Title
Three Applied Economic Studies on Issues Facing African Youth

Author
Herskowitz, Sylvan R

Publication Date
2017

Peer reviewed|Thesis/dissertation

Three Applied Economic Studies on Issues Facing African Youth

by

Sylvan René Herskowitz

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jeremy Magruder, Chair
Professor Aprajit Mahajan
Professor Edward Miguel
Professor Elisabeth Sadoulet

Spring 2017

Three Applied Economic Studies on Issues Facing African Youth

Abstract

Three Applied Economic Studies on Issues Facing African Youth

by

Sylvan René Herskowitz

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Jeremy Magruder, Chair

This manuscript is comprised of three independent essays in applied microeconomics. Each essay examines a distinct challenge or trend impacting African youth. First, I look at the sports betting industry which has exploded across many African countries over the past decade, targeting young men as its primary consumers. In Kampala, Uganda, the site of my study, nearly one in three young men age 18 to 40 participate in sports betting, spending 8-12% of their weekly income on these activities. The second chapter presents results from an impact evaluation of a sports and youth group program in Monrovia, Liberia. The use of sports and youth programs has become extremely popular for many multilateral and international non-profit organizations' efforts to engage and improve the conditions of marginalized or at-risk youth in developing countries. The final chapter looks at the impact of rainfall shocks on marriage decisions among youth from Burkina Faso, a setting where marriages are determined by households and influenced by financial constraints and economic shocks.

In the first essay, "Betting, Saving, and Lumpy Expenditures: Sports Betting in Uganda", I present evidence from my field study on one of the causes of high betting demand among young Ugandan men. In my paper, I show that financial constraints push many men towards betting in the hopes of payouts desired for unmet liquidity needs. I use a range of experimental and quasi-experimental methods and present four pieces of evidence in support of this claim. Winnings from betting increase the size and likelihood of making large, lumpy purchases, with strongest effects among those with limited ability to save. I then use a randomized intervention to improve an alternative strategy of liquidity generation to betting by distributing a basic commitment savings technology and find that it lowers recipients' demand for betting. Next, I use an experimental prime or nudge, increasing the salience of a desired large expenditure, and find that it increases demand for betting tickets, particularly among those with low ability to save. And finally, I use a randomized budgeting exercise and find that it lowers demand for betting among those who learn that their ability to save is better than previously believed. For people whose ability to save and borrow is inhibited or

costly, betting provides an enticing alternative way to generate meaningful sums of liquidity despite imposing considerable losses.

My second essay, "Do Sports Change Lives? Evidence from a Randomized Control Trial", is the product of a joint project with Lori Beaman, Niall Keleher, and Jeremy Magruder. In this project we assess the impact of a popular form of international development program that uses sports both as a direct intervention and also as a point of entry to facilitate engagement with vulnerable youth. The stated ambitions of sports for development programs are typically both admirable and lofty, but we find only limited evidence that the impacts of these programs match their promises. We find evidence of some modest impacts on psychosocial behaviors for young men with moderate improvements in measures of self-esteem and aggressive behavior. We also see an increase in labor force participation for both men and women, although earnings among those working remain unchanged. We also explore whether the research structure required by randomization may have hindered program efficacy. However, it appears that program effects were likely stronger among late registrants who are more likely to have been excluded in the absence of the study. And finally, permitting endogenously formed sports groups may have increased pre-existing social network connections on sports teams. We find that presence of friends does impact program attendance but does not significantly impact program outcomes.

In my third essay, "Marriage Markets and Rainfall Shocks: Evidence from Burkina Faso", I analyze the impact of rainfall shocks on marriage outcomes for young men and women in Burkina Faso. In particular, I use quasi-random variation in rainfall to provide evidence that low rainfall from two years ago causes an increase in the likelihood of marriage for young women by 15%. These effects are strongest among women aged 13-16. In addition, I present evidence that closed marriage markets, where womens' partners are most likely to come from the same local geographic area, respond less to rainfall shocks than areas where marriage markets are likely to be more geographically integrated. This finding is consistent with basic trade theory on market structure and response to simultaneous shifts of supply and demand in opposite directions.

Contents

Dedication

To the hundreds of people who I have leaned on, laughed with, and learned from over the last six years. And to the thousands more I hope are still to come.

In the hope that my work can make a positive contribution in the world's struggle against poverty.

Acknowledgments

First and foremost, I would like to thank the chairs of my dissertation and orals committees. Thank you to Jeremy Magruder for teaching me to trust my instincts and intuition, but also for teaching me to treat them with all the skepticism that untested theories deserve. And thank you to Betty Sadoulet for tirelessly pushing me to make my ideas more refined and more precise. Both of you continuously challenged and encouraged me, holding me accountable to high standards of rigor and quality that I hope to carry with me for the rest of my career. I am deeply grateful to both of you for your patience and support throughout this program.

I would also like to thank the other members of my dissertation committee, Aprajit Mahajan and Edward Miguel. Both of you gave me much needed positivity, encouragement, and feedback through all phases of my dissertation: brainstorming, technical advice, refining, and framing. I also thank the other members of my orals committee, David Levine and Michael Anderson, for your early contributions at the formative stages of my field work.

I would like to thank Alain de Janvry and Steven Buck for the opportunity to teach with them. The classes you led were excellent and I'm grateful to you for empowering me as a Graduate Student Instructor to help teach them. Thank you also to Jeff Perloff, Leo Simon, Sofia Vias-Boas, Peter Berck, Larry Karp, Ethan Ligon, and Brian Wright whose courses, office hours, informal hallway chats, and passing encouragement were central to my studies.

Thank you to Diana and Carmen for always being available for a quick snippet of passing banter, a shared laugh, or some desperate last-minute help filling out an overdue reimbursement request. I can't count the number of times I would have been unintentionally dis-enrolled without your help over the last six years.

With all due respect to my professors, I may have learned even more from my classmates: Aluma, Becca, Daniel, Fiona, Jordan, Josh, Kaushik, Ken, Lauren, Marieke, Marion, Niall, Patrick, Sandile, Seth, and Tarek. I thank you all for your support and help as we battled through (and somehow survived) our course work. Of course, your contributions go far beyond the academic. I'm also immensely grateful for your friendship and support. It would have been impossible without all of you.

Thank you to all of the research partners and implementers I have worked with in both Uganda and Liberia: Amon Natukwatsa, Jackline Namubiru, Stephanie de Mel, and Marwolo T. Wordplee among so many others. It was a joy working with all of you.

And finally, a huge debt of gratitude to my non-Berkeley friends and family as well. To keep things simple, here I'll forego a comprehensive list and just thank Laurel, Dad, and Mom. Your unwavering optimism may have sometimes irked me, but I'm not sure I would have made it this far without it. Thank you all for all of it.

Chapter 1

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

Chapter 1:

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

Sylvan Herskowitz*

November 2016

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.

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. 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

[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)

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

[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

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, 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.

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

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

[13] European Gaming and Betting Association (nodate)

[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.

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

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 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

[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.

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 Ariyabud-

[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.

dhiphongs (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 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:

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

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 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 $Y < \tilde{Y} - B^*$

[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.

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 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.

[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.

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^*:

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

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

[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.

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," "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.

[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.

[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.

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 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.

[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.

[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.

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. 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,

[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.

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

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 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 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.

[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.

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 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]

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

[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.

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 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

[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 absence of the intervention, both respond similarly to the treatment. When estimating 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 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

[39]Table A.10 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.11. 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.

results are shown in Appendix Table A.9.[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 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

[41] Appendix Table A.13 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.13 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.12 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.14). 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.15, 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.

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 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

[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.

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.

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]

[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.16.

[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

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 discouraging 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]

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.19.

[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

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. 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

positive, negative, or zero update. These robustness checks are shown in Appendix Table A.18.

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 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

[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.

[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.

bettors themselves or may be close enough that direct enjoyment of betting pushes them over the edge.

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 (n.d.). *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

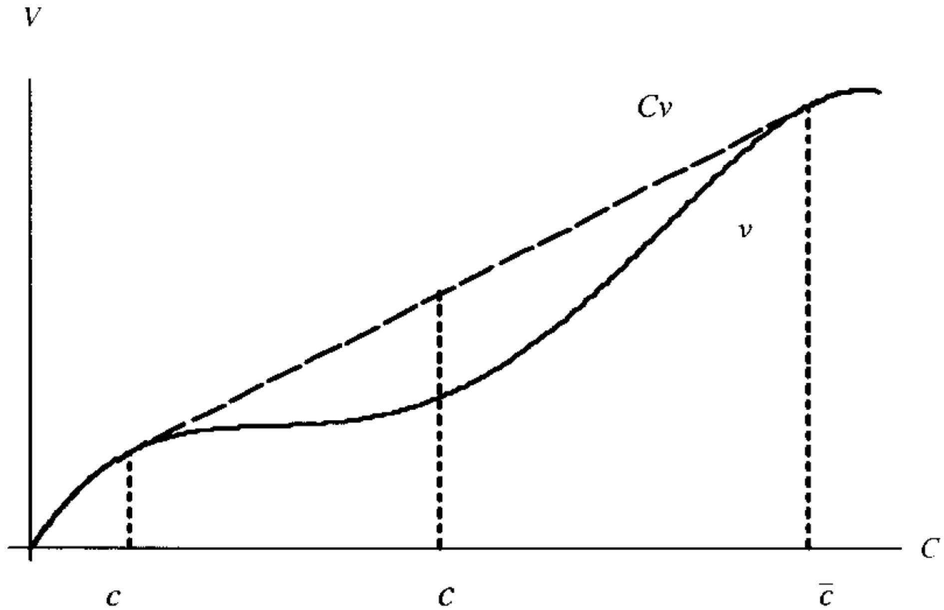

Notes: This figure is from Friedman and Savage's original 1948 paper showing a non-concave indirect utility function where an individual with income level c would be willing to accept a gamble with a 50% chance of income level $\underline{c}$ and a 50% chance of $\bar{c}$.

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

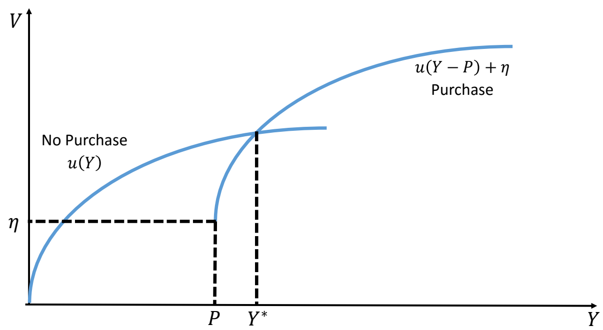

(a) Indirect utility with a lumpy good

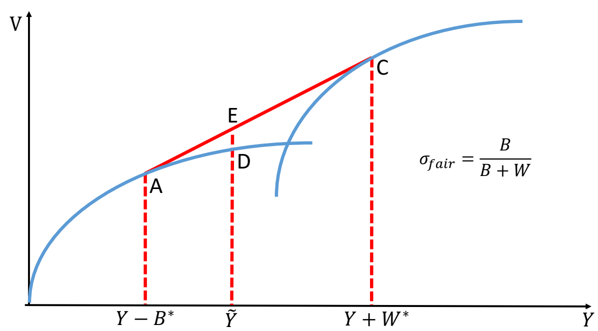

(b) Fair Gambles

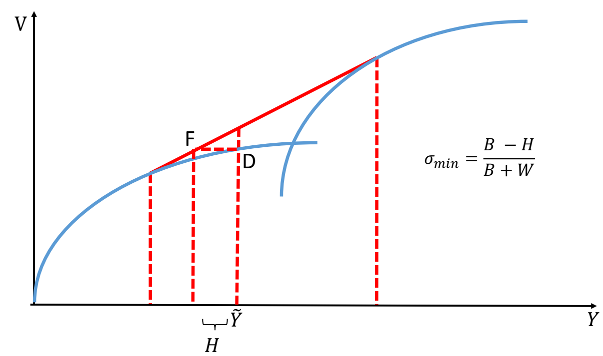

(c) Unfair Gambles

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

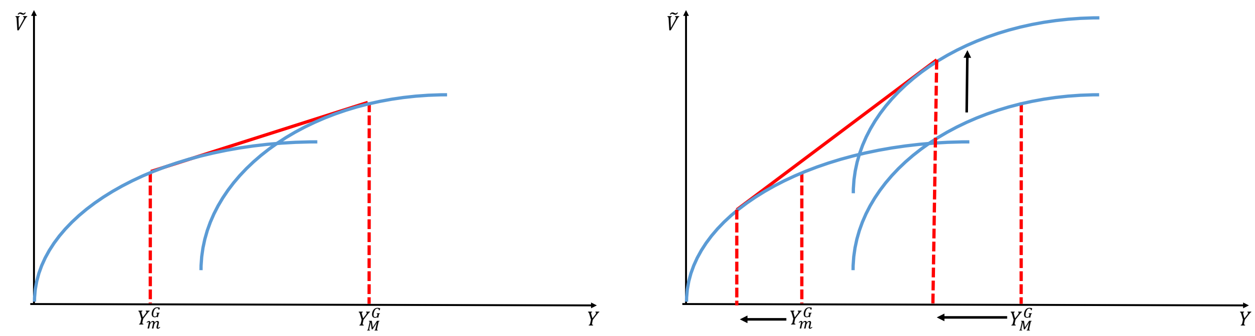

(a) Regular Salience

(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

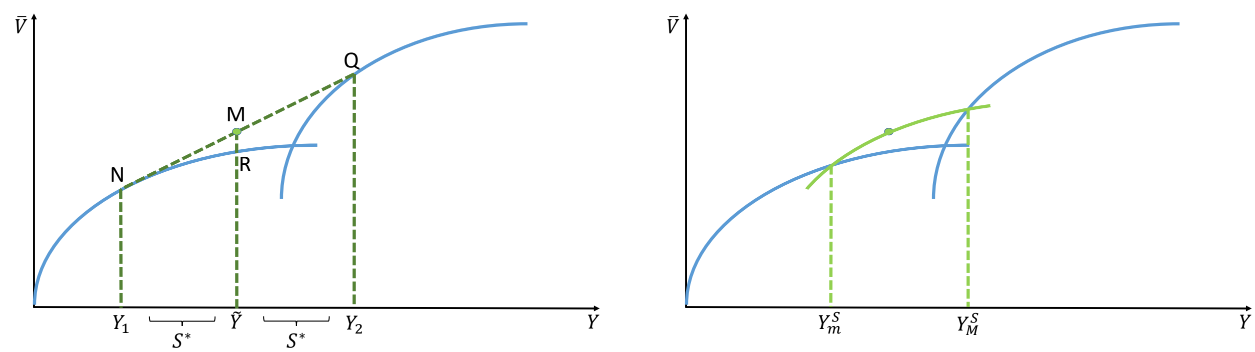

(a) Saving Decision for Y

(b) Optimal Saving Utilities

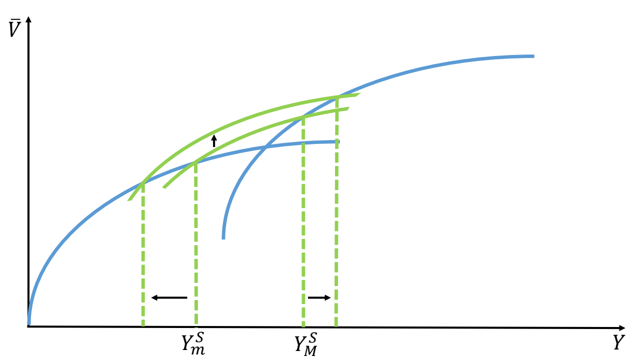

(c) Utility with improved saving ability

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

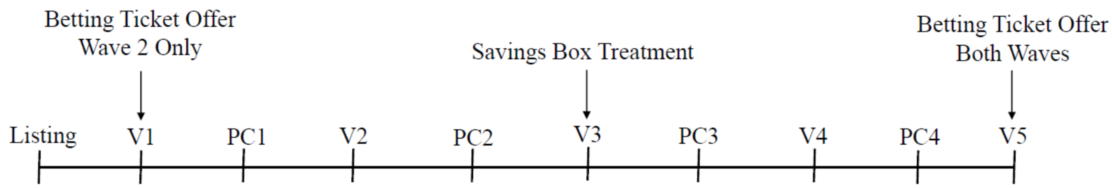

Notes: The figure above illustrates the study timeline for the 1,003 participants in the full study.

Figure 6: Biggest Win During Study Period

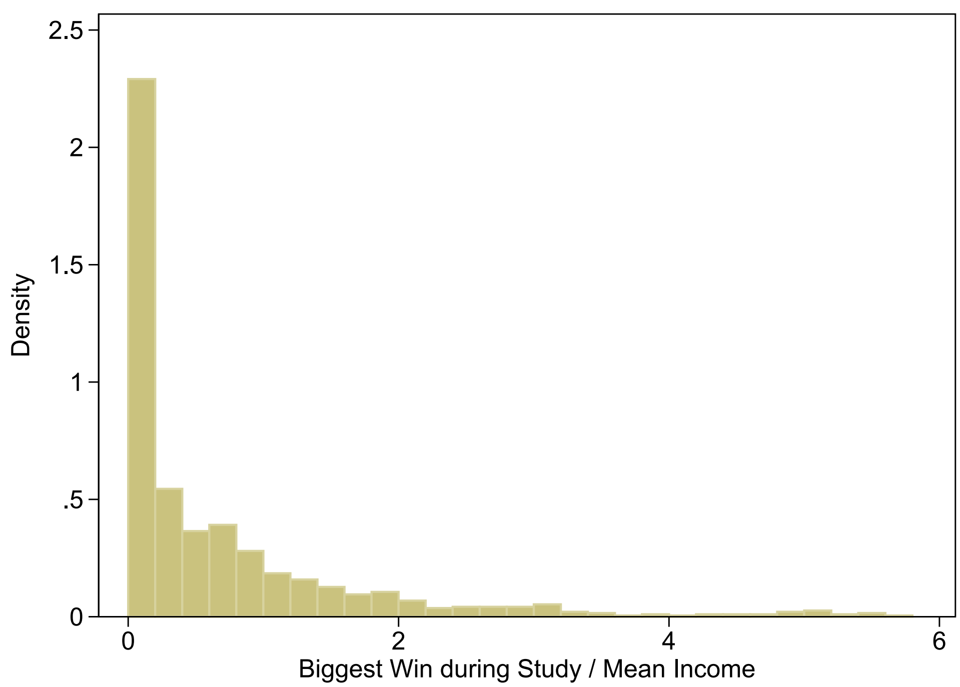

Notes: This figure shows the biggest recorded win for each respondent in the full study over the course of the nine weeks of participation. 10 people had wins bigger than six times their weekly income.

Figure 7: Effect of Winnings on Expenditure Thresholds

(a) All - Regression Raw Coefficients

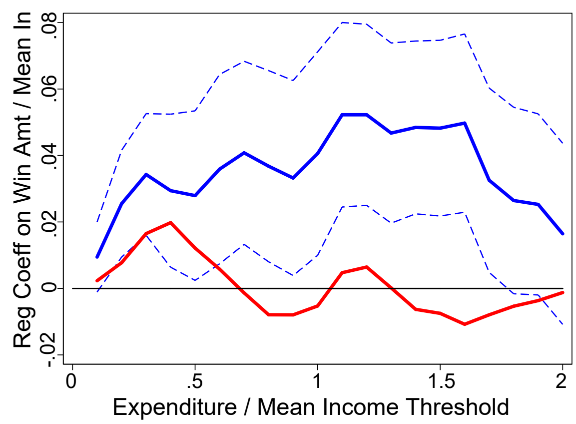

(b) Split - Regression Raw Coefficients

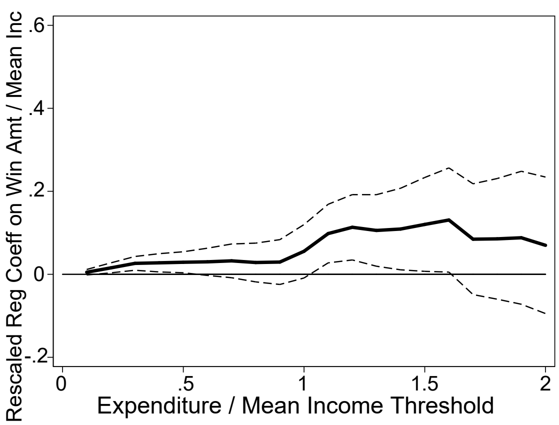

(c) All - Rescaled Regression Raw Coefficients

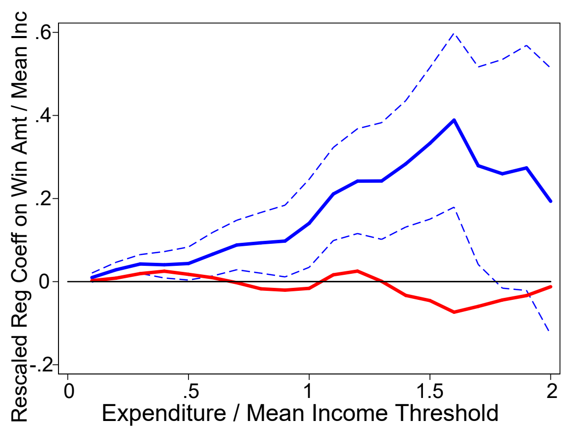

(d) Split - Rescaled Regression Raw Coefficients

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

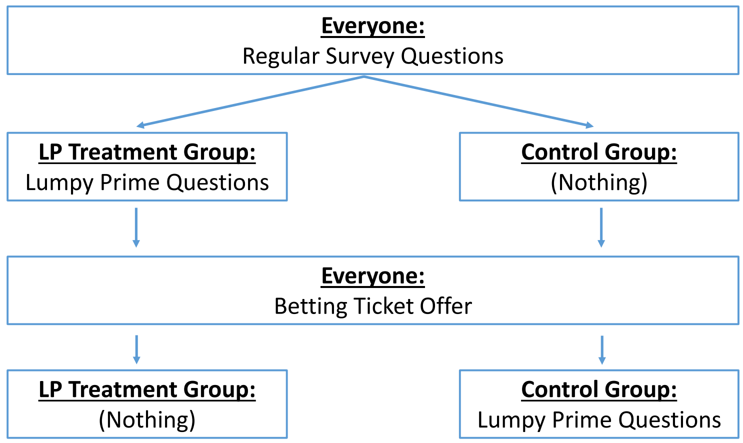

Notes: This figure shows the setup of the lumpy expenditure priming experiment. At the end of the survey, randomly selected respondents from all phases of the project (full and condensed study participants) were led through a priming dialog designed to increase the salience of a desired lumpy expenditure. 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

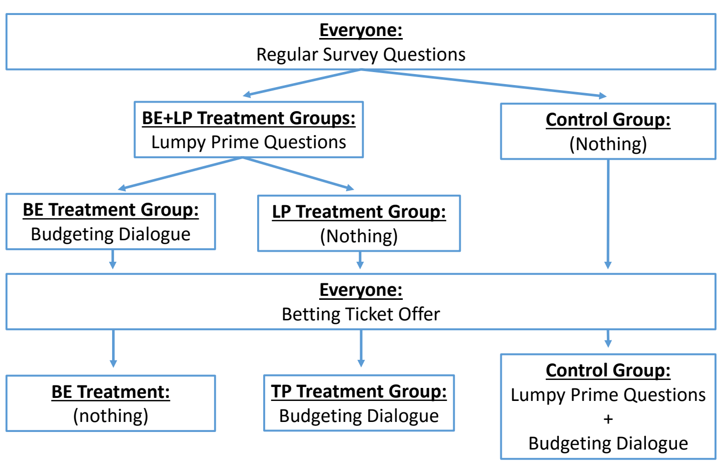

Notes: This figure shows the setup of the budgeting exercise experiment. This was only conducted with participants in the condensed study and the budgeting exercise was nested within the lumpy expenditure prime. People selected for the budgeting exercise treatment group were guided through the respondents' 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

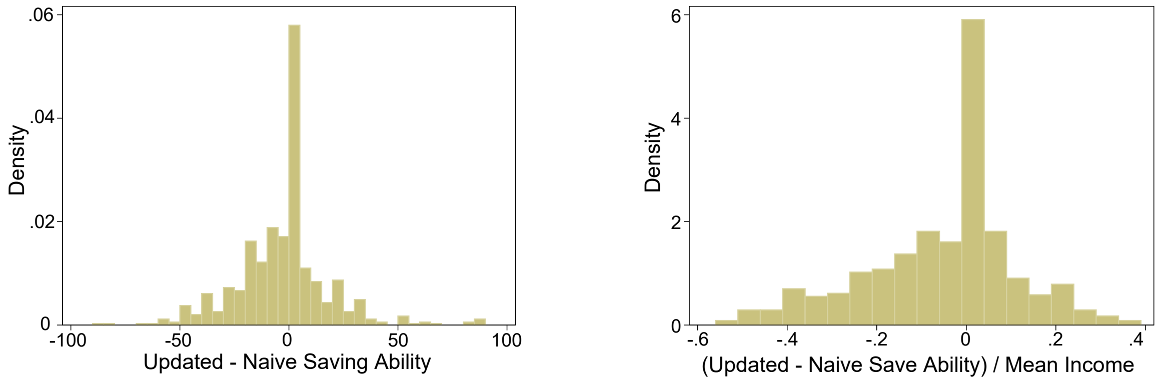

(a) Raw Saving Update (b) Saving Update / Mean Income

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

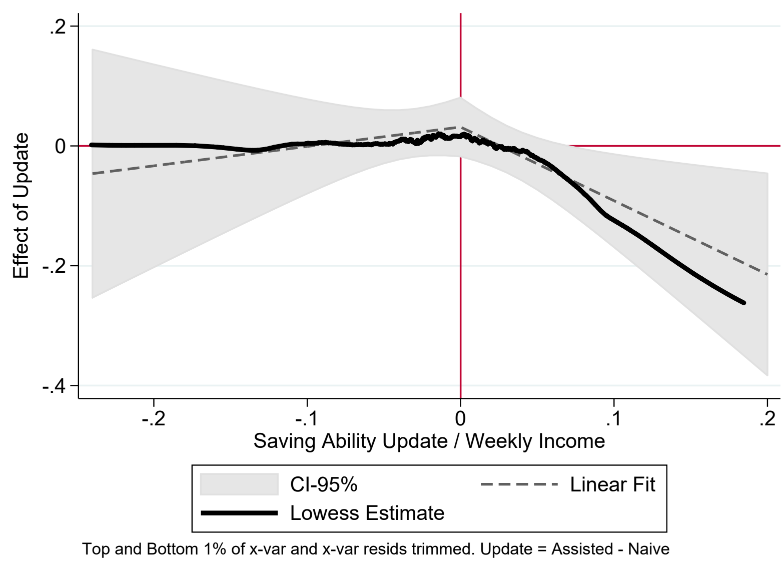

Notes: This figure shows the non-parametric estimate of the effect of the update on respondents' ability to save resulting from the budgeting exercise. The update is the difference between the newly estimated amount an individual can save minus their original naive estimate scaled relative to mean income. 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

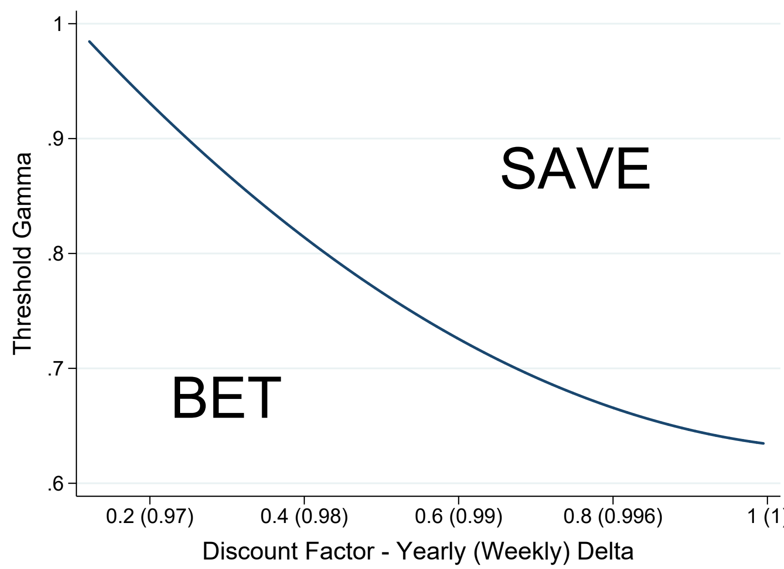

Notes: This graph shows the threshold level of saving ability needed to sustain a saving strategy for each level of patience, based on calculations from the model developed in Section 3 and extended in Section 7. People with a given level of future discounting, δ, with a return on saving above the traced 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)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	FE	FE	LSA	HSA	FE	IV	IV-LSA	IV-HSA	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.

Chapter 2

Can Sports Change Lives? Evidence from a Randomized Control Trial in Monrovia, Liberia

Chapter 2:

Do Sports Transform Lives? Evidence from a Randomized Control Trial in Liberia

Lori Beaman+, Sylvan Herskowitz‡, Niall Keleher*‡, and Jeremy Magruder‡

+Northwestern University
‡UC-Berkeley

April 2017

Abstract

Sports and life skills training are widely utilized interventions among international organizations aiming to engage with and improve the circumstances of vulnerable youth in developing countries. Despite little existing evidence, these programs are motivated with claims of psychosocial benefits and financial empowerment. Using a randomized control trial in partnership with an international non-governmental organization, we assess the impact of a sports and life skills program for vulnerable youth in Monrovia, Liberia. We find limited evidence that the program has a meaningful impact on either psychosocial behaviors or a range of financial outcomes with only modest improvements in self-esteem, aggressive behaviors, and labor force participation. We also test whether the requirements of randomized selection and group assignment of participants limited the impact of the program. To the contrary, we find that the presence of friends does not affect outcomes and that late registrants who may have otherwise been excluded, benefit relatively more than those who enrolled early.

*Send correspondence to: nkeleher@berkeley.edu. We thank Stephanie De Mel for her research assistance and T. Wordplee Marwolo, Joseph Kamara, Dackermue Dolo, Abel Welwean, and other field research staff of Innovations for Poverty Action for their hard work and dedication in carrying out data collection. We are grateful to Innovations for Poverty Action for assistance running the study. We thank the study participants for generously giving their time. This research was funded by the International Initiative for Impact Evaluation [grant number OW4/1094], the Swedish International Development Cooperation Agency [award number 90954-S-001], the International Growth Centre [grant number 1-VCH-VSLE-VXXXX-51300], and Mercy Corps.

"Sport has the power to change the world. It has the power to inspire. It has the power to unite in a way that little else does. It speaks to youth in a language they understand. Sport can create hope where once there was only despair." - Nelson Mandela

1 Introduction

In post-conflict contexts like Liberia, growing populations of unemployed youth represent both an opportunity to develop and enhance the economy as well as a potential source of instability. Conflicts are often fueled by youth facing limited job opportunities.[1] Youth with few job prospects are particularly vulnerable to influence by external forces—violence, drugs, crime, and political elites. Sports-based youth development programs are increasingly viewed both as an important direct intervention and also as an effective vehicle and entry point for complementary programs targeting at-risk youth. Implementers and proponents claim that these programs can increase social capital among participants and improve pro-social behaviors. Despite high levels of investment in these types of programs in recent years, there is almost no rigorous evidence linking sports and youth group participation to the economic and psychosocial outcomes that motivate them. This study uses a randomized control trial to test the impacts of a youth sports group on its beneficiaries.

The form and aims of both sports for development and life skills programs vary greatly. Some programs' goals are modest, such as simply distributing free soccer balls to impoverished communities.[2] Others are more ambitious, with formally organized and coached sports teams linked to highly developed life skills curricula.[3] Without a centralized clearinghouse of sports for development type programs, it is hard to estimate the full scope and resources designated to these activities. However, an important role of sports in the development programing of many large multilateral, bilateral, and non-profit organizations is widely advertised alongside lofty expectations of the claimed impacts of these programs.

The United Nations' designated Office of Sport for Development and Peace asserts that sport can positively impact a wide range of development related outcomes including: gender equality, social integration, development of social capital, peace building and conflict prevention and resolution, trauma relief, and economic development.[4] Similar beliefs are shared by many other international organizations with extensive sport related programming including multilateral organiations such as UNESCO, UNICEF, and the Commonwealth of Nations as well as bilateral aid organizations such as USAID, GIZ, the Swedish International Development Agency, and Australia Aid. An even greater number of international NGOs have focused their entire programming around sports under a belief that sports are an effective tool for achieving important development related out-

[1] See Herbst (2000). And see Collier and Hoeffler (1998) for related economic theory on limited opportunity costs as a driver of conflict.

[2] Example: Ball to All `http://www.balltoall.org/`

[3] Examples: Peace Players International, `https://www.peaceplayersintl.org/`, and Grassroots Soccer, `http://www.grassrootsoccer.org/`

[4] `https://www.un.org/sport/content/why-sport/overview`

comes.

This study uses a randomized control trial to examine the impact of one such program conducted by a large international NGO, Mercy Corps. In its extensive experience working in international development, Mercy Corps has developed a "Sports for Change" method now used in over 25 countries, establishing youth groups and engaging them with opportunities to play on competitive soccer and handball teams while implementing a complementary life skills curriculum. Their program in Monrovia, Liberia, followed a similar format, engaging 1,200 participants across the city and organizing them into 30 youth sports clubs.

We find limited evidence of economically meaningful impacts on psychosocial outcomes. The modest benefits we do see are limited to men who exhibit a reduction in self-reported aggressive behaviors by 0.14 standard deviations as well as a statistically significant 0.15 standard deviation improvement in self-esteem among male participants, though this latter result is not robust to different estimation strategies. For women, we do not find evidence of any significant impacts on psychosocial outcomes. The evidence for economic benefits is even more limited. Cross-sectional estimation shows a statistically significant increase of 7.8% in labor force participation among study participants of both genders. However, this magnitude falls and loses significance in the panel estimation of the same outcome. Other labor market outcomes have positive estimated coefficients with potentially economically meaningful magnitudes but are not statistically distinguishable from zero.

Finally, the implementing partner expressed concern that the structure and imposition of a randomized control trial research design may have hindered benefits of the program. We explore this in two ways. First, without randomization to treatment and control groups, we expect sports teams to be endogenously formed with a greater likelihood of friends being placed in the same group. We find that the presence of a friend on one's sports team increases the likelihood of participating in the program by 7.1 percentage points, a reduction in non-participation of nearly 30%. However, we do not find any impacts on financial or psychosocial outcomes resulting from the presence of friends on randomly assigned groups. Lastly, the study structure imposed a need for Mercy Corps to recruit an excess of eligible participants in order to randomly assign eligible beneficiaries to both treatment and control groups. This expansion of the pool of potential participants and subsequent randomization of actual participants may have led to a less eager set of program beneficiaries than would have been selected in the absence of the study. To assess the credibility of this concern, we look at heterogeneous impacts of the program depending on the time of an applicant's registration under the assumption that those who registered earlier on registration days would have been more likely to be included in the program in the absence of the study. We see no significant differences between early and late registrants on psychosocial impacts. However, we do see some significant differences on financial outcomes whereby those who registered *later* on registration days have more positive outcomes than early registrants. This finding suggests that this method of community mobilization and recruitment utilized by Mercy Corps may unintentionally exclude participants who would benefit the most from their programs.

2 Related Work

With the rise of sports programs in the retinue of development organizations, a growing body of research has concentrated on the design of youth sport programs and the relationship between sports and psychosocial development. However, this body of research has been limited to the fields of psychology and sociology. Perkins and Noam (2007) define sports-based youth development programs as out-of-school-time programs that use a particular sport to facilitate learning and life skill development in youth. The authors posit that sports-based youth development programs provide structure through clear rules, expectations, and responsibilities which, in turn, encourage youth psychosocial development. Rookwood (2008), in discussing the role of soccer as a peace-building activity, suggest that soccer builds trust, respect, and self-discipline. Henley et al. (2007) highlight the role that sports play in supporting and encouraging youth resiliency - successful adaptation despite challenging or threatening circumstances. Jeanes (2013) suggest that peer-led education through sports programs are combined with multi-layered interventions directed at all levels of communities. Further research on sports-based youth development programs stress the positive relationship between sports and psychosocial development (Petitpas et al., 2005; Berlin et al., 2007).

While the above cited studies provide an important psychological and sociological grounding of hypothesized effects and potential benefits, these studies are systematically underpowered in sample size and fail to address endogeneity concerns of participating beneficiaries and program outcomes. With a burgeoning interest and funding dedicated to sports-based youth programs, a higher standard of emprical evidence should be applied in order to establish a causal connection between sports and psychosocial outcomes. Recent reviews of this literature come to the same conclusion, acknowledging the empirical shortfalls of existing research that fail to address the issues of endogenous selection into the program that confound any credible causal claims (Coalter, 2010; Holt and Jones, 2007).

To our knowledge no randomized trial of a sports-based youth development program had been conducted. In order to test the assumptions of the existing literature on sports programs, this paper presents an investigation of the causal relationship between sports-based youth development and both psychosocial and economic outcomes. Through a randomized trial, we set out to address the self-selection bias by conducting a randomized control trial that randomly assigns program applicants to a sports-based youth development program.

3 Program Design

In 2012, Mercy Corps launched its initiative, Promoting Sustainable Partnerships for Economic Transformation (PROSPECTS). PROSPECTS contained several interventions designed to provide vulnerable young Liberians with the psychosocial and pre-employment skills necessary for future formal or self-employment. This study focuses on the evaluation of one initiative within PROSPECTS, Sports for Change.

Sports for Change (SFC) targeted vulnerable youth between the ages of 18 and 25, who were no longer in school, and with minimal formal working experience. This population was broadly considered to be unskilled and "unemployable." The program in Liberia, adapted from the Mercy Corps method of engaging youth in other post-conflict settings, aimed to improve participant psychosocial outcomes, create greater resiliency to negative shocks among youth, and to prepare and facilitate entry of participants either into formal employment or entrepreneurial endeavors.[5] Using sports groups as a means of attracting and engaging with vulnerable youth, the program bundles sports team practices with life skills activities.

The Sports for Change program, sought to organize 1,200 participants into 30 sports clubs of approximately 40 members. One coach/mentor was assigned to each sports club and individual youth were randomly assigned to a sports group within their community. Mercy Corps trained coaches on a curriculum of life-skills lessons designed to complement the regular soccer and handball training sessions. These sessions integrated five core life skills: constructive communication, self-esteem, resilience and problem solving, teamwork and trust building, and strategy making and planning (see Appendix Figure A.1 for an example of a session schedule). The designers of the SFC program felt that these core life skills were not only essential to the participants' psychosocial well-being but also that they provided an essential foundation on which participants could enter the work force either formally or through their own entrepreneurial activities.

Coaches organized two sports group meetings per week for a total of 16 sessions over eight weeks. The typical meeting schedule consisted of one hour of introduction and warm ups, one hour of instructional activities, and one hour of sports. Appendix Table A.1 presents the topics covered in each of the sixteen Sports for Change sessions along with the targeted SFC skills to be emphasized in each session. In addition, participants received USD $2 for each session that they attended.

4 Experimental Design

Recruitment and Random Assignment to Sports for Change

Mercy Corps worked with Innovations for Poverty Action (IPA) to randomly assign a pool of eligible youth to its programs. The random assignment, facilitated through a public lottery, was conducted on an individual-by-individual basis and stratified by gender in order to ensure balance of men and women in treatment and control groups. Randomized assignment allows us to analyze the impact of the Sports for Change program by comparing key outcome variables after the program's conclusion across individuals included in Sports for Change groups and those assigned to the control group who were not included. For the 2,400 participants included in the analysis of the SFC program, 50% were assigned to a sports for change youth group and 50% were assigned to the control group.[6]

[5] Mercy Corps uses the Sports for Change method in over 25 countries.

[6] The initial research and program design of PROSPECTS also included a "cash for work" (CFW) program that involved an opportunity for participating youth to earn money by collecting recyclable materials. However, due to implementation challenges in this component, interest in the program was significantly

Mercy Corps and its implementing partners raised awareness of the registration date in each community by visiting the area a few days before the event. In these visits they publicized the registration day by circulating fliers, posting large banners throughout the community, and informing local authorities. In addition, on the day of the registration itself, Mercy Corps used a truck with a large amplification system to drive through the streets of the community, broadcasting information about the recruitment and encouraging youth to go to the registration location to sign up. This broadcasting approach began early in the morning of the registration day and extended until the targeted number of registrants had been reached, often late in the afternoon or early evening.

At each registration center, typically a local school or community center, interested youths queued in line, waiting for admission in the order in which they arrived. They were then allowed into the registration room in sets of approximately 30 and explained the details of the potential lottery outcomes. After completing registration, individuals then chose their assignment ticket which included whether or not they were in a SFC group and, if so, which team they were assigned to. The tickets were chosen from a covered bucket so that they could not influence their selected outcome.[7] The result of their draw was then recorded by IPA staff before they were then given additional information dependent on their chosen program and group.

As demonstrated in Table 1, the proportion of men and women in each group is stable across all communities. In total, 2400 individuals were deemed eligible for the program and assigned to either the control group or an SFC team. The number of applicants per community varied from 200 to 600. Although Mercy Corps targeted 50% women in each community, the overall average was ultimately 52% of total applicants leading to some community level variation in female participation. However, gender balance across treatments within each community was preserved by the stratification. In total, 1200 individuals (574 men and 626 women) were assigned to an SFC group and 1200 applicants (579 men and 621 women) were assigned to the control group.

Data and Research Timeline

Data for this study was derived from four different sources: registration data, baseline in-person interviews, programmatic data from Mercy Corps collected during the program, and endline phone interviews.

lower than anticipated by the program designers so that very few youths assigned to the CFW group ever actually participated in the program. As a result, the analysis in this study focuses solely on the Sports for Change component of PROSPECTS. The randomized allocation of the full set of 3,000 participants in the original research design was 20% of participants to the CFW program only, 20% to the SFC program only, 20% to both programs, and 40% to the control group. Finally, because of these implementation challenges as well as the budgetary and logistical challenges of conducting follow-up surveys amid the outbreak of Ebola, it was decided to exclude respondents in the CFW only treatment from the endline. We control for randomized assignment to the CFW program in all regressions (for those also assigned to SFC) but otherwise exclude them from further analysis or discussion in this paper.

[7] Group assignment tickets were made prior to each recruitment drive in line with the targeted number of participants for that community and in the proportions of the desired stratification. Participants drew their assignment cards from a covered box so that she could not see her potential assignment while choosing her ticket.

With the assistance of a field officer from IPA, registration forms were completed by the entire pool of eligible applicants. This data included basic demographic information: age, gender, and schooling. In addition, field officers recorded extensive tracking and contact information for future identification in anticipation of the baseline interviews. Following selection of their program and team assignment, group assignment was added to the applicant's record.

In the days immediately following the registration event in each community, IPA conducted an in-person baseline survey. The registration lottery and baseline were conducted according to the schedule shown in Table 2. Initial activities began on July 24, 2013 with the West Point community and concluded with Logan Town on February 2014. In each community, baseline interviews were initiated within a week of the registration date. Between completion of the registration/lottery days and administration of the baseline survey, there was very low attrition with only five out of 2,400 registrants refusing to be interviewed or unable to be traced.

The program was implemented on a phased schedule one community at a time following the registration and baseline schedule. However, implementation start dates varied by community. In all communities, the full set of 16 SFC sessions had been completed before April 2014. Programmatic data on attendance and payments to beneficiaries was recorded and provided by Mercy Corps.

Across all nine PROSPECTS communities, participation in SFC was high. The average attendance in eight of the nine communities was above 50%. On average, the 1,200 individuals assigned to SFC attended 10.35 SFC sessions. In only one community, Peace Island, was average attendance below the 50% threshold. We also find that the SFC program was successful in retaining participants. Among individuals who attended at least any SFC sessions, 70.1% attended at least half of the SFC sessions and 67.1% attended at least three-quarters of the SFC sessions, suggesting the importance of early mobilization and attendance. These high conditional attendance measures suggest that the SFC program was desirable in the eyes of participants and adds credibility to the view that this program was well implemented.

Unfortunately, due to risks associated with travel and restrictions in mobilizing survey teams during the Ebola crisis in Liberia, the endline survey was conducted through computer assisted telephone interviews instead of in-person as initially intended. The endline survey was conducted simultaneously, over the phone, for participants in all nine communities.[8] Endline interviews began on April 3, 2015 and continued through May 9, 2015. A total of 2,081 individuals were successfully interviewed over the phone for the endline survey, a follow-up success rate of 87%.

Descriptive Statistics and Baseline Balance

Table 3 shows summary statistics and balance of participants at baseline by program treatment status and by gender. The targeting of participants was effective in identify-

[8] Stratification of treatment by community alleviates the concern that inconsistency in time between the program and the follow-up survey may bias the assessment of the program's impact. Differences in this timing prevent us from making credible comparisons of treatment effects across communities.

ing youths of both genders. Average age was 21 years old with 83% having completed primary school and slightly more than one out of four having completed secondary. Just over 43% had some form of employment at baseline with an average of 14 hours of work per week and earnings of 21 USD per week among those working.

While most baseline variables are well balanced across treatment status, we note an imbalance in self-esteem whereby those in the control group had higher measures of baseline self-esteem for both women and men. In analysis of this outcome we prefer a panel analysis where we can appropriately control for these baseline imbalance over the cross-sectional analysis.

Additionally, we see a strong and highly significant imbalance in the measure of three month retrospective income by treatment status with respondents in the treatment group reporting considerably higher incomes than those in the control group. In general, we distrust recall of a highly variant measure over such a long window of time, given this imbalance and considerations, exclude it from our analysis on financial outcomes emphasizing the shorter seven day income recall instead.

With 13% attrition from the sample, largely attributable to the needed shift in strategy from in-person to phone interviews, we also look at balance of baseline characteristics in the sample of participants included in the endline in Appendix Table A.2. The general patters look similar. The self-esteem balance looks less severe, while the imbalance in three month income appears worse.

Psychosocial Well-Being and Labor Measures

The PROSPECTS program was built on a presumption that psychosocial well-being and financial empowerment and outcomes are closely linked. In addition to their direct benefits, it is believed that improvements in psychosocial well-being directly impact employment and work force readiness and potentially earnings among youth as they gain confidence and adopt more pro-social behaviors.

While causality between these measures is difficult to establish, we explore these correlations in our baseline data in Table 4. Column (1) shows strong positive correlations between measures of welfare, self-esteem, and locus of change (empowerment) and workforce participation. By contrast, we see that participants with negative measures on the aggression index (more aggressive behaviors) are also correlated with higher levels of workforce participation. We see no clear patterns however on the number of hours worked or income, among those working, over the last seven days. We do, however see some significant and positive correlations between subjective welfare and amount of income earned over the preceding three months in column (6). Overall, the descriptive evidence suggests that if better psychsocial well-being does lead to better financial outcomes, this relationship is confounded by other factors in the cross-sectional correlations.

5 Empirical Results

The sports for change program had two main objectives. First, the implementing partner saw the program as a way to reach vulnerable youth in order to improve their psychosocial well-being and resilience in the face of negative shocks. And second, they believed that the program would improve workforce "preparedness" and therefore positively impact labor related outcomes.

We estimate the direct effects of the program on psychosocial and financial outcomes using two complementary regression methods: cross-sectional estimation and panel estimation. For cross-sectional regressions we use the following regression equation:

$$Y_i = \beta_0 + \beta_1 SFC_i + \lambda \mathbf{X}_i + \delta_c + \epsilon_i \tag{1}$$

Where Y_i is an outcome of interest for individual, i, measured at the endline, and SFC_i is an indicator for whether individual i was assigned to the sports for change program. $\mathbf{X}_i$ is a set of time invariant covariates, in most specifications these include age, age-squared, female, and grade level attained. We also include a set of community fixed effects, δ_c. We use robust standard errors to adjust for heteroskedasticity of the error term.

Panel estimation is conducted using the following regression equation:

$$Y_{i,t} = \beta_0 + \beta_1 SFC_{i,t} + \delta_i + \gamma_t + \epsilon_{i,t} \tag{2}$$

$Y_{i,t}$ is an outcome of interest for individual, i, at time, t, and $SFC_{i,t}$ is an indicator for whether individual i was assigned to the sports for change program in time period t. We also include a set of individual and time fixed effects. Standard errors are clustered at the individual level.

For both estimation strategies, we can interpret β_1 as the causal effect of the program on the outcome variable because SFC status was randomly assigned and therefore should be orthogonal to the error term. For regressions showing the instrumental variables results, we use random assignment to the SFC program as an instrument for having ever participated in a sports for change session according to Mercy Corps' administrative records.

Psychosocial Impacts

Table 5 shows the direct impact of the sports for change program on a set of five psychosocial indices: self-esteem, locus of control, risky behaviors, aggression, and subjective welfare assessments, splitting the sample by gender. Panel (a) shows the cross-sectional estimation using age, age-squared, and grades of schooling completed while Panel (b) shows results from the panel estimation with individual and time fixed effects.

For men, both estimation strategies suggest a positive and significant impact on aggressive behaviors. Assignment to the SFC program improved aggressive behaviors by 0.14 standard deviations relative to those in the control group. This effect is significant at the 95% confidence level in the cross-sectional estimation and at the 90% confidence level in the panel estimation.

We noted in the baseline balance tables that the control group had a higher initial measure of self-esteem for both men and women, statistically significant at the 95% confidence level. Looking at the cross-sectional estimate in Panel (a) on self-esteem we see a positive coefficient for men of 0.09 standard deviations but no statistical significance.

For women the coefficient is negative but also insignificant. Looking instead at the panel estimation in Panel (b) that accounts for the lower starting level of self-esteem for respondents in the treatment group, we now estimate that the program had a positive impact of 0.15 standard deviations for men, significant at the 95% confidence level. The estimated impact for women is 0.09 standard deviations although the estimated coefficient is not distinguishable from zero. None of the other estimated program impacts are greater that 0.1 standard deviations in either direction and none are statistically distinguishable from zero.

Overall, we view this as evidence that the sports for change program had some moderate psychosocial benefits, in particular, by improving aggressive behaviors and self-esteem among male participants. However, the program did not have any discernible impact on female beneficiaries.

Resiliency to Negative Shocks

We implement a different empirical strategy in order to assess the impact of the sports for change program on resiliency. In particular, we want to test whether participation in the sports for change program mitigates the impact of unexpected negative shocks on peoples' well-being. Because we do not have data on shocks preceding the baseline survey, we use a cross-sectional estimation of the following form:

$$Y_i = \beta_0 + \beta_1 SFC_i + \beta_2 (SFC \times LifeEventIndex)_i + \beta_3 LifeEventIndex_i + \lambda \mathbf{X}_i + \delta_c + \epsilon_i \quad (3)$$

Where Y_i is a psychosocial outcome of interest and SFC_i is an indicator for whether individual i was assigned to the sports for change program. The Life Event Index is an index of negative personal and household shocks reported to have occurred in the thirty days prior to the endline survey. $\mathbf{X}_i$ is a vector of individual covariates and δ_c is a set of community level dummies. We use robust standard errors.

Table 6 shows these results, again split by gender. First, the second row shows that greater numbers of negative life events do negatively affect our measures of psychosocial well-being among respondents. Self-esteem and aggressive behaviors for both men and women are negatively impacted by negative life events all statistically significant at either the 95 or 99% confidence level. In addition, welfare for men is significantly negatively impacted at the 90% confidence level. However, the third row of Table 6 shows the interaction term of life events and assignment to a sports for change group. If the sports for change program made participants more resilient in the face of these negative life event shocks, we would expect positive coefficients in this row. The signs of estimates on this interaction term are inconsistent across the row. The only statistically significant coefficient is on risky behaviors in column (10) for women. However, this shows *increased* sensitivity to life events for people who were assigned to a sports for change group, the opposite of increased resiliency. We conclude that there is no evidence of improved resilience resulting from the sports for change program.

Labor Market and Financial Impacts

The results in the previous section show some moderate impacts on psychosocial outcomes, primarily among men. In Table 4 we saw that self-esteem is associated positively with working on the extensive margin, although these correlations are less clear for other labor market outcomes. Aggression has a puzzling strong and negative correlation with working whereby those with most aggressive behaviors are more likely to work although those with less aggressive behavior may earn slightly more. Either through these channels, or directly through its life skills components, the sports for change youth groups may have impacted financial outcomes.

Table 7 shows these results. Each column has a different financial or labor market outcome including whether a respondent was working in the week prior to the interview, the number of hours worked, income over the last seven days (with different functional forms and for different samples of respondents), and an income coping index. Panel (a) again shows the cross-sectional results whereas Panel (b) shows results from the panel estimation. Columns (1)-(3) show the average treatment effects for both, male, and female samples respectively whereas Columns (4)-(6) show the treatment on the treated effect for people who participated in the SFC program at least once. The first row of Panel (a) suggests that people assigned to SFC were 4.6 percentage points more likely to be working at the endline off of a base of 67.6%, significant at the 95% confidence level. The effect sizes appear similar for both men and women. Looking in column (4) we see that the magnitude of this effect increases to 6.3 percentage points among those who showed up to at least one SFC session. Controlling for individual baseline employment levels in the panel estimation, these estimates remain positive but lose about 20% of their magnitude and are no longer statistically distinguishable from zero. We view this as positive, albeit weak, evidence that the SFC program increased labor force participation on the extensive margin.

However, the remainder of the results suggest very little impact of the SFC program on financial outcomes. Starting in Panel (a) the second row shows a positive coefficient but no statistical significance for the number of hours worked among those working. In the second row of Panel (b), restricted to those working in both baseline and endline, we see a negative coefficient but, again, no statistical significance. Looking at another measure of labor force participation, we see a positive coefficient for earning any income over the preceding seven days, but neither specification shows any statistical significance.

Looking instead on the intensive margin of financial outcomes, the coefficients for "Inc 7 Days" in the fourth rows of both panels show the effect of the program on total earnings over the last seven days. The estimates are all positive but insignificant. In order to disentangle whether the positive effect is the result of greater participation on the extensive margin, as already shown, or improved performance on the intensive margin, we trim the sample to only those with positive incomes in the fifth row labeled "I7D ($>$ 0)" of Panel (a). This maintains a similar positive magnitude of just under one USD or roughly 8% of the control mean, but still lacking any statistical significance. In Panel (b) we try a different method and first restrict the sample to respondents who had positive earnings during the baseline (BL>0). We see that the estimates now drop in column (1) for the pooled estimates although they increase for women. Regardless, they remain statistically insignificant. We then constrain the sample further to only people who had positive earnings in both the baseline and endline, (BL/EL, > 0), but still find no significant impact of the treatment on earnings. Similarly we look at the logged measure of income but, again, find a positive coefficient without any statistical significance. Altogether, we do not find any strong evidence of an impact of SFC on the intensive margins of labor force participation or earnings.

Further casting doubt on a meaningful positive financial impact of the program, the final row of Panel (a) shows the effect of the program on a coping index which incorporates a number of financial coping behaviors resulting from financial shortfalls such as restricting consumption or only consuming inferior food types. The variable is coded so that an increase in the coping index measure suggests more reported coping behaviors. Here, we see a positive and significant increase on the prevalence of coping behaviors. The need for these alternative coping strategies contradicts the possibility of meaningful financial benefits for participants.[9] Overall, we find only modest evidence of an increase in labor force participation, but no further evidence of significant financial benefits.

6 Group Formation

The structure of the study randomly assigned eligible and interested youths to SFC youth groups and the control group. This was essential in order to assess the causal impact of the program on the psychosocial and financial outcomes of its participants presented in the preceding section. As discussed in Section 4, this research approach required the recruitment of a larger pool of potential participants than the number ultimately included in the program, thus likely altering the selection of who would have been included in the absence of the study. And second, the study required random assignment among selected participants to their respective sports teams, an approach that likely differed from a more endogenous method of group sorting that would have allowed or even encouraged greater numbers of friends to join the same group.

In this section, we look at whether randomization of participants and teams negatively affected the measured impacts of the program. First, we look at whether estimated impacts of the program were dampened by the inclusion of participants who registered later on recruitment days. Second, if teams were allowed to form endogenously, we would expect a greater likelihood of friends being on the same team. We therefore look at the impact of additional friends in ones group on participation and both financial and psychosocial outcomes.

Composition of Participants

The composition of program participants was likely altered by the requirement of the research design to identify eligible but not-included potential beneficiaries. Needing to form a control group effectively doubled the number of eligible benificiaries that needed

[9] Although we hesitate to make this strong claim, a coherent interpretation of SFC's positive impact on labor force participation and increased coping strategies could result if the program damaged respondents' financial security and thus increased the need both for more income *and* more restrictions on consumption.

to be mobilized in each community. Reaching these recruitment targets proved to be a challenging task for Mercy Corps and its partners. On recruitment days this could be seen in many communities where early in the registration day long lines of eager participants waited to sign up, but later in the day Mercy Corps was forced to do a second round of mobilization in order to achieve community and gender participation targets. If Mercy Corps had not needed to increase their overall numbers in order to fill a control group, they would have stopped mobilization efforts earlier in the day once the available youth group slots had been filled, at approximately the half way mark of recruitment during the study. We can therefore divide the pool of study participants into early and late registrants by community and gender to see if program impacts differed by these relatively eager or reticent registrants.

We estimate these heterogeneous effects using the same cross-sectional and panel estimation strategies as utilized in Section 5 with the inclusion of an interaction term of "early" and SFC treatment status in all specifications and additionally adding the "early" variable itself for the cross-sectional estimation. Table 8 shows the heterogeneous effects of the SFC program on psychosocial outcomes by timing of registration. Panel (a) shows the cross-sectional results. While the positive coefficient on the interaction term may show a stronger positive effect of the program on the early registrants, none are statistically significant and are often mitigated by a negative sign on the overall treatment effect. The results on the interaction term in Panel (b) are even less encouraging with the benefits on aggression for males appearing significant and stronger among late registrants and the sign of the interaction term more often negative than positive.

Table 9 looks instead at the heterogeneity of impacts by registration timing on financial outcomes. Again, Panel (a) shows the cross-sectional estimates. Notably, we see positive significance for a number of outcome variables for late registrants and see *negative* interaction terms in the estimates of each of these equations (although only the coefficient on number of hours worked for women in Column (8) is statistically significant). The panel regression results are similar, though only the coefficients on working are statistically significant among women in Column (7).[10]

Together, these results suggest that the demands of the research design to recruit enough potential respondents to have a control group did not damage the measured impacts of the program. If anything, it appears that the participants who benefited most from the program were those who showed up later on recruitment days and who may have otherwise been excluded by Mercy Corps' recruitment strategy.

Social Linkages on Teams

A second possibility is that the research design may have limited positive benefits by disrupting naturally occurring social networks and linkages that could be beneficial to program participation and outcomes. We therefore look at the impacts of having friends in

[10]It may be that early registrants were, in fact, more eager to participate in the program to start. Or, it may be that early and late registrants differ along some other dimension. Appendix Table A.6 shows that early and late registrants differ along a number of observable characteristics as well. Regardless of what is driving this heterogeneity, the inclusion of these later registrants in the program does not appear to be driving the limited impacts of the program we find in our analysis.

respondents' randomly assigned groups. First, having more friends in ones group could affect participant outcomes if individuals are more likely to show up and participate in the sports for change sessions. And second, the number of friends in your group could improve the value of these sessions and have an effect on either psychosocial or financial outcomes. To estimate the impact of having friends in one's group we use the following regression equation among respondents assigned to an SFC youth group:

$$Y_{i,c} = \beta_0 + \beta_1 Friends_i + \gamma_k + \lambda \mathbf{X}_i + \delta_c + \epsilon_i \tag{4}$$

Where Y_i is an outcome of interest, either measures of SFC attendance, psychosocial outcomes, or financial outcomes. $Friends_i$ is either an indicator for having any friends assigned to their SFC group or the number of friends they have assigned to their group. γ_k is a set of fixed effects for an individual's number of matched friends identified at baseline. $\mathbf{X}_i$ is a set of time invariant covariates including age, age-squared, female, and grade level attained. We also include a set of community fixed effects, δ_c. Robust standard errors are used.

Table 10 looks at the effect of having any friends in your SFC group on program participation. The left panel shows the impact of having any friends in your group on the extensive margin of SFC participation: whether an individual showed up to any team sessions. We see a large and statistically significant increase in program participation by this measure of 7.1 percentage points as the result of having a friend in ones group. This effect is significant at the 95% confidence level and constitutes a reduction of more than 30% in the number of youths who never show up. While the point estimate is slightly stronger for women than for men, the difference in the two groups is not statistically significant. The right panel isolates the intensive margin, restricting the sample to respondents who participated in at least one sports for change session. Here, we see much smaller magnitudes of the effect of friends on participation relative to the outcome mean and none of the estimates are significantly different from zero.[11]

While friends do appear to influence attendance, we do not find any evidence that friends improve program outcomes. Using a similar estimation strategy we also look at whether the presence of friends affects either psychosocial or financial outcomes. For added precision, we also include the lagged dependent variable to equation (4). We see no significant impact of the presence of friends on psychosocial outcomes. Results are included in Appendix Table A.4. We see significant impacts of friends on financial outcomes under only one specification, where total number of friends may boost the number of hours worked, but this marginal significance on two out of eighteen regressions is similar to what we would expect to see from chance. These results are contained in Appendix Table A.5. We therefore do not find any evidence that lack of pre-existing friends in SFC groups meaningfully reduced the program's impact.

[11] Appendix Table A.3 repeats this analysis for the total number of friends assigned to ones sports group. These results, while still of economically meaningful magnitude, are no longer statistically significant.

7 Summary and Conclusion

Using sports as a method of intervention and vehicle for other pro-social programming has come increasingly into fashion. They are viewed as a potentially transformative approach to engaging and positively effecting the lives of vulnerable youth. This study presents results on the impact of one of these programs. The program was well implemented and well attended by participants, and yet the impacts appear to be modest. Males xhibit some modest improvements in pro-social behavior as captured by measures of aggression and self-esteem. And respondents of both genders were more likely to be participating in the work force as a result of program participation, although earnings were not significantly impacted.

These limited results do not appear to have been damaged by impositions of the research design. Having friends increases likelihood of participation, but does not improve program outcomes. And distortions of selection of beneficiaries by requiring the creation of a control group may have improved the measured outcomes of the program by including late registrants who benefitted more than those who signed up more readily.

References

Berlin, Richard A, Aaron Dworkin, Ned Eames, Arn Menconi, and Daniel F Perkins (2007) "Examples of sports-based youth development programs," *New Directions for Student Leadership*, Vol. 2007, pp. 85–106.

Coalter, Fred (2010) "The politics of sport-for-development: Limited focus programmes and broad gauge problems?" *International review for the sociology of sport*, Vol. 45, pp. 295–314.

Collier, Paul and Anke Hoeffler (1998) "On economic causes of civil war," *Oxford Economic Papers*, Vol. 50(4), pp. 53–573.

Henley, Robert, Ivo Schweizer, Francesco de Gara, and Stefan Vetter (2007) "How Psychosocial Sport & Play Programs Help Youth Manage Adversity: A Review of What We Know & What We Should Research.," *International Journal of Psychosocial Rehabilitation*, Vol. 12.

Herbst, Jeffrey (2000) *States and Power in Africa: Comparative Lessons in Authority and Control*, Princeton NJ: Princeton University Press.

Holt, Nicholas L and Martin I Jones (2007) "10 Future directions for positive youth development and sport research," *Positive youth development through sport*, p. 122.

Jeanes, Ruth (2013) "Educating through sport? Examining HIV/AIDS education and sport-for-development through the perspectives of Zambian young people," *Sport, Education and Society*, Vol. 18, pp. 388–406.

Perkins, Daniel F and Gil G Noam (2007) "Characteristics of sports-based youth development programs," *New Directions for Student Leadership*, Vol. 2007, pp. 75–84.

Petitpas, Albert J, Allen E Cornelius, Judy L Van Raalte, and Tiffany Jones (2005) "A framework for planning youth sport programs that foster psychosocial development," *The sport psychologist*, Vol. 19, pp. 63–80.

Rookwood, Joel (2008) "Soccer for peace and social development," *Peace Review*, Vol. 20, pp. 471–479.

8 Tables

Table 1: Treatment Assignment

	SFC		Control		Total
	Male	Female	Male	Female	
West Point	90	70	91	69	320
New Kru Town	116	124	117	123	400
Peace Island	39	41	41	39	160
Buzzy Quarter	38	42	40	40	160
Clara Town	76	84	78	82	320
Dry Rice Market	40	40	40	40	160
Banjor	38	42	40	40	160
Chicken Soup Factory	68	92	67	93	320
Logan Town	69	91	65	95	320
Total	574	626	579	621	2400

Table 2: Registration and Baseline Survey Dates

Community	Registration Date	Baseline Data Collection
West Point	24 Jul 2013	29 Jul - 8 Aug 2013
New Kru Town	7 Sep 2013	11 Sep - 24 Sep 2013
Peace Island	31 Oct 2013	3 Nov - 8 Nov 2013
Buzzy Quarter	7 Nov 2013	9 Nov - 16 Nov 2013
Clara Town	18 Nov 2013	22 Nov - 5 Dec 2013
Dry Rice Market	12 Dec 2013	15 Dec - 19 Dec 2013
Banjor	11 Jan 2014	13 Jan - 16 Jan 2014
Chicken Soup Factory	18 Jan 2014	23 Jan - 30 Jan 2014
Logan Town	1 Feb 2014	4 Feb - 12 Feb 2014

Table 3: Random Assignment Balance and Summary Statistics

	All (N=2395)			Male (N=1150)			Female (N=1245)		
	CTRL (1197)	TRT (1198)	P-val	CTRL (577)	TRT (573)	P-val	CTRL (620)	TRT (625)	P-val
Basic Demographics									
Age	20.95	20.80	0.26	20.67	20.52	0.36	21.16	21.05	0.47
Head of Household	0.131	0.143	0.42	0.156	0.174	0.42	0.110	0.116	0.75
Household Size	6.762	6.598	0.26	6.704	6.410	0.18	6.816	6.770	0.81
Mother Living	0.868	0.870	0.91	0.868	0.892	0.21	0.869	0.849	0.33
Father Living	0.713	0.717	0.84	0.720	0.726	0.83	0.707	0.709	0.94
Has Children	0.454	0.432	0.28	0.224	0.211	0.61	0.669	0.635	0.21
Christian	0.876	0.863	0.37	0.858	0.832	0.23	0.892	0.891	0.97
Muslim	0.111	0.125	0.29	0.127	0.159	0.12	0.097	0.094	0.89
Matched Friends	2.390	2.419	0.72	2.263	2.366	0.37	2.508	2.467	0.71
Education and Cognitive Ability									
Primary School	0.835	0.836	0.95	0.912	0.895	0.35	0.765	0.782	0.45
Secondary School	0.276	0.265	0.57	0.343	0.326	0.55	0.213	0.210	0.89
Grades Completed	11.39	11.35	0.81	12.39	12.26	0.48	10.46	10.52	0.79
Numeracy (0-10)	5.951	5.784	0.04**	6.293	6.286	0.95	5.632	5.323	0.01***
Digits Forward (0-8)	5.176	5.124	0.42	5.295	5.293	0.99	5.066	4.968	0.27
Digits Backward (0-8)	1.981	1.924	0.24	2.102	1.988	0.10	1.868	1.866	0.97
Word Recall 1 (0-10)	3.059	3.072	0.91	3.201	3.190	0.95	2.927	2.963	0.82
Word Recall 2 (0-10)	2.578	2.558	0.85	2.662	2.600	0.66	2.500	2.520	0.89
Ravens Score (0-3)	1.753	1.726	0.51	1.912	1.914	0.96	1.605	1.554	0.34
Risk Aversion (0-6)	3.785	3.758	0.79	3.910	3.775	0.37	3.669	3.742	0.62
Psycho-Social Measures/Indices									
Subjective Welfare	2.292	2.321	0.59	2.234	2.251	0.81	2.347	2.384	0.62
Self-Esteem	20.89	20.39	0.00***	20.78	20.27	0.03**	20.98	20.50	0.04**
Locus of Control	24.09	24.03	0.63	24.24	24.16	0.60	23.95	23.92	0.86
Aggression	2.603	2.546	0.63	2.445	2.401	0.80	2.750	2.678	0.65
Risky Behavior	1.501	1.468	0.67	2.106	2.051	0.68	0.939	0.934	0.95
Depression	21.99	23.28	0.02**	20.54	21.76	0.13	23.34	24.67	0.09*
Labor Force Participation and Earnings									
Working	0.431	0.436	0.83	0.499	0.484	0.61	0.368	0.391	0.40
Hours Worked (7D)	13.92	14.49	0.61	16.03	16.77	0.66	11.96	12.40	0.76
Inc 7 Days	6.423	6.647	0.69	8.254	8.124	0.89	4.710	5.294	0.37
Inc 7 Days (>0)	20.78	24.26	0.43	20.84	29.65	0.24	20.69	18.47	0.57
Log(Inc 7 Days)	2.324	2.384	0.43	2.397	2.538	0.17	2.233	2.218	0.89
Inc 3 Months	47.56	57.84	0.01***	54.85	65.49	0.07*	40.85	50.84	0.05**
Inc 3 Months (>0)	125.0	113.7	0.69	164.7	123.8	0.44	79.2	103.3	0.02**
Log(Inc 3M)	3.716	3.930	0.00***	3.770	4.029	0.01**	3.652	3.828	0.08*

Table 4: Labor and Psychosocial Associations

	(1) Working	(2) Hrs Work	(3) 7dI	(4) log(7dI)	(5) 3mI	(6) log(3mI)
WF	0.023**	0.676	0.788	0.051	2.473	0.086**
	(0.010)	(0.662)	(0.581)	(0.032)	(2.964)	(0.034)
SE	0.043***	-1.391*	-0.099	-0.015	-1.745	-0.028
	(0.010)	(0.801)	(0.618)	(0.033)	(3.089)	(0.037)
LOC	0.027***	-0.363	0.474	0.006	2.266	0.013
	(0.010)	(0.820)	(0.618)	(0.033)	(2.771)	(0.035)
Agg	-0.032***	0.437	0.885	0.064*	-0.133	0.015
	(0.010)	(0.807)	(0.603)	(0.033)	(3.032)	(0.037)
RB	0.002	0.764	-0.710	-0.049	-6.444	-0.043
	(0.011)	(0.870)	(0.702)	(0.037)	(4.488)	(0.041)
N	2368	864	796	796	1230	1230
Mean Y	0.44	28.92	16.49	2.44	87.36	3.82
R2	0.079	0.024	0.049	0.059	0.027	0.039

Notes: Positive earnings of last seven days (7d) or three months (3m) of income. Hrs Work=Hours worked among people working. WF=Subjective Welfare, SE=Self-Esteem Index, LOC=Locus of Control Index, Agg=Aggression Index, RB=Risky Behavior Index. Psychosocial measures normalized with positive set such that it is "better" behavior. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Impact of SFC on Pychosocial Outcomes

Panel (a): Cross-Sectional Estimation										
	Male (N=969)					**Female (N=1086)**				
	(1) WF	(2) SE	(3) LOC	(4) Agg	(5) RB	(6) WF	(7) SE	(8) LOC	(9) Agg	(10) RB
SFC	-0.04	0.07	0.03	0.14**	0.01	0.09	-0.03	-0.06	0.02	-0.01
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
R2	0.017	0.020	0.031	0.026	0.030	0.030	0.010	0.020	0.032	0.015

Panel (a) Notes: Covariates: age, age^2, and school grade completed.

Panel (b): Panel Estimation										
	Male (N=980)					**Female (N=1098)**				
	(1) WF	(2) SE	(3) LOC	(4) Agg	(5) RB	(6) WF	(7) SE	(8) LOC	(9) Agg	(10) RB
SFC	-0.06	0.15**	0.01	0.14*	0.03	0.04	0.09	-0.08	-0.02	0.02
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
R2	0.584	0.645	0.568	0.597	0.644	0.567	0.591	0.557	0.610	0.583

Panel (b) Notes: Regressions include individual and time fixed effects.

Table Notes: SFC=Assigned to a Sports for Change youth group. Hrs Work=Hours worked among people working. WF=Subjective Welfare, SE=Self-Esteem Index, LOC=Locus of Control Index, Agg=Aggression Index, RB=Risky Behavior Index. Psychosocial measures normalized with positive set such that it is "better" behavior. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Effect of SFC on Resiliency to Negative Life Events - Cross-Sectional Estimation

	Male (N=961)					**Female (N=1077)**				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	WF	SE	LOC	Agg	RB	WF	SE	LOC	Agg	RB
SFC	-0.05	0.07	0.04	0.14**	0.01	0.09	-0.02	-0.04	0.02	-0.03
	(0.07)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
LEI	-0.04*	-0.07***	-0.03	-0.06***	0.01	0.01	-0.06***	0.02	-0.05**	0.02
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
SFC x LEI	0.04	-0.04	0.03	0.02	-0.00	-0.01	-0.02	-0.03	0.01	-0.10***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
R2	0.021	0.068	0.035	0.040	0.031	0.030	0.037	0.019	0.039	0.028

Notes: LEI= Life Event Index. SFC=Assigned to Sports for Change youth group. Psychosocial measures normalized with positive set such that it is "better" behavior. WF=Subjective Welfare, SE=Self-Esteem Index, LOC=Locus of Control Index, Agg=Aggression Index, RB=Risky Behavior Index. Covariates: age, age2, and years of education * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Effect of SFC on Financial Outcomes

Panel (a): Cross-Sectional Estimation						ATE			TOT	
					Both	Male	Fem	Both	Male	Fem
Outcome:	Both	Male	Fem	R2	(1)	(2)	(3)	(4)	(5)	(6)
Working	2058	971	1087	0.04	0.046**	0.040	0.049*	0.063**	0.056	0.067*
Mean Control:	*0.676*	*0.741*	*0.617*		(0.020)	(0.027)	(0.029)	(0.027)	(0.038)	(0.039)
Hours Worked	1412	724	688	0.03	0.792	0.023	1.460	1.088	0.033	1.921
Mean Control:	*18.53*	*18.98*	*18.02*		(0.972)	(1.376)	(1.373)	(1.336)	(1.979)	(1.810)
Any Income	2057	970	1087	0.03	0.018	0.005	0.029	0.024	0.006	0.039
Mean Control:	*0.769*	*0.830*	*0.714*		(0.018)	(0.024)	(0.027)	(0.025)	(0.033)	(0.036)
Inc 7 Days (I7d)	2033	958	1075	0.03	0.999	1.058	0.892	1.363	1.456	1.207
Mean Control:	*10.52*	*11.10*	*9.99*		(0.768)	(1.008)	(1.149)	(1.048)	(1.390)	(1.555)
I7D (>0)	1577	792	785	0.03	0.982	1.245	0.611	1.323	1.732	0.797
Mean Control:	*13.74*	*13.42*	*14.06*		(0.933)	(1.132)	(1.471)	(1.257)	(1.578)	(1.920)
Log(I7D)	1577	792	785	0.07	0.071	0.079	0.054	0.095	0.109	0.071
Mean Control:	*2.059*	*2.085*	*2.032*		(0.055)	(0.077)	(0.080)	(0.075)	(0.107)	(0.104)
Coping Index	2055	969	1086	0.02	0.808*	0.544	1.043*	1.101*	0.748	1.407*
Mean Control:	*11.58*	*11.72*	*11.46*		(0.426)	(0.614)	(0.591)	(0.579)	(0.846)	(0.796)

Panel (a) Notes: Covariates gender, age, age-squared, and years education.

Panel (b): Panel Estimation						ATE			TOT	
					Both	Male	Fem	Both	Male	Fem
Outcome:	Both	Male	Fem	R2	(1)	(2)	(3)	(4)	(5)	(6)
Working	2076	974	1087	0.56	0.038	0.044	0.033	0.052	0.060	0.045
Mean Control:	*0.676*	*0.741*	*0.617*		(0.029)	(0.042)	(0.040)	(0.040)	(0.058)	(0.054)
Hours Worked	620	344	276	0.56	-0.150	-1.628	1.431	-0.207	-2.295	1.916
Mean Control:	*20.67*	*21.66*	*19.26*		(2.287)	(3.195)	(3.256)	(3.146)	(4.510)	(4.356)
Any Income	2079	981	1098	0.60	0.016	0.026	0.007	0.022	0.036	0.010
Mean Control:	*0.769*	*0.830*	*0.714*		(0.027)	(0.038)	(0.038)	(0.037)	(0.053)	(0.052)
Inc 7 Days (I7D)	2035	959	1076	0.56	0.546	0.843	0.255	0.744	1.160	0.345
Mean Control:	*10.56*	*11.14*	*10.03*		(0.930)	(1.334)	(1.295)	(1.267)	(1.836)	(1.750)
I7D (BL>0)	797	418	379	0.57	0.284	-1.311	1.866	0.390	-1.784	2.583
Mean Control:	*11.54*	*13.08*	*9.67*		(1.691)	(2.450)	(2.302)	(2.320)	(3.336)	(3.185)
I7D (BL/EL>0)	647	359	288	0.58	1.335	-0.070	2.408	1.807	-0.095	3.252
Mean Control:	*14.30*	*14.86*	*13.47*		(1.906)	(2.688)	(2.660)	(2.581)	(3.650)	(3.593)
Log(I7D)	647	359	288	0.58	0.103	-0.024	0.207	0.139	-0.032	0.280
Mean Control:	*2.093*	*2.119*	*2.054*		(0.112)	(0.152)	(0.163)	(0.151)	(0.207)	(0.220)

Panel (b) Notes: Regressions include individual and time fixed effects. BL>0 looks only at respondents with positive income at baseline. BL/EL>0 looks only at respondents with positive income in both survey rounds.

Table Notes: ATE=Average Treatment Effect. TOT=Treatment on the Treated with random assignment used as instrument for ever having participated in the SFC program. Each row represents a different financial outcome. Columns 1-6 show the regression estimate for the coefficient on being assigned to a sports for change group. Reported R2 is for the regressions including both men and women. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Impact of SFC on Psychosocial Impacts by Time of Registration

Panel (a): Cross-Sectional Estimation

	Males (N=971)					**Females (N=1087)**				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	WF	SE	LOC	Agg	RB	WF	SE	LOC	Agg	RB
SFC	-0.08	-0.01	-0.03	0.14	-0.02	0.04	-0.08	-0.10	-0.02	0.01
	(0.10)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.08)	(0.09)	(0.09)	(0.09)
Early x SFC	0.08	0.16	0.12	0.01	0.06	0.09	0.11	0.09	0.07	-0.04
	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)
Early	-0.06	-0.07	-0.15	-0.08	-0.15*	-0.05	-0.05	0.06	-0.03	0.02
	(0.09)	(0.09)	(0.09)	(0.10)	(0.09)	(0.08)	(0.08)	(0.08)	(0.08)	(0.09)
N	969	971	971	971	971	1086	1087	1087	1087	1087
R2	0.018	0.022	0.033	0.028	0.034	0.030	0.010	0.022	0.033	0.015

Notes Panel (a): Covariates include age age^2 and grades of schooling completed.

Panel (b): Panel Estimation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	WF	SE	LOC	Agg	RB	WF	SE	LOC	Agg	RB
SFC	-0.03	0.14	0.04	0.18**	-0.03	0.07	0.15	-0.16	-0.05	0.03
	(0.10)	(0.09)	(0.10)	(0.09)	(0.09)	(0.10)	(0.10)	(0.10)	(0.09)	(0.09)
Early x SFC	-0.06	0.02	-0.04	-0.08	0.11	-0.05	-0.12	0.15	0.06	-0.02
	(0.11)	(0.10)	(0.12)	(0.11)	(0.11)	(0.12)	(0.11)	(0.12)	(0.11)	(0.11)
N	1960	1964	1964	1964	1964	2194	2196	2196	2196	2196
R2	0.584	0.645	0.568	0.597	0.645	0.567	0.591	0.557	0.610	0.584

Notes Panel (b): Regressions include individual and time fixed effects.

Notes: Early registrants are categorized as being in the first half of registrants for a respondents gender within their community. Psychosocial measures normalized with positive set such that it is "better" behavior. WF=Subjective Welfare, SE=Self-Esteem Index, LOC=Locus of Control Index, Agg=Aggression Index, RB=Risky Behavior Index. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Impact of SFC on Labor and Financial Outcomes by Time of Registration

Panel (a): Cross-Sectional Estimation

	Males						**Females**					
	Working	Hours	Income Last Week			CI	Working	Hours	Income Last Week			CI
	(0/1)	(>0)	Raw	(>0)	Log		(0/1)	(>0)	Raw	(>0)	Log	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
SFC	0.056	-0.191	2.251*	2.822*	0.141	-0.052	0.044	3.719**	1.055	0.731	0.067	0.011
	(0.040)	(1.930)	(1.309)	(1.506)	(0.114)	(0.090)	(0.040)	(1.824)	(1.607)	(1.988)	(0.115)	(0.086)
Early x SFC	-0.029	0.413	-2.320	-2.962	-0.118	0.209*	0.009	-4.687*	-0.317	-0.270	-0.027	0.191
	(0.055)	(2.751)	(2.034)	(2.347)	(0.155)	(0.127)	(0.057)	(2.701)	(2.272)	(2.919)	(0.161)	(0.121)
Early	0.017	-0.105	2.639**	2.573*	0.113	-0.069	-0.069*	5.034***	0.866	1.920	0.189*	-0.101
	(0.040)	(1.882)	(1.310)	(1.488)	(0.108)	(0.094)	(0.041)	(1.876)	(1.619)	(2.108)	(0.114)	(0.086)
N	971	724	958	792	792	969	1087	688	1075	785	785	1086
Control Mean	0.741	18.978	11.102	13.422	2.085	-0.027	0.617	18.021	9.993	14.063	2.032	-0.055
R2	0.023	0.020	0.036	0.039	0.075	0.020	0.037	0.057	0.043	0.042	0.076	0.041

Notes Panel (a): Covariates include age, age^2, and grades of schooling completed. CI=Standardized Coping Index where positive is more financial coping behaviors.

Panel (b): Panel Estimation

	Males						**Females**					
	Working	Hours	Income Last Week				Working	Hours	Income Last Week			
	(0/1)	(>0)	Raw	(EL>0)	(BL/EL>0)	Log	(0/1)	(>0)	Raw	(EL>0)	(BL/EL>0)	Log
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
SFC	0.021	-1.442	0.787	0.121	0.384	0.024	0.108**	-0.202	0.866	2.988	2.745	0.169
	(0.052)	(4.039)	(1.638)	(2.934)	(3.372)	(0.192)	(0.048)	(4.226)	(1.585)	(3.008)	(3.249)	(0.206)
Early x SFC	0.043	-0.372	0.109	-2.865	-0.903	-0.096	-0.147***	2.891	-1.204	-1.933	-0.612	0.071
	(0.059)	(4.493)	(1.961)	(3.662)	(4.041)	(0.215)	(0.054)	(4.383)	(1.831)	(3.521)	(3.806)	(0.213)
N	979	344	959	418	359	359	1097	276	1076	379	288	288
Control Mean	0.741	21.663	11.142	13.084	14.855	2.119	0.617	19.262	10.027	9.674	13.47	2.054
R2	0.534	0.534	0.568	0.567	0.566	0.580	0.568	0.587	0.547	0.571	0.593	0.588

Notes Panel (b): Regressions include individual and time fixed effects. BL>0 looks only at respondents with positive income at baseline. BL/EL>0 looks only at respondents with positive income in both survey rounds.

Table Notes: Early registrants are categorized as being in the first half of registrants for a respondents gender within their community. Logged income variables drop zero values. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Effects of Any Friends on SFC Participation

	Any SFC Participation				Days of SFC Participation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	All	Male	Female	All	All	Male	Female	All
Any Friends	0.071**	0.050	0.074**	0.038	-0.037	-0.055	0.016	-0.025
	(0.031)	(0.045)	(0.035)	(0.047)	(0.112)	(0.216)	(0.163)	(0.193)
Female x Any Friends				0.062				-0.021
				(0.053)				(0.265)
Female	0.012			-0.010	0.086			0.094
	(0.033)			(0.039)	(0.093)			(0.138)
N	1039	486	552	1039	760	351	407	760
Mean Dep Var	0.73	0.73	0.74	0.73	14.71	14.63	14.78	14.71
Mean Any Friends	0.365	0.344	0.383	0.365	0.365	0.344	0.383	0.365
R2	0.077	0.100	0.138	0.078	0.130	0.203	0.137	0.130

Notes: Any Friends refers to whether individual had any friends identified during baseline assigned to the same youth group. Days of participation are restricted to non-zero responses in order to isolate the intensive margin for the effect of friends on participation. Regressions include dummies for amount of potential matches in community and baseline outcome value. Covariates: age, age^2, education attainment, and gender. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Chapter 3

Marriage Markets and Rainfall Shocks: Evidence from Burkina Faso

Chapter 3:

Marriage Markets and Rainfall Shocks: Evidence from Burkina Faso

Sylvan Herskowitz*

April 2013

Abstract

Selection of a marriage partner is among the most important choices in a person's life. Families form the central unit of almost every imaginable development related outcome. Economic shocks may influence household marriage decisions. I develop a simple model whereby supply of and demand for brides differentially respond to income shocks due to marriage transfers made by the husband to the wife's family. In predominantly agricultural societies, rainfall may proxy for these shocks. Using a historical panel of rainfall from the University of Delaware and Demographic and Household Survey data from Burkina Faso, I find evidence that likelihood of marriage for both women and men may be influenced by rainfall in preceding years. Low rainfall increases women's likelihood of marriage in subsequent years. Rainfall that is one standard deviation below the historical mean increases a woman's likelihood of marrige two years later by 2.89 percentage points, or roughly 15%. Results are strongests among young women between 13 and 16 years old. I also present tentative evidence consistent with trade theory that more integrated or "open" marriage markets respond more strongly to rainfall shocks than those that are closed or autarchic. Understanding the influence of economic shocks on marriage decisions and family formation is an important topic in need of future study.

*Department of Agricultural and Resource Economics, University of California at Berkeley. Email: sherskowitz@berkeley.edu. I would like to express my deepest gratitude to Elisabeth Sadoulet and Jeremy Magruder for their invaluable guidance, advice, and patience. I would also like to thank Michael Anderson, Larry Karp, Ethan Ligon, Wolfram Schlenker, Sofia Villas-Boas, Patrick Baylis, Josh Blonz, Marshall Burke, Kyle Emerick, Lauren Falcao, Seth Garz, Ken Lee, and Severin Ouedraogo for their feedback and assistance. All errors are my own.

1 Introduction

Selection of a marriage partner is among the most important choices in a person's life. It is with this person that a new family is started, income is shared, and children are born. Almost every possible development related outcome stems from choices made within the household. It is therefore critical to understand how marriages, the central building block of a family, are formed.

Choice of a marriage partner may be influenced by many factors. While personal affection may play a role, in many settings other more instrumental considerations can matter as well. In particular, where marriages are accompanied by significant costs and transfers, marriage decisions may be influenced by a household's surrounding economic conditions. Given the importance and size of these transfers, I hypothesize that supply and demand for brides in the marriage market may be influenced by household income shocks in the form of rainfall in preceding years.

The existing economic literature on marriage decisions is limited. Mark Rosenzweig and Oded Stark produced a seminal paper on the motivations of marriage contracts in 1989. Their study looked at rural India and presented evidence of risk-mitigating network formation between families as a motivation for marriage.

The topic of marriage decision-making has gained some renewed interest in recent years. Jensen (2012) and Mbiti (2008) both show that increases in returns to female labor lead to delay of marriage. Fafchamps and Quisumbing (2004) look at strategic gift giving by parents to improve children's marriage prospects. Additionally, there is a potentially congruous literature on buffer stocks and consumption smoothing.[1]

A recent working paper from 2011 by Hans Hoogeveen, Bas van der Klaauw, and Gijsbert van Lomwel, analyzes the relationship between marriage timing and income shocks. Looking at data from Zimbabwe, they create a discrete-time dynamic programming model. Hoogeveen et al. find little evidence of any effect of rainfall shocks on marriage and shift their focus to an analysis of idiosyncratic shocks where they do find that loss of cattle leads to increased marriage likelihood for women.

This analysis is most similar to this work by Hoogeveen et al. as I will also look at the impacts of correlated rainfall shocks on likelihood of marriage as evidence of consumption smoothing. However, their population of relatively well off farmers in Zimbabwe may have been too wealthy to have been susceptible to year to year rainfall fluctuations and is considerably wealthier than the population of rural farmers from Burkina Faso where I focus my study. Additionally, in my paper I seek to conduct a more thorough analysis of how large, correlated income shocks may affect the general equilibrium of a marriage market. First, I will look at both the supply and demand side of the marriage market for brides. Second, I look for heterogeneous effects in order to better understand who is more or less likely to be affected by correlated income shocks. Third, I develop and test a trade framework to better understand the broader context under which correlated shocks may be either mitigated or exacerbated.

Using data from the University of Delaware's historical rainfall project along with three rounds of Demographic and Household Surveys in Burkina Faso, I find consid-

[1] See Fafchamps et al. (1998) and Kazianga and Udry (2006).

erable, if mixed, evidence that rainfall shocks affect marriage decisions. In Section 2, I present some relevant background about Burkina Faso. Section 3 establishes a simple model of bride supply and demand and fits it into a trade framework. Section 4 details the empirical strategy of the paper. Section 5 presents the paper's findings. In Section 6, I detail two robustness checks to confirm my findings. Section 7 concludes.

2 Burkina Faso: Setting and Background

Burkina Faso is one of the poorest countries in the world with a GNI per capita of just $570 USD. In 2011 it was ranked 181st out of 187 countries in the United Nations Development Programme's Human Development Index. Agriculture contributes just 35% of national GDP, though it is the primary source of income for 90% of the population (World Bank 2013). Given this dependence on agriculture, Burkinabes are particularly vulnerable to fluctuations in year to year rainfall. Situated in West Africa's Sahel, a region known for its recurrent droughts, these negative shocks to household income are common, and the source of frequent duress.

Meanwhile, weddings in Burkina Faso are the central focus of most young peoples' lives, marking passage from childhood to adulthood. The choice of a spouse is taken seriously, not just by the future bride and groom, but also by their parents and extended families. Traditionally, wedding decisions are heavily influenced, if not entirely controlled, by a young woman's parents. Even for men, family considerations play a significant role in identifying a proper spouse and deciding when a young man has his family's blessing to get married. It is common for the man and woman to have never met before steps for formal engagement have begun (Pacere 1998).

Traditional Burkinabe culture includes the payment of a bride price from the husband's family to the bride's. The quantity and form of this payment vary considerably depending on the wealth of the participating families. Typically bride prices will include some mixture of cattle, goats, rams, poultry, fabrics, jewelry, or simply cash. The value of these gifts and payments is considerable and viewed as compensation for the family's loss of a daughter as she leaves her parents' village, or at least household, in order to move in with her new husband (D'Eveil Pugsada 2000).

Marriage age is young in Burkina Faso with considerable recent attention on the prevalence of child brides. Young women are highly valued in the marriage market for two reasons: 1) they are likely to have more years of reproductive ability and 2) they are less likely to have lost their virginity. Given their high value in the marriage market and particularly given the uncomfortable proximity of many Burkinabe families to minimum subsistence levels, young women are said to be particularly vulnerable to early marriage in times of financial stress (IRIN IRIN 2013).

3 Model

In this section, I develop a basic model showing how present day rainfall impacting household income can lead to changes in the marriage market. First, I show that rainfall may shift demand for and supply of brides in the marriage market. Second, I outline potential sources of heterogeneity stemming from the model. Finally, I fit these supply and demand adjustments into a trade framework.

Bride Supply and Demand Response to Rainfall

In predominantly agricultural areas, rainfall may affect the supply of and demand for brides. Rainfall impacts overall farm yields which, in turn, affect household income. In many settings, marriages, along with their accompanying ceremonies and negotiations require costly expenses as well as significant transfers between parties. Shifts in supply of and demand for brides will depend on the burden of payment for these contributions between the bride and groom's side of the marriage transaction.

Consistent with norms in Burkina Faso, I assume that bride prices are paid by the husband's family to that of the bride. A bride price (p) is paid from the husband's family to that of the bride if they agree to her marriage. Household production is based on a concave agricultural production function with a single input, rainfall (R). Total yearly income is therefore comprised of agricultural production (normalized to a price of one) plus or minus any bride price transfer. A household's yearly income is thus defined as: $f(R) \pm p$ where p will be negative for men or positive for women in years they choose to get married and equal to zero otherwise.[2]

Because marriages generally only take place once (or a limited number of times for polygynous men), timing of marriage is a first order concern, and whether to get married this year or to postpone until later. For simplicity, I assume a two time-period utility maximization problem where an individual must be married in either the first or the second time period and can marry only once. It is assumed that while rainfall and bride price are known for t=1, families must take expectations of next year's rainfall and bride price.

For women's families, the supply side of the market for brides, the payoffs to marriage or postponement are defined as:

$$\begin{aligned} w(postpone) &= v(f(R_1)) + \beta E[v(f(R_2) + P_2)] \\ w(marry) &= v(f(R_1) + P_1) + \beta E[v(f(R_2))] \end{aligned}$$

β is the familiy's future discounting parameter and $v()$ is the household's indirect utility function. Facing these expected payoffs, a daughter will enter the marriage market if:

$$v(y(R_1) + P_1) + \beta E[v(f(R_2)] \geq v(f(R_1)) + \beta E[v(f(R_2) + P_2)]$$

Reshuffling:

$$v(f(R_1) + P_1) - v(f(R_1)) \geq \beta E[v(f(R_2) + P_2)] - \beta E[v(f(R_2))]$$

[2] It is possible that bride prices could go in the opposite direction, as seen in other parts of the world. This would result in opposite signs for the transfer for men and women. However, this is inconsistent with the local setting and local norms that transfers flow from groom to bride.

Holding other factors fixed, the right-hand side of the entry condition is constant if expectations of future rainfall and future bride prices are considered stable.[3] A rise in R_1 leads to a fall in the left hand side because of the concavity of the indirect utility function making the marginal utility from P_1 lower. Conversely, for low levels of rainfall, the additional utility from gaining P_1 becomes more significant. This can be interpreted as a relaxation of her marriage condition.

On aggregate, thinking across many households, some will have their thresholds satisfied while others will still not. I interpret this response as a rise in overall supply of brides following a low realization of rainfall in the first period. Therefore, $\frac{\partial S}{\partial R_1} < 0$.

Switching to the men's side, we follow a parallel construction and analysis. Men's families have the following payoffs:

$$\begin{aligned} m(postpone) &= v(f(R_1)) + \beta E[v(f(R_2) - P_2)] \\ m(marry) &= v(f(R_1) - P_1) + \beta E[v(f(R_2))] \end{aligned}$$

Note that the only difference between payoffs to men from those of women is that the bride price is subtracted from the family's income in the year that he gets married, whereas for women this term was added. Facing these expected payoffs, a son will enter the marriage market if:

$$v(y(R_1) - P_1) + \beta E[v(f(R_2)] \geq v(f(R_1)) + \beta E[v(f(R_2) - P_2)]$$

Reshuffling:

$$v(f(R_1) - P_1) - v(f(R_1)) \geq \beta E[v(f(R_2) - P_2)] - \beta E[v(f(R_2))]$$

Now, we can see that when a houehold experiences stronger than typical rainfall in the first period, and thus has higher income, marriage will be a more attractive option in order to avoid the loss of the bride price in the next round when income is still unknown but expected to be lower. The men's comparative statics are thus opposite from those of women so that positive rainfall increases bride demand, $\frac{\partial D}{\partial R_1} > 0$.[4]

In practice, P may be different for different men or women. Since both husbands and wives vary significantly across a range of different characteristics (such as looks, wealth, family prestige, etc), not all matches will negotiate the same price. This model abstracts from that acknowledgment and suggests how men and women respond to different income shocks on average given a certain bride price, remaining agnostic about which char-

[3] Particularly in the case of expected bride price in the next round, it could and should be questioned whether or not $E[P_2]$ is independent of P_1. For simplicity, I assume that it is. However, if instead we assumed that P_2 is also dependent on R_1 the most likely direction of this bias would be to be negatively correlated with P_1. This would serve to make any resulting inequality from R_1 rainfall even more out of balance and amplify the response to a first period shock.

[4] It should additionally be noted that this result depends on an assumption that men cannot save their income easily between time periods. This is a more difficult assumption to defend than that of credit unavailability. However, where families are in persistent debt, or demands from other family members make savings difficult, there may still be an urge to spend the money when it is available and not to assume that it will still be available in later years.

acteristics determine different price levels.

Supply and Demand Response Heterogeneity

The model readily incorporates heterogeneous effects by levels of wealth. The model predicts that households with higher baseline wealth, additional alternative revenue streams, assets, or buffer stocks would respond less dramatically to rainfall fluctuations. Adding in a baseline income level, A, to both components of yearly income, the entry decision for women becomes:

$$v(f(R_1)+A+P_1)-v(f(R_1)+A) \quad \geq \quad \beta E[v(f(R_2)+A+P_2)]-\beta E[v(f(R_2)+A)]$$

As A gets bigger the effect of R_1 on supply of brides is smaller in magnitude. Therefore if A is larger for the wealthy than the poor such that $A^w > A^p$, then $\frac{\partial S^p}{\partial R_1} < \frac{\partial S^w}{\partial R_1} < 0$. And on the demand side, $\frac{\partial D^p}{\partial R_1} > \frac{\partial D^w}{\partial R_1} > 0$.

A second source of heterogeneity results from women (or men) who, due to some personal characteristic or characteristics, face different price levels. As mentioned, young girls may have particularly high value in the marriage market. The ensuing shift in supply of brides, resulting from income shocks through rainfall, are therefore likely to be greater for households with younger girls than those of older women. [5]

The Trade Framework

The model thus far suggests a shift in demand for and supply of brides in response to rainfall shocks. In short, I assume that there will be an increase in demand and a fall in supply of brides with greater rainfall. However, this does not necessarily mean that there will be a rise in men's marriages and a fall in women's marriages. Because rainfall is a correlated shock affecting both sides of the marriage market, and because every bride needs a husband (and vice-versa), we do not know whether total marriages will go up or down. However, in a trade framework, we can better understand the implications of this asymmetric response to rainfall shocks.

Under Autarchy

If marriage matches are entirely contained within areas experiencing the same rainfall shocks, we can consider a marriage market to be autarchic. This would be the case if search costs for a spouse from anywhere outside of one's immediate vicinity were prohibitively large. Or, if rainfall is so strongly correlated across space, that extending one's

[5] Another way of conceptualizing a higher bride price for younger women would be to incorporate an element of risk into her family's valuation of her future value. The risk of a young girl having an out-of-marriage relationship, getting pregnant, and losing her value on the marriage market is also often cited as a reason motivating parents to marry their daughters at a young age. This effect may not directly impact P_1, but households who just endured a negative shock to household income through low R_1 may become more risk averse to this possible negative outcome. This model does not account for risk-aversion, but a more developed model could and should include this as a consideration.

search outside it's relevant area is practically impossible, we might similarly consider the marriage market to be autarchic.

Under these circumstances, rainfall shocks lead to ambiguous outcomes for the likelihood of a bride or groom getting married. We can imagine that, following a weak rainfall, there is a rise in supply of brides but a simultaneous fall in demand. Figure 1, Panel A shows what these shifts would look like graphically. Bride price will unambiguously fall. However, with no options for marriage except with each other, the net effect on total marriages will depend on the relative elasticities of demand and supply with respect to rainfall. Panel A is drawn showing a larger supply than demand response and a resulting increase in the overall quantity of marriages from Q to Q'. However, it could have been drawn such that the demand response dominates the supply response leading to an overall fall in marriages.

With Free Trade

The opposite extreme would be where either rainfall shocks are extremely localized or where search costs are very low. Under these conditions, similar to those facing a small open economy in a trade model, the divergent responses of bride supply and demand will not be constrained by one another. Instead of the rise in demand for brides being counteracted by the corresponding fall in local supply, a limitless world supply of and demand for brides at an exogenously set world price means that shifts in men's marriage can occur independent of corresponding shifts among local women sand vice versa. In order to preserve an overall balance, it must be that $S(P) = D(P) + EX$. With rainfall acting as a shifter of both supply and demand, local supply of brides must be equal to local demand plus net export of brides. This should make clear that if supply falls and demand rises with rainfall, net exports must fall, meaning *imports* of brides must increase. We expect that $\frac{\partial EXP}{\partial R_1} < 0$. Figure 1, Panel B illustrates these shifts following weak rains.

By comparing these two extreme cases, we can imagine a continuum of marriage market "openness" and infer that an area's marriage response to correlated shocks should more closely resemble the predicted partial equilibrium responses in areas that have more open or integrated marriage markets, holding other factors equal.

Figure 1: Marriage Market Responses Following Low Rainfall

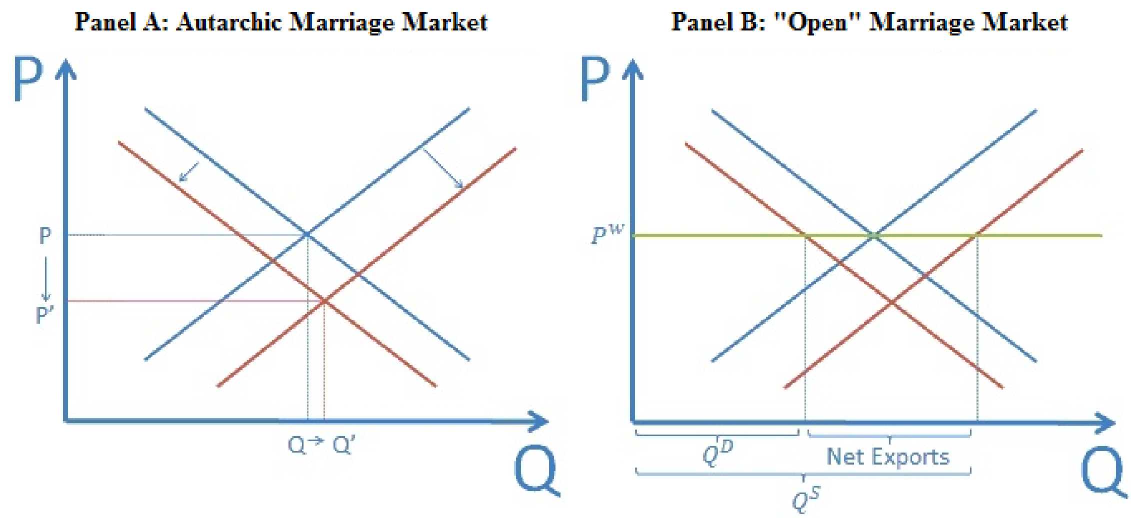

4 Empirical Strategy

Data Sources

This paper uses two main data sources. First, researchers from the University of Delaware have constructed a GPS coded panel of monthly historical rainfall data from 1900 to 2010.[6] Second, I utilize three rounds of the Burkina Faso Demographic and Household Survey (DHS) from 1993, 1998, and 2003.[7] Each survey round was designed to be nationally representative, containing individual interviews with thousands of women, men, and families. Using the GPS locations in the DHS data set, I can link the historical rainfall conditions from the (UD) data to each survey cluster in the DHS surveys.

These data sets have some notable limitations which should be acknowledged at the outset. The UD weather data are known to be extensively "smoothed". This is due to the necessity of extending best guesses of historical rainfall from a limited number of weather observation stations. This smoothing may have eliminated a considerable amount of meaningful geographic variation, however it still represents the best available estimate and contains considerable year to year variation which will serve as my source of year to year income shocks.

While the DHS is rich in data, it has its own limitations. The biggest shortcoming in these data for the present analysis is that while questions about women's duration in the village of the interview are asked, we do not know the specific location they arrived from. Since we can not link rainfall data to an unknown location and since we can also not track women who have left a village following marriage, all analyses of women's behavior are on women who married locally. For men, we do not have the same problem since men typically remain stationary following marriage, with their wives relocating into their households.

[6]See Willmott and Matsuura (2001).
[7]See http://dhsprogram.com/Data/.

Finally, mapping the GPS locations of the DHS surveys into a GIS system and superimposing a map of main roads in Burkina Faso constructed by DIVA-GIS,[8] I calculate the distance from each DHS survey cluster to the nearest major road. However, while proximity to a road is an important variable, there are two issues that should be noted. First, the map is designed to be mark Burkina Faso's most important roads, but there is undoubtedly considerable heterogeneity among these roads while many other important roads are likely to have fallen below the mapping threshold criteria. Second, DHS GPS coordinates are intentionally randomly displaced by 5 kilometers in order to preserve the anonymity of its respondents. However, because the displacement of these coordinates is done randomly, this perturbation will not systematically bias the results, though it means that they are a noisy measure.

Sample and Summary Statistics

Because the analysis focuses on the impact of rainfall on household outcomes through a hypothesized channel of agricultural income, the sample is restricted to respondents in rural locations. To mitigate recall error, observations are further restricted to men and women married for the first time within ten years of their DHS interview. Finally, in order to link these respondents to rainfall shocks, individuals who moved following marriage are dropped. The analysis is thus conducted on marriages that took place between 1983 and 2003. Across the three rounds of the DHS, this constitutes interviews with roughly 2,531 women and 885 men. The interviews are primarily focused on family health but cover a wide range of issues including household characteristics, fertility, and income sources.

Most data included in the DHS surveys was consistent between each round. However, some other variables, such as wealth indices, were not included in all rounds and were excluded from 1993 and 1998 survey rounds. Sample sizes in different specifications may therefore fluctuate depending on the availability of different variables.

In Burkina Faso, the mean rainfall between 1955-2010 is 112.5cm for the critical growing months of May-September with a standard deviation of 62.8cm. As calculated, 19.14% of year-location observations have rainfall of more than a standard deviation below their local mean and 17.88% experience rainfall one standard deviation or more above their local mean. Given that the Sahel is a notoriously drought stricken region, and that I saw no reports of Burkina Faso ever having experienced excessive or harmful rain, I consider higher rainfall to be unambiguously positive for farm output (and household income) while lower rainfall is interpreted as being worse for household output and income.

Table 1 presents summary statistics for women and men in the sample. I provide the number of observations for whom these data were collected in the right most column (revealing the inconsistency of survey content across rounds). Summary statistics confirm what we may have initially expected from anecdotal accounts of Burkina Faso. Women are married young, at an age of barely over 17, whereas men typically wait until they are over 24. Neither men or women seem particularly mobile in the years leading up to their marriage with 89% of women and 85% of men having always lived in their village.

[8] http://diva-gis.org

Polygamy is commonplace across the sample with 42% of women claiming to have other co-wives living in the household with them while only 14% of men in the sample report to be in polygamous households.[9] Average distance to the nearest road is 3.4 kilometers with distances ranging from zero to 18 kilometers. Education levels are strikingly low with only 14% of men and 8.5% of women reporting to have attended even a single year of school. Looking at actual primary school completion 6.6% of men completed primary while a mere 3.6% of women did so.

Identification and Data Treatment

I claim that the primary channel by which rainfall affects households is through its impact on year to year income. It is, however, still possible that rainfall affects local conditions in other significant ways, such as local infrastructure, that also impact local marriage markets. Additionally, without household income data from the actual years where the shock or shocks were experienced, rainfall can not be used as an instrument for exogenous changes in household income. However, given the dependence of the vast majority of Burkinabe on agriculture for their livelihoods, and limiting my analysis to rural populations, it is likely that the income effect has a first-order impact on peoples' livelihoods. Given these considerations and limitations, the following is a reduced form analysis where the interpretation will be the effect of *rainfall* on a marriage market and the ensuing likelihood of marriage following these shocks.

Valid identification relies on the randomness of locally experienced rainfall across time relative to its historical mean and variance in that location. Rainfall itself is coded in a number of ways in the analysis using the UD data. First, historical means and variances are calculated for the rainfall at each survey cluster location for the months that are critical to agricultural production in Burkina Faso, from May through September. Rainfall is coded as a positive shock if it is a standard deviation or more above the historical mean, and similarly as a negative shock when rainfall is more than one standard deviation below the location's historical mean. A continuous measure of rainfall is also used after being demeaned and normalized from local historical averages.

To look at marriage likelihoods, I use data from the DHS. In the DHS data, respondents report their year of first marriage, regardless of that marriage's ultimate outcome or whether additional brides were taken later. This provides information not only on the year that they *did* first get married, but also on all of the preceding years during which they *did not*. Therefore, I expand the data for each individual into a separate observation for each age of the respondent's life with a marriage outcome of zero until the year in which they are married, when the marriage indicator becomes one. In years after marriage, the respondent is no longer included in the data set since he or she is no longer

[9] This seeming mismatch between men and women's reports of household polygamy is likely due to two factors. First, for each polygamous man there will be at least two polygamous women, so this imbalance may be less incongruous than it initially sounds. Additionally, because I am restricting the sample to respondents whose first marriage was within 10 years of the interview in order to limit recall error, the sample is skewed towards men who are closer to their first marriage than the overall population. Since men may add additional wives as they get older, this imbalance is exacerbated when directly comparing this more recently married sample.

eligible for first marriage, creating an unbalanced panel of dates and marriage outcomes. The data is trimmed to remove observations with ages below the youngest 1% or above the top 1% of marriage ages observed in the data. For women, the relevant age range becomes 13-25, while for men it is 16-34.[10]

In addition to age, year, and rainfall patterns, the full specification includes location fixed effects, education, religion, and ethnicity controls. For men, I also include a wealth index calculated by the DHS project, placing men in wealth quintiles. To use these wealth indices as an independent variable in my regressions, one must be willing to assume that a man's household wealth at the time of the survey is a good proxy for his household wealth at the time of marriage. Because data is already trimmed to exclude marriages more than 10 years before the interview, we may be concerned if we expect large changes in income that take place within this medium-length window. Alternatively, it could be that the relevant wealth measure is that of a man's parents (or grandparents). Typically, I expect and make the assumption that wealth levels within 10 years or between generations between a man and his parents are highly correlated, or at least sufficiently well correlated to give an approximation of prior wealth quintiles.

While this may be reasonable assumption for men, a similar assumption of steady wealth levels is less appropriate for women. This is because women leave their homes following marriage in Burkina Faso, and assuming that her new household's income matches that of her parents is even more tenuous than it is for men. As a result, I only include the wealth index for men and am only able to look at heterogeneous wealth effects on the men's (demand) side of the market.

The outcome variable is the binary marriage indicator. The full specification is defined as:

$$Y_{ijt} = \alpha + \beta_0 R_{jt-1} + \mathbf{X'}_i \beta_1 + \beta_2 \gamma_{it} + \beta_3 \mu_j + \beta_4 \delta_t + \varepsilon_{ijt}$$

Where Y_{ijt} is a binary marriage indicator Y, of individual i, in village j, that took place in year t. R_{jt-1} is rainfall that took place prior to the year of that observation. In some regressions, multiple lagged rainfall measures are used simultaneously. $\mathbf{X}_i$ are relevant control variables of the individual such as religion, ethnicity, wealth, and education. γ_{it} is an age fixed effect for individual i at time t, μ_j is a set of location fixed effects, δ_t is a set of year fixed effects, and ε_{ijt} is an error term. All standard errors are clustered at the survey cluster level as this is also the level of observation for the rainfall data (the treatment).

[10] As an example of the data expansion process, imagine a woman in the data set who reports that she first got married in the year 1985, when she was 18. This woman will have one observation in the year 1985, with the age 18, and a binary marriage outcome of 1, because she got married in that year. However, we also have another observation for her, when she was 17, in 1984, and she did not get married. In this way, we expand each individual into a set of observations going back to the lower bound of ages (13 for women or 16 for men), when they can reasonably be thought to have entered the pool of eligible marriage market entrants. In this example we would have a separate observation for this same woman in the data six times, one for each age from 13 through 18. Each observation is then linked with rainfall patterns in the years preceding that date, regardless of whether he or she did or did not get married in that year.

5 Results

Base Specification

Table 2 shows the development of the full base specification for women. In this example the rainfall shock being estimated is for low rainfall that occurred two years prior to the observation. The outcome, as always, is a binary marriage indicator. Column (1) shows the simple regression of getting married in that observation year, regressed on having experienced low rainfall two years earlier in that location, and an intercept term. The sign is in the direction we might expect from the supply side partial equilibrium response: a fall in rainfall increases the likelihood of getting married. However, while the sign is as expected, the standard errors are relatively large and it is not statistically significant. We also note that this specification explains almost none of the variation in the data.

In column (2) geographic sample cluster fixed effects and observation year fixed effects are included in the regression. Year fixed effects are important because norms for age at first marriage have likely increased over the past few decades. Sample cluster fixed effects are needed to control for differences in marriage traditions or rainfall patterns across geographical areas. We see that now, the point estimate has increased to 0.0289 and is significant at the 95% confidence level. Additionally, we are now able to explain over 14% of the variation in the data. Columns (3) and (4) include variables to control for a woman's age in that observation. Column (3) uses both a linear and quadratic age term whereas column (4) uses age fixed effects. The point estimates are very similar to those in column (2), though we see that we are now explaining more of the variation in the data. Column (5) adds fixed effects for religion and ethnicity. The point estimate rises slightly, but is similar to those found earlier.

Finally, column (6) is the full base specification used for all women's regressions going forward, representing the bride supply side of the marriage market. In this specification low rainfall two years prior causes a 2.89 percentage point increase in a woman's likelihood of getting married from a base of 13.95%. This is equivalent to a 20.7% increase in a woman's likelihood of marriage, a considerable magnitude.

As discussed, the sample only includes women who were in the village prior to marriage. Therefore, it does not account for women who left the village following or for marriage. If the model is to be taken seriously, and low rainfall leads to an imbalance with excess local supply of wives, this number may be underestimating the impact of rainfall on marriage likelihood as women leave the village to marry men from other areas. It could also be that women who marry locally are distinct from women who marry outside their village of origin. A plausible correlation would be if poorer women tend to marry locally. If this is the case, and poorer women are more responsive to income shocks than wealthier ones, then this estimate will still be valid for the population in the sample, but may be upwardly biased relative to the national average. Data limitations prevent me from isolating these effects. These concerns are less critical for men who typically remain stationary following marriage.

The full specification for men is the same as that for women with two main differences. First, in addition to the covariates included for women, men's regressions include dummies for each of the wealth quintiles calculated from the DHS's wealth index. And second, because the DHS prioritized interviewing women, there is a much greater number of women in each survey cluster allowing me to use survey cluster fixed effects. However, given the relative sparseness of data for men, I instead use region fixed effects. The estimation results are similar regardless of which set of geographic fixed effects I include.

Table 3 shows the results from three separate regressions using five distributed lags of rainfall “treatments”. Column (1) contains demeaned normalized rainfall deviations for the preceding five years. Column (2) includes indicators for high rainfall from the preceding five years. Column (3) includes indicators for low rainfall shocks in the preceding five years. All regressions also include the full set of base covariates detailed earlier. At the bottom of each column we also see the sum of the five rainfall coefficients along with the standard error for this sum. All three look very close to zero and none are statistically significant. Looking at the individual lag coefficients for all three types of shocks, only one shows significance: the estimate for low rainfall from two years prior. Given our story about marriage decisions responding to recent income shocks, it is reassuring that we are not seeing high magnitudes or high significance in other years considerably further back in the past. If we were, we would be concerned that something else could be driving the result. However, a marriage response to a shock two years earlier seems potentially consistent with a story of supply response given the time a household needs in order to respond.

For the remainder of the paper, I focus on running regressions with a single rainfall measure (and any interactions) at a time.

Pooled Results

Next, Table 4 shows the coefficient estimates of the binary marriage outcome indicator regressed on a single rainfall shock and the full set of covariates, with results for women and men in Panels A and B respectively. Each column of each panel is a separately estimated regression where the heading of the column lists the “treatment”, or rainfall shock of interest. The binary marriage indicator is regressed on this rainfall measure along with the full base specification of covariates, as shown above in Column (6) of Table 2 (with wealth quintile fixed effects included for men). The estimates of these additional covariates are suppressed in order to economize space and to facilitate readability. “Devs” is the continuous measure of rainfall as demeaned standard deviations of rainfall for a given location in the previous year. “Devs 2” is the same measure, but for observed rainfall two years ago and “Devs 3” is this same shock for rainfall experienced three years earlier. “Low” and “High” are binary indicators for a standard deviation of rainfall below or above the historical mean.

Panel A focuses on women. The only coefficient that is statistically significant is from low rainfall two years prior. We can also confirm that the point estimates are similar to those we saw with the distributed lags in Table 3. Although they are not statistically significant, for the deviation and negative rainfall measures the sign of the estimated effects are consistent with the model’s partial equilibrium predictions. For the positive rainfall measures, the results, while not significant are opposite of the expected partial equilibrium response.

Panel B shows mixed results for men. I find a significant effect for men on rainfall that occurred three years earlier on the coefficients for both the low rainfall indicator and the deviated rainfall measure. These coefficients are consistent with the men's demand partial equilibrium response predictions. The binary low rainfall indicator, Low 3, suggests a 2.5 percentage point decrease in the likelihood of a man getting married whereas an additional standard deviation of rainfall increases his likelihood of getting married by just over one percentage point. These are very large treatment effects off of a mean outcome of 10.3%. However, while the signs of the other coefficients are mostly consistent with the predicted partial equilibrium demand response, the coefficients on Low and Low 2 have the opposite effect, where Low 2 has a fairly large magnitude of a 1.9 percentage point *increase* in marriage likelihood, significant at the 90% confidence level.

Interpretation of these findings is unclear. Men may begin saving for marriage early or their side of the bride search process may take multiple years, whereas women, who do not need to save in anticipation of marriage, are more influenced by more recent shocks. If this were the case, then we may see the men's partial equilibrium response dominate for longer lags while the women's response dominates for more recent shocks. Without further information about the savings behaviors and credit access of the population, it is not possible to explain the complete dynamic story of observed effects.

Age Heterogeneity

The impact of rainfall shocks on marriage timing is likely to have considerable heterogeneity by a daughter or son's age. Girls and boys have evolving roles in their respective households and different cultural norms may also influence when marriage is "appropriate". Or, women of different ages may demand different bride prices on the marriage market. To check for differential effects of rainfall by age, observations are divided into age groups. I create age groupings by gender to be balanced in the number of marriages in each group. For women, age brackets are 13-15, 16-17, 18-19, and 20-25 while for men they are 17-20, 21-24, 25-28, and 29-34.

Analysis of the men's sample does not show any significant heterogeneous effects (results are not included but can be requested from the author). Table 5 shows the results for women. The oldest age group, 20-25, is omitted. In order to avoid clutter, the table focuses on the results from the set of low rainfall shocks. For older women, low rainfall in any of the previous three years leads to lower marriage likelihood. Column (1) shows that low rainfall in the previous year lowers the likelihood of marriage by 4.71 percentage points significant at the 99% confidence level. Column (2) shows a negative coefficient for the two year lag with a point estimate of 1.09 percentage points, while Column (3) estimates a 3.34 percentage point decrease for the three year lag, again significant at the 99% confidence level. By contrast, the interacted coefficients for the youngest age bracket with the low rainfall shock is positive for all three lags, with significance at the 99% confidence level for the two and three year lags. These magnitudes more than counteract the negative effect estimated for the oldest age bracket, leading to a positive impact on marriage likelihood. The low rainfall shock from two years earlier results in a more than ten percentage point increase in marriage likelihood for women between the ages of 13 and 15, an immense jump. It is additionally noteworthy that for all three lags the coefficients for each age group get progressively smaller (and eventually becomes negative) as women get older. Finally, it should be noted that there is a large jump in estimated magnitude even between the second oldest and oldest age brackets.

These findings suggest that women who marry later may be different in many unobserved ways from those who marry at a more "typical" age, from 13-19. This may be because women of different ages serve different roles in their households. Younger girls may contribute less to household production and be more valuable to the household as a marriage market asset, while older women may be important for their productive value in the household and needed for their contribution to household labor following worse income shocks. Alternatively, it could be that instead of having dynamic differences in value strictly tied to age (both on the marriage market and within a woman's household), different women may have different inherent value in the marriage market. Richer data on women's characteristics could improve the analysis in order to better understand whether either or both of these effects are present.

Wealth Heterogeneity

Table 6 shows heterogeneous wealth effects for men using the continuous deviated rainfall measures. Instead of using all five wealth quintiles, I use a single binary indicator for men who are in the two bottom wealth quintiles, representing the poorer households in the sample. The results are puzzling. Theory would have suggested that demand should increase most for poorer people following positive rainfall. However, this is not what I find. It is the higher income group that responds more positively while the interaction with being a poor farmer has the opposite sign from what was anticipated.

There are many possible explanations for this surprising result. First, it may be that bride prices for marriages between wealthy families are considerably larger relative to wealthy family income than bride prices are among poorer households. For example, if bride prices demand large amounts of livestock and cash to be transferred among wealthy families whereas marriages among the poorest families only require a token payment or nothing at all, then we could see the observed larger response among wealthier families. This is not, however, consistent with the anthropological literature on Burkina Faso. A second possibility is that poorer households live in more isolated areas than those who are wealthier and we are seeing autarchic outcomes for these families whereby the supply response to these shocks is counteracting the demand response. Unfortunately, without more trustworthy income data for women, this can not be confirmed by checking for similar heterogeneous responses on the women's side.

Marriage Market Integration

The model outlined in the trade section suggests that we expect to see heterogeneous effects dependent on whether a given marriage market is more or less integrated well integrated with areas experiencing different rainfall or income shocks. As we saw in the open trade model, areas that have more integrated marriage markets will have a clearer separation of the local demand and supply partial equilibrium responses. Finding a good indicator of a local area's degree of marriage market "openness" or "closedness" is a challenge. I proxy for this missing parameter in two ways. First, I create a simple dummy for each survey cluster that indicates whether any women in that sample cluster moved into their village for marriage. This can be thought of as a revealed indicator of an open marriage market. Because this is a survey cluster level variable, I can no longer use survey cluster fixed effects and therefore replace them with regional fixed effects. These regressions are run for both men and women.[11]

Table 7 shows results for the women's and men's sides of the marriage market that are expected to be in relative surplus following a given shock. That is to say, I focus attention on women following low rainfall shocks under the hypothesis that there is excess supply, while for men, focus is on response to positive rainfall shocks. Under this hypothesis, we should see coefficients for interaction terms with statistical significance that are positive for both women (on low shocks) and men (on high shocks). In Panel A, for women, this is generally what we see. In particular, low rainfall from three years earlier seems to have a very large and negative effect for women. However, as predicted, this effect is entirely counteracted by the interaction term for low rainfall in open marriage markets. For men, in Panel B, the result runs counter to predictions. Men in closed marriage markets respond significantly and positively to positive income shocks. Though not significant, the sign and magnitude of the interaction between rainfall and being in a more open location seems to counteract this response instead of augmenting it as we might have expected.

These puzzling results may simply be the outcome of the crudeness of my marriage market integration measure. I only know whether a wife came from another village and not necessarily from outside the area experiencing the same rainfall shock. If small remote villages have a lot of marriages between them, but are in fact all close to one another and experience the same shocks, then my measure may be capturing the opposite of what I am attempting to estimate.

In a second attempt to proxy for marriage market integration, I use distance to the nearest road as a proxy for market "closedness". More remote areas, with less access to developed infrastructure will have higher travel costs. Where travel is more expensive, search costs for a spouse are likely to be higher as well, leading to greater marriage market segmentation. Because these measures are proxies for closedness, we now expect the sign on the interaction terms to be negative for low rainfall shocks experienced by women and high rainfall shocks experienced by men. Table 8 shows these results. All six of these regressions have the predicted sign. The coefficient for women in Panel A on the "Low" shock is significant at the 99% confidence level and suggests a decrease of 0.73 percentage points in the likelihood of marriage for each kilometer that a respondent is further from a main road. Panel B shows a very similar magnitude for men in Panel B on the "High" shock interaction term with a 0.71 percentage point decrease in marriage likelihood with every kilometer further from a main road, significant at the 95% confidence level.

Distance to a road is likely correlated with many things. In particular, proximity to roads is likely to correlate positively with wealth. However, if poorer communities are more responsive to shocks and also further from roads (as the model suggests), then es-

[11] Using a continuous measure of the proportion of women who moved into the village for marriage out of all marriages leads to similar results.

timates on road distance are biased, *underestimating* the magnitude of the observed interaction effects. Alternatively, living close to a road may correlate positively with quality financial institutions and available credit. However, this would, again, work against the identified effect whereby those closer to roads (and with more access to credit) should be less likely to respond to rainfall shocks. The potential presence of these factors make detection of a highly significant effect even more surprising.

It should also be noted that the shocks on the most recent rainfall become significant for both sides with closer proximity to a road. This estimated effect could be the result of the time lag between rainfall and market response being smaller in areas closer to roads where you have a higher volume of interpersonal interactions. Areas closer to a road may therefore lower the amount of search time needed to find a suitable partner. Of course there remain many other factors that could correlate with road distance and might have a strong effect. These results can therefore only be thought of as suggestive preliminary findings that marriage market integration and travel costs may affect marriage market responses to income shocks.

6 Robustness Checks

Because rainfall may have a considerable impact on many facets of a developing country's economy and populations' lives, it is important to try and refine and confirm that my hypothesized channels of impact are in fact those responding to rainfall in ways consistent with my story. As a first robustness check, I re-estimate my shock estimates for women but this time include all women in the sample from both urban and rural settings. Given that rural areas are likely to be more dependent on agriculture for their livelihood, we would expect to see a bigger response from interactions of rainfall with rural location than those in the urban settings. Table 9 shows these results. Most notably, the coefficient on Low 2, which had been significant in our initial regressions has fallen in magnitude from 2.89 to 2.18 and is no longer statistically significant. It is comforting that the signs are still in the same direction as before, suggesting that rural agricultural populations were driving the results. It is perhaps unsurprising that we have lost significance given that, even in urban areas, many households depend heavily on agriculture, either as a primary source of income or at least as a significant secondary source. It would have been concerning had we seen the opposite sign from what we expected.

As a second robustness check, national rainfall levels may be a good proxy for agricultural output price levels. If this is the case, then price levels in agricultural products might be lower following high average national rainfall due to increased supply and thus have a negative effect on household income. Local rainfall should still have the same predicted outcomes, even after controlling for this national rainfall impact on price levels. In Table 10, I include national yearly rainfall deviations as an additional regressor. I can no longer include year fixed effects because these dummies would absorb all the yearly variation in rainfall, so I instead include a linear and quadratic year term so that I can control for some of the time trends that are likely present. The coefficients on national rainfall are positive for the two and three year lags, as expected. With lower price levels, household yields are less valuable and women may be more inclined to enter the marriage market.

It is additionally encouraging that the estimate on the Low 2 local rainfall shock is still positive and significant. In fact, it is now indicating a 4.77 percentage point increase in marriage likelihood significant at the 99% confidence level. The Low 3 lag is now also significant at the 99% confidence level with an estimated effect of a 3.28 percentage point increase in marriage likelihood.

However, there are two puzzling findings. First, the High 3 coefficient is large and opposite of predictions. And second, the national rainfall deviations from the previous year are negative. However, the inclusion of national rainfall averages is potentially problematic. First, national rainfall averages are undoubtedly highly correlated with local rainfall measures making it unclear how much of the local income effect is now being re-categorized as a price effect. Second, national rainfall almost certainly affects overall national income, which, under my theory would also affect national bride demand. Following positive national rain, higher demand for brides in the form of a higher volume of suitors could be inducing more women to get married, a mechanism that would act in the same direction as a fall in agricultural produce price. It is thus ultimately unclear whether these results are confirming or refuting the mechanisms I have proposed in my model. Ideally, having local agricultural price data would be a better method of capturing this household income effect than using national rainfall levels.

7 Conclusion

This paper has presented evidence that rainfall patterns can affect the likelihood of a man or woman getting married in subsequent years. The model showed how supply of and demand for brides may react differently in response to rainfall shocks. There are a number of important dimensions of heterogeneity in this response that appear evident in the data such as wealth, age, or the degree of marriage market integration in that area. There is evidence that lagged rainfall affects both the supply and demand side of the marriage market. I find particularly strong evidence that women of different age groups respond in different ways to rainfall shocks, whereby young women are much more likely to get married following poor rainfall then older ones. However, heterogenous wealth effects among men seem to show the opposite effect from that predicted, whereby wealthier men are more influenced by year to year rainfall fluctuations. Finally, I find tentative evidence that higher marriage market integration may lead to more extreme responses to local rainfall shocks.

A better understanding of marriage decisions is important for many reasons. Marriages, at the foundation of most families, are the critical unit of analysis for most development outcomes. The conditions under which they are formed may have lasting implications for the marriage partners, the bride and groom's respective families of origin, as well as for future generations.

Marriage timing is just one component of these decisions, but it is particularly important in a setting like Burkina Faso where a large portion of women get married at very young ages. Early marriage leads to many long-term negative outcomes including higher levels of maternal and child mortality as well as fistula among young mothers. In this sample, less than 10% of women had ever used modern methods of contraception, shedding immediate light on Burkina Faso's national fertility rates of just under six children per woman on average. Exploding population growth is expected to be one of Burkina Faso's greatest development challenges, particularly given that the country's agricultural productivity and infrastructure are notoriously low. Knowing who is most vulnerable to early marriage and having a better sense of when and where this vulnerability is highest would be an important first step in targeting information campaigns or possibly organizing incentive programs to dissuade families from marrying their daughters at such young ages. This paper constitutes a useful framework for thinking about how marriage timing may be influenced by large correlated shocks and for who, where, and when rainfall shocks may have the most influence. Overall, my results raise more questions than they answer.

8 References

Bank, World (2013). "Burkina Faso: Country Brief". www.worldbank.org/U6XAAQ44R0.

D'Eveil Pugsada, Association D'Appui et (2000). *Le Mariage Précoce et Forcé de la Jeune Fille au Burkina Faso*.

Fafchamps, Marcel and Agnes Quisumbing (2004). "Assets at marriage in Rural Ethiopia". In: *FCND Discussion Paper* 185.

Fafchamps, Marcel, Christopher Udry, and Katherine Czukas (1998). "Drought and saving in West Africa: are livestock a buffer stock?"" In: *Journal of Development Economics* 98(55), pp. 273–305.

Hoogeveen, Hans, Bas van der Klaauw, and Gijsbert van Lomwel (2011). "On the Timing of Marriage, Cattle and Shocks". In: *Economic Development and Cultural Change* 60.1, pp. 121–154.

IRIN (2013). "Burkina Faso: Child Marriage Worsens Population Pressure". In: *Humanitarian News and Analysis.* www.irinnews.org/Report/83505/.

Jensen, Robert (2012). "Do Labor Market Opportunities Affect Young Women's Work and Family Decisions? Experimental Evidence from India". In: *Quarterly Journal of Economics* (127), pp. 753–792.

Kazianga, Harounan and Chris Udry (2006). "Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso". In: *Journal of Development Economics* (79), pp. 413–446.

Mbiti, Isaac (2008). "Monsoon Wedding? The Effect of Female Labor Demand on Marriage Markets in India".

Pacere, Maître Titinga Frédéric (1998). *"Pog-Poussoum et Marriage Coutumier chez les Mossé". Fondation PACERE*.

Rosenzweig, Mark and Oded Stark (1989). "Consumption Smoothing, Migration, and Marriage: Evidence from Rural India". In: *Journal of Political Economics* (97(4)), pp. 905–926.

Willmott, C.J. and K. Matsuura (2001). "Terrestrial Air Temperature and Precipitation: Monthly and Annual Time Series (1950 - 2010)". http://climate.geog.udel.edu.

Table 1: Summary Statistics

Panel A: Women's Sample

	Mean	SD	Min	Max	Obs
Age at First Marriage	17.239	2.558	10	40	2531
Husband's Age at Marriage	32.549	24.638	3	87	1667
Moved to Village for Marriage	0.000	0.000	0	0	2531
Always Lived in Village	0.891	0.312	0	1	2531
First or Only Wife at Marriage	0.673	0.469	0	1	2468
Polygamous Household	0.421	0.494	0	1	2468
Attended Any School	0.085	0.278	0	1	2531
Completed Primary	0.036	0.186	0	1	2531
Husband - Any School	0.088	0.283	0	1	2457
Distance to Nearest Road (KM)	3.390	3.162	0	18	2531
Mossi	0.441	0.497	0	1	2531
Islam	0.287	0.452	0	1	2531
Catholic	0.099	0.299	0	1	2531
Traditional	0.476	0.500	0	1	2531

Panel A: Men's Sample

	Mean	SD	Min	Max	Obs
Age at First Marriage	24.372	5.305	14	51	885
Always Lived in Village	0.849	0.359	0	1	885
Polygamous Household	0.139	0.347	0	1	861
Number of Wives	1.148	0.377	1	3	861
Attended Any School	0.141	0.348	0	1	885
Completed Primary	0.066	0.248	0	1	885
Distance to Nearest Road (KM)	4.265	3.448	0	18	884
Mossi	0.470	0.499	0	1	885
Islam	0.199	0.399	0	1	885
Catholic	0.043	0.203	0	1	885
Traditional	0.573	0.495	0	1	885

Table 2: Women's Likelihood of Marriage after Low Rainfall Shock

	(1)	(2)	(3)	(4)	(5)	(6)
Low Rainfall Two Years Ago	0.0104	0.0289**	0.0279***	0.0284***	0.0292***	0.0289***
	(0.0101)	(0.0115)	(0.0108)	(0.0106)	(0.0107)	(0.0107)
Age			0.2540***			
			(0.0162)			
Age Squared			-0.0058***			
			(0.0005)			
Attended Any School						-0.0353***
						(0.0133)
Muslim					-0.0321**	-0.0328**
					(0.0152)	(0.0151)
Catholic					0.0026	0.0014
					(0.0116)	(0.0115)
Traditional					-0.0116	-0.0155
					(0.0180)	(0.0181)
Intercept	0.1858***	-0.0340***	-2.2487***	0.0629***	0.1312***	0.1395***
	(0.0030)	(0.0097)	(0.1296)	(0.0089)	(0.0468)	(0.0441)
Sample Cluster FEs	No	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	No	Yes	Yes	Yes	Yes	Yes
Age Fixed Effects	No	No	No	Yes	Yes	Yes
Ethnicity Fixed Effects	No	No	No	No	Yes	Yes
Standard Errors	Cluster	Cluster	Cluster	Cluster	Cluster	Cluster
Mean of Dep Variable	0.1878	0.1878	0.1878	0.1878	0.1877	0.1877
Number of Observations	13129	13129	13129	13129	13083	13083
Sample	2499	2499	2499	2499	2499	2499
R Squared	0.0001	0.1432	0.2446	0.2713	0.2731	0.2735

Notes: Standard errors in parentheses and clustered at the survey cluster level. The dependent variable is a binary marriage outcome indicator. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Women's Likelihood of Marriage with Multiple Lags

	(1) Deviations	(2) High Shocks	(3) Low Shocks
1 Year Lag	-0.0015	-0.0036	-0.0109
	(0.0044)	(0.0117)	(0.0103)
2 Year Lag	-0.0011	0.0078	0.0285***
	(0.0041)	(0.0111)	(0.0108)
3 Year Lag	0.0000	0.0006	0.0011
	(0.0041)	(0.0108)	(0.0099)
4 Year Lag	-0.0073	-0.0161	-0.0087
	(0.0045)	(0.0122)	(0.0107)
5 Year Lag	0.0031	0.0046	-0.0025
	(0.0042)	(0.0108)	(0.0113)
Intercept	0.1341***	0.1447***	0.1435***
	(0.0462)	(0.0445)	(0.0443)
Mean of Dep Variable	0.1877	0.1877	0.1877
Sum of Coefficients	-0.0069	-0.0068	0.0075
Coeffs SE	0.0086	0.0268	0.0244
Observations	13083	13083	13083
Sample	2499	2499	2499
R Squared	0.2732	0.2732	0.2736

Notes: Standard errors in parentheses. The dependent variable is a binary marriage outcome indicator. Each column represents a separate regression that includes five lagged rainfall measures of the type listed at the top of the column. "Deviations" are locally demeaned and normalized continuous rainfall measures. "High" is a binary indicator for rainfall that was more than a standard deviation over its historical mean and "Low" is a binary indicator for rainfall that was more than a standard deviation below its historical mean. All regressions also include, religion, survey cluster, ethnicity, year, and age fixed effects and schooling variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Likelihood of Marriage after Different Rainfall Shocks

Panel A: Women's Supply Side Response

	(1) Devs	(2) Devs 2	(3) Devs 3	(4) High	(5) High 2	(6) High 3	(7) Low	(8) Low 2	(9) Low 3
Rainfall Measure	-0.0026	-0.0009	-0.0010	-0.0038	0.0083	0.0007	-0.0117	0.0289***	0.0001
	(0.0044)	(0.0040)	(0.0041)	(0.0118)	(0.0111)	(0.0108)	(0.0103)	(0.0107)	(0.0099)
Mean of Dep Variable	0.1877	0.1877	0.1877	0.1877	0.1877	0.1877	0.1877	0.1877	0.1877
Observations	13083	13083	13083	13083	13083	13083	13083	13083	13083
Sample	2499	2499	2499	2499	2499	2499	2499	2499	2499
R Squared	0.2730	0.2730	0.2730	0.2730	0.2730	0.2730	0.2731	0.2735	0.2730

Panel B: Men's Demand Side Response

	(1) Devs	(2) Devs 2	(3) Devs 3	(4) High	(5) High 2	(6) High 3	(7) Low	(8) Low 2	(9) Low 3
Rainfall Measure	0.0075	0.0008	0.0104**	0.0111	0.0165	0.0144	0.0029	0.0192*	-0.0251**
	(0.0047)	(0.0049)	(0.0048)	(0.0145)	(0.0146)	(0.0154)	(0.0106)	(0.0108)	(0.0100)
Mean of Dep Variable	0.1036	0.1036	0.1036	0.1036	0.1036	0.1036	0.1036	0.1036	0.1036
Observations	6744	6744	6744	6744	6744	6744	6744	6744	6744
Sample	735	735	735	735	735	735	735	735	735
R Squared	0.1957	0.1954	0.1961	0.1955	0.1956	0.1955	0.1954	0.1958	0.1961

Notes: Standard errors in parentheses. Each column of each panel represents a separate regression where the dependent variable is a binary marriage outcome indicator. "Rainfall Measure" is the independent variable listed at the top of the column. "Devs" is a locally demeaned, normalized measure of rainfall. "High" and "Low" are binary indicators of rainfall having been at least a full standard deviation above or below the historical mean, respectively. The number following the type of rainfall indicates that it is a lag from two or three years earlier. All regressions also include, religion, ethnicity, year, and age fixed effects and schooling variables. Men's regressions include region fixed effects and wealth quintile dummies while women's regressions include survey cluster fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Women's Marriage by Age

	(1) Low	(2) Low 2	(3) Low 3
Rainfall	-0.0471***	-0.0109	-0.0334*
	(0.0160)	(0.0169)	(0.0184)
Rainfall*13-15	0.0520	0.1117***	0.1327***
	(0.0386)	(0.0409)	(0.0424)
Rainfall*16-17	0.0442**	0.0305	0.0266
	(0.0202)	(0.0203)	(0.0219)
Rainfall*18-19	0.0443**	0.0569***	0.0298
	(0.0203)	(0.0204)	(0.0232)
Age 13-15	0.1389***	0.1297***	0.1272***
	(0.0145)	(0.0143)	(0.0138)
Age 16-17	0.0028	0.0056	0.0061
	(0.0105)	(0.0104)	(0.0100)
Age 18-19	-0.0390***	-0.0413***	-0.0366***
	(0.0110)	(0.0109)	(0.0103)
Intercept	0.0534	0.0457	0.0539
	(0.0455)	(0.0462)	(0.0455)
Mean of Dep Variable	0.1877	0.1877	0.1877
Number of Observations	13083	13083	13083
Sample	2499	2499	2499
R Squared	0.1566	0.1576	0.1573

Note: Standard errors in parentheses. Each column represents a separate regression where the dependent variable is a binary marriage outcome indicator. "Rainfall" is the independent variable listed at the top of the column. "Rainfall" is then interacted with different age brackets with 20-25 as the ommitted age category. All regressions also include, religion, ethnicity, year, survey cluster and age fixed effects and schooling variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Men's Marriage by Poorest Quintile Interaction

	(1) Devs	(2) Devs 2	(3) Devs 3
Rainfall	0.0061	0.0046	0.0193**
	(0.0066)	(0.0068)	(0.0077)
Poor*Rainfall	0.0059	-0.0142*	-0.0201**
	(0.0082)	(0.0080)	(0.0089)
Poor	-0.0160*	-0.0163*	-0.0175*
	(0.0096)	(0.0097)	(0.0100)
Intercept	0.1041*	0.1057*	0.0865
	(0.0561)	(0.0566)	(0.0568)
R Squared	0.2501	0.2502	0.2512
Mean of Dep Variable	0.1028	0.1028	0.1028
Number of Observations	3815	3815	3815
Sample	411	411	411

Note: Standard errors in parentheses. Each column of each panel represents a separate regression where the dependent variable is a binary marriage outcome indicator. "Rainfall" is the independent variable listed at the top of the column. "Devs" is a locally demeaned, normalized measure of rainfall. "Poor*Rainfall" is an interaction term between the rainfall measure of interest and an indicator for a respondent being in either of the bottom to wealth quintiles. All regressions also include, religion, ethnicity, year, region and age fixed effects and schooling variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Excess Supply/Demand Response with "Openness" by Past Marriages

	Panel A: Women (Supply)			Panel B: Men (Demand)		
	(1) Low	(2) Low 2	(3) Low 3	(4) High	(5) High 2	(6) High 3
Rainfall	-0.0613	-0.0022	-0.0958**	0.0837**	0.1005	0.0976
	(0.0403)	(0.0349)	(0.0427)	(0.0396)	(0.0620)	(0.0836)
Rainfall*Open	0.0602	0.0323	0.0992**	-0.0673	-0.1041	-0.0755
	(0.0412)	(0.0357)	(0.0426)	(0.0424)	(0.0634)	(0.0842)
Any Wives from Outside	-0.0259*	-0.0242	-0.0303**	-0.0214	-0.0165	-0.0252
	(0.0145)	(0.0152)	(0.0146)	(0.0296)	(0.0265)	(0.0266)
R Squared	0.2278	0.2283	0.2280	0.2509	0.2509	0.2510
Mean of Dep Variable	0.1877	0.1877	0.1877	0.1028	0.1028	0.1028
Number of Observations	13083	13083	13083	3815	3815	3815
Sample	2499	2499	2499	738	738	738

Notes: Standard errors in parentheses. Each column represents a separate regression where the dependent variable is a binary marriage indicator. "Rainfall" is the independent variable listed at the top of the column. "Rainfall*Open" is the interaction of that rainfall type with a binary indicator for whether any brides in the sample from that cluster came from outside of their village. All regressions also include, religion, ethnicity, year, region, and age fixed effects and schooling variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Excess Supply/Demand Response with "Closedness" by Road Distance

	Panel A: Women (Supply)			Panel B: Men (Demand)		
	(1) Low	(2) Low 2	(3) Low 3	(4) High	(5) High 2	(6) High 3
Rainfall	0.0212	0.0344^{**}	0.0036	0.0521^{**}	0.0140	0.0311
	(0.0144)	(0.0163)	(0.0134)	(0.0222)	(0.0242)	(0.0290)
Rainfall*Road Distance	-0.0073^{***}	-0.0016	-0.0010	-0.0071^{**}	-0.0025	-0.0010
	(0.0028)	(0.0037)	(0.0027)	(0.0036)	(0.0037)	(0.0037)
Dist to Road (KM)	-0.0018	-0.0028^{**}	-0.0028^{**}	-0.0003	-0.0014	-0.0018
	(0.0013)	(0.0012)	(0.0013)	(0.0015)	(0.0013)	(0.0014)
R Squared	0.2285	0.2286	0.2280	0.2511	0.2498	0.2502
Mean of Dep Variable	0.1877	0.1877	0.1877	0.1028	0.1028	0.1028
Number of Observations	13083	13083	13083	3815	3815	3815
Sample	2499	2499	2499	738	738	738

Notes: Standard errors in parentheses. Each column represents a separate regression where the dependent variable is a binary marriage outcome indicator. "Rainfall" is the independent variable listed at the top of the column. "Rainfall*Road Distance" ' is the interaction of that rainfall type with that cluster's distance to the nearest road. All regressions also include, religion, ethnicity, year, region, and age fixed effects and schooling variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Women's Marriage with Rural Interaction

	(1) Dev	(2) Dev2	(3) Dev3	(4) High	(5) High2	(6) High3	(7) Low	(8) Low2	(9) Low3
Rainfall	0.0033	0.0032	0.0027	0.0024	0.0033	-0.0011	-0.0066	-0.0020	0.0009
	(0.0045)	(0.0041)	(0.0038)	(0.0125)	(0.0112)	(0.0123)	(0.0121)	(0.0110)	(0.0109)
ruralshock	0.0001	-0.0014	-0.0025	0.0072	0.0073	0.0034	-0.0052	0.0216	-0.0011
	(0.0053)	(0.0049)	(0.0048)	(0.0154)	(0.0143)	(0.0150)	(0.0144)	(0.0135)	(0.0125)
rural	0.0116	0.0116	0.0114	0.0096	0.0098	0.0102	0.0118	0.0078	0.0109
	(0.0078)	(0.0077)	(0.0076)	(0.0082)	(0.0081)	(0.0080)	(0.0077)	(0.0077)	(0.0079)
Intercept	0.1339***	0.1375***	0.1371***	0.1364***	0.1363***	0.1368***	0.1363***	0.1380***	0.1369***
	(0.0154)	(0.0152)	(0.0151)	(0.0151)	(0.0151)	(0.0151)	(0.0150)	(0.0152)	(0.0151)
R Squared	0.1952	0.1951	0.1951	0.1952	0.1952	0.1951	0.1952	0.1953	0.1951
Mean Dep Var	0.1684	0.1684	0.1684	0.1684	0.1684	0.1684	0.1684	0.1684	0.1684
Num Obs	21560	21560	21560	21560	21560	21560	21560	21560	21560
Sample	3759	3759	3759	3759	3759	3759	3759	3759	3759

Notes: Standard errors in parentheses. Each column of each panel represents a separate regression where the dependent variable is a binary marriage outcome indicator. "Rainfall Measure" is the independent variable listed at the top of the column. "Devs" is a locally demeaned, normalized measure of rainfall. "High" and "Low" are binary indicators of rainfall having been at least a full standard deviation above or below the historical mean, respectively. The number following the type of rainfall indicates that it is a lag from two or three years earlier. All regressions also include, religion, ethnicity, year, survey cluster, and age fixed effects and schooling variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Women's Likelihood of Marriage after Different Rainfall Shocks with National Rainfall

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Devs	Devs 2	Devs 3	High	High 2	High 3	Low	Low 2	Low 3
Rainfall Measure	-0.0032	-0.0008	-0.0002	-0.0100	0.0136	0.0250**	-0.0119	0.0477***	0.0328***
	(0.0043)	(0.0045)	(0.0046)	(0.0111)	(0.0115)	(0.0115)	(0.0100)	(0.0107)	(0.0094)
National Rain Deviations	-0.0115***	0.0213***	0.0199***	-0.0119***	0.0186***	0.0159***	-0.0151***	0.0273***	0.0244***
	(0.0043)	(0.0045)	(0.0044)	(0.0039)	(0.0039)	(0.0039)	(0.0037)	(0.0038)	(0.0039)
Intercept	3232.7325***	-0.0155	-0.0137	3213.1913***	-0.0163	-0.0155	3185.2360***	-0.0217	-0.0177
	(317.9027)	(0.0511)	(0.0517)	(316.5548)	(0.0509)	(0.0514)	(318.6834)	(0.0504)	(0.0513)
Mean of Dep Variable	0.1877	0.1877	0.1877	0.1877	0.1877	0.1877	0.1877	0.1877	0.1877
Number of Observations	13083	13083	13083	13083	13083	13083	13083	13083	13083
Sample	2465	2465	2465	2465	2465	2465	2465	2465	2465
R Squared	0.2574	0.2110	0.2105	0.2574	0.2111	0.2108	0.2574	0.2125	0.2112

Notes: Standard errors in parentheses. Each column of each panel represents a separate regression where the dependent variable is a binary marriage outcome indicator. "Rainfall Measure" is the independent variable listed at the top of the column. "Devs" is a locally demeaned, normalized measure of rainfall. "High" and "Low" are binary indicators of rainfall having been at least a full standard deviation above or below the historical mean, respectively. The number following the type of rainfall indicates that it is a lag from two or three years earlier. All regressions also include, religion, ethnicity, year, survey cluster, and age fixed effects and schooling variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Year fixed effects can no longer be included because of the inclusion of yearly national rainfall so linear and quadratic year terms are included instead.

Appendix A

Additional Appendices, Tables, and Figures for Chapter 1

Appendix A: Additional Tables and Figures

Figure A.1: Correlation of Win Amounts and Biggest Expenditure

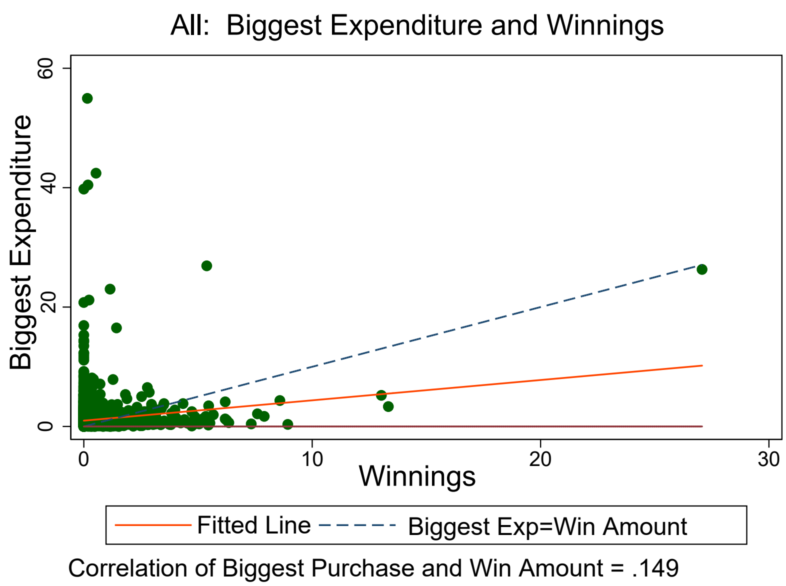

(a)

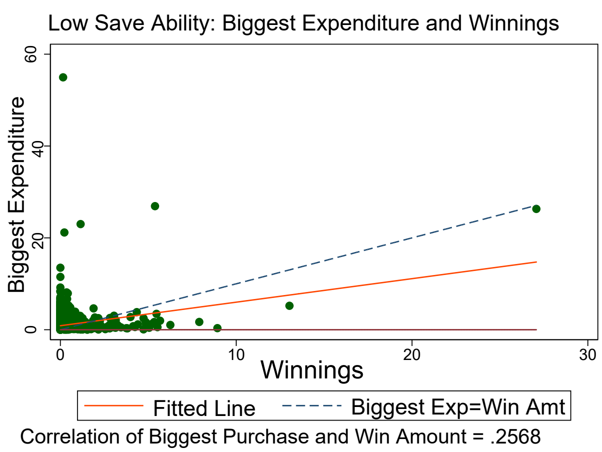

(b)

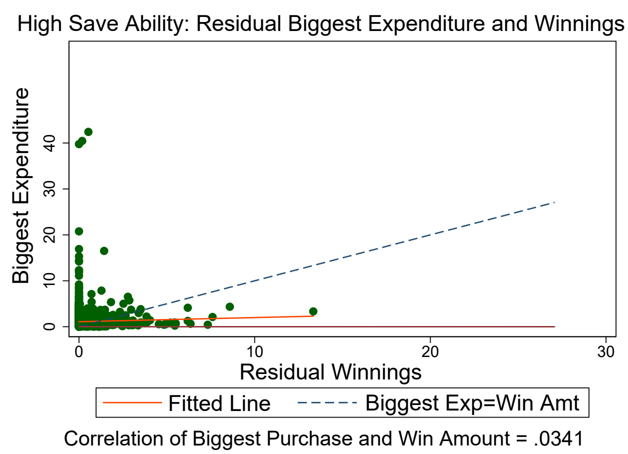

(c)

Notes: This figures shows a set of scatter plots and fitted regression lines of each respondent's win amount in each period and their largest expenditure in that time period, scaled by their mean income amount. In each panel, the dotted blue line has a slope of one, indicating what the relationship would be if the biggest expenditure in each time period were always equal to the amount won in that period. The orange line indicates a fitted regression line between biggest expenditure and the amount won. Panel (a) shows the data for all participants. The correlation is 0.149. Panel (b) shows this relationship only for people with a "low save ability", people who, relative to others at their income level, reported they could save only smal portions of their income. We notice that the correlation is considerably higher, at 0.2568. By contrast, Panel (c) shows this relationship for people with relatively high ability to save and the correlation now falls to just 0.0341.

Figure A.2: Robustness Check on Saving Return Thresholds and Patience for Different Payoffs

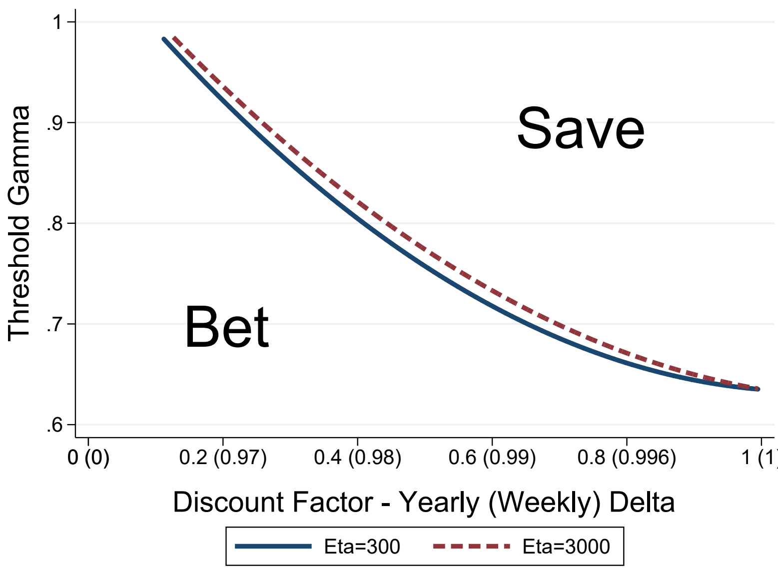

Notes: This graph shows that the results of the relationship between saving ability and patience are not sensitive to large differences in the relative size of foregone utility, ψ, and payoff to the large expenditure, η. Larger payoffs to the lumpy expenditure make people slightly less willing to tolerate expected losses of savings (shifts the treshold upwards) but this difference is minimal, $< 3\%$.

Table A.1: Validation of Revealed Preference Measure of Betting Demand

	Dependent Variable - Proportion of Tickets Demanded (0-1)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mean Bet Exps	0.006**				0.006**	0.005*		
	(0.003)				(0.003)	(0.003)		
Mean Bet/Mean Inc		0.216**						
		(0.090)						
Log(Mean Bet)			0.034***				0.036***	0.034***
			(0.011)				(0.011)	(0.011)
Log(Bet/Inc)				0.028***				
				(0.011)				
Mean Income					-0.000		-0.000	
					(0.000)		(0.000)	
Log(Mean Inc)						0.008		0.002
						(0.019)		(0.019)
Mean Dep Var	0.5662	0.5662	0.5662	0.5662	0.5662	0.5662	0.5662	0.5662
Week FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Num Obs	1980	1980	1979	1979	1980	1980	1979	1979
R2	0.0185	0.0186	0.0212	0.0194	0.0186	0.0186	0.0214	0.0212

Notes: This table examines the relationship between peoples' reported levels of betting and their revealed preference measure of betting demand solicited during the betting ticket offer. Regardless of functional form and whether or not income is included as an additional covariate, people who report to bet more also requested more betting tickets.

Trimmed top and bottom 1% of mean income and top 1% of saving ability and betting prop. Mean Bet Exps = Mean weekly betting expenditures during study or "typical" weekly betting expenditures. for respondents in condensed study. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.2: Listing Results

	Wave 1	Wave 2	Overall
Respondents	2,587	2,935	5,522
Ever Bet	59.3%	60.5%	59.9%
Bet Regularly	33.0%	31.1%	32.0%
Weekly Income (USD) - Mean	21.81	21.12	21.44
Weekly Income (USD) - Median	17.14	17.14	17.14
Weekly Bet Expenditures (USD) - Mean	4.09	4.47	4.28
Weekly Bet Expenditures (USD) - Median	2.00	2.86	2.86
Portion of Income Spent on Betting - Mean	25.5%	23.2%	24.4%
Portion of Income Spent on Betting - Median	14.0%	15.0%	14.3%

Table A.3: Summary Statistics - Financial Constraints

	Full	Condensed
Feel pressure to spend any extra money	27.0%	34.1%
Concerned family members may use money stored at home	23.7%	NA
Concerned thieves may take money from home	48.6%	58.6%
Have had money stolen from home	29.7%	NA
Have a bank account	41.5%	40.7%
Enrolled in mobile money	90.9%	88.8%
Could get a loan from a bank	46.2%	49.5%
Currently in debt	42.6%	23.0%
Median Outstanding Debt / Mean Income	1.4	1.4
Mean Outstanding Debt / Mean Income	3.9	3.7

Notes: "NA" indicates questions that were not included in the mini study survey. "Full" indicates responses of the 1,003 participants in the full study while "Condensed" indicates resonses from the 713 participants in the single-visit study.

Table A.4: Analysis Samples

Analysis	Samples	Reason
1) Usage of Winnings	Full Only	Panel betting and consumption data not part of Mini.
2) Lumpy Good Prime	All	Randomly chosen respondents in all groups received primes.
3) Prime on Savings	Condensed Only	Saving prime designed after full study completed.
4) Savings Boxes	Full / Wave 2	Two ways: cross-sectional (Full) and within (Wave 2).

Notes: Summary of empirical analyses and the sample of participants used along with a reason for this choice.

Table A.5: Effect of Winnings on Biggest Expense - Split Sample by Savings Ability

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	LSA-Med	HSA-Med	LSA-Big	HSA-Big	LSA-Bigger	HSA-Bigger	LSA-Huge	HSA-Huge
Win Amt / Mean Income	0.0279**	0.0121	0.0406***	-0.0053	0.0164	-0.0012	0.0156	-0.0023
	(0.0130)	(0.0132)	(0.0156)	(0.0142)	(0.0139)	(0.0104)	(0.0107)	(0.0086)
Income	0.0604***	0.0329**	0.0184	0.0542***	0.0034	0.0188	-0.0091	0.0069
	(0.0163)	(0.0131)	(0.0147)	(0.0140)	(0.0132)	(0.0120)	(0.0090)	(0.0095)
Bet Amt / Mean Inc	-0.1348*	0.0309	-0.0308	0.1182	0.0027	0.1496*	0.0321	0.0360
	(0.0788)	(0.0724)	(0.1006)	(0.0918)	(0.0762)	(0.0777)	(0.0380)	(0.0457)
Mean Y	0.6408	0.6912	0.2882	0.3271	0.0849	0.1015	0.0244	0.0306
Num Obs	2297	2354	2297	2354	2297	2354	2297	2354
Num Inds	476	478	476	478	476	478	476	478
R2	0.468	0.432	0.461	0.438	0.343	0.434	0.332	0.358

All specifications include betting portfolio moments, time, individual, and survey round fixed effects.
LSA=Low Saving Ability, HSA=High Saving Ability. Biggest expense thresholds: Med=.5*Inc, Big=Inc, Bigger=2*Inc, Huge=4*Inc

Notes: This table shows results from the regression: $Y_{i,t} = \beta_0 + \beta_1 W_{i,t} + \sum_{b=1}^{3} BetMoments_{i,t}^{b} + \lambda X_{i,t} + \gamma_i + \delta_t + \psi_s + \epsilon_{i,t}$. $Y_{i,t}$ 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 all also included. Standard errors are clustered at the individual level. Odd numbered columns (labeled "LSA") look only at respondents who are below the median in the amount of income they can save relative to others of similar income levels while the even numbered columns (labeled "HSA") look at those with relatively higher ability to accumulate savings.

Table A.6: Effect of Winnings on Value of Biggest Expenditure

	(1) All	(2) LSA	(3) HSA	(4) All	(5) All	(6) LSA	(7) HSA	(8) All
Win Amount / Income	0.330**	0.531***	0.034	0.040	0.220**	0.320**	-0.020	-0.014
	(0.160)	(0.187)	(0.059)	(0.057)	(0.088)	(0.127)	(0.049)	(0.049)
LSA * (Win/Income)				0.489**				0.419***
				(0.195)				(0.132)
Mean Y	1.0522	0.9897	1.1134	1.0522	1.0526	0.9889	1.1150	1.0526
Bet Moments	No	No	No	No	Yes	Yes	Yes	Yes
Income	No	No	No	No	Yes	Yes	Yes	Yes
Standard Errors	Indiv	Indiv	Indiv	Indiv	Indiv	Indiv	Indiv	Indiv
Num Obs	4669	2304	2363	4669	4653	2297	2354	4653
Num Inds	955	476	479	955	954	476	478	954
R2	0.319	0.361	0.302	0.326	0.328	0.380	0.309	0.332

LSA=Low Save Ability, HSA=High Save Ability.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows the un-trimmed regression results of the largest expenditure in a given period on the size of an individual's biggest win in that period. All specifications include individual, time, and survey round fixed effects. In columns 1-4 these regressions are unadjusted. Column (1) includes the full population. Column (2) is only people with relatively low ability to save relative to others of the same income level while Column (3) is those with high saving ability. Column (4) tests the difference between these two groups. Columns 5-8 repeat the same regressions again, but now include time varying income as well as the calculated "betting moments" which are the mean, variance, skewness, kurtosis, and higher order terms of a respondent's bets in that time period. We see that the regressions maintain their significance even while the magnitude falls by roughly 40%.

Table A.7: Effect of Winnings on Value of Biggest Expenditure - Drop People with Wins > 5 x Mean Inc

	(1) All	(2) LSA	(3) HSA	(4) All	(5) All	(6) LSA	(7) HSA	(8) All
Win Amount / Mean Inc	0.078	0.108**	0.037	0.043	0.063	0.111**	0.018	0.024
	(0.048)	(0.055)	(0.081)	(0.079)	(0.043)	(0.052)	(0.068)	(0.072)
LSA * (Win Amt/Inc)				0.066				0.073
				(0.097)				(0.096)
Mean Y	1.0342	0.9606	1.1064	1.0342	1.0345	0.9597	1.1080	1.0345
Bet Moments	No	No	No	No	Yes	Yes	Yes	Yes
Income	No	No	No	No	Yes	Yes	Yes	Yes
Standard Errors	Indiv	Indiv	Indiv	Indiv	Indiv	Indiv	Indiv	Indiv
Num Obs	4572	2261	2309	4572	4556	2254	2300	4556
Num Inds	935	467	468	935	934	467	467	934
R2	0.309	0.328	0.302	0.309	0.312	0.330	0.310	0.312

LSA=Low Save Ability, HSA=High Save Ability.
All specifications include individual, time, and survey round fixed effects.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.8: Effect of Winnings on Biggest Expenditure Size - Inverse Hyperbolic Sine Transformation

	(1) All	(2) LSA	(3) HSA	(4) All
Win Amt / Mean Inc (IHST)	0.052**	0.112***	-0.009	-0.001
	(0.024)	(0.036)	(0.031)	(0.030)
Win Amt / Mean Inc (IHST) x LSA				0.106**
				(0.045)
Two Week Income	0.042**	0.026	0.058**	0.042**
	(0.016)	(0.023)	(0.023)	(0.016)
Bet Amt	0.222***	0.016	0.183	0.208***
	(0.084)	(0.096)	(0.121)	(0.079)
Mean Y	0.7562	0.7234	0.7883	0.7562
Standard Errors	Indiv	Indiv	Indiv	Indiv
Num Obs	4653	2297	2354	4653
Num Inds	954	476	478	954
R2	0.477	0.482	0.485	0.478

LSA=Low Save Ability HSA=High Save Ability, $IHST(x) = log(x + (x^2 + 1)^{.5})$ All specifications include individual, survey round, and time fixed effects as well as moments of an individual's betting profile in that time period and higher order terms. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

This table shows an additional robustness check by using an inverse hyperbolic sine transformation. This allows for the inclusion of observations with winning values of zero without over-weighting observations with particularly large winnings. Column (1) shows the pooled results. Column (2) shows the results only among those with relatively low ability to save compared to others with similar income levels and column (3) looks only at those with high saving ability. Column (4) tests for difference in response beween these two groups. Again, we see that there is a significant effect for those with low saving ability but no discernible effect for those with higher ability to save.

Table A.9: Effect of Saving Box Treatment on Betting Demand

	(1) Max	(2) Max	(3) Max	(4) LSA	(5) HSA	(6) IV	(7) IV-LSA	(8) IV-HSA
Savings Box	-0.0706**	-0.0773**	-0.0727**	-0.0799	-0.0723	-0.1475**	-0.1551	-0.1516
	(0.0329)	(0.0346)	(0.0351)	(0.0498)	(0.0506)	(0.0698)	(0.0945)	(0.1028)
Lumpy Prime		0.0991***	0.0920**	0.1484***	0.0356	0.0931**	0.1556***	0.0336
		(0.0364)	(0.0368)	(0.0521)	(0.0530)	(0.0365)	(0.0509)	(0.0516)
Mean Dep Var	0.4185	0.4185	0.4185	0.4188	0.4179	0.4185	0.4179	0.4191
Control Group Mean	0.4443	0.4443	0.4443	0.4433	0.4451	0.4443	0.4433	0.4451
Time FEs	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Treats	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Covariates	No	No	Yes	Yes	Yes	Yes	No	No
Standard Errors	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust
Num Obs	939	939	939	468	469	939	469	470
R2	0.0048	0.0276	0.0495	0.1111	0.0824	0.0394	0.1013	0.0782
Adj R2	0.0037	0.0140	0.0159	0.0480	0.0173	0.0044	0.0332	0.0085

Notes: Dependent variable is an indicator for demanding the maximum number of tickets in the betting ticket offer. Results from regression of $B_i = \beta_0 + \beta_1 SaveBox_i + \lambda X_i + \delta_t + \epsilon_i$. $SaveBox_i$ is an indicator for having been randomly assigned to receive the saving box one month prior to the endline betting demand measure. 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. 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 to have used a saving box over the preceding four weeks. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.10: Takeup of Saving Box (IV-First Stage)

	(1) Both	(2) LSA	(3) HSA
Savings Box Treatment	0.5292***	0.5185***	0.5339***
	(0.0289)	(0.0419)	(0.0402)
Constant	0.1610***	0.1382***	0.1857***
	(0.0152)	(0.0198)	(0.0233)
Control Mean	.161	.138	.186
Mean Dep Var	0.3610	0.3213	0.4009
Standard Errors	Robust	Robust	Robust
Num Obs	939	470	469
R2	0.2854	0.2816	0.2855

Notes: Outcome variable is whether, at endline, respondent said that he had used a saving box over the previous month. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.11: Savings Box Treatment: Total Tickets Demanded

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	FE	FE	LSA	HSA	FE	IV	IV-LSA	IV-HSA	IV
Savings Box	-0.2785	-0.2765	-0.5758**	0.0662	-0.5081**	-0.5329	-1.1702**	0.1264	-0.8307**
	(0.1715)	(0.1721)	(0.2500)	(0.2456)	(0.2038)	(0.3327)	(0.5187)	(0.4694)	(0.3510)
Sav Box x High Sav Abil					0.4989*				0.6657*
					(0.2591)				(0.3705)
Log Income		-0.0561	0.0007	-0.1176	-0.0632	-0.0560	0.0089	-0.1171	-0.0620
		(0.0740)	(0.0988)	(0.1145)	(0.0737)	(0.0747)	(0.0989)	(0.1141)	(0.0744)
Urgent Expense		0.1104	-0.0316	0.1974	0.1099	0.1306	0.0166	0.1931	0.1318
		(0.1971)	(0.3470)	(0.2173)	(0.1932)	(0.2034)	(0.3756)	(0.2176)	(0.2032)
Cost of Needed Expense		-0.0968***	-0.0821***	-0.5484	-0.0949***	-0.1040***	-0.1003***	-0.5552	-0.1015***
		(0.0213)	(0.0235)	(0.8816)	(0.0210)	(0.0231)	(0.0288)	(0.8837)	(0.0225)
Mean Dep Var	2.5342	2.5342	2.5918	2.4840	2.5342	2.5342	2.5918	2.4840	2.5342
Adjusted Control Mean	2.4956	2.4956	2.7468	2.2258	2.4956	2.4956	2.7468	2.2258	2.4956
Num Obs	994	994	490	500	994	994	490	500	994
Num Indivs	497	497	245	250	497	497	245	250	497
R2	0.6509	0.6520	0.6604	0.6494	0.6549	0.6491	0.6490	0.6498	0.6521
Adj R2	0.2897	0.2875	0.2903	0.2681	0.2919	0.2816	0.2665	0.2688	0.2862

Notes: This table provides shows the same main results presented in the main paper while switching the outcome variable from maximum tickets demanded to the raw number (1-4). With a lower powered outcome, significance is reduced and lost on the pooled estimate in Column (1) although it remains highly significant for respondents with low ability to save.

LSA=Low Save Ability, HSA=High Save Ability, Adjusted control mean = baseline mean from treatment group plus time trend from control. All specifications include time and individual fixed effects and control for other treatments. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.12: Effect of Saving Box Treatment on Total Tickets Demanded (Cross-Sectional)

	(1) Max	(2) Max	(3) Max	(4) LSA	(5) HSA	(6) IV	(7) IV-LSA	(8) IV-HSA
Savings Box	-0.0775	-0.0915	-0.0922	-0.0915	-0.1009	-0.1965	-0.1738	-0.2256
	(0.1096)	(0.1129)	(0.1142)	(0.1652)	(0.1617)	(0.2258)	(0.3110)	(0.3276)
Lumpy Prime		0.1733	0.1474	0.3400*	-0.0452	0.1532	0.3627**	-0.0453
		(0.1182)	(0.1178)	(0.1745)	(0.1643)	(0.1162)	(0.1715)	(0.1587)
Mean Dep Var	2.4675	2.4675	2.4675	2.4509	2.4861	2.4675	2.4456	2.4894
Control Group Mean	2.5374	2.5374	2.5374	2.5142	2.558	2.5374	2.5142	2.558
Time FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Treats	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Covariates	No	No	Yes	Yes	Yes	Yes	No	No
Standard Errors	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust
Num Obs	939	939	939	468	469	939	469	470
R2	0.0249	0.0356	0.0742	0.0995	0.1231	0.0714	0.1045	0.1211
Adj R2	0.0197	0.0221	0.0415	0.0355	0.0609	0.0375	0.0365	0.0546

Notes: This table provides shows the same main results presented in the main paper while switching the outcome variable from maximum tickets demanded to the raw number (1-4). With a lower powered outcome, the results are no longer statistically significant although the sign of the effects is preserved.

LSA=Low Save Ability, HSA=High Save Ability, Adjusted control mean = baseline mean from treatment group plus time trend from control. All specifications include time and individual fixed effects and control for other treatments. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.13: Effect of Saving Box on Maximum Ticket Demand w/ Lumpy Prime Interactions

	(1) Max	(2) Max	(3) Low Save Ability	(4) High Save Ability	(5) Max
Savings Box	-0.0727**	-0.0966**	-0.0525	-0.1556**	-0.1280**
	(0.0351)	(0.0438)	(0.0622)	(0.0626)	(0.0587)
Lumpy Prime (LP)	0.0920**	0.0698	0.1780**	-0.0465	0.0706
	(0.0368)	(0.0475)	(0.0700)	(0.0658)	(0.0476)
SB x LP		0.0591	-0.0659	0.2126**	0.0932
		(0.0698)	(0.0965)	(0.1013)	(0.0902)
SB x LP x LSA					-0.0749
					(0.1054)
SB x LSA					0.0656
					(0.0786)
Low Save Ability					-0.0162
					(0.0475)
Mean Dep Var	0.4185	0.4185	0.4188	0.4179	0.4185
Control Group Mean	.4443	.4443	.4433	.4451	.4443
Standard Errors	Robust	Robust	Robust	Robust	Robust
Number of Obs	939	939	468	469	939
R2	0.0495	0.0479	0.1121	0.0917	0.0509
Adj R2	0.0159	0.0153	0.0468	0.0251	0.0130

Notes: Serves as a robustness check on the cross-sectional results to confirm that the effect of the saving box is not driven by the positive correlation of this treatment with receipt of the lumpy expenditure prime. The dependent variable is maximum tickets demanded (0/1). All specifications include time fixed effects, full set of covariates, and control for other treatments.

Table A.14: Effect of Saving Box on Log Liquidity Available

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	FE	FE	LSA	HSA	FE	IV	IV-LSA	IV-HSA	IV
Savings Box	-0.0215	-0.0246	0.1358	-0.2044	0.1370	-0.0486	0.2808	-0.3984	0.1773
	(0.1213)	(0.1218)	(0.1797)	(0.1693)	(0.1455)	(0.2402)	(0.3703)	(0.3285)	(0.2524)
Sav Box x High Sav Abil					-0.3461**				-0.4964**
					(0.1750)				(0.2486)
Log Income		-0.0503	-0.0121	-0.0642	-0.0552	-0.0515	0.0026	-0.0686	-0.0523
		(0.0613)	(0.1613)	(0.0597)	(0.0625)	(0.0624)	(0.1609)	(0.0615)	(0.0640)
Mean Dep Var	3.3823	3.3823	3.2375	3.5159	3.3823	3.3823	3.2375	3.5159	3.3823
Adjusted Control Mean	3.3205	3.3205	3.0864	3.583	3.3205	3.3205	3.0864	3.583	3.3205
Num Obs	1008	1008	496	508	1008	1008	496	508	1008
Num Indivs	504	504	248	254	504	504	248	254	504
R2	0.7560	0.7562	0.7484	0.7639	0.7582	0.7561	0.7497	0.7658	0.7595
Adj R2	0.5047	0.5041	0.4811	0.5134	0.5071	0.5039	0.4838	0.5173	0.5097

LSA=Low Save Ability, HSA=High Save Ability, Adjusted control mean = baseline treatment group mean + time trend.
All specifications control for time, individula, and survey round fixed effects as well as other treatments.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table tests for the effect of the saving box on liquidity available. If greater available liquidity reduces the appeal of betting, this could explain part of the negative effect of the saving box on demand for betting.

Table A.15: Effect of Saving Box on On-Hand Liquidity

	Panel A: Effect on Having 5USD On-Hand								
	FE	FE	LSA	HSA	FE	IV	IV-LSA	IV-HSA	IV
Savings Box	-0.0314	-0.0332	-0.0833	0.0047	-0.0035	-0.0655	-0.1722	0.0092	-0.0161
	(0.0468)	(0.0468)	(0.0661)	(0.0676)	(0.0534)	(0.0922)	(0.1370)	(0.1318)	(0.0933)
Sav Box x High Sav Abil					-0.0637				-0.1087
					(0.0630)				(0.0912)
Log Income		-0.0287	-0.0095	-0.0281	-0.0296	-0.0303	-0.0186	-0.0280	-0.0305
		(0.0322)	(0.0534)	(0.0398)	(0.0327)	(0.0324)	(0.0541)	(0.0399)	(0.0334)
Mean Dep Var	0.7272	0.7272	0.6794	0.7717	0.7272	0.7272	0.6794	0.7717	0.7272
Adjusted Control Mean	0.7579	0.7579	0.79	0.7426	0.7579	0.7579	0.79	0.7426	0.7579
Num Obs	1008	1008	496	508	1008	1008	496	508	1008
Num Indivs	504	504	248	254	504	504	248	254	504
R2	0.6617	0.6624	0.7072	0.6129	0.6630	0.6627	0.7062	0.6127	0.6634
Adj R2	0.3132	0.3131	0.3961	0.2022	0.3130	0.3139	0.3941	0.2017	0.3139

	Panel B: Effect on Having 15 USD On-Hand								
	FE	FE	LSA	HSA	FE	IV	IV-LSA	IV-HSA	IV
Savings Box	0.0143	0.0175	-0.0257	0.0660	-0.0189	0.0345	-0.0531	0.1287	-0.0180
	(0.0417)	(0.0417)	(0.0558)	(0.0615)	(0.0515)	(0.0824)	(0.1158)	(0.1221)	(0.0896)
Sav Box x High Sav Abil					0.0780				0.1154
					(0.0596)				(0.0843)
Log Income		0.0509*	0.0766*	0.0434	0.0520**	0.0517*	0.0738	0.0448	0.0519**
		(0.0260)	(0.0458)	(0.0297)	(0.0257)	(0.0263)	(0.0457)	(0.0301)	(0.0258)
Mean Dep Var	0.1716	0.1716	0.1653	0.1791	0.1716	0.1716	0.1653	0.1791	0.1716
Adjusted Control Mean	0.1865	0.1865	0.1893	0.1838	0.1865	0.1865	0.1893	0.1838	0.1865
Num Obs	1008	1008	496	508	1008	1008	496	508	1008
Num Indivs	504	504	248	254	504	504	248	254	504
R2	0.6185	0.6212	0.6483	0.6032	0.6225	0.6194	0.6465	0.5867	0.6162
Adj R2	0.2254	0.2295	0.2746	0.1822	0.2306	0.2257	0.2709	0.1482	0.2177

LSA=Low Save Ability, HSA=High Save Ability, Adjusted control mean = baseline treatment group mean + time trend.
All specifications include time, individual, and survey round fixed effects and control for other treatmets.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table tests for the effect of the saving box on on-hand liquidity. By reducing the amount of liquidity carried by those treated with the saving box, the betting ticket offer will have a relatively high opportunity cost of daily expenditures. That is to say, the amount of cash offered may have greater appeal for people carrying less cash. Although not statistically significant, this reduction in holdings of cash at the low threshold level may suggest a mechansim through which the effect of the saving box reduces demand for betting.

Table A.16: Effect of Lumpy Prime on Proportion of Tickets Demanded

	(1) All	(2) All	(3) All	(4) All	(5) LSA	(6) HSA	(7) All
Lumpy Good Prime	0.042**	0.042**	0.042**	0.040**	0.079***	0.001	-0.001
	(0.019)	(0.019)	(0.019)	(0.019)	(0.028)	(0.027)	(0.027)
Prime x Low Save Ability							0.083**
							(0.038)
Low Saving Ability							-0.019
							(0.031)
Mean Week Bet			0.135	0.110	0.082	0.187*	0.109
			(0.085)	(0.082)	(0.118)	(0.110)	(0.083)
Liquidity Available			-0.000	-0.000	0.004	-0.004	-0.000
			(0.003)	(0.003)	(0.004)	(0.004)	(0.003)
Save Ability / Mean Inc			-0.011	-0.026	0.183	0.003	0.003
			(0.040)	(0.040)	(0.147)	(0.061)	(0.052)
Mean Week Income		0.001**	0.001**	0.002**	0.002**	0.002*	0.002***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Mean Dep Var	0.6211	0.6211	0.6211	0.6211	0.6292	0.6132	0.6211
Mean Y-Control	0.6008	0.6008	0.6008	0.6008	0.5928	0.6083	0.6008
Full Set of Covariates	No	No	No	Yes	Yes	Yes	Yes
Num Obs	1701	1701	1701	1701	842	859	1701
R2	0.0504	0.0527	0.0572	0.0824	0.1223	0.0944	0.0855
Adj R2	0.0373	0.0391	0.0409	0.0603	0.0796	0.0501	0.0623

All specifications control for other treatments and include price of ticket and time fixed effects.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table provides a robustness check for the lumpy prime treatment, switching to a measure of the proportion of betting tickets demanded (relative to those offered). The effect of the lumpy good prime on betting demand is still positive and significant at the 95% confidence level in all specifications, except for those with high saving ability (as in the original results). The test for difference between people with low and high saving ability remains signficant at the 95% confidence level.

Table A.17: Effect of Lumpy Good Prime on Max Ticket Demand with Saving Box Interaction

	(1) All	(2) All	(3) All	(4) All	(5) All	(6) LSA	(7) HSA	(8) All
Lumpy Prime	0.063**	0.064**	0.066**	0.065**	0.064**	0.121***	0.009	0.013
	(0.027)	(0.027)	(0.027)	(0.027)	(0.027)	(0.038)	(0.037)	(0.035)
Lumpy Prime x Saving Box	0.056	0.051	0.043	0.046	0.043	0.003	0.071	0.041
	(0.061)	(0.061)	(0.062)	(0.062)	(0.062)	(0.088)	(0.089)	(0.062)
Prime x Low Save Ability								0.103**
								(0.048)
Mean Income (USD)		0.002***	0.002***	0.002***	0.002***	0.002	0.004***	0.002***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Save Abil. / Mean Inc			-0.060	-0.061	-0.071	0.189	0.051	0.006
			(0.052)	(0.052)	(0.053)	(0.190)	(0.080)	(0.068)
Mean Y	0.453	0.453	0.453	0.453	0.453	0.476	0.430	0.453
Mean Y-Control Grp	0.418	0.418	0.418	0.418	0.418	0.420	0.416	0.418
Background/Educ Controls	No	No	No	Yes	Yes	Yes	Yes	Yes
Preference Controls	No	No	No	No	Yes	Yes	Yes	Yes
Num Obs	1703	1703	1703	1703	1703	844	859	1703
R2	0.032	0.036	0.041	0.042	0.053	0.088	0.72	0.057
Adj R2	0.018	0.022	0.024	0.022	0.029	0.042	0.025	0.033

Notes: This table provides a robustness check to ensure that the effect of the prime is not driven by positive correlation with the saving box treatment. Even while including an interaction with the saving box treatment, the effect of the lumpy prime remains positive and significant. All specifications also control for all other treatments as well as time and ticket price fixed effects. That is to say, the amount of cash offered may have greater appeal for people carrying less cash. Although not statistically significant, this reduction in holdings of cash at the low threshold level may suggest a mechansim through which the effect of the saving box reduces demand for betting. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.18: Robustness: Effect of Betting Ticket Offer on Saving Update

	(1) Raw	(2) Prop	(3) Positive	(4) Negative	(5) None
Saving Prime Before Ticket Offer	1.157	0.015	-0.006	0.032	-0.026
	(2.135)	(0.015)	(0.041)	(0.037)	(0.035)
N	708	708	712	712	712
Mean Dep Var	-4.5466	-0.0648	0.4775	0.2697	0.2528
R2	0.0003	0.0010	0.0000	0.0011	0.0007

Notes: This table provices a robustness check for the budgeting exercise. If conducting the budgeting exercise after the betting ticket offer affected the reported updates, the identification strategy would no longer be valid. These results show no evidence of the timing of the betting ticket offer affecting reported updates. Outcome variables:
Column (1) Raw Update = Assisted Estimate of Save Ability − Naive Save Ability
Column (2) Prop Update = Raw Update / Weekly Income
Column (3) Positive Update = Update > 0
Column (4) Negative Update = Update < 0
Column (5) None = Update equal to 0.
$p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.19: Effect of Saving Prime Time Update on Betting Demand

	(1)	(2)	(3)	(4)	(5)	(6)
	Max	Max	Max	Max	Max	Max
Lumpy Good Prime	0.115***	0.114***	0.114***	0.114***	0.115***	0.114***
	(0.043)	(0.043)	(0.043)	(0.043)	(0.043)	(0.043)
Savings Prime (SP)	-0.039	-0.180*	-0.190**	-0.217**	-0.183**	-0.286**
	(0.045)	(0.092)	(0.092)	(0.103)	(0.093)	(0.137)
SP x Time Update		0.049*	0.036	0.043	0.038	0.067
		(0.028)	(0.032)	(0.034)	(0.033)	(0.044)
Time Update		0.006	0.039*	0.037*	0.021	0.011
		(0.016)	(0.021)	(0.022)	(0.025)	(0.029)
SP x Time Update x Beta			0.030	0.017		
			(0.033)	(0.044)		
Time Update x Beta			-0.086**	-0.082**		
			(0.038)	(0.040)		
SP x Beta				0.070		
				(0.150)		
SP x Time Update x Delta					0.022	-0.034
					(0.034)	(0.076)
Time Update x Delta					-0.028	-0.010
					(0.037)	(0.045)
SP x Delta						0.198
						(0.236)
Beta	-0.023	-0.023	0.226*	0.195	-0.025	-0.025
	(0.077)	(0.077)	(0.135)	(0.154)	(0.077)	(0.077)
Delta	0.056	0.059	0.038	0.048	0.127	0.062
	(0.074)	(0.074)	(0.079)	(0.082)	(0.129)	(0.158)
N	680	680	680	680	680	680
Mean Dep Var	0.4221	0.4221	0.4221	0.4221	0.4221	0.4221
Full Set of Covariates	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.1457	0.1521	0.1585	0.1588	0.1532	0.1539
Adj R2	0.1158	0.1197	0.1236	0.1226	0.1181	0.1176

Time update is calculated based on the difference between their original guess of how long they would need to save for the lumpy expenditure and the amount of time it would take them to save after going through the budgeting exercise and making a more realistic assessment of how much they could save per week. These time update values are then converted to logs for this analysis so that a coefficient (such as on row (3)) suggests the effect a 100% increase in the amount of time needed to save.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendix B: Sports Betting Details

B1: Odds, Payouts, and Betting Structure

A ticket's payout depends on the choices made by the bettor. Each predicted outcome included on a ticket is associated with a payoff multiplier, such that less likely outcomes are rewarded with higher multipliers. In order to win a ticket, every outcome it includes must have been accurately predicted. Even a single incorrect guess causes the entire ticket to fail. If all predicted outcomes are correct, a bettor wins the product of the selected multipliers times the amount of money wagered on that ticket.

For example, a bettor could bet on specific outcomes (win, loss, or tie) for each of four different matches. If these predicted outcomes had associated multipliers of 1.5, 2, 2, and 5 then his total multiplier is 1.5 x 2 x 2 x 5 = 30. If he bets 2 USD on this ticket and all four outcomes occur, he can redeem his winning ticket for 60 USD. If any of his four predictions do not occur, his ticket becomes worthless.

B2: Estimating the Rate of Return

Estimating the rate of return is complicated and somewhat subtle. It ultimately depends on the number of games you put on your ticket and how much you are targeting. The likelihood that a given ticket will win goes down as the payout rises. However, a more subtle point is that the expected value of a ticket also goes down as tickets are added to the game and the target amount goes up. This is because the betting companies only offer payoff multipliers that are beneath "fair" payoffs. In Uganda, this is estimated as a reduction of roughly 7-8% off of each multiplier. This was estimated from offered match multipliers on over 1,000 bets in Uganda. Because odds are offered for all possible outcomes of an event, it is possible to back out the company's per-match expected earnings.

Betting companies invest a lot of money in having as good an estimate of what the "true" odds of different outcomes are.

This is best illustrated with another example. Imagine a match with associated multipliers of 1.8 if Team A wins, 5.5 if Team B wins, and 2.7 if they tie. If betting companies were offering what they thought to be fair multipliers this would imply that they think the real likelihood that Team A wins is $\frac{1}{1.8} = .5556$. In other words, they would believe that for a dollar spent on this single game wager, they have a 44.44% chance of keeping the dollar and a they have a 55.56% chance that they have to return the dollar and add \$0.8 yielding an expected return of zero. The other two multipliers imply likelihoods of 0.1818 and 0.37. Together, these three probabilities add up to 1.1074. Of course, this is impossible since one and only one of the three outcomes can occur. Therefore, they must be lowering the payouts they offer in order to build in their profit margin. Without knowing the true odds, I need to make the assumption that they shade different sides of the bet similarly. This example suggests that they offered multipliers that they knew should have been 10.74% bigger on average implying an average expected rate of return (no matter what side of the bet is chosen) of 0.903. On average across the bets in my data set, the expected payout is approximately 0.925 per dollar spent *for a single match prediction*.

However, in order to get a larger payout, bettors typically add multiple games onto their

Figure A.3

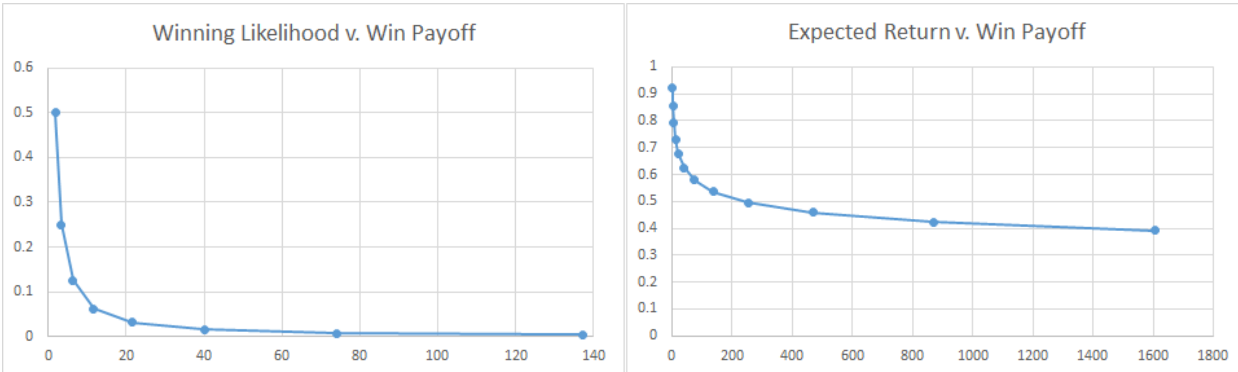

tickets. In fact, most betting companies require a minimum of three matches per ticket. Imagine an individual puts five games on his ticket each of which have a true winning likelihood of 50%. The betting company will shade down its offered multiplier on that outcome by 7.5% and offer 1.85 on these outcomes. Spending one dollar on a ticket and placing five of these matches on a single ticket means that an individual could win \$21.67 if all matches occur as predicted. The true likelihood of the ticket winning is $0.5^5 = 3.125\%$. The rate of return will be $21.67 * 0.03125 = 0.677$. This could have also been calculated with the per ticket adjustment factor of 0.925, so that the expected value of a ticket is a function of the number of games on the ticket and in this example equal to $0.925^5 = 0.677$.

If a bettor exclusively predicted outcomes with "true" likelihoods of 50% a fair bet would offer a multiplier of 2 for each of these outcomes. However, shading down, each offered multiplier is only 1.85. Figure A.3 The graphs below show what happens to the likelihood of winning (on the left) and to the expected return (on the right) for bets as targeted payoffs are increased by adding more games to the ticket. Each dot represents an additional game added to the ticket.

B3: Characterizing Weekly Betting Profiles

Being able to characterize these bi-weekly betting profiles requires an understanding of the structure of the betting system in Uganda. In particular, each betting ticket is defined by three main characteristics: a stake (the price of the ticket), a payout (the amount you can potentially win), and the number of matches you include on the same ticket (number of match multipliers). By knowing these three values for a ticket, I can infer a close approximation for the likelihood that a given ticket wins. Characterizing groups of bets (or weekly betting) with these descriptors allows me to similarly characterize the full distribution of possible outcomes for this set of bets.

First, I calculate what a "fair" payout would be based on the number of matches and the targeted payout of a ticket. As discussed above, I know that multipliers (per match payouts) are shaded downwards by 7.5% on average. This means that the total payout offered (the unfair payout) is equal to:

$$unfair\ \ payout = stake \times \Pi_i^G \underbrace{(fair\ \ multiplier_i \times 0.925)}_{offered\ \ \text{``}unfair'' \ \ multiplier}$$

G is the total number of games included on a single ticket. The product operator includes the offered multipliers for each prediction on the ticket. Bettors only observe this unfair multiplier which is scaled down on average by a factor of 0.925 from what a fair multiplier would be. To estimate what a fair payout would be I have to divide the unfair payout by $(1/0.925^G)$. Once I have calculated the fair payout for a ticket, I can then infer the "true" probability that a ticket will win as $probability = 1/fair\ \ multiplier$.

Although I do not know these values for every individual ticket that a respondent buys, respondents characterized their bets each week. With this information, I calculate each of the moments of a bettors' betting portfolio in a given period of observation. The moments are calculated for a set of n bets with likelihood, p, and a payout amount W.

- **Mean:** $np * W$
- **Variance:** $np(1-p) * W$
- **Skewness:** $(1-2p)/\sqrt{np(1-p} * W$
- **Kurtosis:** $(1-6p(1-p))/(np(1-p))$

The validity of this strategy assumes that given the characteristics of an individual's betting portfolio, the actual outcome is random. This relies on the assumption that bettors in Uganda do not know more than the international betting markets that set the associated odds and payoffs with the games being offered. Recent research by Goddard (2013) confirms that, although some arbitrage opportunities existed 12 years ago, they have nearly vanished today. With so much money at stake in the global sports betting industry, the advancement of data driven analytics and expansion of data availability mean that when arbitrage opportunities exist, they get quickly bid away by those with better information. This then resets the odds which are passed on to the local markets in Uganda. For someone to be able to perform consistently better than expectation would require them to know more about the offered bets than the market and in a way that is not captured by the types of bets they are placing and controlled for in the analysis.

An additional reason to be skeptical that local bettors are identifying subtle arbitrage opportunities is because bettors do not appear to be able to accurately understand the components of calculating joint probabilities or inferring what the offered probabilities imply about the relative likelihood of events. Just 27% of people in the sample understood that the likelihood of flipping two coins and getting double heads was 25%. Just 23% gave a correct answer that flipping three coins and getting all heads was between 10 and 15%. These results are essentially the same when the example given is linked to football teams on a betting ticket where you believe the outcome is 50% likely. 35% did not understand that adding a third coin must lower the overall likelihood of the joint outcome. Finally, just 28% (essentially no better than random with a choice set of three items) understood how to infer win probabilities from the odds as they are displayed and offered by betting shops. This does not rule out that there are a small set of highly sophisticated bettors, however it seems unlikely that very many are able to identify arbitrage opportunities missed by the international markets. In future work I am planning to look more closely on how bettors form their beliefs about different matches and how this may drive them towards certain types of bets.

Appendix C: Contrasting Saving and Betting

This section details the assumptions and derivation of the back of the envelope calculations discussed in Section 7. Section C1 explains the extension and setup of the model originally presented in Section 3. Section C2 details the approach to estimating return on saving.

C1: Balancing Patience and Return on Saving

Imagine that an individual is trying to raise money for a lumpy expenditure that will cost him P_L. The lumpy expenditure is assumed to be consumed immediately (such as a wedding) and the individual has an infinite time horizon.[1] He is trying to decide whether to devote a portion of his weekly income, S, to either saving or betting in order to try and make the expenditure. He will pursue either strategy until he is able to make his purchase, at which point he immediately enjoys its utility payoff, η and then spends the rest of his life consuming only divisible goods. As in the original model, his single period utility is $u(D) + L\eta$.

If he saves, he expects it will take him $K = \frac{P_L}{S\gamma}$ weeks to accrue the liquidity needed to purchase L, where γ is the rate of return on savings. Note that as γ falls, there is slippage in savings and the amount of time required to accrue the needed liquidity increases. For each period that he saves, he sacrifices some amount of consumption of the divisible good, D, and forgoes current utility of ψ such that $\psi \equiv u(Y) - u(Y - S)$. He anticipates a discrete utility payoff of η once he is able to make the lumpy expenditure and δ is his future discounting factor. Expected net utility from saving is therefore:

$$-\sum_{0}^{K} \delta^t \psi + \delta^K \eta = \frac{-\psi(1-\delta^K)}{1-\delta} + \delta^K \eta$$

If instead he chooses to bet, winnings from betting are assumed to be immediately used in their entirety on the lumpy good. The likelihood of winning his bet is p. If he wins, he enjoys η. If he does not, then he will bet again in the next (discounted) period where, again, the likelihood of winning is p. This suggests that expected utility from betting is therefore:

$$-\psi + p\eta + (1-p)\delta[-\psi + p\eta + \delta(1-p)[-\psi \quad ... \quad ... \quad = \frac{-\psi + p\eta}{1-(1-p)\delta}$$

This strategy becomes a geometric sum with the base value of $\psi + p\eta$ and the discount factor of $\delta(1-p)$.

[1] This could also be modeled with a durable good and generate similar results. Similarly, the assumption that he lives for infinite time periods could be relaxed and generate similar results as well.

Setting saving equal to betting becomes:

$$\frac{-\psi(1-\delta^K)}{1-\delta}+\delta^K\eta=\frac{-\psi+p\eta}{1-(1-p)\delta}$$

Next, I want to parameterize this equation in order to solve for the minimum return on saving, γ^*, given an individual's level of patience, δ, before he switches from saving to a betting strategy. Imagine an individual is trying to raise 200,000 shillings for a desired lumpy expenditure (approximately 65 USD or about 2.3 times the median weekly income in the sample just above the median reported winning targets). He can either set aside 10,000 shillings per week for saving or spend them on betting. Based on the types of bets that bettors typically make from the data, I can estimate the likelihood that a bettor would win a 10,000 UGX bet targeting a payout of 200,000 UGX. This likelihood is $p = 0.03125$.[2] The likelihood of losing is therefore $1-p = 0.96875$. As a simplifying assumption, I set the recurrent cost of sacrificing 10,000 shillings in divisible expenditures equal to ten, and assume that the payoff from the lumpy expenditure is large and higher in total value than the marginal utility lost across each time period (with perfect saving this would be 200 thousand shillings). I try low value for η of 300 and a high value of 3,000.

Next, I take a reasonable estimate for peoples' weekly discount factors from research by Mbiti and Weil (2013) in neighboring Kenya. In their population, the authors estimate a yearly discount factor of 0.64 suggesting a weekly discount factor of approximately 0.9915. For each η-δ pair, there is a value of γ, γ^* that will make individuals indifferent between saving and betting. High values of γ^* suggest that only very small expected losses of saving can be tolerated before betting is preferred while a lower γ^* suggests substantial expected losses can be endured. Figure 12 in the main paper shows these results in graphical form for weekly and yearly discount factors with a value of $\eta = 1000$. Appendix Figure A.2 shows this same graph with low and high values for η and shows that the results are not sensitive to the choice of η affecting the threshold magnitude by less than 3%. For Mbiti and Weil's estimated discount factor, threshold gamma will be 0.716. If it falls below 0.63, even the most patient people will prefer to bet.

C2: Estimating Return on Saving

Having established a range of threshold returns to saving that will separate bettors from savers, we need a way to estimate reasonable values for γ before we can say whether betting is likely to be a rational utility maximizing strategy. To do this, I propose that γ should be the product of a set of four contributing factors: inflation, temptation, social-pressure, and loss or theft. None of these are easy to estimate and, to my knowledge, there is little existing research examining this. In the absence of existing estimates, I try and use data

[2]Bettors typically target bets with median multipliers of 1.85. This would require just under five games to be stacked on a bet in order to win at least 200K. Because of how the betting companies systematically lower the payouts on a per game basis. A multiplier of 1.85 that has been shaded down by 7.5% (as appears on average in the data), should have been a fair multiplier of 2. A fair multiplier of 2 implies a fair "true" probability of 50%. We can therefore infer that the true likelihood that a bet with five games on it, each with an offered multiplier of 1.85, has a winning probability of $p = 0.5^5 = 0.03125$.

and information about the local context to create reasonable approximations.

First, inflation in Uganda over the past five years has ranged from 4-24%. This corresponds with a weekly discount factor between 0.9947 and 0.9992. Conservatively estimating that it will take at least 20 weeks to save the needed liquidity for the expenditure, these weekly discount factors correspond with 20 week discount factor between 0.8998-0.9844.

Second, I look to the data for insight on losses from temptation. The survey included consumption modules including a number of temptation goods including weekly expenditures on video clubs, jewelry, soda, alcohol, and gifts for girlfriends. The median portion of income spent on these goods was 4.8%. In data from condensed study, I can show that, after removing essential recurrent expenditures on food, transport, and rent, only 42% of weekly income remained. This suggests that temptation goods use up $4.8/42 = 11.4\%$ of discretionary income. Because consumption of temptation goods is likely a source of enjoyment, all temptation expenditures should not be considered "losses", but we might suspect that some portion of them are not valued. I set a range whereby between 25 and 50% of these expenditures are non-valued (or regretted), constituting losses of 2.85-5.7% of total weekly income or 6.7-13.4% of discretionary income.[3]

In addition, betting itself likely (partially) a temptation good. Assuming that bettors regret the expected losses of bets (but are comfortable with the amount won back) then the median bettor who spends 8.6% of his weekly income on betting regrets the expected losses of 40% of his betting expenditures. This results in additional expected losses equal 3.44% of weekly income or 8.2% of discretionary income. Taking the more conservative estimates of temptation expenditures and comparing them to total weekly income creates a lower bound of $2.85 + 3.44 = 6.29\%$ expected loss to temptation or a discount factor of 0.9361. The higher range of estimates while assuming that money for saving is mixed in with other discretionary income results in expected losses to temptation of $13.4 + 8.2 = 21.6\%$ of discretionary funds or a discount factor of 0.784.

Next, 27% of people in the full sample said that they feel considerable pressure to spend their earnings on other people so this is clearly a concern for many people in the population. Recent work has highlighted the importance of both inter and intra-household pressures on individuals' income (Ashraf 2009; Baland et al. 2011; Goldberg 2011). A recent paper by Jakiela and Ozier (2015) uses a set of lab-in-the-field experiments to identify participants' willingness to sacrifice income in order for the size of their payouts to be hidden from others in the study. With these measures, they are able to back out the level of "Kin Tax" people expect to face from others in the study. Sizes of this kin tax vary considerably and are dependent on a number of factors and ranged from essentially zero up to eight percent, suggesting a discount factor of 0.92 - 1.

[3]It may be unfair to assume that the same proportion of expenditures on temptation goods relative to total income or discretionary income would be taken out of money intended for saving. However, the saving box was a soft commitment saving device that did very little for the actual security of their money, aside from removing it from their own pockets and shielding it from exposure to the temptations they face throughout the day. Take up rates of over 50 percentage points in the saving box treatment are suggestive that people are indeed concerned about exposure of their own money to temptation. It is also possible that saving ultimately ends up being the residual of what people bring home at the end of the day. If this is the case, then the losses from temptation expenditures may cut *first* into saving before they affect less flexible expenditures.

Finally, there is theft and loss. 30% of respondents in the study claimed that they had had their house robbed and money stolen. It is hard to know what rate of theft this implies. Assuming that these robberies happened within the last ten years, it would suggest that every year about 3% of people lose their money stored at home. In the minimum 20 week time window we are presuming in order to save successfully, this would be just 1.15% of participants. It might be that even low incidence of a salient outcome get overweighted in peoples' decision-making (Bordalo et al. 2012), and pick-pocketing or losses of income when out of the house would not be captured in this survey measure, but to be conservative I will use discount factors between 0.9885-1.

A fifth possible factor could be transaction costs. In this example of cash-savings, transaction costs are considered to be zero. However, other saving technologies such as bank accounts or roscas could impose significantly higher transaction costs either through formal fees or through demands on time and effort.

Together, I can construct a low and high range for γ keeping in mind that people will start to switch to betting if γ dips below 0.8. For the upper bound of gamma I take the product of the high γ_i estimates so that $\gamma_{high} = 0.9844 \times 0.9361 \times 1 \times 1 = 0.9215$. For all reasonable levels of patience, these people should clearly prefer saving to betting as a mode of liquidity generation and are likely betting primarily for non-financial motivation. On the other end of the spectrum are those using the lower estimates of each saving discount factor such that $\gamma_{low} = 0.8998 \times 0.784 \times 0.92 \times 0.9885 = 0.6415$. This estimate is below the floor at which we assumed most people would abandon saving and prefer to bet.

Appendix D: Additional Field Targeting Protocols

Listing for Wave 1 was launched in September 2015 and took two weeks. The full set of 91 parishes in Kampala were identified. Because we planned to find and interview respondents at their place of work, we removed 39 parishes without significant commercial activity. Thirteen parishes were then randomly selected for Wave 1.

The field team visited one parish at a time, beginning at its commercial center. The main roads of the community were identified and divided among team members. These field officers were then instructed to walk along their designated route, looking for men who satisfied the targeting criteria. When they found someone suitable, they would approach him and include him in the listing by asking a few short questions to determine his eligibility and interest in participating in the study. After having identified a suitable respondent and including him in the listing, enumerators were instructed to skip the next eligible participant they saw before approaching another potential respondent. This was an effort to mitigate the possibility of spillovers resulting from respondents who knew other people in the study.

Wave 2 was launched in March 2016. Participants were identified following identical protocols and targeting criteria in selected parishes as those in the initial group. The only change was that parishes were no longer randomly selected from the eligible list of locations. Logistic and budgeting considerations influenced the decision to work in areas that were more accessible to the field teams. Ultimately, an additional 23 communities were included in Wave 2.

The same targeting criteria were employed for people in the condensed study. Therefore the populations are very similar (though not identical) between full and condensed study participants. However, inclusion in the condensed study was no longer implemented in a two-stage approach with a listing followed by invitation into the study. Instead, if a suitable respondent was identified, he was invited to participate and, if willing, interviewed immediately. Field team members were instructed to exclude anyone with pre-existing knowledge of the research project.

Appendix E: Comparative Statics

In this appendix I derive the comparative statics resulting from a change in saving ability as described in the model in Section 3. I will show that an improvement in saving ability will lead to an increase in the amount saved and also an increase in the range of incomes for which saving is welfare improving will expand.

As before, utility is $v(Y) = u(D) + L\eta$ where $u(D)$ is conventional concave utility with $u'(D) > 0$ and $u''(D) < 0$ from consumption of divisible goods. L is an indicator of whether or not the lumpy expenditure was purchased such that $L \in \{0, 1\}$ and yields a discrete utility payoff of η. γ is the return on saving which, in this setting, is assumed to be below one. P is the price of the lumpy expenditure and income is Y which is assumed to be the same in both periods.

An individual will only save if after saving he purchases the lumpy good L. This is because of the concavity of $u()$, the assumption that $\gamma \leq 1$, and because income is stable so that the added consumption of the divisible good in the second period will result in less utility in the second period than that sacrificed from period one consumption. The optimal saving choice will therefore be:

$$\max_S u(Y - S) + \delta[u(Y + \gamma S - P) + \eta] \qquad\qquad s.t. \quad S \leq Y$$

The first order condition is:

$$-u'(Y - S^*) + \delta\gamma u'(Y + \gamma S^* - P) = 0$$

I then differentiate again with respect to γ and S in order to show the effect of improved saving ability on amount saved.

$$ds[u''(Y - S^*) + \delta\gamma^2 u''(Y + \gamma S^* - P)] + d\gamma[\delta u'(Y + \gamma S^* - P) + \delta\gamma S^* u'(Y + \gamma S^* - P)] = 0$$

Reorganizing:

$$\frac{ds}{d\gamma} = -\left[\frac{\overbrace{\delta u'(Y + \gamma S^* - P)}^{>0} + \overbrace{\delta\gamma S^* u'(Y + \gamma S^* - P)}^{>0}}{\underbrace{u''(Y - S^*)}_{<0} + \underbrace{\delta\gamma u''(Y + \gamma S^* - P)}_{<0}}\right]$$

Both terms in the numerator are first derivatives of the utility function and therefore positive. Both terms in the denominator are second derivatives of the utility function and therefore negative. Therefore, $\frac{dS^*}{d\gamma} > 0$.

An increase in S^* results in a decrease in period one consumption of D because $D = Y - S$. It is worth nothing that the part of betting that is a normal consumption good will be included in D.

Although the optimal saving amount will increase for everyone, people will only actually save if the utility from saving is bigger than the utility from either purchasing the lumpy expenditure in both time periods or purchasing it in neither without saving.

Next, I show that the increase in saving ability also increases the upper bound of people who save and decreases the lower bound, leading to an expansion of the range of incomes for people who save.

First I show that the upper bound will increase. There are now two equilibrium conditions, optimal savings and equality of saving and direct purchase[4]:

$$\begin{aligned} -u'(Y-S) + \delta\gamma u'(Y + S\gamma - P) = 0 \\ u(Y-S) + \delta u(Y + S\gamma - P) + \eta - (1+\delta)[u(Y-P) + \eta] = 0 \end{aligned}$$

I take derivatives and put them into matrix form omitting the "*" from optimal saving to avoid clutter.

$$\begin{bmatrix} u''(Y-S) + \delta\gamma^2 u''(Y+\gamma S-P) & -u''(Y-S) + \delta\gamma u''(Y+\gamma S-P) \\ -u'(Y-S) + \delta\gamma u'(Y+\gamma S-P) & u'(Y-S) + \delta u'(Y+\gamma S-P) - (1+\delta)u'(Y-P) \end{bmatrix} \begin{bmatrix} dS \\ dY \end{bmatrix} = \begin{bmatrix} -\delta u'(Y+\gamma S-P) - \delta\gamma S u''(Y+\gamma S-P) \\ -\delta S u'(Y+\gamma S-P) \end{bmatrix} d\gamma$$

The determinant of the coefficient matrix is positive because the top left term is negative (second derivatives of utility maximization), the bottom right is negative (the last term is bigger in magnitude than the first two because $P > S$), and the bottom left term is equal to zero. Following Cramer's rule I replace the second column with the column from $d\gamma$. This will be positive as well and therefore $\frac{dY^{Max}}{d\gamma} > 0$.

Following a similar approach as for the upper bound of saving, I look at the effect of a change in saving ability on the lower bound of saving. There are again two equilibrium conditions. Now, they are optimal savings and equality of utility from saving and direct purchase of the good[5]:

$$\begin{aligned} -u'(Y-S) + \delta\gamma u'(Y + S\gamma - P) = 0 \\ u(Y-S) + \delta[u(Y + S\gamma - P) + \eta] - (1+\delta)u(Y) = 0 \end{aligned}$$

[4] If there is a crossing, it will be a single crossing. This is because the marginal utility of $((1+\delta)u'(Y-P) > u'(Y-S) + \delta u'(Y+\gamma S-P)$ so that the direct purchasing utility cuts up through the optimal saving utility and will not cross again.

[5] If there is a crossing, it will be a single crossing. This is because the marginal utility of $((1+\delta)u'(Y) < u'(Y-S) + \delta u'(Y+\gamma S-P)$ so that the saving utility cuts up from below and will not cross again.

I take the derivatives and re-organize in matrix form:

$$\begin{bmatrix} u''(Y-S)+\delta\gamma^2 u''(Y+\gamma S-P) & -u''(Y-S)+\delta\gamma u''(Y+\gamma S-P) \\ -u'(Y-S)+\delta\gamma u'(Y+\gamma S-P) & u'(Y-S)+\delta u'(Y+\gamma S-P)-(1+\delta)u'(Y) \end{bmatrix} \begin{bmatrix} dS \\ dY \end{bmatrix} = \begin{bmatrix} -\delta u'(Y+\gamma S-P)-\delta\gamma S u''(Y+\gamma S-P) \\ -\delta S u'(Y+\gamma S-P) \end{bmatrix} d\gamma$$

The derivative of the coefficient matrix is now negative (because $(1+\delta)u'(Y) < u'(Y-S) + \delta u'(Y+\gamma S-P)$). Replacing the second column of the coefficient matrix with the column for $d\gamma$ and taking the determinant now results in a positive. Therefore $\frac{dY^{min}}{d\gamma} < 0$.

Together, these results show that an increase in saving ability a) increases the amount that people save, S^* which also reduces the amount spent on current divisible good consumption, $D = Y - S^*$, b) the maximum income level for people to prefer saving, Y^M has increased, and c) the minimum income level for people to prefer saving, Y^m has fallen so that the overall range of incomes where saving is preferred has increased.

Appendix F: Effect of Winnings on Biggest Expenditure Size

Section 6.1 of the paper showed the effect of winnings on the likelihood of making expenditures above a given threshold. Alternatively, I look at the continuous effect of winnings on the size of individuals' biggest expenditure in those periods. Figure A.1a shows a scatter plot of winnings in each two-week period and the value of an individual's biggest expense. The dotted line indicates where winning amount is equal to biggest purchase amount. It is clear from the figure that not all winnings are used directly on a large expense and that many large purchases happen in time periods without big wins. People likely pursue many strategies for liquidity generation simultaneously and betting is just one of them. However, there is a positive correlation between win amount and the size of peoples' biggest expenditure. Figures A.1b and A.1c split the sample by people with low and high saving ability respectively and show that this correlation climbs to 0.2568 for people with low saving ability but falls to just 0.0341 for those with high saving ability.

Table A.6 shows these results. Columns (1)-(4) run the raw correlations of biggest expenditure on win amount (both scaled to mean income). More winnings correlate with larger expenditures. In column (1) we see that for each additional dollar won, 33 cents gets put towards the biggest large expense that in that time period. Splitting the sample by high and low saving ability we see that this effect nearly goes to zero for people with high saving ability and is insignificant, whereas 53 cents per additional dollar of winnings for someone with low saving ability get spent on a large lumpy expenditure as shown in column (2). This is a statistically significant difference. Columns (5)-(8) again include the time varying covariates as well as the betting moments needed to establish valid identification. The magnitudes shrink by about a third, suggesting that the exclusion of these additional endogenous covariates were biasing the naive estimates upwards, but we see that the effects are still large and significant when testing between high and low saving ability.

A vulnerability of this approach is that a small number of large wins (and large expenditures) can drive the overall average treatment effect. First, I try trimming the sample to exclude people with particularly high reported winnings (more than five times weekly income). Unsurprisingly, trimming people with large winnings weakens the result since I have thrown away a large amount of valid variation in the treatment. However, the sign of the effects persist and still show highly significant positive effects for people with low saving ability in the full specification such that each additional dollar won increases the biggest purchase made in that period by an average of 0.11 dollars, significant at the 95% confidence level. These results are contained in Table A.7.

I implement an inverse hyperbolic sine transformation on winnings in order to mitigate the weight given to a few people with large winnings without needing to trim them from the sample. Using logs as a way to address skewness in winnings is not possible because I would be forced to drop all observations in which an individual did not win. A similar approach that avoids this problem is to use the inverse hyperbolic sine transformation of winnings. This transformation can still be interpreted similarly to logs (Burbidge et al. 1988; MacKinnon and Magee 1990). These results are in Table A.8 and show similar differentiation between those with high and low saving ability. Winnings have a significant effect on the size of biggest expenditure for those with low saving ability and no clear impact on those with high saving ability.

Having tried a number of different specifications and robustness checks, we consistently see that more winnings, perhaps unsurprisingly, leads to larger purchases in that time period. In addition, we see that after separating the sample between people with high and low saving ability, the effect is typically indistinguishable from zero for those who are likely to be able to use saving as a way to generate liquidity, suggesting that they may not have had many unmet liquidity needs, whereas the effect is stronger and significant for those with low saving ability.

Appendix B

Additional Appendices, Tables, and Figures for Chapter 2

Appendices

Table A.1: Sports for Change Life Skills Sessions: Topics and Skills Addressed

	Sports for Change Skills				
Session	Resilience	Planning	Trust	Self-Esteem	Communication
1: Introduction			X		
2: Me and Others			X	X	X
3: Understanding Emotions	X	X	X	X	X
4: Communication			X		X
5: Relationships	X		X		X
6: Cooperation			X	X	X
7: Believing in Me		X	X		X
8: Conflict and Violence	X	X	X	X	X
9: Collaboration			X		X
10: Motivation	X	X		X	
11: Dealing with Problems	X		X	X	X
12: Making Strategies		X	X	X	
13: Applying SFC Skills in My Life	X	X	X	X	X
14: Planning Graduation Event		X			
15: What Have We Learned?	X	X	X	X	X
16: Review and Closing	X	X	X	X	X

Table A.2: Random Assignment Balance and Summary Statistics for Respondents in Endline

	All (N=2072)			Male (N=980)			Female (N=1092)		
	CTRL (1036)	TRT (1036)	P-val	CTRL (493)	TRT (487)	P-val	CTRL (543)	TRT (549)	P-val
Basic Demographics									
Age	20.99	20.90	0.42	20.75	20.61	0.41	21.21	21.15	0.72
Head of Household	0.134	0.151	0.29	0.156	0.183	0.28	0.115	0.124	0.68
Household Size	6.825	6.675	0.34	6.814	6.576	0.32	6.835	6.763	0.73
Mother Living	0.873	0.871	0.88	0.868	0.893	0.22	0.878	0.851	0.20
Father Living	0.711	0.718	0.73	0.721	0.727	0.82	0.702	0.709	0.79
Has Children	0.453	0.447	0.76	0.227	0.217	0.72	0.659	0.649	0.74
Christian	0.881	0.862	0.20	0.864	0.828	0.11	0.895	0.892	0.83
Muslim	0.107	0.126	0.18	0.119	0.162	0.06*	0.095	0.094	0.94
Matched Friends	2.403	2.436	0.71	2.279	2.400	0.34	2.516	2.468	0.69
Education and Cognitive Ability									
Primary School	0.849	0.854	0.74	0.921	0.908	0.46	0.783	0.807	0.35
Secondary School	0.294	0.290	0.86	0.364	0.357	0.80	0.229	0.231	0.93
Grades Completed	11.58	11.61	0.87	12.53	12.50	0.89	10.73	10.83	0.69
Numeracy (0-10)	6.019	5.847	0.05**	6.314	6.377	0.58	5.752	5.380	0.00***
Digits Forward (0-8)	5.226	5.162	0.35	5.354	5.369	0.88	5.110	4.980	0.16
Digits Backward (0-8)	1.994	1.962	0.53	2.113	2.043	0.36	1.886	1.890	0.96
Word Recall 1 (0-10)	3.060	3.093	0.79	3.186	3.213	0.88	2.945	2.987	0.80
Word Recall 2 (0-10)	2.557	2.573	0.88	2.630	2.619	0.94	2.492	2.533	0.79
Ravens Score (0-3)	1.752	1.735	0.69	1.897	1.932	0.57	1.620	1.561	0.30
Risk Aversion (0-6)	3.789	3.742	0.67	3.860	3.791	0.67	3.725	3.698	0.86
Psychosocial Measures/Indices									
Subjective Welfare	2.315	2.357	0.45	2.253	2.281	0.73	2.371	2.425	0.49
Self-Esteem	20.91	20.49	0.02**	20.76	20.45	0.23	21.04	20.54	0.04**
Locus of Control	24.06	24.13	0.60	24.17	24.27	0.56	23.96	24.00	0.83
Aggression	2.570	2.514	0.65	2.407	2.385	0.90	2.717	2.627	0.60
Risky Behavior	1.466	1.488	0.79	2.091	2.098	0.96	0.899	0.949	0.52
Depression	22.15	22.98	0.17	20.52	21.39	0.32	23.63	24.39	0.37
Labor Force Participation and Earnings									
Working	0.428	0.435	0.73	0.493	0.483	0.77	0.368	0.392	0.41
Hours Worked (7D)	13.42	14.25	0.46	16.25	16.95	0.71	10.86	11.86	0.45
Inc 7 Days	6.301	6.542	0.69	8.083	8.089	1.00	4.677	5.177	0.47
Inc 7 Days ((>0)	21.11	24.65	0.49	21.38	31.77	0.24	20.78	17.32	0.41
Log(Inc 7D)	2.322	2.381	0.47	2.406	2.581	0.12	2.220	2.175	0.71
Inc 3 Months	46.77	58.36	0.00***	53.45	67.96	0.02**	40.80	49.92	0.08*
Inc 3 Months (>0)	130.5	112.5	0.57	177.2	129.3	0.43	78.8	95.3	0.10*
Log(Inc 3M)	3.708	3.919	0.01***	3.761	4.050	0.01**	3.649	3.785	0.20

Table A.3: Effects of Number Friends on SFC Participation

	Any SFC Participation				Days of SFC Participation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	All	Male	Female	All	All	Male	Female	All
Total Friends	0.034	0.035	0.023	0.028	0.013	0.012	0.026	0.051
	(0.021)	(0.031)	(0.024)	(0.030)	(0.066)	(0.112)	(0.083)	(0.089)
Female x Total Friends				0.012				-0.069
				(0.028)				(0.123)
Female	0.012			0.006	0.084			0.121
	(0.033)			(0.037)	(0.092)			(0.115)
N	1039	486	552	1039	760	351	407	760
Mean Dep Var	0.73	0.73	0.74	0.73	14.71	14.63	14.78	14.71
Mean Total Friends	0.522	0.484	0.555	0.522	0.522	0.484	0.555	0.522
R2	0.075	0.100	0.135	0.075	0.130	0.203	0.137	0.131

Notes: "Total Friends" refers to the number of friends identified at baseline assigned to an individual's youth group. Regressions include dummies for amount of potential matches in community and baseline outcome value. Covariates: age, age^2, education attainment, and gender.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.4: Effect of the Presence of Friends in SFC Group on Psychosocial Outcomes

Panel (a): Effect of Any Friends in SFC Group on Psychosocial Outcomes										
	Male (N=971)					Female (N=1087)				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	RB	SE	Loc	Agg	WF	RB	SE	Loc	Agg	WF
Any	0.002	0.035	-0.019	-0.019	0.069	-0.122	0.010	-0.126	-0.052	-0.042
	(0.099)	(0.098)	(0.101)	(0.091)	(0.100)	(0.095)	(0.100)	(0.101)	(0.097)	(0.095)
SFC	0.011	0.071	0.027	0.165**	-0.075	0.042	-0.002	-0.011	0.032	0.105
	(0.071)	(0.069)	(0.070)	(0.073)	(0.073)	(0.067)	(0.070)	(0.069)	(0.077)	(0.071)
N	971	971	971	971	969	1087	1087	1087	1087	1086
R2	0.105	0.114	0.074	0.088	0.051	0.051	0.058	0.040	0.085	0.050
Panel (b): Effect of Total Friends in SFC Group on Psychosocial Outcomes										
	Male (N=971)					Female (N=1087)				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	RB	SE	Loc	Agg	WF	RB	SE	Loc	Agg	WF
Total	-0.016	0.036	0.033	0.031	0.020	-0.003	0.015	-0.063	0.041	-0.077
	(0.094)	(0.059)	(0.058)	(0.051)	(0.053)	(0.039)	(0.059)	(0.055)	(0.049)	(0.055)
SFC	0.020	0.067	0.005	0.143**	-0.061	-0.003	-0.007	-0.024	-0.010	0.131*
	(0.073)	(0.067)	(0.069)	(0.071)	(0.070)	(0.065)	(0.067)	(0.067)	(0.072)	(0.068)
N	971	971	971	971	969	1087	1087	1087	1087	1086
R2	0.105	0.114	0.074	0.088	0.050	0.049	0.058	0.040	0.085	0.051

Outcomes: WF=Subjective Welfare, SE=Self-Esteem Index, LOC=Locus of Control Index, Agg=Aggression Index, RB=Risky Behavior Index. Psychosocial measures standardized so that positive is better. Covariates: Lagged dependent variable, age age^2 and grades of schooling. All regressions include community fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.5: Effect of Friends in Youth Group on Financial Outcomes

Panel (a): Any Friends		Both			Male			Female	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Working	Hours	7dI	Working	Hours	7dI	Working	Hours	7dI
Any	-0.041	1.332	-0.068	-0.044	0.075	-0.237	-0.038	2.052	0.052
	(0.031)	(1.279)	(0.135)	(0.043)	(1.881)	(0.209)	(0.044)	(1.741)	(0.178)
SFC	0.060***	0.848	0.112	0.057*	0.740	0.100	0.060*	1.030	0.114
	(0.023)	(0.873)	(0.094)	(0.031)	(1.340)	(0.133)	(0.033)	(1.141)	(0.131)
R2	0.048	0.051	0.115	0.039	0.053	0.131	0.060	0.086	0.135

Panel (b): Total Friends		Both			Male			Female	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Working	Hours	7dI	Working	Hours	7dI	Working	Hours	7dI
Total	-0.019	1.411*	-0.008	-0.023	0.157	-0.111	-0.016	2.189*	0.046
	(0.018)	(0.836)	(0.083)	(0.025)	(1.231)	(0.114)	(0.026)	(1.196)	(0.124)
SFC	0.056**	0.597	0.093	0.054*	0.691	0.071	0.055*	0.595	0.108
	(0.022)	(0.856)	(0.091)	(0.030)	(1.317)	(0.127)	(0.032)	(1.129)	(0.130)
R2	0.048	0.052	0.115	0.039	0.053	0.129	0.059	0.089	0.135
N	2054	2009	656	968	943	364	1086	1066	292
Mean Dep Var	0.698	13.135	2.135	0.756	14.269	2.141	0.645	12.131	2.126

Notes: Regressions include dummies for amount of potential matches in community and baseline outcome value. Hours=Hours worked over previous seven days. 7dI=Income over last seven days among earners. Covariates: Lagged dependent variable, age, age^2, education attainment, and gender. N and Mean Dep Var are same for both panels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.6: Balance and Summary Statistics at Assignment by Registration Timing

	All (N=2395)			**Male (N=1150)**			**Female (N=1245)**		
	Late (1208)	Early (1187)	P-val	Late (570)	Early (580)	P-val	Late (620)	Early (625)	P-val
Basic Demographics									
Age	20.72	20.10	0.02**	20.38	20.81	0.01***	21.16	21.05	0.47
Head of Household	0.117	0.157	0.01***	0.151	0.179	0.23	0.110	0.116	0.75
Household Size	6.756	6.605	0.30	6.565	6.550	0.95	6.816	6.770	0.81
Mother Living	0.869	0.869	0.98	0.889	0.870	0.31	0.869	0.849	0.33
Father Living	0.711	0.719	0.69	0.712	0.734	0.41	0.707	0.709	0.94
Has Children	0.425	0.462	0.07*	0.195	0.240	0.06*	0.669	0.635	0.21
Christian	0.881	0.858	0.09*	0.858	0.833	0.24	0.892	0.891	0.97
Muslim	0.107	0.129	0.09*	0.132	0.153	0.29	0.097	0.094	0.89
Matched Friends	2.307	2.501	0.02**	2.175	2.452	0.02**	2.508	2.467	0.71
Education and Cognitive Ability									
Primary School	0.826	0.845	0.22	0.891	0.916	0.16	0.765	0.782	0.45
Secondary School	0.245	0.296	0.01***	0.288	0.381	0.00***	0.213	0.210	0.89
Grades Completed	11.23	11.51	0.09*	12.06	12.58	0.01***	10.46	10.52	0.79
Numeracy (0-10)	5.822	5.911	0.28	6.195	6.383	0.08*	5.632	5.323	0.01***
Digits Forward (0-8)	5.137	5.162	0.70	5.295	5.293	0.99	5.066	4.968	0.27
Digits Backward (0-8)	1.944	1.960	0.74	2.021	2.069	0.49	1.868	1.866	0.97
Word Recall 1 (0-10)	3.066	3.065	1.00	3.212	3.179	0.84	2.927	2.963	0.82
Word Recall 2 (0-10)	2.578	2.559	0.85	2.639	2.624	0.92	2.500	2.520	0.89
Ravens Score (0-3)	1.721	1.757	0.37	1.870	1.955	0.14	1.605	1.554	0.34
Risk Aversion (0-6)	3.698	3.844	0.17	3.749	3.934	0.21	3.669	3.742	0.62
Psycho-Social Measures/Indices									
Subjective Welfare	2.256	2.356	0.06*	2.195	2.290	0.19	2.347	2.384	0.62
Self-Esteem	20.51	20.76	0.13	20.43	20.62	0.43	20.98	20.50	0.04**
Locus of Control	23.98	24.14	0.16	24.07	24.32	0.13	23.95	23.92	0.86
Aggression	2.572	2.577	0.97	2.463	2.384	0.65	2.750	2.678	0.65
Risky Behavior	1.461	1.508	0.54	1.975	2.179	0.13	0.939	0.934	0.95
Depression	23.27	22.02	0.03**	21.82	20.49	0.10	23.34	24.67	0.09*
Labor Force Participation and Earnings									
Working	0.429	0.437	0.68	0.512	0.472	0.17	0.368	0.391	0.40
Hours Worked (7D)	14.23	14.18	0.96	17.34	15.48	0.27	11.96	12.40	0.76
Inc 7 Days	6.114	6.945	0.14	8.172	8.206	0.97	4.710	5.294	0.37
Inc 7 Days (>0)	24.85	20.28	0.30	26.35	23.72	0.73	20.69	18.47	0.57
Log(Inc 7D)	2.335	2.371	0.64	2.414	2.519	0.30	2.233	2.218	0.89
Inc 3 Months	54.77	50.65	0.28	68.07	52.31	0.01***	40.85	50.84	0.05**
Inc 3 Months (>0)	130.1	109.1	0.45	158.8	130.4	0.59	79.2	103.3	0.02**
Log(Inc 3M)	3.864	3.783	0.26	3.939	3.853	0.41	3.652	3.828	0.08*

Figure A.1: Sample SFC Session Outline

Session 8: Conflict and Violence

Topics to be covered:

- Why do conflicts happen?
- How do we avoid conflict?
- Fair play
- How do we deal with conflicts when they arise?

Equipment

- Ropes
- Footballs – 6
- Bibs
- Cones
- Other sporting equipment as required by the sport

Session Outline

0-5mins	Energizer	Opening Chant Travelling to new places
5-20mins	Introduction	Review of last session Feedback from last session Introduce topic for the day Outline activities
20-35mins	Warm up	Stretches and exercise
35mins-1hr	Activity 1	Knee Fight
1hr-1hr20mins	Activity 2	Find Similarities and Move On
1hrs20mins-1hr45mins	Activity 3	Line Push and Pull
1hrs45mins-2hrs	Break	Including registration
2hrs-2hrs40mins	Sport	Football (2 varieties to promote conflict resolution!)
2hrs40mins-3hrs	Conclusion	Review main topics discussed Discussion Plan for next week Closing Chant