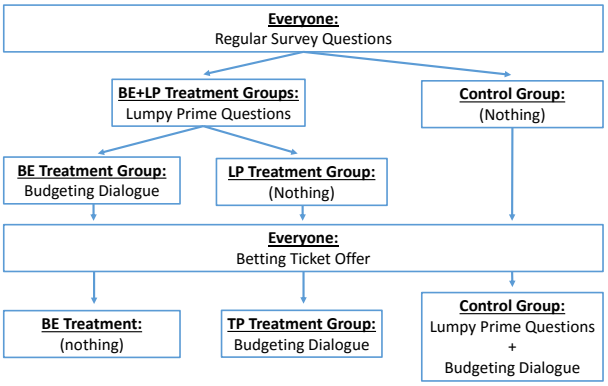

Notes: This figure shows the setup of the budgeting exercise experiment. This was only conducted with participants in the condensed study and the budgeting exercise was nested within the lumpy expenditure prime. People selected for the budgeting exercise treatment group were guided through the respondents' own estimates of their typical weekly earnings and critical recurrent expenditures in order to make a more accurate estimate of their own saving ability. People in the lumpy prime treatment group and the pure control group were guided through the questions after the betting ticket offer was used to elicit a revealed preference measure of betting demand.

Figure 10: Budgeting Exercise Saving Updates

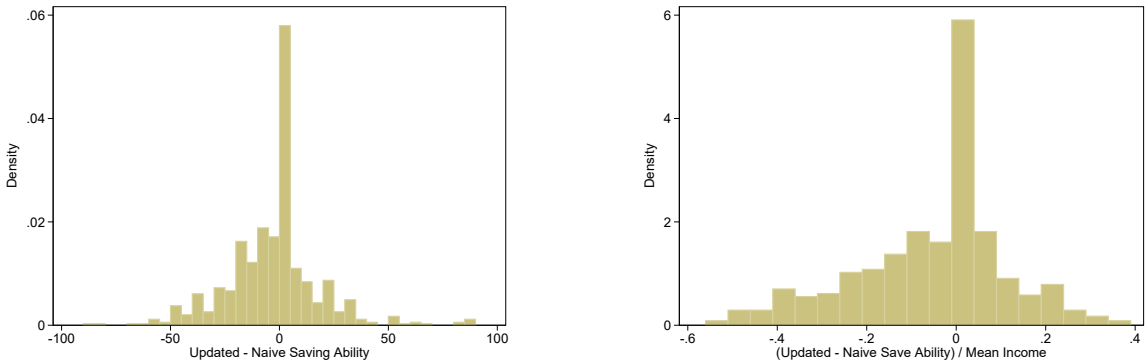

(a) Raw Saving Update

(b) Saving Update / Mean Income

Notes: This figure shows the update size resulting from the budgeting exercise calculated as the amount participants felt they could save in a typical week after the budgeting exercise, minus the amount they estimated naively at the beginning of the survey. Panel (a) is the raw update size in thousands of Ugandan Shillings (3,500UGX $\approx$ 1 USD). Panel (b) converts this update size relative to mean income.

Figure 11: Effect of Savings Ability Update on Max Ticket Demand - Lowess

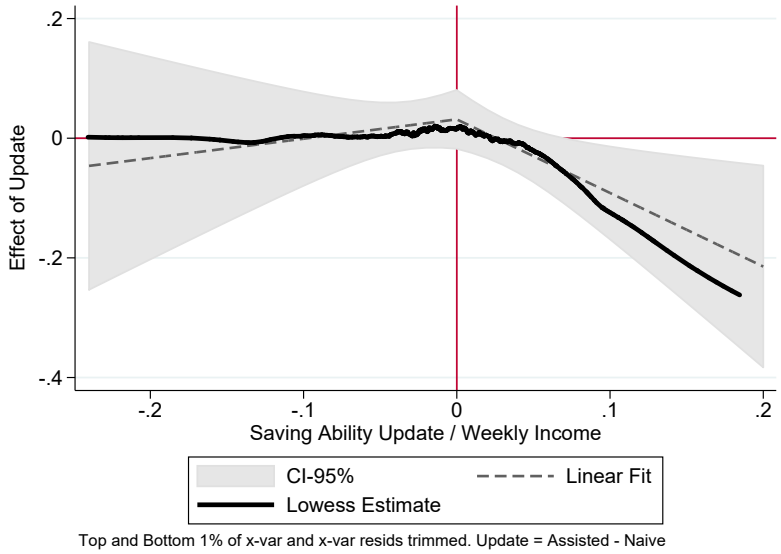

Notes: This figure shows the non-parametric estimate of the effect of the update on respondents' ability to save resulting from the budgeting exercise. The update is the difference between the newly estimated amount an individual can save minus their original naive estimate scaled relative to mean income. The median negative update was -0.15 and the median positive update was 0.1. The y-axis is the likelihood of demanding the maximum number of betting tickets offered during the revealed preference measure of betting demand following the budgeting activity for people in the treatment group.

Figure 12: Separating Saving and Betting Preference by Saving Return Thresholds and Patience

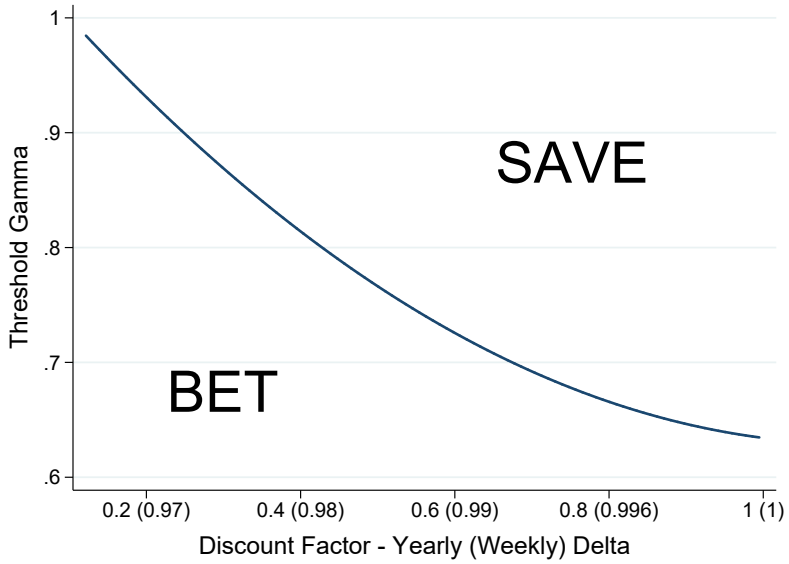

Notes: This graph shows the threshold level of saving ability needed to sustain a saving strategy for each level of patience, based on calculations from the model developed in Section 3 and extended in Section 7. People with a given level of future discounting, δ, with a return on saving above the traced locus will be willing to save in pursuit of a lumpy expenditure. People whose saving ability is below that threshold level of γ will switch to betting. People will begin switching to betting as saving ability worsens until expected losses from saving reach 35%, at which point even the most impatient people will switch to betting. Mbiti and Weil (2013) find a reasonable yearly discount factor of 0.64 in neighboring Kenya. This would imply a threshold γ of 0.716 at which people facing expected losses above 29% for money set aside to saving will switch to betting.

Tables

Table 1: Summary Statistics: Background and Household

	Full Study (N=1,003)				Condensed Study, (N=712)			
Panel (a)	Mean	p10	p50	p90	Mean	p10	p50	p90
Weekly Income (USD)	29.75	11.5	24.5	50.8	37.58	17.1	31.4	60.0
Betting Expenditures (USD)	3.02	0.7	2.0	6.5	4.17	0.9	2.9	8.6
% of Income Spent on Betting	12.73	2.6	8.6	25.2	12.59	2.4	8.3	28.0
Household Size	3.09	1.0	3.0	6.0	2.86	1.0	3.0	5.0
% Contribution of HH Finances	75.00	25.0	100.0	100.0	75.66	25.0	100.0	100.0
Weekly HH Inc. Per Cap. (USD)	15.46	4.1	10.9	30.7	16.92	0.8	10.7	40.3
Weekly HH Inc. Per Cap. Adj. (USD)	19.32	5.3	13.8	42.7	NA	NA	NA	NA
Age	27.11	20.0	27.0	35.0	26.60	20.0	26.0	35.0
Primary	0.84	-	-	-	0.84	-	-	-
Junior Secondary (O-Level)	0.46	-	-	-	0.44	-	-	-
Senior Secondary (A-Level)	0.17	-	-	-	0.19	-	-	-
Panel (b)	Mean	p10	p50	p90	Mean	p10	p50	p90
Available Liquidity (USD)	98.31	5.7	28.6	228.6	66.85	4.3	14.3	142.9
Available Liquidity/Mean Inc	3.76	0.3	1.3	8.4	1.78	0.1	0.6	4.0
Saving Ability Per Week (USD)	12.20	1.4	8.6	20.0	8.81	1.4	6.0	17.1
Saving Ability/Mean Inc	0.51	0.06	0.31	0.87	0.25	0.05	0.21	0.50
Win Target (USD)	143.23	14.4	46.5	265.7	352.60	11.4	57.1	571.4
Win Target / Mean Income	28.87	0.4	2.3	32.1	11.59	0.3	1.9	15.0
Win Target / Available Liquidity	28.10	0.2	1.9	33.3	41.42	0.3	3.0	40.0
Winning Item Cost (USD)	380.47	11.4	57.1	857.1	NA	NA	NA	NA

Notes: HH income only calculated for 97% in full and 92% in condensed study who contributed to household expenses. Condensed study did not ask about targeted item for winnings or number of children and adjusted per capita income could not be calculated.

Table 2: Lumpy Expenditures and Source of Liquidity

Panel (a): Desired lumpy expenditures by category.

Good	Business	Household	Personal
#1	Working capital-19%	Furniture-17%	Clothes-31%
#2	Improve worksite-13%	Entertainment-17%	Phone-11%
#3	Motorcycle-13%	Build/Repair-9%	Vehicle-4%
#4	Tools-12%	Appliance-5%	Entertainment-4%
#5	New venture-2%	School fees-5%	Jewelry-3%
Other	10%	20%	9%
None	33%	27%	38%
Price	$285.6	$114.3	$42.8
$\frac{Price}{Mean\ Inc}$	12.9	4.1	1.8

Panel (b): Likely source of liquidity for desired expenditure.

Source	Most Likely	Likely
Saving	85.4%	95.9%
Betting	6.6%	25.3%
Credit from Family/Friend	2.5%	14.4%
Credit from Bank/Loan Organization	4.9%	11.6%
Credit from Money Lender	0.2%	2.1%
Any Credit Source	7.6%	26.2%

Notes: Panel (a) shows responses to the question "Is there a large expenditure that you are hoping to make in the next few months?" They were asked to name something in each of the three categories. Interviewers were instructed to ensure that the item or expense named was in fact non-divisible (working capital would mean a bulk purchase) and they were additionally instructed to make sure that these expenditures were realistic and not simply something they would like to have as a dream. Panel (b) shows the follow-up question conducted during the condensed study, typically following the identification of a business expense. Panel (b) suggests that betting is considered the second most likely source of money for their desired expense after saving and was cited almost as often as all different sources of credit combined.

Table 3: Impact of Winnings on Biggest Expense Thresholds

	(1) Med	(2) Big	(3) Bigger	(4) Huge
Win Amount / Income	0.0201**	0.0181*	0.0071	0.0084
	(0.0089)	(0.0107)	(0.0085)	(0.0072)
Income	0.0429***	0.0393***	0.0114	-0.0005
	(0.0102)	(0.0098)	(0.0089)	(0.0067)
Bet Amt	-0.0256	0.0792	0.1017**	0.0648**
	(0.0418)	(0.0492)	(0.0439)	(0.0316)
Mean Y	0.6665	0.3078	0.0933	0.0275
Num Obs	4653	4653	4653	4653
Num Inds	954	954	954	954
R2	0.446	0.443	0.387	0.336

Notes: Dependent variable for each column is an indicator for having made an expenditure above a given threshold in that time period. Expenditure thresholds are Med=.5*Inc, Big=Inc, Bigger=2*Inc, Huge=4*Inc. The full regression specification is $Y_{i,t} = \beta_0 + \beta_1 W_{i,t} + \sum_{b=1}^{3} BetMoments_{i,t}^{b} + \lambda X_{i,t} + \gamma_i + \delta_t + \psi_s + \epsilon_{i,t}$. $BetMoments_{i,t}$ are the moments and higher order terms of an individual's bets for that time period. $X_{i,t}$ are time-varying individual covariates. Time, survey round, and individual fixed effects are included in all regressions. Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Treatment Assignment Balance Tables

	Saving Box											
	Panel (a) Wave 1: N=465			Panel (b) Wave 2: N=504			Panel (c) Lumpy Good Prime			Panel (d) Budgeting Exercise		
	Control	Treat	P-Val	Control	Treat	P-Val	Control	Treat	P-Val	Control	Treat	P-Val
Mean Weekly Inc (USD)	29.415	27.967	0.52	31.560	28.285	0.17	31.230	31.068	0.85	36.761	39.301	0.14
Bet Expenditures / Normal Inc	0.160	0.138	0.30	0.127	0.150	0.12	0.119	0.117	0.83	0.124	0.124	0.99
Saving Update (USD)	NA	NA	NA	NA	NA	NA	NA	NA	NA	-0.969	-0.955	0.98
Save Update / Weekly Inc	NA	NA	NA	NA	NA	NA	NA	NA	NA	-0.057	-0.052	0.74
Save Ability (USD)	12.339	12.186	0.88	7.293	7.013	0.60	8.700	7.864	0.14	8.340	7.531	0.10
Save Ability / Mean Inc	0.521	0.520	0.99	0.321	0.309	0.68	0.311	0.296	0.25	0.249	0.218	0.02**
Liquidity Available / Mean Inc	3.390	3.311	0.88	2.540	2.579	0.90	2.446	2.342	0.549	2.634	1.441	0.45
Lumpy Good Purchased	NA	NA	NA	NA	NA	NA	0.115	0.101	0.34	NA	NA	NA
Cost of Lumpy Good / Mean Inc	NA	NA	NA	NA	NA	NA	11.136	11.194	0.97	4.531	5.431	0.46
Cash Offered (UGX)	500	500	1.00	600	600	1.00	616	612	0.451	700	706	0.69
Completed Primary	0.819	0.836	0.68	0.833	0.833	1.00	0.834	0.839	0.78	0.841	0.837	0.90
Completed O-Level	0.424	0.457	0.54	0.472	0.480	0.86	0.460	0.471	0.65	0.448	0.473	0.55
Completed A-Level	0.158	0.129	0.46	0.183	0.194	0.73	0.189	0.186	0.88	0.209	0.167	0.21
Math Score	7.897	7.922	0.90	8.345	8.075	0.07*	0.978	1.029	0.54	3.498	3.488	0.92
Live Alone	0.266	0.267	0.99	0.246	0.266	0.61	0.288	0.281	0.75	0.355	0.360	0.92
Beta Discounting	0.486	0.476	0.75	0.498	0.475	0.43	0.467	0.458	0.58	0.398	0.404	0.83
Delta Discounting	0.641	0.630	0.74	0.604	0.566	0.17	0.589	0.581	0.59	0.545	0.549	0.91
Risk Aversion. 0=Low 6=High	2.100	2.190	0.66	2.159	2.107	0.74	2.300	2.270	0.74	2.655	2.803	0.44
Upside Seeker - Low	0.169	0.164	0.90	0.175	0.171	0.91	0.069	0.070	0.97	0.068	0.059	0.66
Upside Seeker - Mid	0.226	0.172	0.22	0.246	0.222	0.53	0.089	0.085	0.80	0.078	0.108	0.20
Upside Seeker - High	0.266	0.224	0.37	0.270	0.230	0.30	0.097	0.098	0.91	0.080	0.089	0.72
Endline Lumpy Prime	0.249	0.267	0.70	0.448	0.528	0.07*	0.00	1.00	NA	0.00	1.00	NA
Treatment - Saving Box	0.00	1.00	NA	0.00	1.00	NA	0.241	0.264	0.37	NA	NA	NA
Treatment - Wallet	0.496	0.517	0.69	0.306	0.306	1.00	0.230	0.208	0.34	NA	NA	NA
Treatment - Bet Info	0.499	0.500	0.98	0.310	0.310	1.00	0.205	0.231	0.26	NA	NA	NA
N	349	116		252	252		869	834		498	203	

Notes: Values marked as "NA" were either not collected for that round or, in the case of the saving box treatment in Panels (a) and (b) was a potential outcome of the treatment. Saving update was only included as part of the budgeting exercise. No alternative treatments are relevant to the budgeting exercise because these participants were exclusively in the condensed study so that they had not been eligible for treatments.

Table 5: Effect of Savings Box Treatment on Demanding Maximum Betting Tickets - Difference in Differences Estimation

	(1) FE	(2) FE	(3) LSA	(4) HSA	(5) FE	(6) IV	(7) IV-LSA	(8) IV-HSA	(9) IV
Savings Box	-0.1310**	-0.1315**	-0.2456***	-0.0112	-0.1965***	-0.2534**	-0.4990***	-0.0214	-0.3294***
	(0.0541)	(0.0545)	(0.0748)	(0.0818)	(0.0617)	(0.1064)	(0.1583)	(0.1564)	(0.1069)
Sav Box x High Sav Abil					0.1401*				0.1700
					(0.0841)				(0.1223)
Log Income		-0.0152	0.0096	-0.0451	-0.0172	-0.0151	0.0131	-0.0452	-0.0166
		(0.0235)	(0.0299)	(0.0367)	(0.0233)	(0.0241)	(0.0298)	(0.0368)	(0.0238)
Urgent Expense		0.0071	-0.0263	0.0007	0.0070	0.0167	-0.0058	0.0015	0.0171
		(0.0564)	(0.0870)	(0.0692)	(0.0553)	(0.0599)	(0.0995)	(0.0695)	(0.0597)
Cost of Needed Expense		-0.0076	0.0025	-0.5935**	-0.0070	-0.0110	-0.0053	-0.5923**	-0.0103
		(0.0147)	(0.0079)	(0.2432)	(0.0145)	(0.0151)	(0.0098)	(0.2439)	(0.0147)
Mean Dep Var	0.4648	0.4648	0.4959	0.4360	0.4648	0.4648	0.4959	0.4360	0.4648
Adjusted Control Mean	0.4988	0.4988	0.6264	0.3666	0.4988	0.4988	0.6264	0.3666	0.4988
Num Obs	994	994	490	500	994	994	490	500	994
Num Indivs	497	497	245	250	497	497	245	250	497
R2	0.6392	0.6395	0.6817	0.6127	0.6419	0.6296	0.6591	0.6119	0.6331
Adj R2	0.2657	0.2619	0.3348	0.1914	0.2652	0.2416	0.2876	0.1897	0.2472

Notes: Dependent variable is an indicator for demanding the maximum number of tickets in the betting ticket offer. Results from regression of $B_{i,t} = \beta_0 + \beta_1 SaveBox_{i,t} + \lambda X_{i,t} + \gamma_i + \delta_t + \epsilon_{i,t}$. $SaveBox_{i,t}$ is an indicator for an individual having received the saving box in time t. Individual fixed effects, time fixed effects, amount of cash offered, and a set of background controls are included, as well as time-varying covariates, $X_{i,t}$. LSA= Low saving ability, HSA= High saving ability. Robust standard errors are used to adjust for heteroskedasticity in the error term. IV estimation in columns (6)-(8) use randomized assignment to treatment as an instrument for respondents reporting that they have used a saving box over the preceding four weeks. Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Effect of Lumpy Prime on Demand of Maximum Tickets Offered

	(1) All	(2) All	(3) All	(4) All	(5) LSA	(6) HSA	(7) All
Lumpy Good Prime	0.073***	0.074***	0.073***	0.072***	0.122***	0.021	0.020
	(0.024)	(0.024)	(0.024)	(0.024)	(0.034)	(0.034)	(0.034)
Prime x Low Save Ability							0.104**
							(0.048)
Low Saving Ability							0.006
							(0.039)
Mean Week Bet			0.178*	0.156	0.182	0.157	0.153
			(0.101)	(0.099)	(0.142)	(0.152)	(0.100)
Liquidity Available			0.005	0.005	0.011**	-0.001	0.005
			(0.003)	(0.003)	(0.005)	(0.005)	(0.003)
Save Ability / Mean Inc			-0.061	-0.072	0.190	0.050	0.006
			(0.052)	(0.053)	(0.190)	(0.080)	(0.068)
Mean Week Income		0.002***	0.002***	0.002***	0.002	0.004***	0.002***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Mean Dep Var	0.4527	0.4527	0.4527	0.4527	0.4763	0.4296	0.4527
Mean Y-Control	0.4177	0.4177	0.4177	0.4177	0.4198	0.4157	0.4177
Full Set of Covariates	No	No	No	Yes	Yes	Yes	Yes
Num Obs	1703	1703	1703	1703	844	859	1703
R2	0.0316	0.0358	0.0410	0.0523	0.0875	0.0709	0.0568
Adj R2	0.0183	0.0220	0.0244	0.0294	0.0432	0.0255	0.0329

Notes: Dependent variable is an indicator for demanding the maximum number of tickets in the betting ticket offer. Results from regression of $B_i = \beta_0 + \beta_1 LumpyPrime_i + \lambda X_i + \epsilon_i$. $LumpyPrime$ is an indicator for going through the lumpy prime dialog prior to the ticket offer. Individual covariates includes background education and preference variables as well as controls for other treatments during the study and the amount of cash offered instead of tickets. LSA= Low saving ability, HSA= High saving ability. Robust standard errors are used to adjust for heteroskedasticity in the error term. Columns (1)-(4) show stability of the estimated treatment effect regardless of specification. Columns (5)-(7) show significant heterogeneity of response between people with low and high saving ability. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Effect of Budgeting Exercise on Demanding Maximum Betting Tickets

	(1) UGX	(2) UGX	(3) Prop	(4) Both	(5) UGX	(6) Prop
Lumpy Good Prime	0.086**	0.085**	0.084**	0.083*	0.084**	0.084**
	(0.042)	(0.042)	(0.042)	(0.043)	(0.042)	(0.042)
Budgeting Exercise (BE)	-0.025	-0.042	-0.056	0.073	-0.047	-0.026
	(0.044)	(0.044)	(0.046)	(0.081)	(0.051)	(0.056)
BE x Update		-0.004***	-0.538**			
		(0.001)	(0.209)			
Update		0.001*	0.124			
		(0.001)	(0.086)			
BE x (Update > 0)				-0.270***		
				(0.104)		
BE x (Update < 0)				-0.044		
				(0.098)		
Update > 0				0.078		
				(0.059)		
Update < 0				0.004		
				(0.051)		
BE x Positive Update Amount					-0.004*	-0.997**
					(0.002)	(0.472)
BE x Negative Update Amount					0.004*	0.365
					(0.002)	(0.289)
Positive Update Amount					0.000	0.148
					(0.001)	(0.209)
Negative Update Amount					-0.001*	-0.118
					(0.001)	(0.103)
N	706	706	706	706	706	706
Mean Dep Var	0.4164	0.4164	0.4164	0.4164	0.4164	0.4164
Full Set of Covariates	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.1373	0.1452	0.1447	0.1472	0.1454	0.1458
Adj R2	0.1082	0.1137	0.1132	0.1133	0.1114	0.1118

Prop columns scaled to respondent income, UGX in 1,000s. Update = Assisted - Naive Estimate
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Dependent variable is an indicator for demanding the maximum number of tickets in the betting ticket offer. Results from regression of $B_i = \beta_0 + \beta_1 LumpyPrime_i + \beta_2 Budget_i + \beta_3 (Budget \times Update)_i + \beta_4 Update_i + \lambda X_i + \epsilon_i$. $LumpyPrime_i$ is an indicator for doing the lumpy prime dialog before the betting ticket offer. $Budget_i$ is an indicator for doing the budgeting exercise before the betting ticket offer. $Update_i$ is the assisted estimate of the amount that an individual can save from the budgeting exercise minus the naive estimate. UGX columns use the raw measure of the update in thousands of shillings. Prop columns rescale this update relative to an individual's mean income. Individual covariates include background education and preference variables as well as controls for other treatments during the study and the amount of cash offered instead of tickets. Robust standard errors are used to adjust for heteroskedasticity in the error term. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Estimating Return on Cash Saving, γ

γ	Source	Estimate
γ_1	Inflation	0.8998 - 0.9844
γ_2	Temptation	0.784 - 0.9361
γ_3	Social Pressure	0.92 - 1
γ_4	Theft/Loss	0.9885 - 1
γ_5	Transaction Costs	1

Notes: γ is an individual's return on saving or the proportion of each dollar set aside for saving that he expects to be converted into expenditure on a desired purchase or expense. In estimating reasonable levels of γ in the population, I break it down into sub-components and take the product. This is approximated for cash savings, although different components would likely shift if other saving instruments or strategies were used. For example, bank accounts could lower losses from social pressure, temptation, and risk of theft, but impose considerable transaction costs in the form of fees and effort.

- γ_1 is based on the range of inflation rates in Uganda over the previous five years, 2011-2016.
- γ_2 is estimated from the consumption modules in the survey, categorizing certain expenditures as temptation goods (alcohol, video hall tickets, betting) and assuming that people regret between 25-50% of these expenses.
- γ_3 captures expenditures that are made out of obligation or as a result of inter- or intra- household pressure. This estimate is from a recent study by Jakiela and Ozier (2015).
- γ_4 was based on survey responses estimating the frequency of theft.
- γ_5 is assumed to be one for cash savings, but would be lower for most other saving technologies that require usage fees, coordination with others, or effort for either deposits or withdrawals.

The overall range for γ is estimated between 0.6415 and 0.9215. This range of reasonable γs falls considerably below the range of threshold γ^*s illustrated in Figure 12, suggesting that there may be a considerable portion of the population for whom betting is a rationally preferred straetgy to saving, given their levels of patience and return on saving.