

THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

Commitment Savings Products

Theory and Evidence

Anett Hofmann

A thesis submitted to the Department of Economics of the London School of Economics and Political Science for the degree of Doctor of Philosophy.
London, July 2014.

*Für meine Eltern Anke und Frank, deren Liebe und Unterstützung bis heute das
Bedingungsloseste ist, das ich in meinem Leben kennenlernen durfte.*

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others. No chapters are co-authored.

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 39,430 words.

Statement of Inclusion of Previous Work

I confirm that Chapter 1 is an excerpt from a previous study for the Master of Research (MRes) degree in Economics at the LSE, awarded in July 2010. It was published as an *Economic Organisation and Public Policy Discussion Paper*, No. 51, in March 2014.

London, 31 July 2014

A handwritten signature in black ink, reading 'Anett Hofmann'. The signature is written in a cursive style with a large, sweeping initial 'A'.

Anett Hofmann

Abstract

Recent literature promotes commitment products as a new remedy for overcoming self-control problems and savings constraints. This thesis argues that the effects of commitment may be very heterogeneous, and highlights the mechanisms under which commitment may reduce welfare, rather than increase it. It also examines a new type of commitment contract: A formal commitment savings account with fixed regular instalments, introduced in a developing-country context.

Chapter 1 proposes that the popularity of costly or inflexible savings mechanisms as well as of high-interest consumption loans may represent a demand for commitment to fixed instalments. Using a newly collected dataset from Bangladesh, it shows that the introduction of a regular-instalment commitment savings product was associated with a large increase in average savings contributions.

The theoretical framework in Chapter 2 highlights the potential heterogeneity behind such positive average effects: Commitment improves welfare when agents have full knowledge of their preferences, including biases and inconsistencies. If agents are imperfectly informed about their preferences, they may choose ill-suited commitment contracts. I formally show that commitment contracts can reduce welfare if the commitment is not strong enough to discipline the agent, resulting in costly default. I further show that such insufficient commitment contracts are likely to be selected by time-inconsistent agents with ‘partially sophisticated’ preferences: Agents who are neither completely unaware nor fully aware of their time-inconsistency, but anywhere in between those two extremes.

Chapter 3 describes a randomised experiment in the Philippines: I designed and introduced a regular-instalment commitment savings product, intended to improve on pure withdrawal-restriction products by mimicking the fixed-instalment nature of loan repayment contracts. Individuals from a general low-income population were randomly offered to take up the product, and were asked to choose the stakes of the contract (in the form of a default penalty) themselves. The result is that a majority appears to choose a harmful contract: While the intent-to-treat effect on bank savings for individuals assigned to the treatment group is four times that of a withdrawal-restriction product (offered as a control treatment), 55 percent of clients default on their savings contract. The explanation most strongly supported by the data is that the chosen stakes were too low (the commitment was too weak) to overcome clients’ self-control problems. Moreover, both take-up and default are *negatively* predicted by measures of sophisticated hyperbolic discounting, suggesting that those who are fully aware of their bias realise the commitment is too weak for them, and avoid the product. The study suggests that research on new commitment products should carefully consider the risk of adverse welfare effects, particularly for naïve and partially sophisticated hyperbolic discounters.

Acknowledgements

This thesis would not have been possible without the extensive support, guidance and encouragement of a great number of people. First and foremost, I thank my supervisors, Oriana Bandiera, Maitreesh Ghatak, and Gharad Bryan. Thanks to Oriana, for tirelessly providing advice and guidance, for mercilessly seeing through any identification strategy, and for never failing to answer an email. She is living proof that it is possible to be an exceptional researcher, a dedicated and committed supervisor, and an inspiring teacher at the same time, all while having a family. If ever I needed a role model, I would know where to look. Gharad, thank you for challenging my thinking, for volunteering for lively hour-long debates after research presentations, and for inspiring many new ideas and directions. Thanks to Maitreesh for helping me think more clearly, and for pushing me to always firmly ground my work in theory. All three have influenced me (even) more than I probably realise.

I am deeply indebted to Ann Mayuga, Faith McCollister, Megan McGuire, Yoeri Suykerbuyk and Eva Ghirmai of IPA Philippines, without whom the field experiment would not have been possible. I was truly blessed to have Ann as my project manager: She was an inexhaustible source of guidance and good advice on everyday fieldwork implementation, at a time when embarking on my own randomised controlled trial without any experience seemed as daunting as it was exciting. Ann has become a trusted friend. I would also like to thank all enumerators, marketers, data encoders and auditors who made my idea of the IPA Regular Savers project become a reality. Thanks to Betty Wilkinson and Michiko Katagami at the Asian Development Bank, and John Owens at the Rural Bankers' Association of the Philippines, for believing in my idea and pushing me to start this project. Last but not least, I am grateful to 1st Valley Bank of Cagayan de Oro, Philippines, for a productive collaboration.

Chapter 2 and 3 of the thesis have greatly benefited from conversations with Dean Karlan, Stefano DellaVigna, Matthew Levy, George Loewenstein, Lance Lochner, Francesco Nava, Johannes Spinnewijn, Pascaline Dupas, and Alessandro Tarozzi, as well as from seminar participants at LSE and at various conferences.

I was the fortunate beneficiary of a scholarship from the LSE Economics department and the ESRC, which funded the first four years of my PhD. For the field experiment in the Philippines, I gratefully acknowledge funding from the Yale Savings and Payments Research Fund at Innovations for Poverty Action. For the Bangladesh study, I am grateful to Stuart Rutherford, Billal Hossain and Mohammad Maniruzzaman of *SafeSave* Dhaka, for allowing me to use their data, and pro-

viding me with their continued help and support. I would also like to thank Greg Fischer for his supervision on this project.

I will miss the stimulating research environment of STICERD, most notably the kitchen conversations. The prevailing intellectual open-mindedness to think through even the most exotic theories and assumptions is an absorbing state, which will be impossible to switch off. I was privileged to have the company of Jonathan de Quidt, Claudia Steinwender, Christian Siegel, Jason Garred, Matthew Skellern, Michael Best, Michele Piffer, Pedro Pinto, Delger Enkhbayar and many others. I am grateful for their insights and their moral support over the past six years, for their willingness to listen to me ramble about my research, and last but not least for giving me a sense of perspective when I needed it most.

A big heartfelt thank you goes out to Nicolas Weinhold, Erina Ytsma, Tim Komischke, and Nicola Bunte, for being close to my heart through so many years, with no regard for geographical distance. While LSE has taught me that counterfactuals are fundamentally unobservable, I am positive that I could not have made it through my time at LSE without their love, friendship, encouragement, and their firm belief in me, even when I occasionally lost it. Thank you for all the laughs and tears, your sense of humour, your patient and ever-ready ears, and the occasional Thai green curry.

I thank my parents for being my biggest supporters and believers throughout my long school career. They have cheered and suffered through all of my exciting as well as my tough moments with me, and occasionally reminded me that life is about more than just work. Thanks to 'Mama' for sending a million and one 'survival parcels' from Germany to London, and for reminding me of my home wherever I was in the world.

And finally, thanks to my wonderful husband Ryan, for embarking on this roller coaster ride with me, for tolerating the endless evenings and weekends I spent working, for cooking for me when my research crowded out any thoughts of food intake, and for embarking on a new life in Paris with me. You really meant it when you said, "for better or for worse"! Thank you.

Contents

Introduction	11
1 Regular Saver Products: Evidence from Bangladesh	17
1.1 Motivation and Context	17
1.2 The Data	21
1.3 Estimation	23
1.3.1 The Differences-in-Differences Approach	23
1.3.2 Endogeneity Concerns	27
1.3.3 Adjusting the Estimation Strategy	32
1.4 Conclusion and Outlook	34
2 Theoretical Framework: Commitment under Partial Sophistication	36
2.1 First Best	38
2.2 Autarky Equilibrium	39
2.3 Equilibrium with a Regular Saver Commitment Product	44
3 When Commitment Fails: A Randomized Experiment in the Philippines	57
3.1 Overview	57
3.2 Experimental Design	60
3.3 The Survey Instrument	64
3.4 Empirical Strategy	69
3.5 Results	71
3.5.1 Average Treatment Effects	71
3.5.2 Heterogeneity: Descriptive Results	75
3.5.3 Heterogeneity: Regressions	79
3.5.3.1 Predicting Take-Up of the Commitment Savings Products	80
3.5.3.2 Predicting Default and Repeat Take-Up	83
3.5.3.3 Heterogeneous Treatment Effects	87
3.6 Alternative Explanations for Default	90

3.6.1	Income Optimism	90
3.6.2	Aggregate Shocks	93
3.6.3	Limited Attention	94
3.7	Conclusion	95
A	Appendix to Chapter 1: Explaining Adoption Timing	103
B	Appendix to Chapter 2: Mathematical Proofs	107
C	Appendix to Chapter 3	113
C.1	Supplementary Figures and Tables	114
C.2	Robustness Checks	118
C.3	Survey Measurement and Marketing Material	124

List of Tables

1.1	Descriptive Statistics	24
1.2	Impact of a Regular Saver Product	26
1.3	Impact of a Regular Saver Product - Multiple Treatment Groups . . .	28
1.4	Predicting LTS Adoption Timing	31
1.5	LTS Contract Choices	31
3.1	Summary Statistics by Treatment Assignment	68
3.2	Savings Outcomes	72
3.3	Account Usage	79
3.4	ASA Pre-Order	79
3.5	Predicting Demand for Commitment	81
3.6	ASA Defaults & Repeat Take-Ups	84
3.7	Heterogeneous Treatment Effects: Change in Bank Savings	88
3.8	Income Optimism	91
C.1	Personal Savings Goals	114
C.2	Take-Up Rates	114
C.3	Quantile Regressions	115
C.4	Average Treatment Effects - Robustness to Unbalanced Covariates .	119
C.5	ASA Take-Up & Default: Robustness (Sophistication Measure) . . .	120
C.6	ASA Take-Up & Default: Robustness (Real vs. Hypoth. Incentives) .	121
C.7	Real vs. Hypothetical Incentives	122

List of Figures

2.1	Autarky Equilibrium with Partial Sophistication	43
2.2	Regular Saver Equilibrium (high autarky savings)	48
2.3	Regular Saver Equilibrium (low autarky savings)	48
2.4	The Regular Saver Equilibrium with Partial Sophistication	53
3.1	Savings Transactions: Defaulted ASA Clients	77
C.1	Distributional Effect of Treatment on Savings, Loans and Expenditures	117
C.2	Termination Fees (Chosen & Charged)	117
C.3	Test of Risk Aversion (Methodology: Binswanger (1980))	125
C.4	Illustration: Test of Cognitive Ability	125
C.5	Personal Savings Plan (All Treatment Groups)	126

Introduction

Commitment is popular. Contrary to predictions of the standard neoclassical model, the last decade has seen a surge of evidence documenting a demand for (self-) commitment contracts - roughly understood as a voluntary restriction of one's future choice set, in order to overcome intrapersonal conflicts.¹ Applications are as broad as the scope of human ambition, and range from gym memberships, diet clubs and pension savings to self-imposed binding deadlines for academic papers.² More informal arrangements include taking only a fixed amount of cash (and no credit cards) when going shopping, not keeping alcohol or chocolate in the house, and putting one's alarm clock at the other side of the room.³ In the context of developing countries, documented demand for commitment devices goes back to the literature on rotating savings and credit organisations (ROSCAs),⁴ the wandering deposit collectors of South Asia and Africa,⁵ and more recent studies on newly introduced commitment savings products.⁶

Why do people self-commit? Commitment entails the voluntary imposition of constraints on future choices, thereby putting a cost on flexibility, which is weakly welfare-reducing from a neoclassical perspective. Three types of models are frequently cited to rationalise the observed demand for commitment: Models of quasi-hyperbolic discounting (Strotz (1955), Laibson (1997) and O'Donoghue and Rabin (1999)), models of temptation and self-control (Gul and Pesendorfer (2001), Gul and Pesendorfer (2004), Banerjee and Mullainathan (2010)), and dual-self models suggesting the existence of a long-run planning self and a short-run doing self (Fudenberg and Levine (2006)). All of these models generate preferences that are inconsistent over time, and generally suggest that agents are more impatient over current trade-offs (now vs. tomorrow) than over future trade-offs (one year vs. one year plus one day). As a result, they procrastinate activities that involve immediate

¹This paper focuses purely on self-commitment. It does not address commitment contracts adopted with strategic motives with respect to others. Furthermore, the paper abstracts from commitments entered into for convenience or other immediate benefits. As an example, the purchase of Christmas gifts in October qualifies as self-commitment if the agent fears not having enough money left in December, but not if the agent's motivation is purely to avoid the Christmas rush.

²See DellaVigna and Malmendier (2006) for gym memberships as a commitment device, Benartzi and Thaler (2004) for 401(k) pension savings, and Ariely and Wertenbroch (2002) for academic assignments.

³For an overview of commitment devices, see Bryan et al. (2010). For a humorous illustration, see popular articles and Internet videos on the 'money-shredding alarm clock'.

⁴See Besley et al. (1993), Anderson and Baland (2002), Ambec and Treich (2007) or Gugerty (2007).

⁵See e.g. Besley (1995) on West Africa's susu collectors.

⁶See Ashraf et al. (2006b), Brune et al. (2011) and Dupas and Robinson (2013) for the use of withdrawal-restriction savings accounts, discussed further in Chapter 1. Also see Duflo et al. (2011) for commitment to fertilizer use via advance purchase.

costs and later rewards (saving for a new TV, going to the gym), and do too much of activities that involve immediate gratification but later costs (using high-interest credit cards, buying temptation goods). If individuals with such preferences realise their own time-inconsistency, they will have a positive willingness to pay for commitment devices which eliminate tempting options from their future choice sets (or make them more expensive), thus allowing them to follow through with their plans (to save, to eat healthily, to exercise). In theory, this will increase their welfare from an ex-ante (or long-run) perspective.

Is commitment a good idea? Especially in the development literature, the answer seems to be yes. Recent years have seen a multitude of papers promoting commitment savings in particular as a remedy for behavioural savings constraints, and thus as a possible way out of (credit-constraint based) poverty traps. Commitment savings have been hailed as increasing savings levels (Ashraf et al. (2006b)), agricultural input use (Brune et al. (2011)), pension contributions (Benartzi and Thaler (2004)), microenterprise investment (Dupas and Robinson (2013)), and chances of successful smoking cessation (Giné et al. (2010)).

The first chapter of this dissertation is concerned with the optimal design of such commitment savings products. Drawing on evidence from Bangladesh, the chapter proposes that the popularity of inflexible ROSCAs, costly deposit collectors and (partly) of microcredit and informal borrowing may simply represent a demand for a commitment savings product with fixed periodic contributions (hereafter: ‘regular saver product’), as it is commonly offered by banks in rich countries. It presents observational evidence from a microfinance institution (MFI) in Dhaka, where the introduction of a commitment savings product with fixed regular instalments was associated with a 180 percent increase in average savings contributions after five months (estimated using a Difference-in-Difference methodology). The product design mimics the fixed instalment structure found in loan repayment contracts, thus providing an incentive to make regular future deposits and smooth consumption. The study is motivated by empirical evidence suggesting that microloans and informal loans are often taken out for consumption purposes, or for recurring business expenditures - rather than as a one-off investment.⁷ With loans that are not directly required for income generation, the question arises why individuals are willing to pay substantial loan interest charges rather than choosing to save. Especially for those who borrow in frequent cycles, the long-term difference between expensive loan cycles and equivalent savings cycles reduces to (i) one initial loan disbursement and (ii) a binding fixed-instalment structure that is

⁷See e.g., Ananth et al. (2007).

rarely available in savings products.⁸ The idea that time-inconsistent agents benefit from commitment to fixed regular instalments has previously been discussed in the context of loans (Fischer and Ghatak (2010), Bauer et al. (2012)). In contrast, the existing literature on formal commitment savings products has largely focused on withdrawal restrictions (see the discussion in Section 1.1).

Shifting the focus from the design to the strength (or the ‘stakes’) of a commitment contract, the remainder of the dissertation asks: Are people good at choosing the ‘right’ commitment contract? And if not, can commitment be harmful? By construction, correctly choosing a welfare-improving contract requires some knowledge about one’s preferences, including possible biases and inconsistencies: To determine whether a contract will enable an agent to follow through with a plan, the agent needs to anticipate how his future selves will behave under the contract. Consequently, if the agent is overconfident, or imperfectly informed about his own future preferences, the contract may result in undesirable behaviour, and the agent may be hurt, rather than helped. Given that the very nature of most commitment contracts is to impose penalties (usually of a monetary or social nature) for undesirable behaviour, adopting a commitment device that is ill-suited to one’s preferences may ‘backfire’ and become a threat to welfare.⁹

The second and third chapter of this dissertation argue that commitment can be harmful if agents select into the wrong commitment contract - and that they frequently do. In Chapter 2, I develop a model of commitment under partial sophistication and stochastic income. The model shows that a savings product with fixed regular instalments can benefit agents by increasing achievable savings levels and smoothing savings contributions, as long as the ‘stakes’ of the contract (the penalty upon default) are sufficiently high to enforce compliance. Given a full set of contracts, ‘fully sophisticated’ agents (i.e., those with perfect knowledge of their time preferences) always choose contracts which are ex-ante welfare-improving. In contrast, commitment contracts can reduce welfare if the commitment is not strong enough to discipline the agent, resulting in costly default. I show that such insufficient commitment contracts are likely to be selected by time-inconsistent agents

⁸Informal arrangements like ROSCAs may constitute an exception, but they are inflexible to an individual member’s needs. Also, ROSCAs were not widely available in the study region. Deposit collectors (see Ashraf et al. (2006a)), if available, provide another alternative - but a deposit collection service does not commit the individual to deposit any particular amount, and the individual may be tempted to deposit the minimum necessary to avoid social sanctions.

⁹Consider any type of commitment contract with front-loaded fees, such as retirement savings products with acquisition or management costs. Fees are generally subtracted from the contributions during the first few years of the contract, generating a ‘J curve’ in the asset value. Canceling or defaulting on the contract during early years generates high negative returns. A similar argument can be made for front-loaded gym membership costs.

with ‘partially sophisticated’ preferences: Agents who are neither completely unaware nor fully aware of their time-inconsistency, but anywhere in between those two extremes. This effect is discontinuous at full sophistication, and may lead to substantial welfare losses from commitment even when the discrepancy between *believed* and *actual* degree of time-inconsistency is small.

In Chapter 3, I conduct a randomised field experiment with 913 individuals in the Philippines. Individuals could sign up for a new commitment savings account with fixed regular instalments, and were given the chance to choose the stakes of the contract (in form of a default penalty) themselves. I find that the *average* effect on bank savings is large and significant: The Intent-to-Treat (ITT) effect on bank savings is roughly four times that of a conventional withdrawal-restriction product that was offered as a control treatment. However, a striking feature of the results is that the *median* client appears to choose a ‘harmful’ contract: 55 percent of clients default on their savings contract, and incur the associated penalty. The magnitude and timing of defaults is difficult to reconcile with rational expectations and idiosyncratic shocks (a ‘bad luck’ scenario). Instead, it is suggestive of individuals making ‘mistakes’ in contract choice. The explanation most strongly supported by the data is that the chosen stakes were too low (the commitment was too weak) to overcome clients’ self-control problems. In addition, both take-up and default are *negatively* predicted by measures of sophisticated hyperbolic discounting, consistent with the notion that those who are fully aware of their bias realise the commitment is too weak for them, and stay away. The results from a subsequent repeat marketing stage with the offer of ‘pre-ordering’ the product for a second round support the impression that a significant share of clients took up the commitment contract by mistake. Alternative explanations for default that find some support in the data are income overoptimism and household conflict. A pure stochastic shock explanation appears unlikely.

This dissertation builds and expands on the literature in three ways. First, it provides the first causal evidence on the effect of a formal regular-instalment commitment savings product:¹⁰ Chapter 1 uses the introduction of a regular-instalment product in Bangladesh as a natural experiment to assess whether such products can help individuals in overcoming savings constraints. Motivated by the encouraging results of this case study (a large increase in savings contributions on *average*), Chapter 3 introduces a regular-instalment product in a randomised setting. If it is true that a significant share of the demand for microcredit and informal borrowing is just a demand for commitment to fixed instalments, then we should expect to

¹⁰Previous literature has studied ROSCAs (see Besley et al. (1993) and the discussion in Section 1.1), but these are informal in nature, and inflexible to an individual member’s needs.

see that the introduction of a fixed-instalment microsavings product will result in (i) substantial increases in saving and (ii) a reduction in the demand for loans. I find strong support for an increase in savings, and a large but statistically insignificant reduction of loan demand. Furthermore, the dissertation provides the first direct comparison of a regular-instalment savings product with a pure withdrawal-restriction product.¹¹ I estimate an average effect on bank savings that is about four times the effect of the withdrawal restriction product. This is consistent with the theoretical work of Amador et al. (2006), who show that when individuals value both commitment and flexibility, the optimal contract involves a minimum (per-period) savings requirement.

Second, to the author's knowledge, it is the first study to explicitly discuss heterogeneous effects (and possible welfare risks) of commitment contracts, and link these to measures of (partially) sophisticated time-inconsistency. This makes it closest in spirit to DellaVigna and Malmendier (2006), who show that U.S. consumers choose gym contracts which are cost-inefficient given their attendance frequency. It also relates to the theoretical work of DellaVigna and Malmendier (2004). In the realm of commitment savings, the literature has largely focused on positive *average* effects, highlighting the promising role that commitment savings could play in overcoming behavioural savings constraints. However, welfare inference critically depends on the distribution of effects in the population. I establish that these effects may be very heterogeneous, including the possibility of a majority being hurt by the product. The results of this study complement previous findings: Ashraf et al. (2006b) find that a withdrawal-restriction product increased savings by 81 percent on average after 12 months, but 50 percent of the 202 clients made no further deposits after the opening balance of 100 pesos. Out of 62 clients who selected an amount goal, only six reached this goal within a year, suggesting that the remainder may have their initial savings tied up indefinitely.¹² Similarly, Dupas and Robinson (2013) document that offering Kenyan women savings accounts with withdrawal restrictions led to a 45 percent increase in daily business investment on *average*, but 43 percent of women made no further deposits after opening the account. Fi-

¹¹The withdrawal-restriction product tested in this study (Gihandom) directly corresponds to the SEED product in Ashraf et al. (2006b): Terms and Conditions are identical, and the study locations are 70km (2h by local bus) apart. The magnitude of estimated effects is comparable, considering that Ashraf et al. (2006b) estimate an ITT of 411 pesos after 12 months in a sample of previous savings account holders, whereas I estimate an ITT of 148 pesos after 4.5 months (on average) in a general low-income population.

¹²Neither Ashraf et al.'s SEED product, nor the Gihandom product used in this study are fool-proof, in the sense that clients could have borrowed the goal amount for five minutes from a friend, deposited it at the bank, and received their savings back. Neither study finds any evidence that this happened.

nally, Giné et al. (2010) offered smokers in the Philippines a commitment contract for smoking cessation, in which smokers would deposit savings into a withdrawal-restriction account, and forfeit their savings to charity if they failed a nicotine test after 6 months. The authors point out that offering the commitment contract increased the likelihood of smoking cessation by 3 percentage points. Looking at heterogeneity, 66 percent of smokers who took up the product failed the nicotine test, forfeiting an average of 277 pesos in savings. Interestingly, those who defaulted on their contract had chosen much lower stakes relative to those who succeeded (successful quitters saved 1,080 pesos on average). While the direction of causality is unclear, this is consistent with the idea that individuals tend to choose commitment products which are too weak to overcome their self-control problems. In summary, a closer look at the heterogeneity behind average treatment effects in the literature reveals that adverse effects of commitment products may be widespread.

Third, Chapter 3 proposes a novel measure of sophistication for time-inconsistent agents. Previous literature has often assumed a one-to-one mapping from the take-up of a commitment product to the presence of (fully) sophisticated time-inconsistency (with the notable exception of Tarozzi and Mahajan (2011), who follow a structural approach). Such a one-to-one mapping does not allow for the possibility that individuals may take up commitment products by mistake. I propose a survey-based measure of sophistication, which relies on the interaction between observed time-inconsistency (as measured by conventional time discounting questions), and measures of self-perceived temptation.

Chapter 1

Regular Saver Products: Evidence from Bangladesh

1.1 Motivation and Context

There are many reasons why even individuals with very low incomes have the need to convert small everyday sums of money into larger lump sums. Among them are life-cycle needs (such as marriage, childbirth, education or home building), insurance against emergencies (like sickness, loss of employment, natural disasters) and business opportunities (e.g. buying land or a new machine). To obtain these lump sums, individuals use a variety of financial services – many of which are very costly. In particular, observational evidence indicates a popularity of borrowing for consumption purposes, as well as for periodic business expenses (e.g., to maintain stock).¹ In both cases, obtaining a loan is not essential to generating the income that will repay it (as opposed to one-off investments, say, purchasing a piece of land): Many business clients go through a loan cycle every month for years, paying back in small instalments every week. A savings cycle with frequent contributions differs from this in one initial loan disbursement, which pales in size compared to annualized interest rates between 100 and 500 percent.² Given such interest rates, the decision to borrow rather than to save seems puzzling.

A similar puzzle is the popularity of costly or inflexible savings mechanisms: Consider the wandering deposit collectors found in South Asia and Africa.³ Rutherford (1999) illustrates their work with the example of a woman in India who collects a fixed amount of 5 or 10 Rupees per day from a client. She returns their savings

¹See e.g., Collins et al. (2009), the Indian vegetable vendors studied in Ananth et al. (2007), or Rutherford (1999).

²See e.g., Rutherford (1999) or Ananth et al. (2007).

³See e.g. Besley (1995) on West Africa's susu collectors.

1.1. MOTIVATION AND CONTEXT

after 220 days, charging a fee of 20 days' deposits for her service. This is equivalent to an annualized interest rate of *minus* 30 percent. Demand for such services is high – but there are few individuals who have built up the trust and reputation needed to provide this service, severely limiting supply.

Finally, a savings mechanism that has been prominently featured in the literature (e.g. Besley et al. (1993), Gugerty (2007)) and is prevalent all across the developing world are rotating savings and credit organizations (ROSCAs). A ROSCA is a group of people who meet regularly to each put a fixed amount of money into a 'pot'. The content of the pot is given to a different member of the group at each meeting. While ROSCAs are costless, they are inflexible to an individual member's needs.

The obvious question to ask is: Why do people not just save at home, instead of relying on such costly devices? An answer suggested in the existing literature are hyperbolic (or quasi-hyperbolic) preferences,⁴ described in the introduction as preferences which generate impatience over current trade-offs and (relative) patience over future trade-offs. Consequently, hyperbolic agents procrastinate saving. This leads either to a failure to reach their savings target, or to a failure to smooth consumption over time by saving too much in the last minute. If hyperbolic individuals become aware of their time-inconsistency, they will be willing to pay for commitment savings products which enable them to overcome their present bias by forcing them to save.

A second factor which receives increasing attention are financial claims from relatives, neighbours and spouses. Slum communities and rural villages are often close-knit social networks. If an individual has managed to save money at home, she may be informally obligated to help out a neighbour who has fallen on hard times, a relative who needs medical treatment, or simply a husband who claims the remains of the household budget for personal expenses.⁵ Awareness of claims on liquid assets may sharply reduce incentives to save in the first place. As in the previous case, such individuals may develop a preference for commitment to 'safeguard' their savings.

In recent years, a burgeoning literature on commitment savings has confirmed the empirical link between hyperbolic discounting, intra-household bargaining and a preference for commitment: Bauer et al. (2012) find that individuals in their sample who have been identified as hyperbolic keep a lower proportion of their savings

⁴Most prominently, Laibson (1997) and O'Donoghue and Rabin (1999).

⁵See the literature on informal insurance and need-based gift exchange, e.g. Morduch (1999).

at home and are more likely to borrow from microcredit organizations.⁶ Ashraf et al. (2006b) find that hyperbolic preferences predict take-up of a commitment savings product which features withdrawal restrictions. Gugerty (2007) finds ROSCA participation patterns to be consistent with both time-inconsistent preferences and low intra-household bargaining power (only husband has formal sector income) – i.e., individuals value public pressure to save. Finally, Anderson and Baland (2002) find an inverted U-shaped relationship between women’s autonomy from their husbands (as proxied by female share of household income) and participation in ROSCAs: Participation in ROSCAs is high when female autonomy is high enough to persuade their husband to let them join, but not so high that the protection which the ROSCA provides is no longer needed.

This case study provides suggestive evidence that the demand for costly or inflexible informal savings mechanisms (such as ROSCAs and deposit collectors), as well as the observed popularity of consumption loans and cyclical business loans, may (partly) reflect a demand for a commitment savings product with fixed regular instalments. Using a model of hyperbolic discounting,⁷ it can be shown that individuals cannot reach their welfare-maximising level of savings when saving at home, and that a regular saver product can increase their achievable savings level, as well as smooth savings contributions over time. Using a novel dataset with administrative savings data for 2006-2010 from a small MFI in Bangladesh, it is found that the introduction of a regular saver product is indeed associated with a large increase in individual’s average savings contributions. The intuition is straightforward: If individuals are present-biased, the ability to tie themselves to a particular savings schedule which they would not be able to stick to otherwise improves their welfare by enabling them to afford lump-sum expenditures (such as school tuition or a house repair). It may further improve their welfare by smoothing savings (and consumption) over time. Regarding relatives or spouses’ claims, a regular saver product will help by restricting withdrawals on past savings. For present savings, if the penalties associated with default are high enough, the individual may be able to argue that this money is ‘accounted for’ and cannot be used for a different purpose.

While this view is neither new nor surprising, the existing literature on formal commitment savings has focused almost exclusively on savings products featuring

⁶They also attribute the preference of present-biased women for microcredit over other sources of borrowing to the structure and support that microcredit provides. However, the central feature of fixed regular repayments is shared by virtually all other common forms of borrowing.

⁷The model originally developed for this case study is omitted here, as it is nested in the theoretical framework in Chapter 2. It corresponds to the special case of full sophistication ($\tilde{\beta} = \beta$) and deterministic income ($\lambda = 0$). The original model is available in Hofmann (2014).

1.1. MOTIVATION AND CONTEXT

withdrawal restrictions, but without an obligation to deposit in the first place. Notable examples include Ashraf et al. (2006b), who find a positive effect on savings levels using a withdrawal-restriction product in the Philippines. Dupas and Robinson (2013) study a similar (though not as strict) withdrawal-restriction product in Kenya and find increased investment in clients' small businesses. Brune et al. (2011) randomly offered withdrawal-restriction accounts to smallholder cash crop farmers, and found a positive effect on agricultural input use, as well as on subsequent crop sales and household expenditures. Other notable savings interventions included text message reminders (Karlan et al. (2010)) and deposit collectors (Ashraf et al. (2006a), Giné et al. (2010)).⁸ However, to the author's knowledge, none of the formal commitment savings accounts which have been studied so far commit clients to make future deposits.

For a hyperbolic discounter, simply giving him the opportunity to restrict his withdrawals will not do - he will be prevented from touching past savings, but will have no added incentive to contribute further. Similarly, someone constrained by relatives' claims can protect past savings, but not current income (before visiting the bank). Benartzi and Thaler (2004)'s 'SMarT' program in the U.S. allows individuals to commit to allocating their future salary increases towards their retirement savings. This comes closest to the objective of this study, in the sense that the program imposes fixed (or even increasing) savings contributions. However, the context of defined-contribution pension plans in the U.S. is very different from the Bangladeshi slum context studied here.

This study takes advantage of a 'natural experiment' following the introduction of a commitment savings product with fixed regular contributions by the MFI *SafeSave* in a slum community in Dhaka, Bangladesh, in June 2009. Using administrative savings data obtained from the MFI for the period 2006-2010, the case study presents suggestive evidence on the effect of the regular saver product on individuals' savings contributions. In a classical differences-in-differences estimation, I find that savings contributions during a five-months period increase by 523Tk (about U.S. \$7) more relative to past year savings than those of a 'control group' formed by later adopters of the product. This represents a 180 percent increase compared to average previous year savings. Allowing for multiple 'treatment groups' depending on month of entry into the product further shows a significant effect for four or more months of treatment.

⁸Also see Beshears et al. (2011) for a comparison of different withdrawal restrictions, Dupas and Robinson (2011) for the effect of informal savings technologies on health investment, and Bryan et al. (2010) for a review of commitment devices beyond savings.

The estimates are to be taken with a pinch of salt: The data comes from a field experiment, where selection into ‘treatment’ with the product is not randomly assigned, but self-selected. The resulting endogeneity is kept to a minimum by using early adopters of the product as a treatment group and late adopters (who had the product for only part of the period of observation) as a control group. The paper takes the remaining endogeneity of adoption timing seriously and rules out as many differences between early and late adopters as the data permits. I find that the estimated treatment effect is robust and even increases when various controls as well as pre-existing time trends for the treatment groups are included. Unfortunately, it is impossible to control for all differences between early and late adopters. Hence, this analysis of the effect of a commitment savings product with fixed regular contributions is best interpreted as a starting point for further research, and as a motivation to improve on the estimates presented here using randomized control methodologies (this is done in Chapter 3).

While it would ultimately be interesting to compare the performance of a savings product with restricted withdrawals to a savings product with fixed regular contributions (also done in Chapter 3), this case study focuses on the question whether regular saver products can increase savings levels (by overcoming savings constraints) relative to autarky. Severe data limitations mean that the analysis is largely focused on average effects: While the dataset is novel and stretches over four years, it is of administrative nature, and individual characteristics are largely unobserved. When discussing time-inconsistency, the case study further abstracts from the issue of sophistication, which will be crucial for the later parts of this dissertation. Both the heterogeneity of treatment effects and the consequences of (partially) sophisticated time-inconsistency will be addressed in the theoretical framework in Chapter 2 as well as the experimental setting in Chapter 3.

1.2 The Data

The previous section suggested that individuals subject to short-term consumption temptations and financial claims from relatives may benefit from a savings product which enables them to commit to fixed contributions every period. The remainder of Chapter 1 studies one such product: The Long Term Savings (LTS) product offered by *SafeSave*, an MFI located in Dhaka, Bangladesh. *SafeSave*, which was founded by microfinance researcher Stuart Rutherford in 1996, aims to provide day-to-day money management services such as basic savings and loan in-

struments to the poor.⁹ In June 2009, *SafeSave* introduced the LTS product, enabling clients to save either 50Tk or a multiple of 100Tk (about U.S. \$1.3) each month for a chosen contract length of 3, 5, 7 or 10 years. A particular feature of *SafeSave* is branchless banking: ‘Field officers’ visit each client’s house every working day to collect general savings, and once a month LTS savings. The only time clients need to visit a bank branch is to open an account (general or LTS), or to take withdrawals or loans in excess of 500Tk. These field officers informed clients about the new LTS product in June 2009.

The analysis focuses on the 333 clients who opened an LTS account at a particular *SafeSave* branch (Bauniaband branch, the first to offer LTS) between June 2009 and February 2010, conditional on having had a general savings account since August 2008 or longer.¹⁰ The variable of interest is the increase in clients’ savings contributions after LTS adoption: Do clients save more when they use LTS? Data is available on clients’ general savings level plus LTS savings as of 31 August and 31 January each year, between August 2006 and January 2010. This makes it possible to compare savings contributions made during the five-month period from August to January, for the years 2006 to 2009.¹¹ During the last of these four periods, from August 2009 to January 2010, LTS was available. This paper explores the variation contained in the timing of LTS adoption: Those who adopted LTS soon after its introduction in June 2009 were ‘treated’ with the product for the full five month period of interest. Those clients who adopted LTS later in 2009 received ‘partial treatment’ in the sense that they used LTS for only a fraction of the period from August 2009 to January 2010. The analysis then compares whether the savings contributions of ‘fully treated’ individuals increased more relative to previous years’ contributions than the savings of ‘partially treated’ or untreated individuals did. The obvious difficulty with approaching LTS as a treatment is that assignment into the treatment group does not occur randomly, as individuals select into the treatment group by adopting early. The effects that can arise from such endogeneity of treatment are discussed in Section 1.3.2 and Appendix A.

Two features of the LTS product design are worth noting for their divergence from the later theory framework: First, interest rates are not perfectly identical across commitment and non-commitment accounts. Interest rates for LTS are 7, 8, 9, and 10 percent for a 3, 5, 7, and 10 year contract, respectively, where the aver-

⁹For more details on *SafeSave*, see Rutherford (1999) and www.safesave.org.

¹⁰Of originally 334 clients, one was excluded as an outlier, as her savings in 2009 deviated from the average in 2009 by -19.6 standard deviations, or -32,709Tk.

¹¹Note that contributions made during January to August periods are not considered, as they span seven months and may contain seasonal effects.

age chosen contract length is 5.3 years. With inflation ranging around 9 percent in 2008 and 2009, the real interest rate on LTS is negative or roughly zero. It seems unlikely that clients adopted LTS to increase the size of the pie, rather than for its commitment features. However, it cannot be ruled out that individuals chose LTS to evade inflation, not because they wanted or needed commitment. In comparison, the (non-commitment) general savings account pays 6 percent interest without any commitment.¹² Given the uncertainty of most slum dwellers' incomes and consumption needs, it seems doubtful whether a client who does not desire commitment would sign a 5-year contract for an additional 2 percent of interest.

Secondly, the penalty upon canceling the LTS contract is complex: If a client runs three months or less behind on payments, a small nominal 'service fee' needs to be paid to catch up on any past missed payments. If he is more than three months behind, the LTS contract is canceled, and the penalty imposed is that he loses all LTS interest payments above the general account interest rate, with retrospective effect since he opened his LTS. For simplicity, the three-month grace period may be thought of as one period within a theoretical framework. The penalty for canceling LTS is more 'lost bonus' than actual penalty, in the sense that agents receive the same returns which they would have received, had they never joined LTS. If agents are rational, this will reduce the efficacy of the product. For behavioural agents who exhibit reference dependence and loss aversion, the penalty would still be effective: While holding an LTS, individuals are likely to adjust their reference point to the higher interest payments they believe to accumulate over time, and will be reluctant to give them up. This ambiguity in the design of the commitment product will be corrected in the experimental setting in Chapter 3.

1.3 Estimation

1.3.1 The Differences-in-Differences Approach

An obvious first step to estimate the effect that LTS has on client's savings is to look at how much early adopters saved between August 2009 and January 2010, i.e., while using LTS, and compare this to their savings during previous years' August to January periods. To isolate time effects, this change in savings contributions can be compared to the change in savings of a control group during the same time period. The resulting differences-in-differences (DD) estimate can be interpreted as a causal treatment effect only under the strong assumption that the treatment group

¹²This interest rate of 6 percent requires a savings balance of at least 1000Tk – a criterion that 285 of 333 clients in our sample satisfied as of August 2009.

would have evolved identically to the control group, had they not been treated (i.e., had they not used LTS). The plausibility of this assumption will be discussed in Section 1.3.2. The first part of the analysis considers all clients who adopted LTS between June 2009 and September 2009 ('early adopters') as 'treated', while those who adopted LTS between October 2009 and February 2010 ('late adopters') are used as a control group. Note that LTS deposits must be made on the 20th of each month, and 'adopting LTS in September' means that clients chose to first contribute to LTS on 20th of September. Hence, clients adopting LTS in September can be considered as 'fully treated' in the sense that they contributed on all LTS payment dates between 31 August 2009 and 31 January 2010.

Table 1.1 presents descriptive statistics of clients' characteristics (as far as available), the LTS contract options they chose, and their average savings contributions by group. Column (1) of Table 1.2 presents the regression estimates of a DD approach generalized to multiple time periods, which at this point should be interpreted as correlations only. The regression equation is

$$S_{it} = \alpha + \gamma T_i + \lambda_t + \beta D_{it} + \epsilon_{it}$$

where S_{it} are an individual's savings contributions (general plus LTS savings) in period t , i.e. the difference between their savings balances in January of year $t + 1$ and August of year t . T_i is a dummy that switches on for individuals who adopted LTS in or before September 2009, λ_t is a vector of year dummies for 2006-2009, and D_{it} is a treatment dummy that switches on for observations of treated individuals in 2009. The estimated $\hat{\beta}$ shows that, on average, early adopters increased their savings contributions by 523 Tk (about \$7) more relative to past year savings than late adopters did during the treatment period. At average contributions of 290Tk in the previous year, this implies an increase in savings of 180 percent. The treatment group coefficient $\hat{\gamma}$ is not significant: Early adopters did not appear to save more (in contribution levels) than late adopters before LTS was introduced.

TABLE 1.1: DESCRIPTIVE STATISTICS

	Early Adopters	Late Adopters	All
Age	30.6	32.9	31.4
Female	89.7	89	89.5
<i>Occupation (percent)</i>			
Housewife	66.5	52.3	61.9
Small Business	6.7	15.6	9.6

continued on next page

1.3. ESTIMATION

	Early Adopters	Late Adopters	All
Private Service	11.6	9.1	10.8
Handicrafts	6.7	12.8	8.7
Student	8.5	10.2	9.0
<i>LTS Entry Month (percent)</i>			
June 2009			27.9
July 2009			17.7
August 2009			11.4
September 2009			10.2
October 2009			7.2
November 2009			4.5
December 2009			3.0
January 2010			11.7
February 2010			6.3
<i>LTS Contract Details</i>			
Contract Length (years)	5.1	5.76	5.3
Monthly Deposit (Taka)	210	235	218
<i>Purpose of LTS Contract</i>			
Education	56	51.4	54.6
Ornaments	1.4	1.8	1.5
Land Purchase	5.8	9.2	6.9
House Construction	1.3	0.9	1.2
Medical Treatment	4.9	0.9	3.6
Marriage	10.7	12.8	11.4
Business	19.9	23	20.8
<i>Mean Savings Contributions</i>			
Aug 2006 - Jan 2007	318	197	277
Aug 2007 - Jan 2008	323	301	316
Aug 2008 - Jan 2009	292	285	290
Aug 2009 - Jan 2010	1064	499	879
Observations	224	109	333

Notes: "Early Adopters" are all clients who adopted an LTS contract at Bauniaband branch between June and September 2009. "Late adopters" are all clients who adopted an LTS contract at Bauniaband branch between October 2009 and February 2010. While income data was not available on client level, bank staff estimate average monthly household income between 5000 and 6000 Tk.

1.3. ESTIMATION

TABLE 1.2: IMPACT OF A REGULAR SAVER PRODUCT
Dependent Variable: Change in Total Savings from August (t) to January (t+1)

	(1)	(2)	(3)
Treatment group*2009	523.1*** (195)	654.6*** (209)	750.9*** (264)
Treatment group	42.45 (58.3)	39.90 (61.9)	148.3 (240)
Treatment group trend			-51.16 (106)
Female		4.513 (174)	4.570 (134)
Age		24.70* (13.0)	24.70* (14.6)
Age ²		-0.308** (0.15)	-0.307* (0.18)
Interaction Area*2009	NO	YES	YES
Time effects	YES	YES	YES
Occupation FE	NO	YES	YES
Robust SEs	YES	YES	NO
R-squared	0.06	0.08	0.08
Observations	1185	1185	1185

*significant at 10 percent; **significant at 5 percent; *** significant at 1 percent. Exchange rate is about 100Tk for US \$1.3.

The cut-off between early and late adopters may seem arbitrary, given that joining LTS was possible every month. Alternatively, the concept of treatment and control group can be generalized to allow for separate treatment groups depending on month of LTS entry. This allows us to test whether savings increase in the number of months that a client used LTS – in other words, whether the coefficient on treatment in the savings equation increases in the number of months that a client was ‘treated’ with LTS.¹³ The regression equation becomes

$$S_{it} = \alpha + \gamma_m T_i + \lambda_t + \beta_m D_{it} + \epsilon_{it}$$

¹³This is similar in concept (but much simpler in implementation) to the approach of Duflo (2001), who showed that a higher degree of exposure to a school construction program in Indonesia had increased effects on wages and education.

where m indexes months of treatment. The equation contains a vector of treatment group effects γ_m (dummy T_i is one if individual i was “treated” m months), a vector of time effects λ_t , and a vector of treatment effects β_m , where D_{it} is a vector of dummies that switches on for observations in 2009 for individuals treated for m months. Treatment varies at the individual level, and the analysis explicitly assumes that observations are independent at the individual level, i.e., there is no group structure to the time-individual specific errors ϵ_{it} (see Section 1.3.3 for concerns regarding the error structure).

Column (1) of Table 1.3 reports the regression estimates of a DD with five treatment groups, ranging from two to six months of treatment. One and zero months of treatment (January and February adopters) have been combined as a control group. The six months group (June-August adopters) has been added to account for possible differences between those who already held an LTS when the period of observation started on 31 August, and those who only joined LTS on the first relevant payment day on 20th September. The regression in column (1) shows an even stronger correlation of savings with treatment than before: While the effects for six, five and four months are highly significant, the coefficient on treatment peaks at 1177Tk (\$15) for five months treatment, i.e., September adopters, and then roughly decreases in level and significance as duration of treatment drops. It is interesting to observe that all group and time effects are insignificant – suggesting that levels of savings contributions were roughly the same across groups before treatment.

1.3.2 Endogeneity Concerns

The previous section pointed out that ‘treatment’ in the form of early LTS adoption is highly correlated with higher savings, but has refrained from making any statements of causal nature. A necessary requirement to make causal statements in an OLS regression is that the coefficient of interest can be estimated consistently, which in turn requires that regressors are uncorrelated with the error term. In our case, this requirement asks that treatment is orthogonal to anything that determines savings. This may not be plausible considering that treatment is self-selected, i.e., there may be omitted variables in the error term which cause individuals to both adopt the product early, and to save more. Much of the appeal of DD estimation comes from the fact that it avoids many of the endogeneity problems commonly associated with comparisons between heterogeneous individuals: In particular, omitted variables that are either time-invariant or group-invariant will not affect the validity of a DD estimation, even if they are correlated with treatment or take-up.

TABLE 1.3: IMPACT OF A REGULAR SAVER PRODUCT - MULTIPLE TREATMENT GROUPS
Dependent Variable: Change in Total Savings from August (t) to January (t+1)

	(1)	(2)	(3)
<i>Treatment Effects</i>			
Six months treatment	759.2*** (262)	863.9*** (275)	816.3** (331)
Five months treatment	1177*** (404)	1265*** (410)	1094** (474)
Four months treatment	963.6*** (334)	885.9*** (319)	-0.476 (533)
Three months treatment	306.3 (416)	263.4 (422)	210.2 (632)
Two months treatment	518.7 (435)	529.7 (445)	851.0 (755)
<i>Group Effects</i>			
Six months	51.44 (61.0)	46.14 (67.4)	-5.095 (298)
Five months	-54.91 (152)	-77.43 (159)	-269.6 (443)
Four months	-58.76 (128)	-96.10 (134)	-1125** (507)
Three months	127.3 (109)	113.5 (111)	53.23 (577)
Two months	-154.0 (147)	-103.9 (160)	286.6 (735)
<i>Group specific trends</i>			
Trend (6 months)			24.75 (132)
Trend (5 months)			90.97 (195)
Trend (4 months)			478.5** (222)
Trend (3 months)			27.98 (257)
Trend (2 months)			-177.2 (318)
Female		-6.061 (174)	-6.042 (135)
Age		23.42* (13.8)	24.00 (15.1)
Age ²		-0.292* (0.16)	-0.299 (0.19)
Interaction Area*2009	NO	YES	YES
Time effects	YES	YES	YES
Occupation FE	NO	YES	YES
Observations	1185	1185	1185
R-squared	0.07	0.09	0.10
Robust SEs	YES	YES	NO

*significant at 10 percent; **significant at 5 percent; *** significant at 1 percent.

Note: Exchange rate is about 100Tk for US \$1.3.

Starting from the equation

$$S_{it} = \alpha_i + \lambda_t + \beta D_{it} + \gamma z_i + \epsilon_{it}$$

where z_i is unobserved, the classic two-period DD estimator $\hat{\beta} = \Delta \bar{S}_t^T - \Delta \bar{S}_t^C$ differences out any effects that are constant either across time or across individuals. As an illustration, suppose that higher income causes early take-up, and higher income also causes higher savings levels, but it does not cause different savings time trends. Then the effect of income is time-invariant, and thus income will not bias the DD estimator. The real problem for the DD approach adopted are unobservables which cause both early adoption *and* a change in the savings time trend. The following decomposition for the two-period case illustrates this: The treatment effect that we are ultimately interested in is

$$E[S_{it}^T - S_{it}^{NT} \mid T]$$

where $E[S_{it}^{NT} \mid T]$ are the expected savings of a treatment group individual in *absence* of treatment. The DD coefficient that we estimate converges to

$$E[S_{it}^T - S_{i,t-1}^{NT} \mid T] - E[S_{it}^{NT} - S_{i,t-1}^{NT} \mid NT]$$

which can be decomposed into

$$\underbrace{E[S_{it}^T - S_{it}^{NT} \mid T]}_{\text{Treatment Effect}} + \underbrace{E[S_{it}^{NT} - S_{i,t-1}^{NT} \mid T] - E[S_{it}^{NT} - S_{i,t-1}^{NT} \mid NT]}_{\text{Selection Bias}}.$$

In other words, the identifying assumption is that

$$E[S_{it}^{NT} - S_{i,t-1}^{NT} \mid T] = E[S_{it}^{NT} - S_{i,t-1}^{NT} \mid NT].$$

The estimator is biased if the (counterfactual) change in savings for the treatment group absent treatment is different from the change in savings for the control group.¹⁴ Given the lack of a suitable instrumental variable, this case study will not be able to fully solve the endogeneity problem. Instead, the remainder of this section will take a closer look at the selection bias by suggesting possible differences between early and late adopters, controlling for them where possible, and discussing their implications for savings behaviour.

¹⁴See Besley and Case (2000) for a clear discussion of endogenous treatment and possible remedies in the case of policy choice.

Shedding Light on LTS Adoption

The most obvious step in analysing take-up is to test whether observable characteristics can predict the timing of LTS adoption. Table 1.4 presents the results of regressing LTS entry month on age, occupation, gender and area of residence. Likewise, it may be interesting to look at correlations of contract choices, i.e., whether early adopters chose higher monthly contributions. Table 1.5 presents a regression of entry month on other choices – which is purely to illustrate correlations and should be interpreted as such.

Table 1.5 suggests that early adopters were more likely to choose shorter contracts, but committed to roughly similar monthly contributions. They also opened their general bank accounts at the same time. Table 1.4 reveals that while age, gender and occupation are not associated with adoption timing, the area where clients live strongly predicts when they took up the LTS product. The importance of area may seem striking, considering that *SafeSave* only operates within a one kilometre radius of a bank branch (in this case, Bauniaband branch), so the areas considered are relatively small and close together. The addition of a deposit-collection service means that transaction costs are independent of area. At first sight, this may look like an omitted variable bias: For instance, area may be correlated with income, which may influence take-up. While individual income data is not available, observational evidence from field officers who visit each client every day suggests that all areas are comparatively poor. Second, insights such as those from Bertrand et al. (2004b) suggest that situational factors such as walking distance to the branch (to sign an LTS contract) may be important for take-up. Walking distances are Block A (5min), Block B (10min), Block C (16min), Block D (12min), Block E (12min), Kurmitola and Mirpur (5-15min). With Block C being the farthest from the branch and at the same time the earliest to adopt LTS, the distance explanation does not seem plausible.

There are various other explanations for a ‘patchwork pattern’ of adoption, such as herding behaviour and other mechanisms of technology adoption.¹⁵ Appendix A discusses these explanations in detail and shows that, in particular with herding behaviour, random factors may lead to a profitable innovation spreading very quickly in one area, and not at all or with considerable delay in another area, even if individuals are completely homogeneous in their savings preferences (implying zero selection bias). Factors which may determine adoption timing independent of area include different degrees of risk aversion and the amount of benefit which individuals derive from the product: Suppose information about the manageability

¹⁵Compare Besley and Case (1994) and Foster and Rosenzweig (1995).

1.3. ESTIMATION

TABLE 1.4
PREDICTING LTS ADOPTION TIMING

Dependent Variable: LTS Entry Month (1=June 2009)	
Female	-0.0859 (0.55)
Age	-0.00975 (0.061)
Age ²	0.000243 (0.00077)
<i>Occupation FE</i>	
Housewife	0.151 (0.65)
Handicrafts	0.972 (0.77)
Small Business	0.462 (0.80)
Private Service	0.0203 (0.75)
<i>Area of Residence</i>	
Bauniaband Block B	-1.302*** (0.43)
Bauniaband Block C	-1.396*** (0.42)
Bauniaband Block D	-0.587 (0.49)
Bauniaband Block E	-0.574 (0.48)
Kurmitola	2.374*** (0.90)
Mirpur	1.617* (0.92)
Constant	4.228*** (0.97)
Observations	333
R-squared	0.12

*significant at 10 percent; **significant at 5 percent; *** significant at 1 percent.

Notes: The omitted category for occupation is student. The omitted category for area is Bauniaband Block A.

TABLE 1.5
LTS CONTRACT CHOICES

Dependent Variable: LTS Entry Month (1=June 2009)	
Quarter of bank entry (1=2000Q1)	-0.00972 (0.016)
Contract length (yrs)	0.161** (0.070)
Monthly deposit (Tk)	0.0000952 (0.00088)
<i>Purpose of LTS contract</i>	
Marriage	-0.573 (0.54)
Education	-0.710* (0.38)
Ornaments	0.172 (1.37)
Land Purchase	0.0270 (0.67)
House Construction	-0.194 (1.38)
Medical Treatment	-1.413* (0.85)
Constant	3.558*** (0.63)
Observations	333
R-squared	0.04

*significant at 10 percent;

**significant at 5 percent;

***significant at 1 percent.

Notes: The omitted category for purpose is Business.

of LTS is a public good. Early adopters will be those who derive the largest potential benefit from the product, and thus have the highest stake in learning about it - for instance, because their time-inconsistency prevents them from achieving a desired savings goal (commonly a lump-sum expenditure) without commitment, but a commitment product with fixed instalments does make their savings goal achievable (anticipating the theoretical framework in Chapter 2, this corresponds to the range $\beta \in [\hat{\beta}_B, \hat{\beta})$). Individuals who could have reached their savings goal without commitment, but who benefit from commitment through smoother savings contributions (in theory terms, $\beta \geq \hat{\beta}$), may still adopt LTS, but wait for others to adopt first. Appendix A elaborates on this argument and shows that the implications for a possible selection bias are ambiguous, instead pointing towards heterogeneous treatment effects.

Unfortunately, the data do not allow a further disentanglement of which factors were at work in adoption timing. Most plausibly, it was a mixture of the above plus some factors which have not been discussed - like shocks that occurred simultaneously with treatment. Ultimately, what matters for the estimated treatment effect is whether the factors which determined adoption also determined savings trends absent treatment. The following subsection will attempt to control for two sources of selection bias.

1.3.3 Adjusting the Estimation Strategy

How can we eliminate or at least reduce a possible selection bias in order to isolate the treatment effect? First, the effect of area can be controlled for by interacting the treatment period (2009) with the area dummies. Suppose time-individual specific effects were such that clients in areas who adopted early (like Block C) would have increased their savings even without treatment. Early adoption leads to disproportional representation of these clients in the treatment group and biases the estimated treatment effect. Including 2009*area as a control will capture this effect. Column (2) of Tables 1.2 and 1.3 show that including these area interactions further *increases* the treatment effect, and suggests that early adoption blocks B, D and E would have *decreased* their savings absent treatment.

A second source of selection bias that can be controlled for are pre-existing trends. In other words, if early adopters are different from late adopters in a way that causes them to follow different savings time trends before treatment, this can be controlled for by including a time trend specifically for the treatment group. Column (3) of Tables 1.2 and 1.3 show that the estimated treatment effects are largely robust to including treatment-specific time trends. Time trends are mostly insignif-

icant, suggesting that savings patterns of the treatment and the control group were comparable before treatment. Clients treated for four months (October entrants) constitute a special case, as their treatment effect is fully absorbed by a large negative group effect and a large positive time trend. This is almost certainly due to the fact that the regressors explain only four periods of data, and would disappear with a longer time series.

An important source of bias remains in time-individual specific shocks which occurred simultaneous with LTS introduction: For instance, if there was a change in circumstance for a particular group of people which caused them to both adopt early and increase their savings (absent treatment), then the estimate of the treatment effect is biased upwards. Unless these shocks target specific age, occupation or area groups, the scarce nature of the data prevents attempts to control for the resulting bias. As a consequence, the treatment regressor is plausibly still correlated with the time-individual specific error. The logical next step is an instrumental variable analysis – but lacking a suitable instrument, this step has to be postponed until more comprehensive data is available.

Concerns about Standard Errors

After discussing at length a possible bias of the treatment effect estimator, it is worth taking a closer look at the accuracy of the standard errors. Bertrand et al. (2004a) point out that serial correlation in the dependent variable severely biases standard errors in many DD analyses: If the dependent variable is positively correlated, OLS overstates the amount of information contained in the data and thus understates standard errors. This bias is particularly severe if the treatment variable is of a ‘switch-on’ nature and thus itself serially correlated. The bias increases in the length of the time series. Applying these factors to the present analysis, serial correlation can likely be ruled out as a significant source of bias: The dependent variable, savings contributions, is the first difference obtained from savings balances. Unless savings balance is an explosive time series, savings contributions will not exhibit positive serial correlation. If savings balance does not have a unit root but exhibits autocorrelation of $0 < \rho < 1$, savings contributions will be *negatively* correlated, leading to an *overstatement* of standard errors. However, the time span of the data is short (4 periods) and ends after the treatment period, preventing autocorrelation of treatment. Therefore, it seems unlikely that serial correlation of any sign had a significant effect on standard errors.

Another frequent concern in the DD literature is the clustering of the error terms. Grouped error terms are a concern when the outcome of interest varies at

the individual level while treatment varies only at the group level.¹⁶ This is not the case here, as both savings and treatment vary for each individual.

A final concern addressed here is heteroscedasticity of regression residuals, which may result from heteroscedastic errors in the true model, or simply from a nonlinear conditional expectation function $E[S_{it} \mid D_{it}, X_{it}]$. While robust standard errors improve on conventional ones asymptotically, Angrist and Pischke (2009) show that when heteroscedasticity is modest, robust standard errors are more biased than OLS errors. They propose a rule-of-thumb to use the maximum of conventional and robust standard errors, which has been applied throughout the analysis: For every regression column in Tables 1.2 and 1.3, the maximum of conventional and robust standard errors was reported, with the choice of errors indicated in the last row of the table.

1.4 Conclusion and Outlook

This case study has presented some first empirical evidence on the effect of a commitment savings product with fixed periodic contributions in a slum in Dhaka, Bangladesh. The analysis indicates that clients who used the product during the entire observation period increased their savings significantly relative to clients who only used the product for a part of the observation period. I also find that the effect of the product is significant for four or more months of treatment.

In attempting to understand the nature of the endogeneity, the study finds a significant effect of area on adoption timing, but no correlation of adoption timing with other savings choices, such as monthly contribution or bank entry timing. I also estimate that early and late adopters follow the same savings trends before treatment. A story of herding behaviour or other technology adoption mechanisms would be consistent with this evidence, explaining differences in adoption timing without necessarily implying a selection bias. However, the endogeneity problem cannot be fully resolved due to possible unobservable time-individual specific effects. A related issue is the potential heterogeneity of treatment effects: While the benefit of the product may consist in reaching a previously unattainable savings goal for some individuals, it may simply consist in smoothing savings contributions for others (this can be motivated theoretically, see Chapter 2).

The objective of the study is to provide a motivation for further research into regular saver products, and to tackle the plethora of questions it does not answer: First, a randomized controlled experiment would help to account for unobserved

¹⁶See Angrist and Pischke (2009), Chapter 8.2.

1.4. CONCLUSION AND OUTLOOK

determinants of participation by giving the product only to a fraction of those who want it, and using the others as a control group. Second, the present analysis lacks data on individuals' time preferences. Surveys and lab experiments are needed to identify individuals' time preferences as well as the strength of the social claims on their savings, in order to link hyperbolic discounting and relatives claims' to a demand for regular saver products. Third, it is desirable to more explicitly compare the performance of different types of commitment savings products – e.g., high frequency versus low frequency deposit products, versus the pure withdrawal-restriction products which have been a subject of academic interest in recent years. The randomized field experiment in Chapter 3 attempts to answer some of these questions.

Chapter 2

Theoretical Framework: Commitment under Partial Sophistication

The following chapter develops a formal understanding of the interaction between commitment and partial sophistication. Focusing specifically on a regular-installment savings product, it sheds light on (i) why sophisticated hyperbolic discounters can benefit from commitment to fixed instalments, (ii) why commitment reduces welfare if it is too weak to be effective, (iii) why partially sophisticated hyperbolic discounters are likely to select into such weak commitment contracts, and (iv) why those with high perceived degrees of time-inconsistency may avoid commitment. The autarky setting (Section 2.2) is based on the autarky savings framework in Basu (2014), generalised to allow for partial sophistication and stochastic income (creating a need for flexibility). Using this setting as a benchmark, I then introduce a commitment savings product with regular instalments (Section 2.3). The commitment design differs from previous models of commitment products by a default penalty that is conditional on per-period contributions, and the simultaneous absence of any withdrawal restrictions. A version of the regular-installment design with full sophistication and deterministic income was developed in Hofmann (2014) (it is nested in the current model).

Consider an agent who chooses whether to save for a nondivisible good which costs the lump-sum $2 < p < 3$ and yields a benefit $b > 3$. The agent lives for 3 periods and receives a per period income of 1 (barring shocks), which he can either consume or save. He cannot borrow. His instantaneous utility is twice differentiable and strictly concave, with $u'(c) > 0$, $u''(c) < 0$, and $u'(0) = \infty$. Throughout, assume the interest rate is $R = 1$ and $\delta = 1$ for simplicity. Define s_t as the amount

of savings that he sends from period t to $t + 1$, so that $c_t = y_t + s_{t-1} - s_t \geq 0$. Lifetime utility as evaluated in each period is given by the discounted sum of the instantaneous utilities:

$$U_t = u(c_t) + \beta \sum_{k=t+1}^3 u(c_k)$$

For $\beta < 1$, the agent is *present-biased*: He exhibits a lower rate of discount over current trade-offs (t vs. $t + 1$) than over future trade-offs ($t + s$ vs. $t + s + 1$, $s > 0$). Following O'Donoghue and Rabin (1999), the agent's degree of sophistication about his present bias is captured in the parameter $\tilde{\beta} \in [\beta, 1]$, which he believes he will use in all future periods. In particular, the agent believes in period t that his utility function in period $t + s$ will be

$$U_{t+s} = u(c_{t+s}) + \tilde{\beta} \sum_{k=t+s+1}^3 u(c_k).$$

For a fully sophisticated agent, $\tilde{\beta} = \beta$. A fully naive agent believes he will behave time-consistently in the future, captured in $\tilde{\beta} = 1$.

A need for flexibility is introduced through stochastic income shocks: With a per-period probability of λ , the agent loses his income in that period. This shock has a variety of interpretations: It can be interpreted directly as a loss of income, e.g., from redundancy, bad business, or illness of an income-earning household member. With a minor modification, it can be interpreted as a reduced-form taste shock: Suppose the sudden illness of a family member changes preferences such that utility stays unchanged if a hospital visit (at cost 1) is consumed and paid for, and drops to $u(c) = -\infty$ without a hospital visit. The implication of a shock is that the agent's lifetime income drops to (at most) 2, which means the nondivisible good can no longer be purchased. When a shock hits, any plans to save for the nondivisible are abandoned, and existing savings are optimally spread over the remaining periods for consumption. This results in a third interpretation: More generally, the shock λ corresponds to the probability that, for any time-consistent reason, the agent no longer finds it optimal to save for the nondivisible.¹

While there is much ambiguity over the definition of welfare for time-inconsistent agents, the paper will follow the convention proposed by O'Donoghue and Rabin (1999): An agent's welfare is evaluated from an ex-ante perspective, and corresponds to the lifetime utility of the period 0 agent:

¹Another time-consistent explanation why an agent may no longer wish to purchase the nondivisible are state-dependent preferences. In contrast to income shocks, this would not necessarily result in a precautionary savings motive.

$$W = E[u(c_1) + u(c_2) + u(c_3)].$$

The advantage of this convention is that no particular period is favoured (since no consumption takes place in period 0).

2.1 First Best

I assume throughout that b, p are such that it is optimal for a time-consistent agent to save for the nondivisible. For $\lambda = 0$, consumption smoothing implies that the agent optimally distributes the required savings burden of $p - 1$ evenly over periods 1 and 2, and uses his period 3 income plus accumulated savings to purchase the good. Denote the implied savings profile as $s_1 = \frac{p-1}{2} \equiv \bar{s}$, $s_2 = p - 1 = 2\bar{s}$. For $\lambda > 0$, there is a precautionary savings motive, even if the agent does not intend to save for the nondivisible good. Denote such precautionary savings s_t^{No} . It can be shown that the optimal savings path is slightly increasing, i.e., $s_1 < \bar{s}$.² Since the present analysis focuses on regular-installment products, I assume that desirability of the nondivisible still holds for fixed equal instalments $\bar{s} \equiv \frac{p-1}{2}$:

$$\begin{aligned} & (1 - \lambda)^2[2u(1 - \bar{s}) + (1 - \lambda)u(b) + \lambda u(p - 1)] \\ & + (1 - \lambda)\lambda[u(1 - \bar{s}) + u(\bar{s} - s_2^{No}) + E(u(y_3 + s_2^{No}))] \\ & + \lambda[u(0) + E(u(y_2 - s_2^{No}) + u(y_3 + s_2^{No}))] \\ & \geq E[u(y_1 - s_1^{No}) + u(y_2 + s_1^{No} - s_2^{No}) + u(y_3 + s_2^{No})] \end{aligned} \quad (2.1)$$

where s_t^{No} is chosen to optimally spread available current assets over the remaining future periods, conditional on not buying the nondivisible (i.e. $s_t^{No} = f(y_t, s_{t-1}, \lambda)$). Inequality 2.1 reflects the ‘survival probabilities’ of the savings plan: With probability $(1 - \lambda)^2$, no shock occurs in period 1 and 2, thus $c_1 = c_2 = 1 - \bar{s}$, and the period 2 self successfully transfers $2\bar{s} = p - 1$ to period 3. If no shock hits in period 3, the agent buys the nondivisible and consumes $c_3 = b$. If a shock does hit, he is unable to buy the good, and consumes $c_3 = p - 1$ instead. The second and third line of inequality 2.1 account for the case that saving fails because of a shock realization in period 2 (probability $(1 - \lambda)\lambda$) or period 1 (probability λ), respectively.

²The probability of remaining shock-free (and thus obtaining the nondivisible) increases over time, from $(1 - \lambda)^3$ ex-ante to $(1 - \lambda)$ once period 2 has been reached without a shock. This makes it optimal to slightly skew the savings burden $p - 1$ towards period 2. To see this formally, note that expected utility decreases in s_1 when evaluated at $s_1 = \bar{s}$: $dU/ds_1 = (1 - \lambda)^2[-u'(1 - s_1) + u'(2 + s_1 - p)] + (1 - \lambda)\lambda[-u'(1 - s_1) + u'(s_1 - s_2^{No})] < 0$ for $s_1 = \frac{p-1}{2} > 0.5$. By the envelope condition, $dU/ds_1 = \frac{\partial U}{\partial s_1} + \frac{\partial U}{\partial s_2} \cdot \frac{\partial s_2}{\partial s_1} = \frac{\partial U}{\partial s_1}$.

2.2 Autarky Equilibrium

The results in this section generalise the no-banking equilibrium from Basu (2014) to allow for partial sophistication and stochastic income.

The analysis proceeds by assuming that no shock has hit up to period t . If a shock does hit (i.e., if $y_t = 0$), the agent immediately gives up any plans to save for the nondivisible, and instead spreads available savings s_{t-1} optimally over the remaining periods. Denote such savings as s_t^{No} . For $\beta = \tilde{\beta} = 1$, the agent will always buy the nondivisible given the above condition (and absent shocks). The savings path will be perfectly smooth ($s_1 = \bar{s}$, $s_2 = 2\bar{s}$) if $\lambda = 0$, and slightly increasing ($s_1 < \bar{s}$, $s_2 = 2\bar{s}$) if $\lambda > 0$. If $\beta \leq \tilde{\beta} \leq 1$ (with at least one inequality strict), the three period selves engage in strategic interaction. Savings behaviour can be analysed by backward induction, taking into account the agent's belief about his future preferences.

Period 3

The agent will buy the nondivisible whenever he can afford it, i.e., whenever there is no shock, and $s_2 \geq p - 1$. Additional savings $s_2 > p - 1$ are simply consumed, as are savings that are not sufficient to buy the good. Since there are no future choices, the sophistication level does not influence behaviour at this stage. The consumption profile is

$$c_3 = \begin{cases} y_3 + s_2 - p + b & \text{if } y_3 + s_2 \geq p \\ y_3 + s_2 & \text{if } y_3 + s_2 < p \end{cases}$$

Period 2

The period 2 self knows the good will be bought if and only if he sends $s_2 \geq p - 1$, and absent shocks. He decides whether to send $s_2 = p - 1$, in which case the good is bought, or less. Due to consumption smoothing motives, it is never optimal to send $s_2 > p - 1 > 1$, which exceeds the magnitude of a shock. If the agent prefers not to save for the good, he will want to smooth s_1 over periods 2 and 3: $s_2^{No}(s_1) = \operatorname{argmax}(u(y_2 + s_1 - s_2) + \beta E[u(y_3 + s_2)])$ subject to $0 \leq s_2 < p - 1$ (note $E[u(y_3 + s_2)] = \lambda u(s_2) + (1 - \lambda)u(1 + s_2)$). This equation also describes his savings behaviour in case of a shock, where $y_2 = 0$, imposing an additional restriction

2.2. AUTARKY EQUILIBRIUM

$0 \leq s_2 < s_1$. He is willing to save $s_2 = p - 1$ if s_1 is such that

$$u(1 + s_1 - (p - 1)) + \beta[(1 - \lambda)u(b) + \lambda u(p - 1)] \geq u(1 + s_1 - s_2^{No}) + \beta E[u(y_3 + s_2^{No})]. \quad (2.2)$$

Lemma 1. (a) *The period 2 agent is willing to save for the nondivisible and transfer $s_2 = p - 1$ if s_1 is bigger than some threshold value, $s_1 \geq s_{min}$. (b) s_{min} is strictly decreasing in the time-inconsistency parameter β . (c) The effect of the shock frequency λ on s_{min} is ambiguous.*

(All proofs are in Appendix B.)

As in period 3, the level of sophistication does not affect the analysis: The period 2 self knows his true current β , but may mistakenly think that his period 3 self will apply $\tilde{\beta} \geq \beta$ to future decisions. As there are no future decisions once period 3 has been reached, this is of no consequence. Also note that the period 2 self conditions his behaviour on the s_1 received from period 1, regardless of the beliefs held by the period 1 self.

Period 1

Analogue to the minimum s_1 threshold for period 2, it is useful to identify the maximum s_1 that period 1 is willing to save, conditional on purchase of the nondivisible. If this maximum is bigger than the minimum required, the agent is theoretically able to purchase the good (whether saving is successful in equilibrium depends on the coordination between the selves, which is discussed in the equilibrium subsection below). In period 1, sophistication first comes into effect: Period 1 anticipates period 2's decisions, but is overconfident that his future self will be more patient than he is, i.e., he believes his future self uses $\tilde{\beta} \geq \beta$. This belief affects not only his perception of s_{min} , but also directly enters his own optimality considerations via $\tilde{s}_2^{No} \equiv s_2^{No}(\tilde{\beta})$, his perception of s_2^{No} . It is easy to show that \tilde{s}_2^{No} increases in $\tilde{\beta}$, and that $\tilde{\beta} \geq \beta$ implies $\tilde{s}_2^{No} \geq s_2^{No}$. The special case of full sophistication is obtained by setting $\tilde{s}_2^{No} = s_2^{No}$. Conditional on the nondivisible *not* being purchased (i.e., period 2 is believed to save $\tilde{s}_2^{No} < p - 1$), period 1 saves only for precautionary purposes: $s_1^{No} = \argmax(u(y_1 - s_1) + \beta E[u(y_2 + s_1 - \tilde{s}_2^{No}) + u(y_3 + \tilde{s}_2^{No})])$ for $s_1^{No} \geq 0$ and $y_t = \{0, 1\}$. The occurrence of a shock implies $y_1 = 0$ and thus $s_1^{No} = 0$.

Assume no shock has hit up to period 1. The expected utility from saving for the nondivisible is then captured in the left-hand side of inequality 2.3, conditional on period 2 saving $s_2 = p - 1$. If no shock hits in any future period (probability $(1 - \lambda)^2$), the nondivisible is consumed in period 3. However, the agent also

2.2. AUTARKY EQUILIBRIUM

takes into account that saving may fail because of shocks in period 3 (probability $(1 - \lambda)\lambda$) or period 2 (probability λ). The maximum that period 1 is willing to pay for the expected purchase of the nondivisible, denoted s_{max} , is the maximum value of s_1 which satisfies

$$\begin{aligned} & u(1 - s_1) + \beta(1 - \lambda)^2(u(1 + s_1 - (p - 1)) + u(b)) \\ & + \beta(1 - \lambda)\lambda(u(1 + s_1 - (p - 1)) + u(p - 1)) \\ & + \beta\lambda(u(s_1 - \tilde{s}_2^{No}) + E[u(y_3 + \tilde{s}_2^{No})]) \\ & \geq u(1 - s_1^{No}) + \beta E[u(y_2 + s_1^{No} - \tilde{s}_2^{No}) + u(y_3 + \tilde{s}_2^{No})]. \end{aligned} \quad (2.3)$$

If there is no such value, let $s_{max} = 0$.

Lemma 2. (a) The maximum that period 1 would be willing to save, denoted s_{max} , is strictly increasing in the time-inconsistency parameter β . (b) s_{max} weakly decreases in the amount of naiveté, $\tilde{\beta} - \beta$.

By assuming desirability of the nondivisible for a time-consistent agent (inequality 2.1), it follows that $s_{max}(\beta = \tilde{\beta} = 1) \geq \frac{p-1}{2}$. We further know that $s_{max}(0) = 0$. In addition to the maximum which period 1 is willing to save in order to purchase the good, consider the *optimal* way in which period 1 would like to allocate the savings burden of $p - 1$ across periods 1 and 2.

Lemma 3. (a) The optimal allocation of savings from period 1's perspective, denoted $s_1 = s_{opt}$, is characterized by

$$u'(1 - s_{opt}) = \beta[(1 - \lambda)u'(2 + s_{opt} - p) + \lambda u'(s_{opt} - \tilde{s}_2^{No})(1 + \frac{\delta \tilde{s}_2^{No}}{\delta s_1} \cdot \frac{1 - \tilde{\beta}}{\tilde{\beta}})].$$

(b) s_{opt} is strictly increasing in β , and always smaller than s_{max} .

The term involving $\delta \tilde{s}_2^{No} / \delta s_1$ is a result of the time-inconsistency (for a time-consistent agent, the envelope condition applies): \tilde{s}_2^{No} is chosen optimally given period 2's preferences (more specifically, period 1's belief thereof), which makes it suboptimal from period 1's perspective for $\tilde{\beta} < 1$. As a result, s_1 has a first-order positive effect on \tilde{s}_2^{No} .

Unfortunately, the effect of sophistication on s_{opt} is ambiguous. Holding β constant and increasing $\tilde{\beta}$ (sophistication falls), period 1 is more confident about period 2 following his interests in the future - in particular with respect to precautionary savings for period 3. As $\tilde{\beta}$ increases, it becomes more attractive to send savings to period 2, as the period 2 self is believed to spread them more equally across periods 2 and 3. On the other hand, period 1 no longer has to overcompensate for period 2's bias, sending excessive savings just to ensure some of them are passed on to period 3. It depends on the specific values of λ , $u''(c)$, $\tilde{\beta}$ and β which effect is stronger.

Autarky Equilibrium with Full Sophistication

Given a decreasing $s_{min}(\beta)$ and an increasing $s_{max}(\beta)$ -function, there is a threshold level $\hat{\beta}$ such that $s_{min}(\beta) \leq s_{max}(\beta)$ for any $\beta \geq \hat{\beta}$. The fact that $\hat{\beta}$ is in the relevant interval $(0, 1]$ follows from $s_{min}(0) > s_{max}(0)$ and $s_{min}(1) \leq s_{max}(1)$: The former follows from $s_{min}(0) > 1$, $s_{max}(0) = 0$. The latter is a consequence of desirability (inequality 2.1), by which a time-consistent agent always purchases the good. Since the different period selves are perfectly able to anticipate each other's behaviour, the nondivisible will be purchased (absent shocks) for all $\beta \in [\hat{\beta}, 1]$. Absent shocks, equilibrium savings are

$$s_1 = \begin{cases} \max(s_{min}, s_{opt}) & \text{if } \beta \in [\hat{\beta}, 1] \\ s_1^{No} & \text{if } \beta \in [0, \hat{\beta}) \end{cases}, \quad s_2 = \begin{cases} p - 1 & \text{if } \beta \in [\hat{\beta}, 1] \\ s_2^{No} & \text{if } \beta \in [0, \hat{\beta}) \end{cases}$$

If a shock occurs in any period, the individual gives up any plans to save for the nondivisible, and merely smoothes available assets $y_t + s_{t-1}$ over future periods, saving $s_t^{No} \geq 0$ for all t after the shock.

Importantly, it is ambiguous whether autarky savings will be above or below $\bar{s} \equiv \frac{p-1}{2}$. This will complicate the later analysis of the regular saver product (e.g., compare Figures 2.2 and 2.3). Considering that a time-consistent agent saves \bar{s} (for $\lambda = 0$) or slightly below \bar{s} (for $\lambda > 0$), this question corresponds to O'Donoghue and Rabin's (1999) *pre-emptive overcontrol*: A sophisticated hyperbolic discounter may both save more or less than a time-consistent agent, depending on the numerical values used for $(b - p)$ and $u''(c)$. In the following, scenarios with $s_{min}(\hat{\beta}) = s_{max}(\hat{\beta}) < \bar{s}$ will be referred to as "low autarky savings", and scenarios with $s_{min}(\hat{\beta}) = s_{max}(\hat{\beta}) > \bar{s}$ will be referred to as "high autarky savings".

Autarky Equilibrium with Partial Sophistication

Allowing for partial sophistication, the period 1 agent overestimates the patience of his future self, $\tilde{\beta} > \beta$. As discussed previously, this affects the s_{max} - and s_{opt} -function used by period 1 via precautionary savings. However, the main effect of partial sophistication is to cause period 1 to underestimate the amount of savings s_1 required to convince period 2 to save $s_2 = p - 1$ (and thus facilitate the purchase of the nondivisible). For ease of graphical illustration, assume $\tilde{\beta} = \beta + \gamma$. Denote the resulting perceived s_{min} -function as $\tilde{s}_{min}(\tilde{\beta}) = \tilde{s}_{min}(\beta + \gamma) = \tilde{s}_{min}(\beta) < s_{min}(\beta)$: For

2.2. AUTARKY EQUILIBRIUM




Figure 2.1: Autarky Equilibrium with Partial Sophistication

a constant sophistication level γ , perceived minimum savings \tilde{s}_{min} can be expressed as a function of β .³ This allows me to define thresholds in terms of β only:

Define $\check{\beta}$ such that $\tilde{s}_{min}(\beta) \leq s_{max}(\beta)$ for any $\beta \geq \check{\beta}$. For $\beta \in [\check{\beta}, 1]$, the period 1 agent will believe that he is able to save for the nondivisible. Note $\tilde{s}_{min}(\beta) < s_{min}(\beta)$ implies that $\check{\beta} < \hat{\beta}$. Furthermore, define β_{max} such that $s_{min}(\beta) \leq s_{opt}(\beta)$ for any $\beta \geq \beta_{max}$. For $\beta \in [\beta_{max}, 1]$, the optimal savings choice from period 1's perspective is more than required given period 2's true preferences. It follows that $\check{\beta} \leq \hat{\beta} \leq \beta_{max}$. Absent shocks in period 1 and 2, the autarky savings outcome is

$$s_1 = \begin{cases} \max\{\tilde{s}_{min}, s_{opt}\} & \text{if } \beta \in [\check{\beta}, 1] \\ s_1^{No}(\beta) & \text{if } \beta \in [0, \check{\beta}) \end{cases}, \quad s_2 = \begin{cases} p-1 & \text{if } \beta \in [\beta_{max}, 1] \\ s_2^{No}(\beta) & \text{if } \beta \in [0, \beta_{max}) \end{cases}$$

The savings path is illustrated by the red dashed line in Figure 2.1. The most remarkable feature of this savings function is that period 2 “eats” period 1's savings for $\beta \in [\check{\beta}, \beta_{max})$.⁴ For $\beta \in [\check{\beta}, \hat{\beta})$, the agent believes he can save for the nondivisible, but is not genuinely able to do so given his true preferences. This is the region

³While a constant sophistication level γ is convenient for the purposes of graphical illustration, the model's results do not depend on assumptions about the functional relationship between $\tilde{\beta}$ and β , other than $\tilde{\beta} \geq \beta$.

⁴This is comparable to the theoretical result in Duflo et al. (2011) for farmers' decision to save for fertilizer.

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

where $\tilde{s}_{min}(\beta) \leq s_{max}(\beta) < s_{min}(\beta)$. Period 1 sends $s_1 = \tilde{s}_{min}(\beta)$, anticipating that this will be enough to incentivise period 2 to save $p - 1$. Period 2 responds by consuming the savings, transferring only $s_2^{No} < p - 1$ to period 3.

Even more paradoxically, for $\beta \in [\hat{\beta}, \beta_{max})$, the agent fails to obtain the non-divisible because of a coordination failure between his different selves: $\tilde{s}_{min}(\beta) < s_{min}(\beta) \leq s_{max}(\beta)$, so the nondivisible could be bought if period 1 saved $s_1 \geq s_{min}(\beta)$. Instead, incorrect beliefs about this future preferences lead him to save $s_1 = \max\{\tilde{s}_{min}, s_{opt}\} < s_{min}(\beta)$, and again period 2 consumes period 1's savings. It is only for $\beta \in [\beta_{max}, 1]$ that the savings sent by period 1 are sufficient for purchasing the nondivisible: As β rises, period 1 becomes sufficiently patient to save more than \tilde{s}_{min} voluntarily, eventually reaching the point where $s_{opt}(\beta)$ becomes larger than the required true $s_{min}(\beta)$. Conditional on the absence of shocks, the nondivisible is purchased for the region $\beta \in [\beta_{max}, 1]$.

2.3 Equilibrium with a Regular Saver Commitment Product

The following section investigates the effect of offering agents a commitment to fixed regular contributions - as commonly found in loan contracts, pension savings, and other forms of regular saving. As pointed out by Fischer and Ghatak (2010) for the case of microloans, small frequent instalments may mediate time-inconsistency problems of hyperbolic discounters. In a savings setting, commitment to fixed instalments may help agents to reach savings goals, and smooth savings contributions.⁵

The Regular Saver product is defined as follows: Consider an agent who can commit in period 0 to deposit a fixed amount $\bar{s} = \frac{p-1}{2}$ in a bank account in both period 1 and 2. He also chooses a default penalty D , subject only to a limited liability constraint which prevents negative consumption. Once the agent fails to deposit \bar{s} in a period, he is charged the default penalty D , but immediately receives back any accumulated savings. In addition, he is free to save at home independently of his bank contributions. His total cumulative savings (in the bank plus at home) which are transferred from period t to $t + 1$ can then be captured as s_t . The penalty D is imposed in period 1 if $s_1 < \bar{s}$, and in period 2 if $s_1 \geq \bar{s}$, $s_2 < 2\bar{s}$. The contract is successfully completed with $s_1 \geq \bar{s}$, $s_2 \geq p - 1$. The assumption that the contract is

⁵Section 2.1 argues that the first-best savings schedule is $s_1 = \bar{s}$ for $\lambda = 0$, and slightly increasing for $\lambda > 0$, i.e., $s_1 < \bar{s}$. For small λ , this effect is likely to be small. Commitment products with increasing savings schedules are possible, but may pose serious challenges to institutional implementation: The first-best schedule will depend on individual values of λ , $u''(c)$, p and b . The present analysis focuses on fixed-instalment products due to their empirical popularity and ease of administration.

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

signed in period 0 simplifies things greatly, as the agent is not subject to temptation in this period.⁶ As before, the savings outcome can be derived using backwards induction, with a contract-signing period 0 discussed at the end.

Period 3

Period 3 behaviour is identical to that in autarky. The agent will buy the nondivisible whenever he can afford it, i.e., whenever $s_2 \geq p - 1$ holds, and absent shocks.

Period 2

Suppose the contract is still active in period 2. In other words, period 1 has not been hit by a shock, and has transferred $s_1 \geq \bar{s}$. Suppose a shock hits in period 2: At an asset level of $s_1 < 1$ and contractual savings of $s_2 = 2\bar{s} = p - 1$, default is unavoidable. The resulting consumption level is $c_2 = s_1 - D - s_2^{N_0} \geq 0$, implying that a penalty of $D \leq s_1$ can be enforced. Absent shocks, period 2 is faced with the decision of whether to send $s_2 = 2\bar{s} = p - 1$ (it is never optimal to send $s_2 > p - 1$). He is willing to do so if he receives an s_1 that satisfies

$$u(2 + s_1 - p) + \beta[(1 - \lambda)u(b) + \lambda u(p - 1)] \geq u(1 + s_1 - D - s_2^{N_0}) + \beta E[u(y_3 + s_2^{N_0})] \quad (2.4)$$

Since the inequality differs from the autarky case only in the penalty D , the same proof can be used to show that the nondivisible is bought for any $s_1 \geq s_{min}^B$. The threshold $s_{min}^B(\beta)$ will be strictly lower than $s_{min}(\beta)$ in the autarky case: The right-hand side of the inequality decreases when D is introduced, while the left-hand side stays unchanged. The effect of the penalty disappears for $s_1 < \bar{s}$: Period 1 has already defaulted on the contract and paid the penalty, so the contract is no longer active in period 2. As a result, $s_{min}^B(\beta) = s_{min}(\beta)$ for $s_1 < \bar{s}$. The two sections of the $s_{min}^B(\beta)$ -function combine with a horizontal line at $s_{min}^B(\beta) = \bar{s} = \frac{p-1}{2}$. In this region, the s_{min} required by period 2 is lower than \bar{s} if he faces the penalty, and higher than \bar{s} if he does not. To keep the contract active and ensure that period 2 faces the penalty, the period 1 agent needs to save $s_1 \geq \bar{s}$. Therefore, the minimum s_1 needed to incentivise period 2 to save is \bar{s} . Finally, in the region where $s_{min}^B(\beta) > \bar{s}$, the period 2 agent is not willing to save for the nondivisible unless period 1 makes additional savings at home.

⁶This assumes that the bank can enforce the penalty, even in the case that the agent defaults before depositing any savings. See Section 2 and footnote 16 on how I dealt with this issue in the study.

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

Period 1

Consider the maximum s_1 that period 1 is willing to save, once subjected to a penalty for $s_1 < \bar{s}$. Denote this as s_{max}^B . Absent shocks in period 1, s_{max}^B is given by the maximum value of s_1 which satisfies

$$\begin{aligned} & u(1 - s_1) + \beta(1 - \lambda)^2(u(1 + s_1 - (p - 1)) + u(b)) \\ & + \beta(1 - \lambda)\lambda(u(1 + s_1 - (p - 1)) + u(p - 1)) \\ & + \beta\lambda(u(s_1 - D - \tilde{s}_2^{No}) + E[u(y_3 + \tilde{s}_2^{No})]) \\ & \geq u(1 - D - s_1^{No}) + \beta E[u(y_2 + s_1^{No} - \tilde{s}_2^{No}) + u(y_3 + \tilde{s}_2^{No})]. \end{aligned} \quad (2.5)$$

If there is no such value, let $s_{max}^B = 0$. The intuition for inequality 2.5 is analogue to that of inequality 2.3: The agent considers the expected utility from saving for the nondivisible, conditional on period 2 saving $s_2 = p - 1$. He takes into account that the savings plan may fail to survive until the purchase of the nondivisible, due to shocks in future periods. As described in Section 2.2, partial sophistication implies that the agent uses $\tilde{s}_2^{No} \equiv s_2^{No}(\tilde{\beta})$ to assess period 2's behaviour.

In contrast to the inequality for s_{min}^B , both sides of the s_{max}^B - inequality are affected by the penalty: Even for a devoted saver, the penalty is unavoidable if a shock hits in period 2, causing the left-hand side to decrease in D (discounted by $\beta\lambda$). On the right-hand side, the penalty is the consequence of a deliberate decision to default in period 1. Note that, due to limited liability, the penalty cannot be enforced if there is a shock in period 1: With no income or previous savings, $c_1 = s_1 = 0$.

Proposition 1. *For small shock frequencies λ , and in the region where savings are skewed towards period 1, $s_1 \geq \bar{s} \equiv \frac{p-1}{2}$, adopting a regular-instalment product increases the maximum the agent is willing to save, i.e., $s_{max}^B > s_{max}$. A sufficient constraint on the shock frequency is $\lambda < \frac{u'(1)}{u'(0.5)}$. In the region $s_1 < \bar{s}$, adopting the regular-instalment product unambiguously decreases s_{max} , i.e., $s_{max}^B < s_{max}$.*

Note that inequality 2.5 is specific to the region $s_1 \geq \bar{s}$: The penalty is not charged in period 1 if the agent saves for the nondivisible. Consider the case where necessary savings are $s_1 < \bar{s}$, i.e., period 1 could ensure the good is bought even if he does not contribute \bar{s} . In this case, he faces a penalty whether or not he saves for the good. The penalty D enters in period 1 on both sides of the inequality (later periods are unaffected by D , as the contract is no longer active). The resulting threshold $s_{max}^B(\beta)$ is strictly lower than the original threshold $s_{max}(\beta)$.

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

Figure 2.2 shows that the two sections of the $s_{max}^B(\beta)$ -curve combine with a vertical line. To see why, extend the lower section of $s_{max}^B(\beta)$ to the \bar{s} -line. For any β in this range, s_{max} is below \bar{s} if the agent is charged the penalty even if he saves for the nondivisible, and s_{max} is above \bar{s} if he is not charged. Since the penalty does not apply for $s_1 \geq \bar{s}$, the maximum that he is willing to pay is given by the upper part of the $s_{max}^B(\beta)$ -curve. Even a high $s_{max}^B \geq \bar{s}$ does not rule out that the period 1 agent may optimally save $s_1 < \bar{s}$ for the nondivisible, deliberately incurring the penalty as a “premium” for procrastinating the savings burden onto period 2. Consider $s_{opt}^B(\beta)$, the optimal way in which period 1 would like to split the savings burden $p - 1$ across periods 1 and 2, when subjected to the penalty. In autarky, $s_{opt} = \bar{s} > 0.5$ holds only for a time-consistent agent, and given $\lambda = 0$. In the presence of time-inconsistency and a positive shock frequency $\lambda > 0$, s_{opt} is strictly below \bar{s} . In consequence, the introduction of a penalty reduces optimal savings further, as the agent needs to pay both s_1 and the penalty D , as a premium for procrastinating savings. Algebraically, $s_{opt}^B < s_{opt}$ follows from Lemma 3, after allowing for the fact that period 1’s consumption is now $c_1 = 1 - s_{opt} - D$.

Alternatively, period 1 may prefer to jump to $s_1 = \bar{s}$, rather than paying the penalty. The vertical part of s_{max}^B illustrates that it is never optimal to choose savings in the region $\hat{s}(\beta) < s_1 < \bar{s}$, where $\hat{s}(\beta)$ denotes the savings level which makes period 1 indifferent between saving \hat{s} plus paying the penalty D , and saving $s_1 = \bar{s}$, thus avoiding the penalty.⁷ Intuitively, if the necessary savings s_1 are such that $s_1 + D > \bar{s}$, then period 1 is trivially better off to save \bar{s} . Furthermore, the threshold \hat{s} is strictly lower than $\bar{s} - D$ for $\beta > 0$: At equal instantaneous cost $s_1 + D = \bar{s}$, it is strictly preferable to save \bar{s} , for the sake of the additional consumption D in the next period. Finally, willingness to jump to $s_1 = \bar{s}$ requires that $s_{max}^B \geq \bar{s}$. As a result, $\hat{s}(\beta)$ is only defined for the range of β such that $s_{max}^B(\beta) \geq \bar{s}$ (see Figure 2.2).

Lemma 4. *The threshold $\hat{s}(\beta)$ weakly decreases in β . Equivalently, as β increases, a larger range $s_1 \in (\hat{s}(\beta), \bar{s})$ is strictly dominated by \bar{s} .*

Equilibrium and Contract Choice: Full Sophistication

With full sophistication, the nondivisible good is purchased whenever $s_{max}^B(\beta) \geq s_{min}^B(\beta)$, which occurs for any $\beta \in [\hat{\beta}_B, 1]$. Equilibrium savings (absent shocks) are

⁷Formally, \hat{s} is the lowest value of s_1 which satisfies

$$\begin{aligned} & u(1 - \hat{s} - D) + \beta(1 - \lambda)(u(2 + \hat{s} - p)) + \beta\lambda(u(\hat{s} - \tilde{s}_2^{No}(\hat{s})) + E[u(y_3 + \tilde{s}_2^{No}(\hat{s}))]) \\ & \leq u(1 - \bar{s}) + \beta(1 - \lambda)(u(2 + \bar{s} - p)) + \beta\lambda(u(\bar{s} - D - \tilde{s}_2^{No}(\bar{s})) + E[u(y_3 + \tilde{s}_2^{No}(\bar{s}))]). \end{aligned}$$

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT




Figure 2.2: Regular Saver Equilibrium (high autarky savings)




Figure 2.3: Regular Saver Equilibrium (low autarky savings)

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

analogue to those for autarky, except for a lower savings threshold $\hat{\beta}_B < \hat{\beta}$, and a dominated region $s_1 \in (\hat{s}(\beta), \bar{s})$:

$$s_1 = \begin{cases} \max(s_{min}^B, s_{opt}^B) & \text{if } \beta \in [\hat{\beta}_B, 1] \text{ and } \max(s_{min}^B, s_{opt}^B) \notin [\hat{s}, \bar{s}) \\ \bar{s} & \text{if } \beta \in [\hat{\beta}_B, 1] \text{ and } \max(s_{min}^B, s_{opt}^B) \in [\hat{s}, \bar{s}) \\ s_1^{No} & \text{if } \beta \in [0, \hat{\beta}_B) \end{cases}$$

$$s_2 = \begin{cases} p - 1 & \text{if } \beta \in [\hat{\beta}_B, 1] \\ s_2^{No} & \text{if } \beta \in [0, \hat{\beta}_B). \end{cases}$$

Figures 2.2 and 2.3 illustrate the Regular Saver equilibrium with full sophistication. Figure 2.2 shows the savings path starting from high autarky savings, while Figure 2.3 starts from low autarky savings. It is critical for the welfare implications of the regular saver product whether $\hat{\beta}_B \leq \hat{\beta}$. In other words, is the nondivisible achievable for a larger range of preferences when the regular saver product is used? The answer is yes, given a sufficiently large penalty. Since the period 0 agent chooses the penalty himself, $\hat{\beta}_B \leq \hat{\beta}$ is guaranteed to hold under full sophistication.⁸ As a result, for $\beta \in [\hat{\beta}_B, \hat{\beta})$, the nondivisible is achievable with the regular-instalment product, but not without it. The threshold $\hat{\beta}_B$ decreases in the size of the chosen penalty D (a corollary of Proposition 2). Furthermore, for the region $\beta \in [\hat{\beta}, 1)$, the Regular Saver product weakly smoothes savings contributions (and thus consumption) towards \bar{s} .

Period 0 Adoption Decision and Penalty Choice In principal, any agent with $\beta \in [0, 1)$ can benefit from commitment. Given a sufficiently large penalty, it makes the nondivisible achievable and smoothes savings: Absent shocks, the contract is trivially enforceable in period 1 if $D > \bar{s}$, and in period 2 if $D > 2\bar{s}$. Even with $\beta = 0$, it is cheaper for the agent to make the contracted-upon savings contribution than to pay the penalty. As a result, the threshold $\hat{\beta}_B$ can be moved to an arbitrarily low β . The downside of commitment is the risk of “rational default” due to shock frequency λ : The penalty not only acts to discipline the agent when income is available, it also needs to be paid when the agent no longer finds it welfare-maximising (or feasible) to save for the nondivisible. Limited liability implies that this risk is limited to shocks in period 2: If a shock hits in period 1, the agent has no assets or

⁸Formally, $\hat{\beta}_B \leq \hat{\beta}$ holds regardless of penalty size if starting from a high-savings autarky scenario with $s_{min}(\hat{\beta}) = s_{max}(\hat{\beta}) \geq \bar{s}$. For scenarios with low autarky savings $s_{min}(\hat{\beta}) = s_{max}(\hat{\beta}) < \bar{s}$, the penalty D needs to be large enough to make an agent with $\beta = \hat{\beta}$ willing to jump to $s_{max}^B(\hat{\beta}) \geq \bar{s}$ to ensure $\hat{\beta}_B \leq \hat{\beta}$.

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

income, thus the penalty cannot be enforced. In period 3, the contract is no longer active. In contrast, if a shock hits in period 2, the agent's savings of $s_1 \geq \bar{s}$ may be lost to the penalty D , leaving the agent worse off than if he had not adopted commitment.

The resulting decision is a two-step problem: The period 0 agent first decides which penalty D offers the optimal trade-off between commitment and flexibility. He then makes a binary choice between adopting the regular saver product with the optimal penalty, or not adopting the product. Unfortunately, the choice of the optimal penalty is non-monotonic in β , and sensitive to the autarky scenario, due to the consumption smoothing motive: Consider starting from a low $\beta < \hat{\beta}$ in Figure 2.2 (high autarky savings). Increasing D will first shift the upper part of the s_{min}^B - and s_{max}^B -curve to the left, until $s_{max}^B(\beta) = s_{min}^B(\beta)$ holds for the agent's β (in other words, until $\hat{\beta}_B = \beta$). The nondivisible is now achievable, but at a skewed savings schedule $s_1 = s_{min}^B(\hat{\beta}_B) > \bar{s}$. The agent may choose to increase D further, in order to decrease s_{min}^B and smooth savings towards \bar{s} . However, the benefit associated with smoother savings contributions is a discrete drop from the benefit associated with achieving the nondivisible, and the agent may not deem it worthwhile to increase D further in the face of shock frequency λ . To see why the optimal penalty is non-monotonic in β , consider starting from a high $\beta \geq \hat{\beta}$ in Figure 2.3 (a scenario with low autarky savings). In autarky, the nondivisible is achievable, and the agent saves $s_1 = \max(s_{min}, s_{opt}) \ll \bar{s}$. While the Regular Saver product is not needed to achieve the nondivisible, it can help to smooth consumption: As D increases, $\hat{s}(\beta)$ falls, and the dominated region $s_1 \in (\hat{s}(\beta), \bar{s})$ becomes larger, until it eventually includes $\max(s_{min}^B, s_{opt}^B)$.⁹ The period 0 agent would like to choose the penalty at the minimum level which will make him jump to $s_1 = \bar{s}$.¹⁰ Thus, he will choose D such that $\hat{s}(\beta) = \max(s_{min}^B, s_{opt}^B)$ holds exactly. As $\max(s_{min}^B, s_{opt}^B)$ first decreases in β along with s_{min}^B , and then increases in β along with s_{opt}^B , the penalty required to make $\hat{s}(\beta) = \max(s_{min}^B, s_{opt}^B)$ first increases and then decreases in β . Depending on $u''(c)$ and λ , the take-up decision for $\beta \in [\hat{\beta}, 1)$ (where the Regular Saver product is exclusively used for consumption smoothing) may be similarly non-monotonic.

⁹Note that the s_{min}^B -curve is unaffected by D in the region $s_1 < \bar{s}$, as the contract is no longer active in period 2. Meanwhile, s_{opt}^B decreases in D , as it factors in the default on the contract. Therefore, the only possibility to smooth consumption via penalty D is through its effect on the dominated region $(\hat{s}(\beta), \bar{s})$.

¹⁰This is a simplification: The first-best is to save slightly below \bar{s} , which is not feasible with a regular-saver product. However, for small λ , this difference is small, and the agent is better off with a smooth savings schedule \bar{s} , compared to leaving the savings decision entirely at the discretion of period 1.

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

For the sake of simplicity, I will abstract from the consumption smoothing motive, and focus on the range of $\beta \in [0, \hat{\beta})$. For this range of β , the nondivisible is not achievable in autarky, and obtaining it constitutes the primary benefit of the Regular Saver product. This focus is empirically meaningful: It restricts the analysis to the part of the population who are not able to save for lump-sum consumption expenditures by themselves, i.e., without the use of commitment. This is consistent with data from my sample population.¹¹ Define $D_{eff}(\beta)$ to be the minimum effective penalty which achieves $s_{max}^B(\beta) \geq s_{min}^B(\beta)$. Given full sophistication, a Regular Saver contract with a penalty D_{eff} will enable the agent to save for the nondivisible (absent shocks). By construction, $D_{eff} = 0$ for $\beta \geq \hat{\beta}$.

Proposition 2. *For a given shock frequency λ , the minimum effective penalty D_{eff} that will enforce saving weakly decreases in the time-inconsistency parameter β .*

Proposition 3. *The optimal Regular Saver contract for a fully sophisticated agent with $\beta < \hat{\beta}$ depends on the effect of the minimum effective penalty, $D = D_{eff}$: Where D_{eff} results in $s_{min}^B(\beta) \leq s_{max}^B(\beta) \leq \bar{s}$ (illustrated in Figure 2.3), equilibrium savings contributions are perfectly smooth, i.e., $s_1 = \bar{s} \equiv \frac{p-1}{2}$ and $s_2 = 2\bar{s} = p - 1$. Consequently, the optimal contract is to choose D_{eff} . Where D_{eff} results in $s_{max}^B(\beta) \geq s_{min}^B(\beta) > \bar{s}$ (illustrated in Figure 2.2), the optimal contract involves $D \geq D_{eff}$, with equilibrium savings weakly skewed towards period 1 ($s_1 \geq \bar{s}$).*

For plausible ranges of the parameters, the case where D_{eff} guarantees perfect consumption smoothing $s_1 = \bar{s}$ (and thus eliminates the need to choose a higher penalty) coincides with the “low autarky savings” scenario.¹² The specific parameter restrictions needed are the subject of current research.

Having determined the optimal penalty for the Regular Saver product, the period 0 agent then faces the binary decision of whether or not to take up the product.

¹¹35 percent of the study population reported zero savings of any form, the median level of liquid assets (bank and home savings) was 500 pesos (U.S.\$ 12), and the most common way to afford lump-sum expenditures was through high-interest borrowing (which includes a commitment to fixed instalments).

¹²As discussed in the autarky section, “low autarky savings” refers to a situation where $s_{min}(\hat{\beta}) = s_{max}(\hat{\beta}) < \bar{s}$. It does not refer to the specific savings made by the agent in autarky.

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

The following inequality is sufficient for take-up to be optimal:

$$\begin{aligned}
& (1 - \lambda)^3 [u(1 - s_1) + u(2 + s_1 - p) + u(b)] \\
& + (1 - \lambda)^2 \lambda [u(1 - s_1) + (u(2 + s_1 - p) + u(p - 1))] \\
& + (1 - \lambda) \lambda [u(1 - s_1) + u(s_1 - D_{eff} - s_2^{No}) + E(u(y_3 + s_2^{No}))] \\
& + \lambda [u(0) + E(u(y_2 - s_2^{No}) + u(y_3 + s_2^{No}))] \\
& \geq E[u(y_1 - s_1^{No}) + u(y_2 + s_1^{No} - s_2^{No}) + u(y_3 + s_2^{No})] \tag{2.6}
\end{aligned}$$

where $s_1 = \max\{\bar{s}, s_{min}^B(\hat{\beta}_B)\}$ and $y_t = \{0, 1\}$ depending on the realisation of shocks. The rows of inequality 2.6 describe the different cases of shock occurrence: The savings plan could be undisturbed by shocks until the end of the agent's life (first row), it could fail in period 3 (second row: period 3 lacks the income to buy the nondivisible), a shock in period 2 could lead to costly default (third row), or a shock in period 1 could prevent saving for the nondivisible altogether (fourth row).

This leads to the following results for period 0's adoption decision: Whether the agent adopts commitment will critically depend on shock frequency λ , nondivisible benefit b , price p , and required penalty $D_{eff}(\beta)$. However, *ceteris paribus*, those with the lowest values of β will require the highest penalties D_{eff} . Consequently, those with the lowest values of β are the *least* likely to adopt the product. To see this, realise that the benefit of an effective commitment contract (obtaining the non-divisible with a smooth schedule $s_1 = \bar{s}$) is independent of β : The period 0 agent bases his adoption decision on the welfare function $W = E[u(c_1) + u(c_2) + u(c_3)]$, which does not directly depend on β . Put simply, the time-inconsistency parameter β determines how difficult it is for the agent to save for the nondivisible, but not how much he benefits from achieving it.¹³ As a result, agents with low β (and therefore a high required penalty D_{eff}) will find that commitment harms them in expectation, and will not adopt it.

The result on welfare is straightforward: Given full sophistication, everyone who adopts the commitment product is made better off in expectation. Agents perfectly anticipate their own behaviour, and assess the required degree of commitment (D_{eff}) correctly. The only reason for contract defaults are shocks: A fraction

¹³Strictly speaking, the benefit from commitment is only independent of β for $\lambda = 0$. With $\lambda > 0$, the period 0 agent has to rely on his future selves to make precautionary savings. The lower is β , the larger is the disagreement between the selves over how much should be saved for shocks. The commitment contract insures the agent against shocks at least in period 3 (no savings are available if a shock hits in period 1, and a shock in period 2 would leave the agent with $s_1 - D \geq 0$). Since precautionary savings decrease in β , the insurance effect of commitment is slightly more valuable for lower β . However, this effect is unlikely to quantitatively dominate the offsetting effect of a higher required penalty D_{eff} .

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT




Figure 2.4: The Regular Saver Equilibrium with Partial Sophistication

λ of adopters defaults each period. In summary, commitment through a regular-instalment product will be weakly welfare-increasing for sophisticated hyperbolic discounters.

Equilibrium and Contract Choice: Partial Sophistication

The derivations for penalty choice and adoption decision for partially sophisticated agents are analogous to those for full sophistication – except that the period 0 agent systematically applies an incorrect belief $\tilde{\beta} > \beta$. This results in a biased perception not only of $\tilde{s}_{min} < s_{min}$ (as for period 1), but also of $\tilde{s}_{max}(\beta) \equiv s_{max}(\tilde{\beta}) > s_{max}(\beta)$, as period 0 is overconfident about the patience he will have in period 1. Since the same belief $\tilde{\beta}$ is used to assess \tilde{s}_{min} and \tilde{s}_{max} , the partial sophistication bias in period 0 graphically corresponds to a shift in the entire schedule by a constant $\tilde{\beta} - \beta \equiv \gamma$ (see Figure 2.4).

The analysis will focus on those agents with $\beta < \tilde{\beta} < \hat{\beta}$: The part of the population who is not only unable to save without commitment, but who is also aware of this fact (for instance, because they have not observed themselves save in the past). Given a large benefit b of the nondivisible good, the primary motivation of such agents for adopting the regular-instalment commitment savings product will be to achieve the nondivisible. The minimum penalty which is perceived to be effective

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

in making the nondivisible achievable (in other words, the penalty which results in $\tilde{s}_{min}^B = \tilde{s}_{max}^B$) is then $D_{eff}(\tilde{\beta})$, denoted \tilde{D}_{eff} . By construction, $\tilde{D}_{eff} = 0$ for $\tilde{\beta} \geq \hat{\beta}$.

The optimal penalty choice for partially sophisticated agents is a corollary of Proposition 3: An agent who believes to have $\tilde{\beta} < \hat{\beta}$ will unambiguously choose the perceived minimum effective penalty, $D = \tilde{D}_{eff}$, whenever he anticipates that this will result in perfect consumption smoothing, i.e., when $\tilde{s}_{min}^B(\beta) \leq \tilde{s}_{max}^B(\beta) \leq \bar{s}$ at \tilde{D}_{eff} . This is likely to happen under low autarky savings scenarios. Where \tilde{D}_{eff} results in $\tilde{s}_{max}^B(\beta) \geq \tilde{s}_{min}^B(\beta) > \bar{s}$ (typically in high autarky savings scenarios), the agent chooses $D \geq \tilde{D}_{eff}$, and anticipates equilibrium savings $s_1 = \tilde{s}_{min}^B \geq \bar{s}$. By Proposition 2, D_{eff} decreases in β . Therefore, $\tilde{\beta} > \beta$ implies $\tilde{D}_{eff} < D_{eff}$.

The take-up decision is determined in the same way as for fully sophisticated agents, and captured in inequality 2.6. The inequality does not (directly) depend on β , but it depends on the (perceived) effective penalty D_{eff} . Partially sophisticated agents differ from fully sophisticated agents precisely in the fact that they perceive a lower $\tilde{D}_{eff} < D_{eff}$ to be sufficient. As a response, for the range of $\beta < \tilde{\beta} < \hat{\beta}$, the regular-instalment product is more attractive to partially sophisticated agents than to fully sophisticated agents: Conditioning on β , and holding λ, p, b and $u''(c)$ constant, those with higher sophistication gaps γ have a lower \tilde{D}_{eff} , and are thus more likely to adopt the product.

The savings outcome is a function of the chosen penalty. In addition, it may critically depend on the degree of *learning* which the agent undergoes during his life: In period 0, he believes he will use the parameter $\tilde{\beta}$ in all future periods. In period 1, he realises his true current β . In a static model where the agent does not update his beliefs after he observes his behaviour, period 1 will continue to believe that he will use $\tilde{\beta}$ in the future (much like a dieter who observes himself eating chocolate, but repeatedly plans to be more disciplined tomorrow). The other extreme is full updating: As period 1 learns his true current β , he updates his belief to $\tilde{\beta} = \beta$ for all future periods. Ali (2011) characterizes conditions under which agents' beliefs converge to full sophistication, presuming Bayesian updating. This paper discusses the two extreme assumptions: The case without updating, and the case of full updating.

Suppose the minimum effective penalty is perceived sufficient to guarantee perfect consumption smoothing, i.e., $\tilde{s}_{min}^B(\beta) \leq \tilde{s}_{max}^B(\beta) \leq \bar{s}$ at \tilde{D}_{eff} . The period 0 agent chooses \tilde{D}_{eff} and expects $s_1 = \bar{s}$. In period 1, the agent learns his true current β , and thus $s_{max}^B < \tilde{s}_{max}^B$. Without updating, period 1 still believes in \tilde{s}_{min}^B . With full updating to $\tilde{\beta} = \beta$, the agent also learns the true $s_{min}^B > \tilde{s}_{min}^B$. In this scenario, updating is of no consequence for the savings outcome: \tilde{D}_{eff} is constructed to make $\tilde{s}_{min}^B = \tilde{s}_{max}^B$

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

hold exactly. The realisation that $s_{max}^B < \tilde{s}_{max}^B$ is sufficient to inform the agent that saving is not feasible: Whether he believes in \tilde{s}_{min}^B or s_{min}^B only determines the size of the gap $s_{max}^B < \tilde{s}_{min}^B < s_{min}^B$ which keeps him from saving (see Figure 2.4). As a response, he abandons his savings plan in period 1, pays the penalty, and saves s_1^{No} .

Starting from a situation where \tilde{D}_{eff} provides incomplete consumption smoothing, i.e., $\tilde{s}_{max}^B(\beta) \geq \tilde{s}_{min}^B(\beta) > \bar{s}$, will generally produce the same result: Agents choose their penalty at \tilde{D}_{eff} or slightly above. For most parameter specifications, period 1's realisation that $s_{max}^B < \tilde{s}_{max}^B$ will result in $s_{max}^B < \tilde{s}_{min}^B < s_{min}^B$, which leads to immediate contract default irrespective of learning behaviour. For illustration, consider an agent with $\beta = \check{\beta}_B$ in Figure 2.4, who believes that his future selves will use $\tilde{\beta} = \hat{\beta}_B$. The agent's β and $\tilde{\beta}$ are at the banking thresholds $\check{\beta}_B$ and $\hat{\beta}_B$ by construction of the penalty \tilde{D}_{eff} and the assumed function $\tilde{\beta} = \beta + \gamma$.

However, the result may differ in cases where the agent has a particularly strong motive for consumption smoothing: The agent may voluntarily increase the penalty beyond \tilde{D}_{eff} in order to reduce $\tilde{s}_{min}^B \geq \bar{s}$ and get closer to \bar{s} . The success of this endeavour depends on the size of the penalty, and on learning: The agent may increase D until $\tilde{s}_{min}^B = \bar{s}$ holds exactly (higher penalties cannot be optimal, since their only effect is to increase the cost of default in case of a shock). In period 1, the agent realises $s_{max}^B < \tilde{s}_{max}^B$. Without updating, if $\tilde{s}_{min} = \bar{s} < s_{max}$, the agent still believes he is able to save, and transfers $s_1 = \tilde{s}_{min} = \bar{s}$ to period 2 (this follows from $s_1 = \max(\tilde{s}_{min}^B, s_{opt}^B)$ and $s_{opt}^B < \bar{s}$). However, in reality, $s_{min}^B > \tilde{s}_{min}^B = \bar{s}$. The penalty that is sufficient to reduce \tilde{s}_{min}^B to \bar{s} is not sufficient to reduce s_{min}^B to \bar{s} . Comparable to the coordination failure in autarky, once period 2 arrives, the agent eats his savings, and fails to save for the nondivisible. The situation is welfare-reducing relative to autarky, as the effect of an uneven consumption path is exacerbated by the loss of the penalty D . Instead, consider the case with full updating: In period 1, he learns that $s_{min}^B > \tilde{s}_{min}^B = \bar{s}$. If the chosen penalty is large enough to guarantee $s_{max}^B \geq s_{min}^B$ for his true preferences, the agent is willing to save $s_1 = s_{min}^B > \bar{s}$. While the agent fails to achieve consumption smoothing, updating his beliefs enables him to avoid contract default, and obtain the nondivisible. When do such cases occur? The motive for consumption smoothing must be large, and the sophistication gap low. Therefore, successful saving under partial sophistication is most likely to occur for high autarky savings, small sophistication gaps γ , small shock frequency λ (so the agent is less averse to big penalties), and large nondivisible prices p (increasing the benefits to consumption smoothing).

2.3. EQUILIBRIUM WITH A REGULAR SAVER COMMITMENT PRODUCT

The resulting welfare implications are discouraging: For $\beta < \tilde{\beta} < \hat{\beta}$, and without updating of beliefs, all partially sophisticated adopters default. Agents are particularly likely to adopt the contract if they have a high $\tilde{\beta}$, as is the case for those with large sophistication gaps γ . Default always occurs in period 1 when choosing $D = \tilde{D}_{eff}$. Welfare is unambiguously reduced: It decreases from $W_A = E[u(y_1 - s_1^{No}) + u(y_2 + s_1^{No} - s_2^{No}) + u(y_3 + s_2^{No})]$ in autarky to $W_{RS} = E[u(y_1 - s_1^{No} - D) + u(y_2 + s_1^{No} - s_2^{No}) + u(y_3 + s_2^{No})]$ with the commitment product. When choosing $D > \tilde{D}_{eff}$, default in period 2 is possible under some parameter specifications. Finally, with full updating of beliefs, the agent may be able to fulfill the contract and obtain the nondivisible under parameter specifications which strongly encourage consumption smoothing.

Chapter 3

When Commitment Fails: A Randomized Experiment in the Philippines

3.1 Overview

The field experiment was designed to test the efficacy of regular-instalment savings products in a randomised setting. I partnered with 1st Valley Bank, a rural bank based in Mindanao (Philippines). The sample population of 913 individuals was obtained by conducting a door-to-door baseline survey in low-income areas in proximity to two selected bank branches. The baseline survey elicited time preferences, with a random half of individuals receiving real monetary rewards. Further elicited measures included perceived time-inconsistency, risk aversion, financial claims from others, cognitive ability, financial literacy, intrahousehold bargaining power, household demographics, and measures of saving, borrowing, and household expenditures. After the baseline survey, all individuals were provided with a marketing treatment, which included a personalised savings plan for an upcoming expenditure and a free non-commitment savings account with 100 pesos (U.S. \$2.50) opening balance.¹ Personal savings plans featured a self-chosen goal date, goal amount, and a weekly or bi-weekly instalment plan (see Figure C.5 for a sample). The idea was to encourage individuals to save for their lump-sum expenses (such as school fees, business capital, or house repairs), rather than following the common practice of borrowing at high informal moneylender rates. At the end of the marketing visit, a randomly chosen 50 percent (the 'Regular Saver' group) were

¹At the time of marketing (October 2012), the exchange rate was 42 Philippine pesos per U.S. dollar.

offered a new commitment product called ‘Achiever’s Savings Account’ (ASA). ASA committed clients to make fixed regular deposits and pay a penalty upon default, which effectively made all features of the personal savings plan binding. The default penalty was chosen by the client upon contract signing, and framed as a charity donation.² There was no compensation for the restrictions, no added help (such as deposit collectors or text message reminders), and a standard market interest rate.³

As a control treatment, an additional 25 percent of the sample (the ‘Withdrawal Restriction’ group) were offered the commitment savings account studied in Ashraf et al. (2006b), Giné et al. (2010), Brune et al. (2011), and Karlan and Zinman (2013): The ‘Gihandom’ savings account (Visayan for ‘dream’) allowed individuals to commit to either the goal date or the goal amount from their savings plan, by restricting withdrawals before the goal had been reached. This account did not include any obligation to make further deposits after the opening balance. The remaining 25 percent of the sample received no further intervention after the marketing treatment, and constitute the control group. For those in the control group (and those who rejected the commitment accounts), none of the savings plan features were binding. Since individuals’ expenditures were due at different times, the outcome variable of interest are individuals’ savings at the time of their goal date. The study concluded with a comprehensive endline survey, a ‘customer satisfaction survey’ for ASA clients, and a repeat marketing stage where ASA clients could opt to ‘pre-order’ the product for a second round.⁴

I find that demand for commitment is high, even in a general low-income population with little previous bank exposure: Take-up rates were 27 percent for ASA and 42 percent for Gihandom, in spite of the fact that all individuals were given a free standard savings account (with 100 pesos) immediately prior to the commitment offer.⁵ Offering ASA was more effective at increasing savings: By the time individuals reached their goal date (an average of 130 days later), bank savings in the Regular Saver group had increased by 585 pesos (U.S. \$14) relative to the control group, whereas bank savings in the Withdrawal Restriction group had increased by 148 pesos (U.S. \$3.50, as measured by the Intent-to-Treat effect) relative to the

²The concept is roughly comparable to the Stickk initiative (www.stickk.com), where people are asked to set their own stakes, but applied to the requirements of a developing country context.

³The bank’s standard interest rate as of September 2012 was 1.5 percent per annum, and decreased to 1 percent in November 2012. This interest rate was the same across all offered accounts. The inflation rate for 2012 was 3.1 percent.

⁴Pre-orders were not legally binding, but involved a cost through substantial paperwork.

⁵The difference in take-up rates may be partly driven by liquidity concerns: ASA required an opening balance equal to the first weekly deposit (minimum 150 or 250 pesos), whereas Gihandom could be opened with 100 pesos.

control.⁶ The control group saved an average of 27 pesos. The scale of effects suggests that a commitment product with fixed regular instalments is highly effective at increasing savings on *average*. However, this average hides a lot of heterogeneity in the case of both products: 55 percent of ASA clients defaulted on their savings contract, incurring penalties (charity donations) between 150 and 300 pesos - the equivalent of a day's wage (the stated treatment effect already accounts for these charges). The penalty for the withdrawal-restriction product Gihandom was less salient, but existent: 79 percent of Gihandom clients made no further deposits after the opening balance. For those who had chosen to make their goal amount binding (45 percent), this meant their initial savings were tied up indefinitely, or until the bank would exhaust their account with dormancy fees.⁷

Using conventional measures of actual time-inconsistency and a novel measure of *perceived* time-inconsistency (sophistication), the data suggest that present-biased preferences by themselves do not predict take-up of a commitment product, but they do predict default. In contrast, sophistication drives both take-up and default: As an agent's degree of sophistication rises, he becomes less likely to adopt commitment, and less likely to default, conditional on take-up. This is consistent with the interpretation that partial sophistication about time-inconsistency leads agents to adopt weak commitment contracts, and subsequently default. Highly sophisticated agents are more cautious about adopting commitment, but have higher chances of success when they do choose to commit. The notion of 'weak commitment' is supported by the observation that 80 percent of ASA clients chose the minimum admissible default penalty (150 or 250 pesos, depending on the savings goal). Finally, the data is strongly bi-modal, in the sense that almost all clients either (i) stop depositing immediately after the opening balance or (ii) complete their savings plan in full. The timing of defaults is inconsistent with idiosyncratic shock patterns, where individuals rationally default following large random shocks to their income or expenditures.

The chapter proceeds as follows. Section 3.2 describes the experimental design employed in more detail. Section 3.3 explains the survey instrument, with particular view to the measurement of time-inconsistency and sophisticated hyperbolic discounting. Section 3.4 outlines the empirical strategy. Section 3.5 presents empirical results. Section 3.6 discusses other explanations. Section 3.7 concludes.

⁶The Intent-to-Treat (ITT) effect measures the effect of being *offered* the product. An increase of 585 pesos (148 pesos) corresponds to 27 percent (7 percent) of median weekly household income in our sample.

⁷Dormancy fees are very common with Philippine banks, and commonly start after two years of inactivity.

3.2 Experimental Design

I designed and implemented the commitment savings product ASA in cooperation with 1st Valley Bank, based in Mindanao, Philippines. 1st Valley Bank is a small rural bank that offers microcredit, agricultural insurance, salary loans, and pension products to its clients. The bank agreed to offer both the regular-instalment product ASA and the withdrawal-restriction product Gihandom in two of their bank branches: Gingoog and Mambajao. Gingoog is a city of 112,000 people in northern Mindanao, and Mambajao is a municipality of 36,000 people on Camiguin Island. For these two bank branches, both ASA and Gihandom constituted new product additions at the time of the study.⁸

The sample was obtained through door-to-door visits in all low and middle income areas in proximity to the bank branches. In each household, the survey team identified the person in charge of managing the household budget (in 94 percent of the cases, this was the mother of the family), and invited them to take part in a household survey. The baseline survey was completed with all individuals who (i) had some form of identification,⁹ (ii) claimed to have a large upcoming expenditure (such as school fees, house repairs, or business expansions) and (iii) agreed to receive a visit from a financial advisor (to talk about how to manage household expenses). After the baseline survey, individuals were randomly assigned to three groups: 50 percent of individuals were assigned to a 'Regular Saver' (R) group, and 25 percent each were assigned to a 'Withdrawal Restriction' (W) and a control (C) group.

Approximately one week after having completed the household survey, individuals received a visit from a bank marketer. Of 913 surveyed individuals, 852 could be re-located.¹⁰ Marketers engaged individuals in a conversation about how to manage large lump-sum expenses, and talked about the benefits of saving. Focusing on one particular expenditure, individuals were encouraged to make a formal 'Personal Savings Plan', which contained a purpose, a goal amount, a goal date, and a fixed equal instalment plan with due dates (see Figure C.5). Median savings goals were 2400 pesos across all groups (roughly comparable to a median household's weekly income of 2125 pesos), with a median weekly instalment of 150 pe-

⁸Gihandom had been previously offered at other 1st Valley Bank branches, but not at the two branches in question.

⁹A valid form of identification was required by the bank to open a savings account. Accepted forms of identification included a birth certificate, tax certificate, voter's ID, barangay clearance, and several other substitutes.

¹⁰A test for equality of means in the probability of being reached by the marketer across treatment groups yields an *F*-statistic with a *p*-value of 0.16. Individuals in group R were as likely to be reached as individuals in group C, but slightly less likely than individuals in group W.

3.2. EXPERIMENTAL DESIGN

sos. Common savings goals were school tuition fees, house repairs, and Christmas gifts (see Table C.1 for an overview of savings plan characteristics). The duration of savings plans was limited to 3–6 months, so that the outcome could be observed by the study. The median duration was 137 days. In addition, everyone was offered a standard non-commitment savings account (henceforth called ‘ordinary savings account’) as a ‘welcome gift’ from the bank, along with an encouragement to use this account to save for the expenditure. This ordinary savings account contained a free 100 pesos opening balance, which also constitutes the minimum maintaining balance.¹¹

At the end of the visit, individuals in group R were asked whether they wanted to commit to the fixed-instalment structure outlined in their Personal Savings Plan by taking up the ASA product, and the product features were explained. In contrast, individuals in group W were offered the option to restrict withdrawals of their savings until they reached either the goal amount or the goal date specified in their Personal Savings Plan, implemented through the use of the Gihandom product. It is to be expected that the marketing treatment itself influenced individuals’ savings behaviour, as evidenced in the literature on mental accounting.¹² However, up to the point of offering the commitment products, the marketing script was identical across all groups, which prevents a bias of the estimated treatment effects. The effect of marketing as such cannot be identified, as there was no marketing-free group. However, the fact that the control group saved an average of 27 pesos until their respective goal dates indicates that, given a non-negativity constraint on bank savings, the effect of marketing was small.

Out of 852 individuals located for the marketing visit in September and October 2012, 788 accepted the free ordinary savings account, and 748 agreed to make a savings plan. In group R, 114 clients (out of 423 offered) accepted the ASA product.¹³ In group W, 92 (out of 219 offered) accepted the Gihandom product. Table C.2 summarizes the take-up results.

The regular-instalment product ASA committed clients to a specific instalment plan with weekly or bi-weekly due dates. An account was considered in default from the day the client fell *three* instalments behind. In case of a default, the account was closed, an ‘Early Termination Fee’ was charged, and any remaining sav-

¹¹Individuals were able to close this account and retrieve funds by visiting the bank, but incurred a 50 peso closing fee. During the period of observation (September 2012 until 15 April 2013), no client closed their account.

¹²Most prominently, see Thaler (1985) and Shefrin and Thaler (1988).

¹³One member of the control group was mistakenly offered the ASA product and accepted, which means a total of 115 ASA accounts were opened. This constitutes a mild contamination of the control group, and means that the estimated ASA treatment effect is a lower bound on the true effect.

3.2. EXPERIMENTAL DESIGN

ings were returned to the client. A termination fee that is directly linked to the instalment structure distinguishes the ASA product from withdrawal-restriction or standard accounts, and represents its key commitment feature. No fee was charged after successful completion. The amount of the fee was chosen by the client upon signing the ASA contract: Each client signed a 'Voluntary Donation Form', which specified a termination fee that would be donated to charity in case of a default. Clients were given a choice of three national (but not locally-based) Philippine charities.¹⁴ While the instalment structure may appear rigid at first sight, a variety of flexibility features were included to allow clients to adapt to changing circumstances: First, clients could fall up to two instalments behind at any given time, making it theoretically possible to miss every other instalment, and pay a double instalment in alternate weeks. To encourage timely depositing, a small 10 peso (\$0.25) admin fee had to be paid upon making up a missed past instalment, but this fee did not accumulate over time. Deposits towards future weeks could be made at any time, as long as they were in increments of the weekly instalment. This was a practical requirement, as the client's progress was monitored by making ticks on a collection card for each successful week (see Figure C.5). The possibility of making future deposits early effectively provided a form of insurance against uneven income streams. Withdrawals during the savings period were only possible by allowing default to occur.

Enforceability of the termination fee was facilitated through the account opening balance: To complete the opening of an ASA, clients had to deposit an opening balance equal to their first weekly instalment, but at least 150 pesos for savings goals below 2500 pesos, and at least 250 pesos for savings goals of 2500 pesos and above. The same threshold applied for the termination fee: Clients could choose a termination fee as high as they liked, but no lower than a minimum of 150 or 250 pesos, respectively. Consequently, the minimum termination fee could always be enforced. Higher termination fees could be enforced only if the client continued to save, or if their opening balance exceeded the minimum. By nature of the contract, all ASAs were either successfully completed or in default by the goal date,¹⁵ and any remaining savings were transferred to clients' ordinary savings accounts.

The withdrawal-restriction account Gihandom was simpler in structure: Clients chose to restrict withdrawals before either their goal date or their goal amount (specified at contract signing) was reached. Out of 92 Gihandom clients, 39 chose

¹⁴ Attitudes towards charities were measured in the baseline, and an indicator for previous charitable contributions is available as a control variable. See Appendix C.3 for a description.

¹⁵ After the goal date, there was a one-week grace period to make any outstanding deposits, but no client made use of this.

3.2. EXPERIMENTAL DESIGN

the amount goal, and 53 chose the date goal. The goal amount can be interpreted as the stronger restriction, since additional deposits need to be made in order to receive savings back. Formally, there was no limit on how long individuals could take to reach the goal amount. However, as is common for Philippine banks, significant dormancy fees were applied after two years of inactivity. While the marketers encouraged individuals to deposit the first weekly instalment from their savings plan as an opening balance, the formal minimum opening balance for Giandom was 100 pesos. The difference of 50–150 pesos (depending on the savings goal) in the mandatory opening balances between ASA and Giandom is a possible explanatory factor in the difference between take-up rates. Finally, two features were common to both ASA and Giandom: First, opening balances for both products were deliberately collected one week *after* contract signing. The practical motivation behind this was to give individuals time to prepare for the expense. The theoretical motivation was to free the decisionmaker from temptation in the contract-signing period – a sophisticated hyperbolic discounter should choose a welfare-maximising contract when asked in period 0, but not necessarily when asked in period 1.¹⁶ Second, both products shared the same emergency provisions: In case of a medical emergency or death in the family, a relocation to an area not served by the bank branch, or a natural disaster (as declared by the government),¹⁷ clients could close their account and access their savings without any penalties. Within the six months of observation, no client exercised this option.

In order to identify the treatment effect of a commitment to fixed regular instalments, individuals were left to themselves during the savings period, without any help from deposit collectors or reminders. After all goal dates had been reached, a comprehensive endline survey was administered. The endline survey focused on all types of savings (including at home and in other banks), outstanding loans, expenditures, changes in income, and various types of shocks experienced since the baseline survey. In addition, existing ASA clients were offered the option to sign up for a 'Pre-Order' of the product: Clients were informed that the bank may decide to offer ASA for a second round, conditional on sufficient demand. While the Pre-Order did not involve a financial commitment, it involved the completion

¹⁶This approach is similar to that in Benartzi and Thaler (2004), who let employees commit to allocate *future* salary increases to their pension plan. It could be argued that the late collection of opening balances effectively just delayed when individuals entered the commitment contract. In a purely financial sense, this is true. However, signing the contract was associated with substantial paperwork, as well as a non-financial commitment to the marketers, who personally collected the opening balance after one week. Out of 159 individuals who initially signed the ASA contract, 45 failed to deposit an opening balance. The corresponding number for Giandom is 24 out of 116 initially signed contracts.

¹⁷Provided appropriate documentation, i.e. a hospital bill, death certificate, or proof of relocation.

of a new savings plan, a new 'Voluntary Donation Form', and a decision on a new termination fee (to deter cheap talk).

3.3 The Survey Instrument

The household survey administered at the beginning of the study had two objectives: First, to measure factors commonly suspected to influence the demand for (commitment) savings products. Second, the survey data on savings, loans, income, and expenditures provides the baseline for the estimation of treatment effects (see Section 3.5.1).

I measured time-inconsistent preferences using the common method of multiple price lists (MPLs): Individuals were asked to choose between a fixed monetary reward in one period and various larger rewards in a later period. A randomly chosen half of the sample received real rewards, whereas for the other half, the questions were hypothetical. After a set of questions using a near time frame (now versus one month), the same set of questions was repeated for a future time frame (one month versus two months). The outcome of interest was the size of the later reward necessary to make the individual switch from preferring the (smaller) earlier reward to the (larger) later reward. For illustration, consider the following sample questions:

1. Would you prefer to receive P200 guaranteed today, or P250 guaranteed in 1 month?
2. Would you prefer to receive P200 guaranteed in 1 month, or P250 guaranteed in 2 months?

The earlier reward was kept constant at 200 pesos, while the later reward gradually increased from 180 to 300 pesos. Individuals whose preferences satisfy standard exponential discounting will be time-consistent – i.e., the amount necessary to make them switch from the earlier reward to the later reward will be the same whether they are asked to choose between now and one month, or between one month and two months. I identify as hyperbolic discounter those who are impatient in the present, but patient in the future, i.e., the reward needed to make them wait for one month is larger in the present than in the future (thus the term 'present biased'). In the opposite direction, individuals who exhibit more patience now than in the future are classed as 'future biased'. An individual who always prefers the earlier reward in all questions (for both near and future time frames) is classified as 'impatient'. The two sets of questions were separated by at least 15 minutes of other

3.3. THE SURVEY INSTRUMENT

survey questions, in order to prevent individuals from anchoring their responses to earlier answers. An ad-hoc randomisation based on individuals' birthdays determined who played the game with real rewards (see Appendix C.3 for details). For those with real rewards, one of their choices was paid out, selected at random by drawing a ping pong ball with a question number from a black bag. To prevent uncertainty about whether future payments would be guaranteed (causing an upward bias of the present-bias measure), both cash and official post-dated bank cheques were presented during the game.

I find 16.6 percent of individuals to be present-biased, and 18.9 percent of individuals to be future-biased.¹⁸ An additional 36 percent of individuals were classified as 'impatient' (they consistently chose the earlier reward, implying a switching point outside of the 180–300 pesos survey range, and thus an inconclusive time-consistency status). These estimates are slightly below comparable estimates in the literature, but show a similar tendency for future bias to be as common as present bias (Ashraf et al. (2006b) find 27.5 percent present-biased and 19.8 percent future-biased, Giné et al. (2012) find 28.5 and 25.7 percent, respectively, Dupas and Robinson (2013) find 22.5 and 22 percent, and both Brune et al. (2011) and Sinn (2012) find 10 percent present-biased and 30 percent future-biased). Explanations that have been proposed for future bias include utility from anticipation (Loewenstein (1987), Ameriks et al. (2007)), varying degrees of future uncertainty (Takeuchi (2011), Sayman and Öncüler (2009)), and survey noise.

In addition to a standard measure of preference reversals, it is vital to the analysis to obtain a measure of sophistication. In particular, this measure should not in itself be derived from a demand for commitment. The approach pursued in this paper relies on a simple idea: Multiple price lists provide a measure of actual time-inconsistency, independent of an individual's awareness of said inconsistency. If *observed* inconsistency could be interacted with a measure of *perceived* inconsistency, a measure of sophisticated hyperbolic discounting would result.

Such an awareness measure exists: The self-control measure proposed by Ameriks et al. (2007), henceforth referred to as ACLT. Using survey questions, the authors elicit how individuals would *optimally* like to allocate a fixed resource over time. They then ask which allocation individuals would be *tempted* to consume (if not exercising self-control), and finally, which allocation they *expect* they would consume in the end. While originally intended to identify the parameters of the Gul

¹⁸At 17.9 percent, present bias was more frequent among those with hypothetical rewards, than among those with real rewards (15.2 percent), but the difference is not significant. This suggests that a bias from uncertainty is unlikely. A detailed comparison of real and hypothetical incentives is beyond the scope of this paper, and will be provided in a separate working paper on the elicitation of time preferences. Appendix C.2 gives a summary.

3.3. THE SURVEY INSTRUMENT

and Pesendorfer (2001) model, the questions require the individual to critically assess future temptations and their (hypothetical) response to them. The resulting measure reflects an individual's *perceived* (rather than actual) self-control problems. This makes the ACLT questions, interacted with a measure of observed time-inconsistency (e.g., through MPLs), a promising candidate to measure sophistication.

The setup is as follows: Respondents were presented with a hypothetical scenario of winning 10 certificates for “dream restaurant nights”. In this scenario, each certificate entitled the holder and a companion to an evening at any local restaurant of their choice, including the best table, an unlimited budget for food and drink, and all gratuities. The certificates could be used starting immediately, and would be valid for two years. Any certificates not used after two years would expire. I presented the ACLT scenario along with an example list of local middle-class restaurants which were chosen to be above what respondents could usually afford, and which were regarded as highly desirable. This was intended to prevent simple substitution of certificates into everyday consumption (given the low income levels in the sample, respondents were used to eat either at home, or in simple street eateries, *carinderias*). In addition, the restaurant framing has the added benefit of being directly linked to consumption, thus avoiding the common concern with cash rewards that money is fungible and does not have to be associated with an immediate consumption shock (cf. Frederick et al. (2002)). In line with the ACLT design, I then asked the following questions:

1. Think about what would be the ideal allocation of these certificates for the first and the second year. From your current perspective, how many of the ten certificates would you ideally like to use in year 1 as opposed to year 2?
2. Some people might be tempted to depart from this ideal allocation. For example, there might be temptation to use up the certificates sooner, and not keep enough for the second year. Or you might be tempted to keep too many for the second year. If you just gave in to your temptation, how many would you use in the first year?
3. Think about both the ideal and the temptation. Based on your most accurate forecast of how you would actually behave, how many of the nights would you end up using in year 1 as opposed to year 2?

The answers to these questions provide two important measures: Perceived self-control (from (3) – (1), *expected* – *ideal*) and perceived temptation (from (2) – (1), *tempted* – *ideal*). However, these measures were designed for the Gul-Pesendorfer

3.3. THE SURVEY INSTRUMENT

model, which does not directly translate into the $\beta\delta$ -model of hyperbolic discounting which underlies this analysis. The models are not nested, and there is no direct equivalent to self-control and temptation in the model of hyperbolic discounting. From the perspective of the $\beta\delta$ -model (where self-control does not exist), we would expect the two measures to be the same - namely the difference between the optimal ex-ante allocation, and (the individual's perception of) the allocation that results in a subgame perfect equilibrium between the different selves. Following this logic, both measures are equally suitable to assess an individual's awareness of their time-inconsistency.

For the purposes of my empirical analysis, I choose to focus on *tempted* – *ideal* as an awareness measure for time-inconsistency. I then interact awareness of time-inconsistency with observed time-inconsistency (in MPLs), and obtain a measure of sophisticated hyperbolic discounting: *tempted* – *ideal* * *presentbias*.¹⁹ The reason for focusing on the temptation measure is as follows: Suppose costly self-control does exist. An individual who exercises full self-control, and thus realises the (ex-ante) ideal allocation, might still have a demand for commitment. While a commitment device would not change the de-facto allocation he consumes, it can increase his utility by removing temptation, and thus the need to exercise costly self-control. Therefore, a low or zero measure of *expected* – *ideal* (i.e., good self-control) might still be associated with a demand for commitment, while time-consistent preferences would not. As a result, for the purposes of analyzing demand for a commitment savings product, the temptation measure provides a better indication of whether individuals feel they could benefit from commitment. In this sense, perceived temptation is closely related to the concept of sophistication.

I observe that 81.6 percent of individuals report strictly positive values of temptation, with a median temptation of two certificates. Given the much lower frequency of observed present bias, the question arises how to interpret temptation without present bias. This paper remains agnostic about the precise theoretical connection between models of self-control and those of hyperbolic discounting, and instead offers a simple intuition: If an individual reports to be tempted, but behaves in a time-consistent fashion, this may be due to the exercise of self-control. This hypothesis is supported by the data: Conditional on a given level of temptation, non-present biased individuals report significantly better self-control than present-biased individuals.

¹⁹I censor values of temptation and self-control at zero. I interpret observed negative values as measuring something other than temptation and self-control – e.g., not having time to go to restaurants as often as individuals would ideally like, or inability to understand the survey question. Negative values occurred in 4 (42) out of 910 cases for temptation (self-control).

3.3. THE SURVEY INSTRUMENT

TABLE 3.1: SUMMARY STATISTICS BY TREATMENT ASSIGNMENT

	R-Group	W-Group	Control	F-stat P-value
Age*	43.8337 (0.6029)	43.4493 (0.8214)	44.25 (0.8412)	0.8039
Female*	0.9409 (0.0110)	0.9430 (0.0154)	0.9430 (0.0154)	0.9912
Education (yrs)	10.55604 (0.1662)	10.39207 (0.2417)	10.56388 (0.2513)	0.8398
HH Income	2890.89 (124.26)	2485.78 (165.13)	3194.43 (272.45)	0.0481
#HH members	5.07221 (0.0909)	5.179825 (0.1399)	5.429825 (0.1398)	0.1081
Real Rewards*	0.5033 (0.0234)	0.5219 (0.0332)	0.5263 (0.0331)	0.8371
Financial Claims*	0.3934 (0.0229)	0.3877 (0.0324)	0.3860 (0.0323)	0.9867
Existing Savings Account	0.4683 (0.0234)	0.4649 (0.0331)	0.4254 (0.0328)	0.5176
Impatience	0.3217 (0.0219)	0.4035 (0.0326)	0.3333 (0.0313)	0.0959
Present Bias*	0.1723 (0.0180)	0.1614 (0.0246)	0.1560 (0.0246)	0.8388
Perceived Temptation	2.3838 (0.0889)	2.1850 (0.1122)	2.4714 (0.1210)	0.2249
Risk Aversion	4.2254 (0.0932)	4.6360 (0.1219)	4.1316 (0.1289)	0.0104
Cognitive Ability	2.9365 (0.0592)	2.8860 (0.0889)	2.9342 (0.0955)	0.8870
Financial Literacy	1.8556 (0.0466)	1.8377 (0.0682)	1.8509 (0.0694)	0.9767
Donates to Charity	0.3961 (0.0229)	0.3860 (0.0323)	0.4518 (0.0330)	0.2836
N	457	228	228	913

Note: A starred variable indicates that the randomisation was stratified on this variable.

3.4. EMPIRICAL STRATEGY

In addition to the measures for present bias and sophistication (*tempted – ideal * presentbias*), the survey obtained measures of financial claims from others, risk aversion, cognitive ability, financial literacy, bargaining power within one’s household, distance to the bank branch (via GPS coordinates), attitudes towards charitable giving, and frequency of income or expenditure shocks, as well as an indicator for having an existing bank account. These measures are discussed in Appendix C.3.

Table 3.1 presents summary statistics for the main observed covariates from the survey. Tests for equality of means across treatment groups were conducted to verify that the randomisation was balanced. Randomisation into treatment groups occurred shortly after the baseline survey, which means that covariates were available at the time of randomisation. A star next to a variable in Table 3.1 indicates that the randomisation was stratified on this variable. In three cases, means were statistically different across treatment groups: Income, impatience and risk aversion. Income and impatience have no predictive power in any of the later regressions. In particular, wealthier individuals are no more likely to take up a commitment product than poorer individuals. Risk aversion does have predictive power for the take-up of Gihandom (W-group). Robustness checks are reported in Appendix C.2.

3.4 Empirical Strategy

The primary objective of the study was to analyse the demand for and the effects of a regular-installment commitment savings product, and to compare its performance to traditional withdrawal-restriction commitment products. Given the heterogeneity of results, a deduced objective is to document possible risks of commitment contracts, in particular with respect to partially sophisticated hyperbolic discounting. The main outcomes of econometric interest are a range of treatment effects (on bank savings, other savings, loan demand, and expenditures), as well as predictors of take-up, contract outcome (successful or default), and the pre-order decision (comparable to repeat take-up).

For the estimation of treatment effects, denote by R_i an indicator variable for assignment to the ‘Regular Saver’ treatment group – all individuals in this group were offered the ASA product. Denote by W_i an indicator variable for assignment to the ‘Withdrawal Restriction’ group – all individuals in this group were offered the Gihandom product. Treatment effects can be estimated using the equation

$$\Delta Y_i = \alpha_0 + \alpha_R R_i + \alpha_W W_i + \epsilon_i \quad (3.1)$$

3.4. EMPIRICAL STRATEGY

where ΔY_i denotes the change in the outcome variable of interest. In Section 3.5.1, I focus on bank savings, but also provide estimates for total savings, loan demand, and expenditures (see Figure C.1). Bank savings refers to the change in savings held at 1st Valley Bank. For ASA clients, this is the sum of their savings in ASA plus their savings in the non-commitment savings account provided to them. For Gihandom clients, it is the sum of their Gihandom savings plus their savings in the ordinary account. For everyone else (i.e., the control group and those who rejected the commitment product offered to them), bank savings refer to their ordinary savings account only (recall that individuals were encouraged to use the ordinary savings account to follow the personal savings plan provided to them). Summing all existing savings accounts per individual means that crowd-out between savings devices at the bank will not impact the analysis. However, individuals could have substituted away from home savings, or savings at other banks. To observe such effects, I also analyse a measure of other savings, obtained from survey data, which includes home savings, money lent out to others or safekept elsewhere, and money at other banks. The time frame for measuring savings runs from the date of the baseline survey visit to the individual's savings goal date – i.e., savings durations vary at the individual level. This is a consequence of focusing the marketing on particular expenditures: If savings were measured at the end of the study, even a successful saver would have a savings balance of zero if he has already paid for the expenditure.

An OLS estimation of equation 3.1 provides $\hat{\alpha}_R$ and $\hat{\alpha}_W$ – estimates of the Intent-to-Treat effects of the regular-instalment product ASA and the withdrawal-restriction product Gihandom. The ITT measures the mean causal effect of having been *offered* the product, which is likely to be an average of the effect of using the product, and of simply feeling encouraged to save because of the product offer. However, as has been outlined in Section 3.2, individuals in all groups received an identical marketing treatment. Only after a personal savings plan for an upcoming expenditure had been made, and an ordinary savings account had been opened, did the marketers offer individuals in groups R and W the possibility to bindingly commit to selected features in their savings plan. Under the assumption that the mere offer of commitment has no effect on savings (other than via encouraging people to use the product), the ITT will be a composite of the Treatment-on-the-Treated effect (TOT) on those who took up the product offered to them, and a zero effect (relative to the control) on those who did not take up the product.²⁰ In this case, the TOT can be estimated by dividing the ITT ($\hat{\alpha}_R$ and $\hat{\alpha}_W$) by the fraction of take-ups. Alter-

²⁰See Imbens and Angrist (1994) and Duflo et al. (2007) for a discussion on ITTs and local average treatment effects.

natively, equation 3.1 can be estimated using an instrumental variables approach, with take-up (ASA_i and $Gihandom_i$) as the regressors and assignment to treatment (R_i and W_i) as orthogonal instruments.

Predictors of the take-up, default and pre-order decision can be summarized in a binary choice equation. I use a probit model to estimate

$$Choice_i = \beta_0 + \beta X_i + \epsilon_i,$$

where $Choice_i$ can be an individual's decision to take up ASA (if in group R), to take up Gihandom (if in group W), to default on an ASA contract, or to pre-order ASA for a second round. The vector X_i contains demographics (age, gender, marital status, income, assets, household size, years of education), as well as all survey-based measures mentioned in Section 3.3. In addition, all binary choice regressions contain marketer fixed effects. This is to filter any noise from differences in marketer ability.

3.5 Results

3.5.1 Average Treatment Effects

Effect on Total Bank Savings

This section presents estimates of the effects of the two commitment treatments on individuals' total savings held at the partner bank. The outcome variable of interest is the change in a client's total savings balance at the partner bank, summed across ordinary savings accounts and any commitment savings products (ASA or Gihandom). The savings period is specific to each individual, starting with the date of the baseline survey, and ending with the goal date specified in an individual's personal savings plan.²¹ The cost of this reliance on the goal date is that it diminishes the sample to those 748 individuals who a) could be located for the marketing visit and b) were willing to make a savings plan with the marketer. This form of attrition is orthogonal to assignment to treatment.

Column (1) in Table 3.2 estimates that assignment to the Regular Saver treatment group increased average bank balances by 585 pesos (U.S.\$14) relative to the control group. This estimate already includes any charged termination fees due

²¹All accounts except for those of existing 1st Valley Bank clients were opened after the marketing stage, implying the observed change in savings is equal to the savings balance. For those 18 clients who had previously existing 1st Valley Bank savings accounts, the existing account was monitored instead of opening a new ordinary savings account. Existing bank clients were still offered commitment savings products, in accordance with their assignment to treatment.

TABLE 3.2: SAVINGS OUTCOMES (OLS, PROBIT)

	(1) Change in Bank Savings	(2) Purchased Savings Goal	(3) Borrowed to Purchase Goal (given purchase)	(4) Change in Other Savings (survey-based)
Regular Saver Treatment (ASA)	585.4652*** (129.2510)	0.1156** (0.0486)	0.0509 (0.0621)	426.8112 (671.8442)
Withdrawal Restr. Treatment (Gihandom)	148.2429*** (40.9269)	0.1322** (0.0545)	0.2109*** (0.0808)	-328.1585 (705.4607)
Constant	27.1600*** (9.3987)			63.4513 (531.0279)
Mean Dep. Variable		0.4992	0.1922	
R ²	0.02			0.00
Observations	748	615	307	603

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Entries in columns (1) and (4) represent OLS coefficients. Entries in columns (2) and (3) represent marginal coefficients of the corresponding probit regressions.

to default. In contrast, individuals assigned to the Withdrawal-Restriction group saved on average 148 pesos more than the control group. Noting that the average duration of savings periods was 130 days (about 4.5 months), this estimate is roughly comparable to the effect estimated in Ashraf et al. (2006b): In a sample of previous savings account holders, the authors find that their withdrawal-restriction product SEED increased average savings by 411 pesos after 12 months. Given that the product design of SEED and Gihandom was identical, the Gihandom estimates presented here also serve to replicate and confirm the results of Ashraf et al. (2006b). Furthermore, the estimates confirm a small but significant increase of 27 pesos in savings for the control group. Two interpretations are possible: First, the marketing treatment could have led to higher savings even in the absence of commitment products. Second, the savings increase could be a result of the monetary rewards received in the baseline survey. Randomisation into treatment groups was stratified on whether individuals had received real rewards, which ensures that the resulting income shock is exogenous to treatment. Note that savings increases are net of the 100 peso opening balance contained in the free ordinary savings account – this amount constituted the minimum account maintaining balance, and no client closed their ordinary savings account during the period of observation.

In addition to the ITT effects reported in Table 3.2, an instrumental variables regression of the change in bank savings on an indicator for take-up of the commitment products (using assignment to groups R and W as orthogonal instruments) provides an estimate of the Treatment-on-the-Treated (TOT) effect (discussed in Section 3.4). The TOT regression suggests that taking up the regular-instalment product ASA increased savings by 1928 pesos, while taking up the withdrawal-restriction product Gihandom increased savings by 324 pesos. Both estimates are conditional on the assumption that being offered a commitment product has no effect on savings, other than through use of the product (equivalently, those who rejected the commitment products on average saved the same as clients in the control group). This assumption is supported by the fact that the marketing treatment was identical across all groups. The increased gap in the TOT effects of ASA and Gihandom relative to their ITT effects is a result of the higher take-up rate for the Gihandom product.

The remainder of Table 3.2 presents the results from a probit estimation of whether individuals purchased the savings goal (see Table C.1) they had been saving for: At the end of the endline survey, individuals were asked whether they had purchased the good specified on their personal savings plan.²² This is a distinct

²²The survey team was informed about this savings purpose, in case individuals had forgotten.

question from whether individuals achieved a certain amount of money in a bank account – they could have saved for the good at home, or found a different way to pay for it. If respondents confirmed having purchased the desired good, they were further asked how they paid for it, and in particular whether they borrowed (from any source, including friends or family). Due to attrition in the endline survey, the sample for this estimation is limited to the 615 individuals who a) had made a savings plan during the marketing stage and were b) reached by the endline survey.²³ Exactly half of the individuals reported to have bought the good, or paid for the expenditure, that was named on their personal savings plan. These 307 individuals constitute the sample for the probit regression in column (3), which estimates the relationship between treatment and the likelihood of borrowing for the purchase (conditional on purchase). Due to sample selection issues, the coefficients in column (3) do not have a causal interpretation, and should be interpreted as correlations. Borrowing was not uncommon: Slightly below 20 percent of individuals chose loans or family borrowing as a means of affording the expenditure.

Column (2) of Table 3.2 confirms that both the Regular Saver treatment and the Withdrawal Restriction treatment increased an individual's chances of purchasing their savings goal. The coefficients for the two treatments are not significantly different from each other. However, column (3) shows that individuals in the Withdrawal Restriction group were significantly more likely to borrow in order to obtain the good: Considering marginal effects, assignment to group W is associated with a 19.6 percentage points higher likelihood of borrowing (of 31 rather than 11.4 percent). In comparison, assignment to group R (being offered ASA) increased the likelihood of obtaining one's savings goal, but seems unrelated to the probability of borrowing for the expenditure. This may suggest that the ASA product indeed helped individuals to purchase a savings goal using their own money, and without the use of loans.

Figure C.1 (Appendix C) shows the impact of the Regular Saver treatment and the Withdrawal-Restriction treatment on the cumulative distribution of changes in bank savings, total savings, outstanding loans, and expenditures.²⁴

Testing for Crowd-Out of Savings

A caveat about the estimation presented above is that it is restricted to savings at the partner bank. During the baseline survey, 46 percent of the sample reported

²³Both 'having a savings plan' and 'being reached by the endline survey' are orthogonal to assignment to treatment.

²⁴The observation period ends with the goal date for bank savings and total savings, and with the date of the endline survey for outstanding loans and expenditures.

to have an existing savings or checking account. This number is partly driven by microentrepreneurs, who are required to hold an existing savings account when obtaining microloans (the pairwise correlation is 0.18). More than one quarter of bank account holders reported not to have used their account in the last 12 months, and dormant accounts were common. The regression in column (4) of Table 3.2 seeks to establish whether the savings increases observed at the partner bank constituted new savings, or whether a simple substitution from other sources of savings (at home, or at other institutions) took place.

The outcome variable in column (4) is the change in an individual's total savings balance outside of the partner bank, as measured by survey data: During the baseline survey, individuals were asked about their savings at home, money lent out or safekept by others, informal savings, and savings at other institutions. An incentive of 30 pesos was paid for showing an existing bank passbook. The same exercise was repeated during the endline survey six months later, except that individuals were now questioned about the savings they kept around the time of their goal date. Unfortunately, the survey data is very noisy, and coefficients are estimated with substantial imprecision.²⁵ The available evidence does not suggest that a substitution took place between savings increases at the partner bank, and savings at home or at other institutions. All coefficients are insignificant. Moreover, the coefficient for being assigned to the Regular Saver treatment is positive – if anything, individuals who were offered ASA may have been encouraged to save even more in other savings vehicles, in addition to deposits made to their ASA accounts. In contrast, the coefficient for the Withdrawal Restriction treatment is negative. While this could easily represent survey noise, it is consistent with the 'safekeeping' explanation discussed earlier: Individuals may decide to shift existing assets into an account where they know other members of their household will not be able to access them.

3.5.2 Heterogeneity: Descriptive Results

The ASA results were very bi-modal: At the time of their goal date (between December 2012 and April 2013), 51 ASA clients (45 percent) had reached their savings goal. They had completed all scheduled instalments with a median of 12 transactions,²⁶ and reached savings goals between 950 and 7150 pesos (U.S.\$170). By design, accounts were closed after completion of the savings plan, and clients could

²⁵To account for some outliers in the stated balances, the savings data has been truncated at 1 percent, reducing the sample from 615 to 603 observations.

²⁶ One transaction can cover several weeks' instalments.

withdraw their savings in order to pay for the planned lump-sum expenditure (any remaining savings were transferred onto clients' ordinary savings account). Many of these clients pro-actively enquired at the bank to roll over their account into a new ASA contract. While rolling over contracts was not an immediate possibility during the study period, the repeat marketing stage included the option to 'pre-order' the product for a second round, should the bank decide to offer the product again. The pre-order contract was not financially binding, but included substantial official paperwork. Two thirds of the successful clients took up this offer (see Table 3.4), devised a new savings plan, and chose a new termination fee. The bank has since decided to offer new ASA contracts to those enquiring about them at the branch.

The situation looked very different for the remaining 63 ASA clients (55 percent) who defaulted on their savings contract. After falling three deposits behind, their accounts were closed, and the initially agreed termination fee charged (and transferred to charity). What happened? Two possibilities emerge:²⁷ (i) Clients had chosen an ASA contract which was optimal for them in *expectation*, and then rationally defaulted upon observing a shock (in other words, a 'bad luck' scenario). Or (ii), clients chose the contract by mistake. If the 'bad luck' explanation is true, the timing of the defaults should depend on the shock distribution: If shocks are independently distributed across individuals and time, and hazard rates are small, the timing of defaults should be roughly uniform over time. In sharp contrast, Figure 3.1 illustrates that clients had a tendency to default either right from the start, or not at all: Out of 63 defaults, 35 clients stopped depositing immediately after the opening balance, 10 clients made one more deposit, another 10 made between three and five deposits, and only 8 clients made more than five deposits (see Figure 3.1). Approximating transactions with weeks (85 percent of clients chose weekly instalments), Figure 3.1 also illustrates the expected default timing given a hazard rate of 0.028 per week. This estimate is obtained from the endline survey: The sample population was questioned about the occurrence of 17 types of common emergencies (sickness, loss of job, bad business, flood damage) including a flexible 'other' category. 45% reported at least one shock within 6 months, with an average of 0.72 shocks, equivalent to 0.028 shocks per week. This hazard rate is neither consistent with the overall frequency of defaults (observed 55 percent versus predicted 29 percent based on a 12 week contract), nor with the steep timing of defaults. A much higher hazard rate of 0.56 shocks per week would be needed to explain that 56 percent of all defaults happened immediately after opening. A hazard rate this

²⁷Strictly speaking, this assumes that an individual would not take up a contract if he knows that default is certain.




Figure 3.1: Savings Transactions: Defaulted ASA Clients

high is problematic: It predicts an overall default frequency of 99 percent within 6 weeks, which contradicts both the observed 45 percent ‘success rate’ on contracts lasting 12-24 weeks, as well as the thick tail of the default distribution (13 percent of defaults occur more than 6 weeks after opening). The observed default timing is difficult to reconcile with the exponential pattern that would be generated by any i.i.d. hazard rate. Unless there was an aggregate shock which affected all defaulting clients immediately after opening their accounts, a pure shock explanation seems unlikely. Evidence of aggregate shocks is discussed in Section 3.6.2. Heterogeneous hazard rates are discussed in Section 3.6.1.

The second possibility requires a deviation from rational expectations: Individuals could have chosen their contract by mistake. Mistakes (defined as choices that are not optimal under rational expectations) can happen if individuals have incorrect beliefs about their future preferences or their income distribution, including the probability of shocks to either of the two. Chapter 2 outlines why a time-inconsistent agent with incorrect beliefs about the degree of his time-inconsistency is likely to select into a commitment contract that is too “weak” to overcome his self-control issues, leading to default. Looking at the data, it is notable that 80 percent of individuals chose the minimum permissible termination fee for their savings goal (P150 for goals below P2500, and P250 for goals of P2500 and above), roughly equivalent to a day’s wage. The observed combination of minimum penal-

ties and high default rates raises the question whether individuals underestimated the amount of commitment it would take to make them save. This is consistent with the observed tendency to default soon after account opening, as individuals start behaving according to their true degree of time-inconsistency upon entering the depositing phase. Could rational expectations about stochastic future time-inconsistency explain the data? If individuals had correct beliefs on average about their future preferences, they should realize which contract (and in particular, which penalty) will be effective for them on average. Moreover, risk-averse preferences imply that individuals who internalize the risk of default should either sign up for stronger commitments (to be on the safe side), or stay away from commitment. This is inconsistent with the frequency of observed defaults, and the tendency to choose the minimum penalty.

Figure C.2 (Appendix C) lists the chosen termination fees of the 114 ASA clients, and contrasts them with how much was charged ('successful' indicates that no fee was charged). Not all chosen fees were enforceable: Whenever clients chose a fee strictly above the minimum and later defaulted on their contract, the charged fee was the lesser of chosen fee and savings balance at the time of default. The minimum fee was always enforceable through the opening balance.

Unfortunately, it is safe to conclude that the ASA contract likely reduced the welfare of a significant share of its adopters. For the 35 clients who defaulted immediately after depositing the opening balance, losing the opening balance (through the termination fee) was the only economic consequence of the contract, thus leaving them worse off. For those who defaulted later during their savings plan (thus making a shock explanation more likely), an argument can be made that the contract helped them to achieve savings which they would not otherwise have been able to achieve, at a negative return of 150 or 250 pesos (which is still less than common interest payments to local moneylenders). A cautious estimate of the frequency of 'mistakes' is provided by the pre-order results: 55 percent of all ASA clients (71 percent of defaulting clients and 35 percent of successful clients) chose not to order the product again (see Table 3.4).

For the Gihandom accounts, both benefits and risks were less pronounced: Out of 92 accounts, only five reached the goal amount specified in their savings plan (three were date-based, two were amount-based). The 53 clients who had opted for a binding date goal received their savings back after the savings period completed. Their median savings were 100 pesos (average 286 pesos), which is equivalent to the minimum opening balance. Out of 39 clients who had opted for binding goal amounts, 35 were still open at the end of the six-month observation period (average

TABLE 3.3: ACCOUNT USAGE

Average # of Deposits (includes opening balance)	Mean	Median	#accounts
ASA (all)	6.76	5	114
_____(successful)	11.98	12	51
_____(default)	2.52	1	63
Gihandom Account (all)	1.68	1	92
_____(date-based)	1.68	1	53
_____(amount-based)	1.69	1	39
Control Savings Account	0.43	0	788

TABLE 3.4: ASA PRE-ORDER

	Yes	No	
Successful	33	18	51
Default	18	43	63
	51	63	114

savings 141 pesos).²⁸ 85 percent of all amount-based Gihandom accounts (and 79 percent of Gihandom accounts overall) had no further deposits after the opening balance. This creates a parallel between Gihandom and ASA defaults: Similarly to the ASA clients who made no further deposits, amount-based Gihandom accounts effectively lose their opening balance if they do not continue to deposit. A difference between the two commitment products is that the penalty for discontinuing to save on an amount-based Gihandom account increases with every deposit, while the default penalty for ASA is fixed.

Finally, 582 clients opened exclusively an ordinary savings account – either because they were assigned to the control group, or because they rejected the commitment product offered to them. Out of these clients, one reached the goal amount specified within their savings plan. Summary statistics for transactions in all accounts can be found in Table 3.3.

3.5.3 Heterogeneity: Regressions

In an attempt to resolve the puzzles presented in the previous section, this section analyses empirical predictors of take-up for the two commitment products, as well as default and pre-order decisions.

²⁸Two accounts were closed after reaching the goal amount, and another two were closed after the bank mistakenly treated them as date-based accounts.

3.5.3.1 Predicting Take-Up of the Commitment Savings Products

Columns (1) to (3) in Table 3.5 present the results of a probit regression of the ASA take-up decision on a number of potential determinants, limiting the sample to the Regular Saver group (R), where ASA was offered.²⁹ The first notable fact is that not a single demographic factor seems to correlate with the take-up decision. Age, gender, income, assets, marital status, education and household size all appear to be insubstantial for the decision to take up the regular-installment product.

The main factors which do predict ASA take-up are the proposed measure of sophisticated hyperbolic discounting (see Section 3.3) and a measure of cognitive ability (see Figure C.4 for a sample question from the cognitive ability test). The positive predictive power of cognitive ability is reassuring: The ASA product is more complex in its rules than traditional savings accounts (but no more complex than a loan contract). The significance of cognitive skills suggests that those clients who were more likely to understand the rules were also more likely to take up the product. This may also be interpreted as evidence against possible manipulation by the bank marketers.

Present bias on its own is not a predictor of take-up, consistent with the intuition that *perceived* time-inconsistency, rather than *actual* time-inconsistency, determines demand for a commitment product. Perhaps more surprisingly, the association of commitment take-up and sophisticated hyperbolic discounting is significant and *negative*. Recall from Section 3.3 that sophistication is measured as the interaction of present bias (from multiple price list questions) and self-reported temptation. In other words, those who exhibit hyperbolic preference reversals, but at the same time report *low* levels of temptation, are more likely to take up the product. In contrast, those who report being strongly tempted tend to stay away from the product. To interpret interaction coefficients, note that present bias is a binary variable, whereas temptation is in the interval $[0, 10]$. A possible explanation for this link comes from theory: Partially sophisticated agents (i.e. those with time-inconsistent preferences and positive but low self-perceived temptation) have a positive demand for commitment. They take up the product and choose a low default penalty, which they anticipate will be sufficient to make them save. In contrast, agents who perceive themselves as strongly tempted have two choices: Either they take up the product with a penalty that is sufficiently large to make them save, or they stay away from the product completely. The latter choice may be optimal if the required

²⁹The sample for the regression is restricted to those clients who could be located for the marketing visit. Inability to locate individuals for marketing is orthogonal to treatment group assignment.

TABLE 3.5: PREDICTING DEMAND FOR COMMITMENT (PROBIT)

Commitment Take-Up	ASA (1)	ASA (2)	ASA (3)	Gihandom (1)	Gihandom (2)	Gihandom (3)
Age	-0.0015 (0.0019)	-0.0004 (0.0020)	-0.0003 (0.0020)	0.0017 (0.0029)	0.0009 (0.0028)	0.0011 (0.0028)
Female	0.0328 (0.0914)	0.0592 (0.0936)	0.0536 (0.0868)	0.2418 (0.1687)	0.2347 (0.1501)	0.2300 (0.1516)
Married	0.0076 (0.0650)	0.0095 (0.0640)	0.0165 (0.0630)	-0.0932 (0.0939)	-0.0848 (0.0872)	-0.0952 (0.0887)
HH Income	0.0018 (0.0081)	0.0013 (0.0080)	-0.0028 (0.0079)	0.0003 (0.0118)	-0.0054 (0.0108)	-0.0052 (0.0109)
Assets	0.0007 (0.0145)	-0.0048 (0.0143)	-0.0084 (0.0143)	0.0220 (0.0224)	0.0253 (0.0213)	0.0258 (0.0216)
HH Members	0.0125 (0.0100)	0.0105 (0.0097)	0.0130 (0.0095)	0.0229 (0.0157)	0.0202 (0.0154)	0.0217 (0.0158)
Education (yrs)	-0.0055 (0.0066)	-0.0075 (0.0064)	-0.0094 (0.0064)	0.0276*** (0.0094)	0.0316*** (0.0093)	0.0315*** (0.0093)
Present Bias	0.0757 (0.0866)	0.0636 (0.0870)	0.0827 (0.0864)	0.0809 (0.1260)	0.1020 (0.1256)	0.1046 (0.1301)
Soph. Present Bias (Pres.Bias*Temptation)	-0.0622** (0.0291)	-0.0579** (0.0290)	-0.0631** (0.0292)	-0.0363 (0.0491)	-0.0524 (0.0510)	-0.0530 (0.0532)
Perceived Temptation	-0.0114 (0.0123)	-0.0067 (0.0124)	-0.0046 (0.0125)	-0.0058 (0.0212)	0.0005 (0.0205)	0.0006 (0.0209)
Impatience	-0.0047 (0.0476)	0.0053 (0.0467)	-0.0008 (0.0464)	-0.0124 (0.0717)	-0.0224 (0.0684)	-0.0219 (0.0688)
Financial Claims	-0.0022 (0.0426)	-0.0076 (0.0418)	-0.0038 (0.0414)	0.1079 (0.0663)	0.1166* (0.0638)	0.1185* (0.0646)
Risk Aversion		-0.0049 (0.0106)	-0.0059 (0.0105)		0.0497*** (0.0167)	0.0500*** (0.0168)
Cognitive Ability		0.0353* (0.0188)	0.0363* (0.0187)		0.0157 (0.0236)	0.0174 (0.0239)
Financial Literacy		0.0425* (0.0246)	0.0328 (0.0250)		-0.0115 (0.0322)	-0.0120 (0.0321)
HH Bargaining Power		0.0063 (0.0113)	0.0053 (0.0113)		0.0444*** (0.0164)	0.0456*** (0.0167)
Distance to Bank			-0.0265 (0.0207)			0.0084 (0.0260)
Exist. Savings Account			0.0998** (0.0444)			-0.0022 (0.0667)
Donates to Charity			0.0221 (0.0424)			-0.0100 (0.0625)
#Emergencies last yr			-0.0161 (0.0277)			-0.0215 (0.0493)
Marketer FE	YES	YES	YES	YES	YES	YES
Mean Dep. Variable	0.2687	0.2687	0.2687	0.4115	0.4115	0.4115
Observations	402	402	402	209	209	209

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Entries in the table represent the marginal coefficients of the probit regressions.

effective penalty is very high: Given a constant probability of ‘rational default’, in which a shock (say, the loss of one’s business) makes it optimal for an individual to discontinue their contract, agents with a higher default penalty have more to lose. As a result, for a fully sophisticated agent with medium to high time-inconsistency, a low penalty may not be effective, and a high penalty may be too risky in the face of uncertainty.

Column (3) looks beyond ‘deep’ individual characteristics and investigates correlations with other choices. While neither distance to the bank branch, attitude towards charities (proxied by having donated any positive amount to charity in the past 12 months) nor estimated shock frequency significantly affect take-up probability, individuals with an existing bank account (at any local bank) were more likely to take up the product. Given a widespread scepticism towards banks in the study area, this may be interpreted as a sign of trust in and familiarity with the banking system.

Columns (4) to (6) present the same regressions applied to take-up for the withdrawal-restriction product Gihandom, limiting the sample to group W (where Gihandom was offered). Most strikingly, there is no overlap in the factors predicting ASA and Gihandom. If the products were perceived as close substitutes, and individuals in need of commitment merely took up whichever commitment product was offered to them, then the empirical analysis should find that the same factors which predict ASA take-up also predict take-up of Gihandom. A look at the data confirms that the sets of determinants for the two products are mutually exclusive, suggesting that individuals perceived ASA and Gihandom rather differently. Specifically, Gihandom take-up is predicted by high education (measured in years of schooling), high risk aversion (choosing a ‘safe’ lottery in Figure C.3), high household bargaining power (measured using questions on who decides what in a household), and strong claims from others on own liquid assets. Considering a 94 percent female sample population, this combination of factors is reminiscent of the evidence presented in Anderson and Baland (2002): In their study, the authors argue that Kenyan women use commitment devices (ROSCAs) to protect their savings against claims from their husbands. They propose an inverted U-shaped relationship between women’s power in their household, and participation in ROSCAs. While the Kenyan context studied by Anderson and Baland (2002) is different from the Philippine context studied here, the evidence in Table 3.5 is consistent with the explanation that women took up Gihandom to ‘safeguard’ their savings from intra-household conflicts. The withdrawal restriction featured in the Gihandom account is well-suited to preventing other household members from ac-

cessing savings, but allows the woman to retain flexibility regarding when to make deposits. The estimated linear relationship of commitment take-up with household bargaining power is unable to capture the proposed inverted U-shape. However, both household bargaining power and female education may be associated with an increased autonomy of the woman in planning to build up savings of her own. Finally, the strong predictive power of risk aversion is consistent with a precautionary savings motive: Those women who are particularly concerned about consumption variance and the possibility of shocks will be more interested in putting savings aside for future hard times.

No evidence currently suggests that demand for the withdrawal-restriction product Gihandom is associated with intra-personal conflicts and time-inconsistency. A reservation must be made with respect to statistical power: The sample of group W is half the size of group R, reducing the precision of estimates. Summing up, the evidence currently available suggests that demand for ASA is related to time-inconsistency and partial sophistication, while demand for Gihandom appears to be related to household bargaining and safekeeping motives.

3.5.3.2 Predicting Default and Repeat Take-Up

Table 3.6 presents marginal coefficients from probit regressions with ASA default as well as the ASA pre-order decision as the dependent variable. A take-up regression (column (3) from Table 3.5) has been added for comparison. In addition to the regressors from the take-up regressions, Table 3.6 also includes the number of emergencies (illness or death of household members, unemployment, damage due to natural disasters, and a range of other income and expenditure shocks) which the household suffered since the baseline survey.

Column (2) predicts ASA default among those randomly assigned to the Regular Saver group (R). Individuals who do not commit cannot default. Thus, the coefficients are best understood as predictors of who took up the commitment product ‘by mistake’, proxied by take-up and subsequent default. Note that this interpretation abstracts from the possibility of rational default following a shock realization. The results provide further support to the partial sophistication hypothesis: Present-biased individuals are significantly more likely to take up the ASA product and then default. This effect is particularly strong for agents who report low levels of temptation, representing naive and partially sophisticated hyperbolics. In contrast, more sophisticated hyperbolics are *less* likely to default: Note that temptation is in $[0, 10]$ with a median of 2. Aggregating the coefficients for present bias (0.11*), sophistication (-0.045**) and temptation (-0.02*) yields a lower likelihood of default

TABLE 3.6: ASA DEFAULTS & REPEAT TAKE-UPS (PROBIT)

Dependent Variable	(1) ASA Take-Up	(2) Default (R-Sample)	(3) Default (takeup-Sample)	(4) Pre-Order (takeup-Sample)
Age	-0.0003 (0.0020)	-0.0024 (0.0016)	-0.0057 (0.0041)	-0.0080* (0.0043)
Female	0.0536 (0.0868)	0.1189 (0.0911)	0.3663* (0.1871)	0.0745 (0.1942)
Married	0.0165 (0.0630)	0.0064 (0.0536)	0.0363 (0.1369)	-0.2433 (0.1490)
HH Income	-0.0028 (0.0079)	0.0034 (0.0057)	0.0206 (0.0159)	0.0124 (0.0177)
Assets	-0.0084 (0.0143)	-0.0093 (0.0115)	-0.0384 (0.0238)	0.0007 (0.0298)
HH Members	0.0130 (0.0095)	0.0123 (0.0077)	0.0159 (0.0143)	0.0015 (0.0179)
Education (yrs)	-0.0094 (0.0064)	-0.0030 (0.0053)	0.0016 (0.0121)	-0.0133 (0.0142)
Present Bias	0.0827 (0.0864)	0.1119* (0.0654)	0.4837* (0.2550)	-0.4860** (0.2353)
Soph. Present Bias (Pres.Bias*Temptation)	-0.0631** (0.0292)	-0.0453** (0.0230)	-0.1718 (0.1299)	0.2375** (0.1170)
Perceived Temptation	-0.0046 (0.0125)	-0.0202* (0.0105)	-0.0655*** (0.0244)	0.0109 (0.0301)
Impatience	-0.0008 (0.0464)	-0.0030 (0.0372)	0.0293 (0.0882)	0.0319 (0.1001)
Financial Claims	-0.0038 (0.0414)	-0.0113 (0.0330)	0.0095 (0.0848)	0.0097 (0.0905)
Risk Aversion	-0.0059 (0.0105)	-0.0181** (0.0084)	-0.0673*** (0.0192)	0.0199 (0.0240)
Cognitive Ability	0.0363* (0.0187)	0.0365** (0.0143)	0.0658* (0.0384)	-0.0385 (0.0428)
Financial Literacy	0.0328 (0.0250)	-0.0168 (0.0204)	-0.1403*** (0.0401)	0.0582 (0.0465)
HH Bargaining Power	0.0053 (0.0113)	-0.0116 (0.0090)	-0.0778*** (0.0236)	0.0792*** (0.0257)
Distance to Bank	-0.0265 (0.0207)	-0.0115 (0.0165)	-0.0365 (0.0536)	0.0643 (0.0615)
Exist. Savings Account	0.0998** (0.0444)	0.0296 (0.0363)	-0.0641 (0.0860)	0.1694* (0.0890)
Donates to Charity	0.0221 (0.0424)	0.0422 (0.0332)	0.0802 (0.0888)	0.2122** (0.0896)
#Emergencies last yr	-0.0161 (0.0277)	0.0005 (0.0213)	0.0558 (0.0601)	-0.0855 (0.0677)
#Emergencies since baseline		-0.0033 (0.0182)	0.1156* (0.0687)	-0.0414 (0.0658)
Marketer FE	YES	YES	YES	YES
Mean Dep. Variable	0.2687	0.1468	0.5463	0.4630
Observations	402	402	108	108

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Entries in the table represent the marginal coefficients of the probit regressions.

for all time-inconsistent agents with perceived temptation values higher than the median. This is in line with the proposed explanation: Sophisticated hyperbolic discounters either do not select into the product (if an effective penalty would be prohibitively high), or they choose a contract which is incentive-compatible for their preferences (adjusted through the size of the weekly instalment, or the size of the penalty).

Columns (3) and (4) restrict the analysis to the sample of clients who took up the ASA product, and should be interpreted as correlational evidence only: The regressions condition on an endogenous variable, and are likely subject to sample selection bias. In predicting default occurrence (column (3)), the marginal coefficient on present bias has quadrupled, and kept its significance. Interestingly, the link between present bias (as proxied by observed time-inconsistency) and default seems to be much stronger than the link between present bias and take-up. This is consistent with the theoretical intuition that an agent's take-up decision should be driven by *perceived* time-inconsistency, as proxied by the sophistication measure. In contrast, once the agent has adopted the contract, actual time-inconsistency will determine the success of the contract (in addition to a sophistication effect). The temptation measure now has strong predictive power on its own, even when not interacted with present bias. Expanding on the discussion from Section 3.3, individuals who report being positively tempted but do not exhibit hyperbolic preference reversals in MPLs could be either one of two things: a) they are time-inconsistent, but incorrectly classified as time-consistent in MPL questions, or b) they are subjectively feeling tempted but behaving time-consistently, possibly due to the exercise of costly self-control. In both cases, higher awareness of temptation will prompt individuals to choose more manageable (incentive-compatible) contracts – either through higher penalties or through lower weekly deposits (conditional on income). The data confirm that ASA clients with higher perceived temptation are indeed more likely to choose a penalty strictly above the minimum. However, due to lack of variation in penalties, this relationship is not significant.

Moving on to the pre-order (repeat take-up) decision, the coefficients on present bias (-0.49^{**}) and sophistication (0.24^{**}) are large and significant. The aggregate coefficient for a present-biased individual with the median value of perceived temptation is approximately zero. This has a convenient interpretation: Relatively naive hyperbolic discounters (those with below-median reported temptation) are unlikely to take up the ASA product again. From the previous analysis, there is a high chance that these individuals defaulted on their contract, and at the same time had not anticipated the default risk. These clients have 'burnt their fingers'. The result

that such individuals do not take up the product again is encouraging, insofar as it suggests learning about their preferences. The reverse holds true for present-biased individuals with above-median reported temptation (sophisticated hyperbolics): The aggregate coefficient on their time preferences is positive, suggesting a higher likelihood to pre-order ASA for a second round. This is consistent with the conjecture that sophisticated hyperbolic discounters choose the ‘right’ contract, which is incentive-compatible with their true preferences, and optimal in expectation. However, it does not imply a one-to-one mapping from successful ASA completion to the decision to pre-order: A sophisticated client who chose a contract that was optimal in expectation, but then rationally defaulted following a shock, might well decide to take up the product again. This is supported by an imperfect mapping from account status to pre-order decision in the data (see Table 3.4).

A number of other factors can help in explaining the observed default rates. The most obvious candidate - the occurrence of shocks during the savings period - finds some support in the take-up sample (‘emergencies since baseline’, column (3)). Shock occurrence was estimated by asking for common income or consumption emergencies during the endline survey.³⁰ The positive correlation of defaults with shocks, in combination with the fact that 45 percent of clients completed their ASA contract successfully, suggests that a significant portion of the take-up sample likely did choose a contract which was optimal for them in expectation. Clients without shocks could complete their plan successfully, while those with shocks rationally defaulted. The theoretical prediction that shock realisation should be irrelevant to the pre-order decision (as it does not affect contract optimality in expectation) is supported by the data (see column (4) of Table 3.6). Other factors predicting default include financial literacy (-), household bargaining power (-), risk aversion (-) and cognitive ability (+). Financial literacy is perhaps the least surprising: Individuals with poor numeracy skills tend to do worse at managing their household finances, and may fail to allocate a portion of the household budget to regular ASA deposits. The positive significance of cognitive ability may be partially explained by the predictive power that cognitive ability has for take-up of the ASA product, as those who struggle to understand the product’s rules don’t select in. Within the take-up sample, cognitive ability may reflect predictability and rationalisability of behaviour: An individual with high cognitive skills is more likely to realise when

³⁰There is a risk that clients who defaulted had a stronger incentive to report shocks, in order to preserve their self-image or reputation. However, the endline survey was framed as coming from a research organisation, with no direct link to the bank. The survey was identical across the sample, and made no reference to ASA or Gihandom. Note that attrition in the endline survey was compensated by imputing the median shock value for those who did not participate.

a contract is no longer optimal (even for time-inconsistent reasons), and default on it, rather than display noisy or unsystematic behaviour.

A similar puzzle arises from the negative correlation of risk aversion with default (but not with take-up). An explanation requires a closer look at how the measure was obtained: The risk aversion measure is a score in $[1, 6]$, indicating which lottery individuals chose from a set of lottery options with increasing expected value and increasing variance (see Figure C.3). If preferences are characterised by reference dependence (with the no-risk lottery A as a reference point) and loss aversion, then the choice of a safe lottery would measure loss aversion rather than risk aversion. A high degree of loss aversion can be associated with a lower likelihood to default on the ASA product (to avoid loss of the penalty). Finally, it is interesting to note that household bargaining power is unrelated to take-up of the regular-installment product ASA, but strongly related to defaulting on it (within the take-up sample). A possible explanation is that intra-household conflicts played no role in individual's motives to take up the product – but that, much like a shock, individuals soon learnt that it caused household conflicts to try and put aside a portion of the household budget every week, beyond the reach of other household members. This can be interpreted as a learning process in adopting a new savings technology (loan repayment is similar in structure, but may be easier to agree on in a household because of the higher penalties involved). Consequently, clients with low bargaining power may have yielded to these disagreements, and defaulted on their contracts. The large positive association of household bargaining power with the pre-order decision provides further support for a learning explanation: Once individuals had learnt about the difficulties of regularly diverting a share of the household budget, only those with sufficient autonomy in their household chose to take up the product again.

3.5.3.3 Heterogeneous Treatment Effects

Table 3.7 examines treatment effect heterogeneity across a number of dimensions of interest. The regression set-up is identical to that in column (1) of Table 3.2: The change in savings held at the partner bank is regressed on indicators for assignment to the treatment groups. In addition, the treatment indicator for the Regular Saver group is interacted with variables which have been shown to predict take-up or default, or which are of interest in themselves. Interaction variables include present-biased preferences, the self-reported sophistication measure, having an existing savings account, household bargaining power, and household income.

3.5. RESULTS

TABLE 3.7: HETEROGENEOUS TREATMENT EFFECTS: CHANGE IN BANK SAVINGS

	(1)	(2)	(3)	(4)	(5)
Regular Saver (R)	707.1748*** (257.1454)	481.4925*** (134.9550)	287.4050*** (63.6388)	455.4759*** (133.2911)	351.9444*** (114.1216)
Withdrawal Rest. (W)	129.4941*** (40.6258)	146.3426*** (39.6980)	147.2989*** (40.7323)	154.5030*** (42.6279)	148.8766*** (41.4302)
R * Present bias	-58.8396 (543.3738)				
Present bias	57.6709 (83.9149)				
R*Soph. Present Bias (R*Pres.Bias*Temptation)	20.7603 (81.9725)				
Soph. Present Bias	-6.0420 (18.1875)				
R * Temptation	-48.8363 (62.2914)				
Temptation	-4.9396 (9.2861)				
R*High Cognitive Ability		261.2952 (291.0735)			
High Cognitive Ability		-30.7091 (36.4527)			
R * Existing SA			621.5350** (269.7281)		
Existing SA			74.9955* (41.3577)		
R * Income				44.6249 (56.0615)	
Income				7.6109 (5.4829)	
R * HH power					88.1604* (50.5904)
HH power					6.9664 (10.3868)
Constant	34.9483 (26.7103)	41.1479** (17.9218)	-7.0002 (18.9202)	2.7796 (17.7092)	8.8140 (29.9048)
R ²	0.02	0.02	0.04	0.02	0.03
Observations	720	748	748	745	748

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Heterogeneity in treatment effects is most pronounced for existing savings account holders. Existing savings account holders increased their savings balances by 622 pesos more than those without an existing account after being offered the Regular Saver product. Put differently, the intent-to-treat effect of the Regular Saver

product was 909 pesos for existing savings account holders, and only 287 pesos for those without existing accounts. This seems particularly surprising in light of the fact that, in *absence* of the Regular Saver treatment, existing account holders saved only 75 pesos more than those without existing accounts. The evidence suggests that existing account holders were not necessarily active savers before the intervention, but felt strongly motivated by the Regular Saver treatment. A possible explanation relates to mistrust and negative preconceptions towards banks, which were common in the population.³¹ Existing account holders were more likely to be familiar with basic bank transactions, and more trusting of the banking system as a whole.

It is worth noting that treatment effects appear to be relatively uniform across measures of present bias and sophistication. Theory predicts that a present-biased agent with a low degree of sophistication is likely to select into a commitment contract that is too weak to be effective given his preferences, resulting in default soon after take-up. After taking into account the default penalty, savings with the commitment product should be weakly smaller than savings without the commitment product. The positive association between (naive) present bias and default is supported empirically by the regressions in Table 3.6. The negative effect of present-biased preferences on savings should be mitigated or even reversed with increasing levels of sophistication: The agent is more likely to choose an incentive-compatible contract, increasing the chances of successfully reaching his savings goal. The sign of the aggregate coefficient on sophisticated present-biased preferences relative to time-consistent behaviour is theoretically ambiguous, as illustrated by O'Donoghue and Rabin (1999). Column (1) of Table 3.7 shows that all estimates of treatment effects with respect to measures of present bias and sophistication are small and insignificant. A likely reason are the composition effects inherent in ITT estimates: Individuals with sophisticated time-inconsistent preferences were much less likely to select into the product (see Table 3.5). Thus, a lower percentage of sophisticated agents were 'treated'. Theoretical arguments in Chapter 2 confirm that a sophisticated agent may choose to stay away from commitment, if the effective default penalty is prohibitively high in the presence of shocks. If his preferences are such that he cannot achieve his savings goal in autarky, he will choose not to save.

Columns (3) and (4) show that the estimated treatment effect is relatively uniform across household income level, as well as an above-median indicator for cog-

³¹It was a common belief that banks were "not for poor people". In addition, some individuals believed that savings deposited at a bank would likely be lost if the bank became insolvent. Deposit insurance does exist in the Philippines, but may be associated with years of waiting time. See e.g., Dupas et al. (2012) on trust-related challenges in banking the poor.

3.6. ALTERNATIVE EXPLANATIONS FOR DEFAULT

nitive ability. Column (6) suggests that successfully maintaining the Regular Saver product ASA was facilitated by having a certain degree of household bargaining power: Individuals who report to be the primary decisionmaker in many aspects of household budgeting respond to the Regular Saver treatment with larger savings increases than those with low bargaining power. Using a score $[0, 5]$, each one-point increase in bargaining power corresponds to an increase of 88 pesos in savings after being offered the Regular Saver product. This effect is consistent with the incidence of household conflicts caused by the weekly ASA instalments (see Section 3.5.3.2). Note that individuals with high bargaining power did not save more absent treatment – it is the interaction of sufficient bargaining power and the Regular Saver treatment which helped individuals to save. This explanation differs markedly from a ‘safekeeping’ motive: If individuals took up ASA to mitigate household bargaining issues, we would expect the interaction coefficient to be negative (as those with low power would benefit more from treatment).

3.6 Alternative Explanations for Default

Previous sections have focused on partially sophisticated hyperbolic preferences in explaining why a majority of individuals who choose to adopt a regular-instalment commitment product will default soon after opening their accounts. This section will consider alternative explanations: Income optimism, aggregate shocks, and limited attention.

3.6.1 Income Optimism

As suggested by Browning and Tobacman (2007), the consumption behaviour of someone who is overoptimistic about his future income distribution cannot be distinguished from someone who is impatient – both will overconsume in the present. Overoptimistic beliefs about future income could explain the observed measure of time-inconsistency (from MPL questions): If individuals expect their future income to be higher than their current income, they may select the smaller, sooner reward when presented with the ‘now vs. 1 month’ frame, but choose the larger, later reward when presented with the ‘1 month vs. 2 months’ frame. As a result, they would be falsely classified as present-biased. Income optimism could further explain default incidence: If people were overoptimistic about their income when they adopted the Regular Saver product, and realised this upon starting their savings plan, default may have become an optimal response.

3.6. ALTERNATIVE EXPLANATIONS FOR DEFAULT

TABLE 3.8: INCOME OPTIMISM

	Not Present-Biased	Present-Biased	All	T-stat P-value
Prediction Gap (growth)	3.2904 (0.6976)	3.6777 (1.2298)	3.3570 (0.6146)	0.81
Prediction Gap (level)	1.2698 (1.1704)	-2.2231 (2.6292)	0.6686 (1.0698)	0.22
Observations	582	121	703	

	No Take-Up	Take-Up	All	T-stat P-value
Prediction Gap (growth)	1.738007 (0.9255)	5.043011 (1.8137)	2.582418 (0.8325)	0.08
Prediction Gap (level)	1.140221 (1.8676)	-2.569892 (2.6941)	0.1923077 (1.552)	0.30
Observations	271	93	364	

	Successful	Default	All	T-stat P-value
Prediction Gap (growth)	4.227273 (2.6477)	5.77551 (2.5106)	5.043011 (1.8137)	0.67
Prediction Gap (level)	-5.318182 (3.7117)	-0.1020408 (3.8799)	-2.569892 (2.6941)	0.34
Observations	44	49	93	

Standard deviations in parentheses. All numbers are group averages.

Using data on predicted and realised incomes, I construct a measure which plausibly captures income optimism for groups. It is impossible to identify optimism on an individual level – an individual who reports to have lower income than predicted may either experience a bad draw from a correct income distribution (the ‘bad luck’ explanation), or he may have systematically biased beliefs about his income distribution (‘optimism’). However, the law of large numbers implies that individuals should correctly predict their income *on average* if their beliefs about income are unbiased. On the other hand, if the present bias measure captures income optimism rather than time-inconsistency, then individuals classified as present-biased should have higher predicted-minus-realised income gaps than those classified as not present-biased. Further, if defaults were caused by individuals systematically misjudging their future income, then defaulting clients should have higher prediction gaps than those who successfully completed their contract.

Table 3.8 presents group averages of prediction gaps across three dimensions: The observed measure of present bias, take-up of the Regular Saver product ASA,

3.6. ALTERNATIVE EXPLANATIONS FOR DEFAULT

and default on ASA. Prediction gaps are measured as follows: During the baseline survey in September and October 2012, individuals were asked to predict their average weekly household income for each month from October 2012 to March 2013. To make this task easier, individuals chose one of 31 income brackets, numbered from 1 for '0-50 pesos per week' to 31 for 'more than 10,000 pesos per week'. Six months later, in late March and April 2013, this exercise was repeated during the endline survey, except that individuals now stated their realised weekly income for the same time period. Two measures of optimism (or bad luck) are obtained: *Prediction Gap (growth)_i* is the difference between predicted income growth and realised income growth, where growth is measured as $Growth_i = \sum_{m=Nov}^{Mar} (bracket_m - bracket_{October})$. In other words, income growth is proxied by the sum of deviations from October income, in units of income brackets. This approach is conservative, in the sense that it is robust to individuals using different income benchmarks for their October income in baseline and endline survey.³² An alternative measure of optimism is *Prediction Gap (level)_i*, obtained by the simple difference between predicted and realised income levels (summed), $Prediction\ Gap\ (level)_i = \sum_{m=Oct}^{Mar} (bracket_m^{pred} - bracket_m^{real})$. Consistent with noise in benchmark income levels, *Prediction Gap (level)_i* exhibits more variation than *Prediction Gap (growth)_i*. Note that these measures cannot be included as covariates in take-up or default regressions – both because they are not meaningful on an individual level, and because they use data from the endline survey, and may thus not be orthogonal to treatment.

The sample for Table 3.8 are those individuals who participated in both the baseline and endline survey. The average prediction gap for income growth across the sample was 3.36 brackets, suggesting that moderate income optimism may be common. However, the average prediction gap is not higher for individuals classified as present-biased – if anything, the level measure suggests they may have been more pessimistic. In contrast, the average prediction gap is significantly higher for individuals who adopted the ASA product compared to those who did not, suggesting that those entering commitment contracts may have been more optimistic about their future income. Finally, individuals who defaulted on ASA did not report significantly higher prediction gaps for income growth than did clients who

³²For instance, individuals might have referred to the household income of their core household in the baseline survey, and their extended household in the endline survey, or vice versa. Clear definitions of what constitutes a household were provided, but some grey areas were unavoidable (e.g., where families lived with uncles or cousins, and shared a common budget for food, but not for other household expenses).

3.6. ALTERNATIVE EXPLANATIONS FOR DEFAULT

successfully completed their contract. However, it is worth noting that the level measure points to a possible pessimism of successful clients.

Summing up, there is mixed evidence that those who adopted ASA were optimistic about the growth of their income, relative to those who rejected the offer. The evidence does not suggest a connection between optimism and the observed measure of present bias. In addition, income optimism alone cannot explain why individuals demand commitment. Further, it does not provide a rationale for the observed link with the sophistication measure (which is based on self-reported temptation).

Similar arguments apply for optimism regarding the shock frequency λ (as defined in Chapter 2): For instance, individuals could have heterogeneous shock frequencies λ_i , where shocks may refer to income shocks, consumption emergencies, and more generally the risk that saving may no longer be optimal. With rational expectations about λ_i , individuals with high shock frequencies are *ceteris paribus* less likely to select into commitment. However, if individuals have biased beliefs about λ_i (such as the belief that one's shock frequency rate corresponds to the average shock frequency in the population), then the consequence of a commitment contract may be a bulk of defaults soon after opening (as those individuals with the highest λ_i are likely to drop out first). Therefore, biased beliefs about the shock frequency provide another potential explanation for default occurrence. Its limitation is similar to that of income optimism: Biased beliefs about λ alone do not predict a demand for commitment. Neither do they explain a correlation with measures of present-bias or sophistication.

Less parsimonious explanations may involve a combination of different factors, such as fully sophisticated hyperbolic preferences in combination with income optimism. This combination may predict both a demand for commitment and subsequent default. However, it fails to explain why measures of sophistication are *negatively* associated with take-up and default. In this sense, partial sophistication provides a parsimonious explanation that is consistent with the evidence.

3.6.2 Aggregate Shocks

Idiosyncratic and independent shocks are unlikely to cause the default timing pattern apparent in Figure 3.1. However, if an aggregate shock hit the sample population around the time of account opening, this may help to explain why 55 percent of clients defaulted shortly after adopting the product. The Philippines is a well-known area for earthquakes and tropical storms, and had recently been hit by tropical storm Washi (Philippine name 'Sendong') in December 2011, causing

3.6. ALTERNATIVE EXPLANATIONS FOR DEFAULT

1,268 casualties (more than half of them in Cagayan de Oro, a city 126km west of the study location).³³ The risk of such shocks was thus well-known at the time of marketing in September 2012, possibly affecting take-up rates. Indeed, tropical storm Bopha (Philippine name 'Pablo') hit the Mindanao region between December 2 and December 9, 2012. As opposed to storm Washi, storm Bopha did not cause flash flooding, and the main effect on the study location was a six-day power outage. While this may have affected on large businesses, power outages of several hours each day were common in the study area even before the storm, and provisions against power outages were widespread. Because of its limited effect on the area, storm Bopha was not locally classified as a natural disaster (which would have invoked both ASA's and Gihandom's emergency provisions). In the endline survey, 20.5 percent of the sample population reported some damage to their house or crops, with a median damage value of 1400 pesos (U.S. \$33, conditional on non-zero damage). Within the sample of defaulting ASA clients, the percentage affected by the storm was 20.4. Asked whether they suffered reductions in income because of the power outages, only 3 out of 732 endline survey respondents answered in the affirmative.

While some negative effects of the storm cannot be ruled out, the timing of the storm does not match the timing of the defaults: The ASA accounts were opened between 20 September and 28 October. Out of 63 defaults, 35 made no further deposit after their opening balance, resulting in contract default upon the third missed deposit, three weeks after opening.³⁴ An additional 15 clients made one or two deposits after opening (see Figure 3.1 for the distribution of transactions). By the time of the storm in early December, most of the contract defaults had already occurred.

3.6.3 Limited Attention

An intuitive explanation for default suggests that clients may have simply forgotten to make their weekly deposits. Limited attention models such as that of Banerjee and Mullainathan (2008) suggest that attention is a scarce resource, which needs to be divided between home and work in order to catch emerging problems before they cause damage. In their model, the amount of damage an individual suffers from problems occurring at home or at work (such as a child's sickness, or running out of stock for one's business) is a function of the attention which the individual

³³Statistics from the Philippine National Disaster Risk Reduction and Management Council (NDR-RMC).

³⁴85 percent of ASA clients opted for weekly deposits, 15 percent opted for bi-weekly deposits.

invests into home life and workplace. Given the relatively low stakes of the Regular Saver account (with default penalties roughly equivalent to a day's household income), it would be understandable if individuals prioritised their attention on their home and work lives, rather than on their bank accounts. However, this explanation predicts that individuals would not take up the Regular Saver product in the first place: During the marketing stage, ASA was clearly presented as attention-intensive: Clients were presented with an explicit savings plan including due dates for each week, and given the instruction to physically deposit their instalments at the bank. Most respondents received their income in cash, and bank transfers were uncommon. In the Banerjee and Mullainathan (2008) model, not investing attention in one aspect of one's life incurred a risk that a costly problem would go unnoticed. In contrast, not investing attention in the ASA product (by adhering to the weekly schedule) resulted in *certain* default. As a result, if the returns to investing attention at work or at home exceeded the returns to investing attention in the ASA schedule, then individuals should not have adopted the product. The data suggests that this indeed reduced take-up: Among the clients who were assigned to the Regular Saver treatment but chose not to adopt the product, being "too busy to go to the bank" was a common reason for rejecting. Among those clients who accepted the offer, "distance to the bank branch" does not predict default (as measured by GPS coordinates, see Table 3.6).

3.7 Conclusion

Commitment devices are receiving increasing attention both in the academic literature and in the public eye, and are generally portrayed as a promising way to overcome intrapersonal conflict. The *SafeSave* case study presented in Chapter 1 of this dissertation illustrates the positive effect of a newly introduced commitment savings product in Bangladesh on individuals' average savings contributions.

Using the example of a commitment savings product in the Philippines, I present evidence that people may fail at choosing commitment contracts which are suitable for their preferences. I argue that an individual's ability to correctly choose a welfare-improving commitment contract depends on his degree of sophistication, i.e., on the individual's awareness of the nature of his time-inconsistency. I observe that a majority of individuals who take up a commitment product choose very low stakes for this commitment, and then default on it. Both take-up and default decisions are systematically linked to low measures of sophisticated time-inconsistency, suggesting that imperfect (or partial) levels of sophistication are widespread. By

3.7. CONCLUSION

the nature of commitment, a tendency to choose unsuitable contracts is costly. Implications reach beyond commitment savings, and may extend to rich country applications such as gym contracts (as shown by DellaVigna and Malmendier (2006)), diet clubs, and long-term pension savings plans.

From a policy perspective, the presented comparison between a (harder) regular instalment commitment and a (softer) withdrawal-restriction commitment may suggest a possible trade-off between efficacy and risk of offering commitment products: Offering stronger commitments with more pressure may provide greater benefits on *average* – as observed by a fourfold effect of the ASA product on average bank savings, and an increased likelihood of purchasing one's savings goal using own funds. However, offering stronger commitments may also involve an increased risk of adverse effects on welfare for partially sophisticated agents. In the present study, a 'softer' commitment contract is exemplified by the date-based Gi-handom account: At the end of the savings period, individuals simply received their savings back, and 'undesirable' behaviour went unpenalized. While the absence of penalties may keep welfare risks to a minimum, beneficial effects of the product may be similarly limited: Offering the account had a comparatively small effect on average savings, and an even smaller effect on the median.

The welfare risks suggested in this study are not singular – a closer look at heterogeneity behind average treatment effects in the literature may reveal that adverse effects of commitment products are widespread. As a consequence, research on new commitment products should carefully consider possible risks to welfare, with particular view to partially sophisticated time-inconsistent agents.

Bibliography

Ali, S Nageeb, “Learning Self-Control,” *The Quarterly Journal of Economics*, 2011, 126(2), 857–893.

Amador, Manuel, Iván Werning, and George-Marios Angeletos, “Commitment vs. Flexibility,” *Econometrica*, 2006, 74(2), 365–396.

Ambec, Stefan and Nicolas Treich, “ROSCAs as Financial Agreements to Cope with Self-Control Problems,” *Journal of Development Economics*, 2007, 82(1), 120–137.

Ameriks, John, Andrew Caplin, John Leahy, and Tom Tyler, “Measuring Self-Control Problems,” *The American Economic Review*, 2007, 97(3), 966–972.

Ananth, Bindu, Dean Karlan, and Sendhil Mullainathan, “Microentrepreneurs and Their Money: Three Anomalies,” Working Paper, Financial Access Initiative 2007.

Anderson, Siwan and Jean-Marie Baland, “The Economics of Roscas and Intra-household Resource Allocation,” *Quarterly Journal of Economics*, 2002, 117(3), 963–995.

Angrist, Joshua D. and Joern-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist’s Companion*, Princeton: Princeton University Press, 2009.

Ariely, Dan and Klaus Wertenbroch, “Procrastination, Deadlines, and Performance: Self-control by Precommitment,” *Psychological Science*, 2002, 13 (3), 219–224.

Ashraf, Nava, Dean Karlan, and Wesley Yin, “Deposit Collectors,” *The B.E. Journal of Economic Analysis & Policy*, 2006, 5(2).

—, —, and —, “Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines,” *Quarterly Journal of Economics*, 2006, 121(2), 635–672.

BIBLIOGRAPHY

- Banerjee, Abhijit and Sendhil Mullainathan**, "The Shape of Temptation: Implications for the Economic Lives of the Poor," Working Paper, National Bureau of Economic Research 2010.
- Banerjee, Abhijit V.**, "A Simple Model of Herding Behavior," *Quarterly Journal of Economics*, 1992, 107(3), 797–817.
- **and Sendhil Mullainathan**, "Limited Attention and Income Distribution," *American Economic Review*, 2008, 98(2), 489–493.
- Basu, Karna**, "Commitment Savings in Informal Banking Markets," *Journal of Development Economics*, 2014, 107, 97–111.
- Bauer, Michal, Julie Chytilova, and Jonathan Morduch**, "Behavioral Foundations of Microcredit: Experimental and Survey Evidence from Rural India," *American Economic Review*, 2012, 102(2), 1118–1139.
- Benartzi, Shlomo and Richard H. Thaler**, "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, 2004, 112(1), 164–187.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, 119(1), 249–275.
- **, Sendhil Mullainathan, and Eldar Shafir**, "A Behavioral Economics View of Poverty," *American Economic Review*, 2004, 94(1), 419–423.
- Beshears, John Leonard, James J. Choi, David Laibson, Brigitte C. Madrian, and Jung Sakong**, "Self Control and Liquidity: How to Design a Commitment Contract," Working Paper Series WR- 895-SSA, RAND November 2011.
- Besley, Timothy**, "Savings, Credit and Insurance," in Jere Behrman and T.N. Srinivasan, eds., *Handbook of Development Economics*, Vol. 3A, Elsevier, 1995.
- **and Anne Case**, "Diffusion as a Learning Process: Evidence from HYV Cotton," Papers 174, Princeton, Woodrow Wilson School - Development Studies 1994.
- **and —**, "Unnatural Experiments? Estimating the Incidence of Endogenous Policies," *Economic Journal*, 2000, 110(467), F672–F694.
- **, Stephen Coate, and Glenn Loury**, "The Economics of Rotating Savings and Credit Associations," *American Economic Review*, 1993, 83(4), 792–810.

BIBLIOGRAPHY

- Binswanger, Hans P**, "Attitudes Toward Risk: Experimental Measurement in Rural India," *American Journal of Agricultural Economics*, 1980, 62(3), 395–407.
- Browning, Martin and Jeremy Tobacman**, "Discounting and Optimism Equivalences," 2007.
- Brune, Lasse, Xavier Giné, Jessica Goldberg, and Dean Yang**, "Commitments to Save: A Field Experiment in Rural Malawi," Policy Research Working Paper Series 5748, World Bank 2011.
- Bryan, Gharad, Dean Karlan, and Scott Nelson**, "Commitment Devices," *Annual Review of Economics*, 2010, 2(1), 671–698.
- Collins, Daryl, Jonathan Morduch, Stuart Rutherford, and Orlanda Ruthven**, *Portfolios of the Poor - How the World's Poor Live on \$2 a Day*, Princeton: Princeton University Press, 2009.
- DellaVigna, Stefano and Ulrike Malmendier**, "Contract Design and Self-Control: Theory and Evidence," *The Quarterly Journal of Economics*, 2004, 119(2), 353–402.
- and —, "Paying Not To Go To The Gym," *The American Economic Review*, 2006, 96(3), 694–719.
- Duflo, Esther**, "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment," *American Economic Review*, 2001, 91(4), 795–813.
- , **Michael Kremer, and Jonathan Robinson**, "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya," *The American Economic Review*, 2011, 101(6), 2350–90.
- , **Rachel Glennerster, and Michael Kremer**, "Using Randomization in Development Economics Research: A Toolkit," *Handbook of Development Economics*, 2007, 4, 3895–3962.
- Dupas, Pascaline and Jonathan Robinson**, "Why Don't the Poor Save More? Evidence from Health Savings Experiments," Working Paper 17255, National Bureau of Economic Research 2011.
- and —, "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya," *American Economic Journal: Applied Economics*, 2013, 5(1), 163–192.

BIBLIOGRAPHY

- , **Sarah Green, Anthony Keats, and Jonathan Robinson**, “Challenges in Banking the Rural Poor: Evidence from Kenya’s Western Province,” Working Paper 17851, National Bureau of Economic Research 2012.
- Fischer, Greg and Maitreesh Ghatak**, “Repayment Frequency in Microfinance Contracts with Present-Biased Borrowers,” Economic Organisation and Public Policy Discussion Paper 21, London School of Economics 2010.
- Foster, Andrew D. and Mark R. Rosenzweig**, “Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture,” *Journal of Political Economy*, 1995, 103(6), 1176–1209.
- Frederick, Shane, George Loewenstein, and Ted O’donoghue**, “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, 2002, 40(2), 351–401.
- Fudenberg, Drew and David K Levine**, “A Dual-Self Model of Impulse Control,” *The American Economic Review*, 2006, 96(5), 1449–1476.
- Giné, Xavier, Dean Karlan, and Jonathan Zinman**, “Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation,” *American Economic Journal: Applied Economics*, 2010, 2(4), 213–235.
- Giné, Xavier, Jessica Goldberg, Dan Silverman, and Dean Yang**, “Revising Commitments: Field Evidence on the Adjustment of Prior Choices,” Working Paper, National Bureau of Economic Research 2012.
- Gugerty, Mary Kay**, “You Can’t Save Alone: Commitment in Rotating Savings and Credit Associations in Kenya,” *Economic Development and Cultural Change*, 2007, 55(2), 251–282.
- Gul, Faruk and Wolfgang Pesendorfer**, “Temptation and Self-Control,” *Econometrica*, 2001, 69(6), 1403–1435.
- and —, “Self-Control and the Theory of Consumption,” *Econometrica*, 2004, 72(1), 119–158.
- Hofmann, Anett**, “Just A Few Cents Each Day: Can Fixed Regular Deposits Overcome Savings Constraints?,” EOPP Discussion Papers 51, London School of Economics 2014.
- Imbens, Guido W and Joshua D Angrist**, “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, 1994, 62(2), 467–475.

BIBLIOGRAPHY

- Johnson, Simon, John McMillan, and Christopher Woodruff**, "Property Rights and Finance," *American Economic Review*, 2002, 92(5), 1335–1356.
- Karlan, Dean and Jonathan Zinman**, "Price and Control Elasticities of Demand for Savings," 2013.
- , **Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman**, "Getting to the Top of Mind: How Reminders Increase Saving," Working Paper 16205, National Bureau of Economic Research July 2010.
- Laibson, David**, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, 112(2), 443–478.
- Loewenstein, George**, "Anticipation and the Valuation of Delayed Consumption," *The Economic Journal*, 1987, 97(387), 666–684.
- Morduch, Jonathan**, "Between the State and the Market: Can Informal Insurance Patch the Safety Net?," *World Bank Research Observer*, 1999, 14(2), 187–207.
- O'Donoghue, Ted and Matthew Rabin**, "Doing It Now or Later," *American Economic Review*, 1999, 89(1), 103–124.
- Rutherford, Stuart**, "The Poor and Their Money," FDRP Working Paper 3/99, University of Manchester, Institute for Development Policy and Management 1999.
- Sayman, Serdar and Ayse Öncüler**, "An Investigation of Time Inconsistency," *Management Science*, 2009, 55(3), 470–482.
- Shefrin, Hersh M and Richard H Thaler**, "The Behavioral Life-Cycle Hypothesis," *Economic Inquiry*, 1988, 26(4), 609–643.
- Sinn, Miriam**, "Risk and Time-Inconsistency: Evidence from a Field Experiment in West Bengal," 2012.
- Strotz, Robert Henry**, "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 1955, 23(3), 165–180.
- Takeuchi, Kan**, "Non-Parametric Test of Time Consistency: Present Bias and Future Bias," *Games and Economic Behavior*, 2011, 71(2), 456–478.
- Tarozzi, Alessandro and Aprajit Mahajan**, "Time Inconsistency, Expectations and Technology Adoption: The Case of Insecticide Treated Nets," ERID Working Paper 105, Duke University 2011.

BIBLIOGRAPHY

Thaler, Richard, "Mental Accounting and Consumer Choice," *Marketing Science*, 1985, 4(3), 199–214.

Appendix A

Appendix to Chapter 1: Explaining Adoption Timing

Section 1.3.2 has pointed out that area of residence strongly influenced when clients took up the product. However, early adoption areas do not seem richer than others and are not closer to the bank branch. One possible factor at work may be a herding behaviour story in the style of Banerjee (1992): Consider homogeneous individuals who decide in a sequential fashion whether or not to adopt LTS. The sequential decision process could be due to reasons exogenous to savings, e.g. because families need to find time to sit down at the kitchen table and solve the optimization problem whether or not they want to commit to a 5 year contract.¹ Assume there is a network structure within each block or area (but not between areas): Communication between neighbours is good enough that agents know whether someone in their block has chosen to i) adopt, ii) not adopt, iii) is yet undecided. Everyone has a private signal whether the LTS product is good or bad for them: “Good” means it will be manageable with their budget, and will improve their situation. “Bad” means it will prove to be unmanageable, taking away slack from their budget, causing transaction cost and withdrawal fees, and constrain them until they decide to cancel it without reaping any returns. Assume further that the true return to the product is the same for everyone and positive, and that a large majority of individuals (say, 90 percent) receive a good signal. Finally, add a tie-breaking rule that someone who is indifferent will not bother to go to the bank to adopt the product.

¹See also the work on limited attention spans by Banerjee and Mullainathan (2008), which advocates the view that the poor may miss or delay investment opportunities because of time and attention they spend on domestic problems.

Suppose the first two individuals who decide receive a good signal. The first individual will see only his signal and adopt. The second individual will observe the first's choice, conclude that his signal must have been positive, and adopt as well. If someone with a bad signal follows, he will see the first two's choice, discard his own signal, and adopt. In a similar fashion, the entire block will adopt the product regardless of their signals.

Suppose now the first individual has a bad signal, and will not adopt. The second individual will observe this, conclude that the first had a bad signal and weigh this against his own signal. Even if he had a good signal himself, he will be indifferent, and will therefore not adopt. The third person sees that the first two did not adopt, and will do the same. Likewise, no one in the block adopts.

The consequence is that even with completely homogeneous individuals, a profitable innovation may spread very quickly in one area, and not at all or with considerable delay in another area, depending on the signals of the first one or two individuals who find the time to optimise. What would this explanation imply for our results on the treatment effect? If random factors determine who chooses first, then early adopter areas need not be different at all from late adopter areas in their savings behaviour, and what we are observing is an actual treatment effect. This explanation is supported empirically by the fact that early adopters did not choose higher monthly contributions than late adopters did, as can be seen from the lack of correlation between entry month and monthly deposit in Table 1.5.

Of course, one may question that the adoption decision is a sequential process in a random order. Instead, we may believe that the order in which individuals optimise may result from heterogeneous savings preferences. This leads us to a technology adoption scenario as presented in Besley and Case (1994): Information (about manageability) is a public good, and those with higher stakes in the public good have bigger incentives to adopt the product first to learn about it. For example, those who have a higher income may have higher potential gains from the product, and may thus adopt early. Areas who have more of such "pioneers" will then adopt the product faster, even if the area is not richer on average. Note that we do not need to assume different degrees of risk-aversion for this explanation. What are the implications for savings? The group of pioneers that causes an area to adopt earlier than another area may be small, so that early adoption areas do not necessarily differ much on average from late adoption areas, keeping the selection bias small. If, however, the group of pioneers is large, and pioneers follow different savings trends (for instance, because their income grows faster) than the

average client, then there is a positive selection bias, and we are overestimating the treatment effect.

Furthermore, area does not predict early adoption perfectly. If it is generally the case that richer, or more educated clients in all areas tend to adopt early, and these groups follow a higher savings trend, then there will be a positive selection bias. Another classic issue in the technology adoption literature is risk aversion: With higher risk aversion, the value of learning is higher. Intuitively, people who are more risk averse will adopt later, waiting deliberately for the information revealed by the actions and experience of others. Regarding LTS savings, risk averse clients are likely to commit to higher amounts and longer contracts the later they enter. Hence, if risk aversion is what drives adoption timing, we may expect monthly contributions and contract length to be *negatively* correlated with early adoption (which the data confirms for contract length but not for monthly deposits). However, we cannot make any statement about the sign of the selection bias: If the most risk averse clients are in the control group (i.e. they are late adopters), then what matters are their savings contributions *before* adopting LTS. Without LTS, we may expect risk averse clients to save more in *levels* than other clients, e.g. as an insurance to income uncertainty. But we cannot conclude anything on the *change* in their savings during the treatment period.

Finally, what insights can be derived from the model outlined in Section 2? As mentioned in the technology adoption scenario, agents who adopt early may be those who have the highest potential gains from LTS, and consequently the largest interest in learning whether or not it is manageable. The model identifies these agents as those with $\beta \in [\hat{\beta}_B, \hat{\beta})$, respectively $p \in (p_{max}, p_{max}^B]$. In other words, theory suggests that early adopters may be those who can reach their desired savings goal p with LTS but not without it. Late adopters could have reached their p in autarky, but at less smooth savings contributions – thus their potential return to LTS is positive, but their stake in learning about its manageability is lower. The main implication of the model are heterogeneous treatment effects: If the variation of β in the population is low and that of p is high, then early adopters are the clients who aim for higher p 's, and who need LTS to achieve them. If p does not vary in the population, but β varies, then everyone aims for the same savings goal, but those with lower β need LTS to achieve it. In both cases, the treatment effect is higher for early adopters than it would be for late adopters. Finally, β and p could be constant across the population, with variation consisting only in the degree of relatives' claims, μ . Since p_{max} increases in μ (for given β), clients with demanding relatives (low μ) may need LTS to achieve a p that clients with higher μ can achieve

without LTS. So clients with low μ adopt early, and again the treatment effect is bigger for early adopters.

What about the selection bias, $E[S_{it}^{NT} - S_{i,t-1}^{NT} \mid T] - E[S_{it}^{NT} - S_{i,t-1}^{NT} \mid NT]$? Absent treatment, early adopters would have saved zero as they could not achieve their desirable p , hence the change in their savings is zero. Late adopters may have been saving for p at home during the treatment period. Whether or not this contribution increased or decreased from t to $t + 1$ depends on their allocation of the savings burden $p - 1$ across periods. In the language of the model, if LTS was introduced in a period 1, and $\beta, p, u(c_t)$ were such that $s_1 > \bar{s}$ in autarky, then their savings would decrease from period 1 to 2, causing a positive selection bias. If $s_1 < \bar{s}$, the selection bias would be negative. Finally, one may argue that late adopters anticipated saving with LTS in the following period, and thus did not save at home, resulting in a selection bias of zero.

This subsection has discussed a small number of factors that may have determined adoption timing – many others remain in the dark. The herding model is consistent with adoption in some areas but not in others, and is further supported by the lack of correlation between monthly contribution and entry timing as well as the non-existence of pre-existing time trends. Classical technology adoption mechanisms are supported by the fact that late adopters choose longer contracts. Finally, timing may be explained by different degrees of hyperbolic discounting or stronger claims from relatives. While this preliminary analysis does not have survey data on client's time preferences, studies like Ashraf et al. (2006b) show that individuals with a lower β and a higher degree of sophistication will be most interested in a commitment savings product.

Appendix B

Appendix to Chapter 2: Mathematical Proofs

Lemma 1. (a) The period 2 agent is willing to save for the nondivisible and transfer $s_2 = p - 1$ if s_1 is bigger than some threshold value, $s_1 \geq s_{min}$. (b) s_{min} is strictly decreasing in the time-inconsistency parameter β . (c) The effect of the shock frequency λ on s_{min} is ambiguous.

Proof. (a) The period 2 agent is willing to save $s_2 = p - 1$ if s_1 is such that

$$u(1 + s_1 - (p - 1)) + \beta[(1 - \lambda)u(b) + \lambda u(p - 1)] \geq u(1 + s_1 - s_2^{No}) + \beta E[u(y_3 + s_2^{No})]$$

It is sufficient to prove that once s_1 is high enough to satisfy the inequality above (i.e., buying the good is optimal), the inequality will also be satisfied for all higher values of s_1 . Consider a value s_1' such that buying the good is optimal, then

$$u(2 + s_1' - p) + \beta[(1 - \lambda)u(b) + \lambda u(p - 1)] \geq u(1 + s_1' - s_2) + \beta E[u(y_2 + s_2)].$$

The inequality holds for all $s_2 < p - 1$, thus it also holds for $s_2^{No}(s_1'')$, the s_2 that is optimal at a higher level $s_1'' > s_1'$, conditional on the nondivisible not being bought. Due to strict concavity of $u(c_t)$,

$$u(1 + s_1' - s_2^{No}(s_1'')) - u(1 + s_1' - (p - 1)) \geq u(1 + s_1'' - s_2^{No}(s_1'')) - u(1 + s_1'' - (p - 1)),$$

i.e., the consumption gain $(p - 1) - s_2^{No}$ from deciding not to save for the good in period 2 gives a higher utility gain when starting from the lower consumption level $1 + s_1'$ than when starting from consumption level $1 + s_1''$. Since

$$\beta[(1 - \lambda)u(b) + \lambda u(p - 1)] - \beta E[u(y_2 + s_2^{No}(s_1''))] \geq u(1 + s_1' - s_2^{No}(s_1'')) - u(2 + s_1' - p)$$

holds by the optimality of buying the good at s_1' , substitution and rearranging yields

$$u(2 + s_1'' - p) + \beta[(1 - \lambda)u(b) + \lambda u(p - 1)] \geq u(1 + s_1'' - s_2^{No}) + \beta E[u(y_2 + s_2^{No})]$$

for all $s_1'' > s_1'$. Therefore, when s_1 has reached some threshold s_{min} , saving for the nondivisible is optimal for all $s_1 \geq s_{min}$.

(b) For a given β , evaluate inequality 2.2 at $s_1 = s_{min}$. If β is increased to $\beta' > \beta$, the inequality still holds: $u(b) > u(1 + s_2)$ and $u(p - 1) > u(s_2)$ for all $s_2 < p - 1$ given $b > p$. Intuitively, the weight of the reward of saving increases relative to the cost. Since $u'(c) > 0$, the inequality becomes more slack, and will still be satisfied for $s_1' = s_{min} - \epsilon$. Therefore, s_{min} decreases in β .

(c) Investigating the sign of $\delta s_{min} / \delta \lambda$, note that an increase in λ makes it less attractive to save for the nondivisible (which will not be obtained in case of a shock), increasing s_{min} . However, a stronger motive for precautionary savings on the right-hand side decreases the savings difference $(p - 1) - s_2^{No}$, which decreases s_{min} . Which effect dominates is a function of $(b - p)$ and $u''(c)$. Formally, both sides of the inequality decrease in λ . As the shock hits, the right-hand side loses 1, at a consumption level $1 + s_2^{No} < p$. The left-hand side loses $(b - p) + 1 > 1$, at a higher consumption level $b > 1 + s_2^{No}$. \square

Lemma 2. (a) The maximum that period 1 would be willing to save, denoted s_{max} , is strictly increasing in the time-inconsistency parameter β . (b) s_{max} weakly decreases in the amount of naiveté, $\tilde{\beta} - \beta$.

Proof. (a) Evaluate inequality 2.3 at $s_1 = s_{max}$. For each side separately, take the derivative w.r.t. β . By the envelope condition, $\frac{dU}{d\beta} = \frac{\delta U}{\delta \beta} + \frac{\delta U}{\delta s_1^{No}} \frac{\delta s_1^{No}}{\delta \beta} = \frac{\delta U}{\delta \beta}$. For a time-inconsistent period 1 agent with $\tilde{\beta} < 1$, only s_1 is a choice variable – \tilde{s}_2^{No} is inferred by backward induction, and depends on his belief $\tilde{\beta}$ (rather than on β). The resulting derivative of the left-hand side is bigger than the derivative of the right-hand side:

$$\begin{aligned} & (1 - \lambda)^2(u(2 + s_1 - p) + u(b)) \\ & + (1 - \lambda)\lambda(u(2 + s_1 - p) + u(p - 1)) \\ & + \lambda(u(s_1 - \tilde{s}_2^{No}) + E[u(y_3 + \tilde{s}_2^{No})]) \\ & > E[u(y_2 + s_1^{No} - \tilde{s}_2^{No}) + u(y_3 + \tilde{s}_2^{No})] \end{aligned}$$

This inequality follows from inequality 2.3, noting that $u(1 - s_{max}) < u(1 - s_1^{No})$ holds by definition of s_{max} . As a result, when s_1 is held constant at s_{max} , and β is

increased, the left-hand side increases more than the right-hand side does, so the original inequality is maintained and becomes more slack. The inequality will still hold for $s_1 = s_{max} + \epsilon$. Thus, s_{max} is strictly increasing in β .

(b) For a given β , an increase in $\tilde{\beta} > \beta$ is associated with a less sophisticated agent. The parameter $\tilde{\beta}$ enters the s_{max} -function through period 1's expectation of period 2's precautionary savings, $\tilde{s}_2^{No}(s_1) = \text{argmax}(u(y_2 + s_1 - s_2) + \tilde{\beta}E[u(y_3 + s_2)])$. An increase in $\tilde{\beta}$ causes expected precautionary savings \tilde{s}_2^{No} to increase. This brings savings closer to period 1's ideal: Since period 1 discounts period 2 and 3 at the same rate, he would like his future self to save more than he actually does. As $\tilde{\beta}$ increases, period 1 is more optimistic that period 2 will follow his preferences. As a result, both sides of inequality 2.3 increase in $\tilde{\beta}$. However, the agent is more dependent on precautionary savings if he does not save for the nondivisible good, since savings for the nondivisible act as an insurance against shocks. Thus, the left-hand side of the inequality increases less than the right-hand side, and the inequality may no longer hold at the original s_{max} . Hence, s_{max} weakly decreases in $\tilde{\beta}$. \square

Lemma 3. (a) The optimal allocation of savings from period 1's perspective, denoted $s_1 = s_{opt}$, is characterized by

$$u'(1 - s_{opt}) = \beta[(1 - \lambda)u'(2 + s_{opt} - p) + \lambda u'(s_{opt} - \tilde{s}_2^{No})(1 + \frac{\delta \tilde{s}_2^{No}}{\delta s_1} \cdot \frac{1 - \tilde{\beta}}{\tilde{\beta}})].$$

(b) s_{opt} is strictly increasing in β , and always smaller than s_{max} .

Proof. (a) Maximising expected lifetime utility from period 1 perspective, conditional on purchase of the nondivisible (i.e., on $s_2 = p - 1$), yields the following first-order condition for $s_1 = s_{opt}$:

$$\begin{aligned} u'(1 - s_{opt}) &= \beta[(1 - \lambda)u'(2 + s_{opt} - p) + \lambda u'(s_{opt} - \tilde{s}_2^{No})] \\ &\quad + \beta \lambda \frac{\delta \tilde{s}_2^{No}}{\delta s_1} [-u'(s_{opt} - \tilde{s}_2^{No}) + Eu'(y_3 + \tilde{s}_2^{No})] \end{aligned}$$

Note that $\delta U_1 / \delta \tilde{s}_2^{No} \neq 0$ given $\tilde{\beta} < 1$: Period 1 self does not expect his future self to share his preferences, thus the envelope condition does not apply for \tilde{s}_2^{No} . The first-order condition for s_{opt} can be simplified using the first-order condition from \tilde{s}_2^{No} : $\beta Eu'(y_3 + \tilde{s}_2^{No}) = u'(s_1 - \tilde{s}_2^{No})$. Substituting this into the above and simplifying yields Lemma 3.

(b) s_{opt} is determined by the equation in part (a). Increasing β unambiguously increases the right-hand side of the equation (note $\delta \tilde{s}_2^{No} / \delta s_1 > 0$). To clear, the marginal utility of period 1 consumption must increase, implying an increase in s_{opt} . Thus, s_{opt} increases in β . Further, $s_{opt} \leq s_{max}$, follows by the definition of s_{max} . \square

Proposition 1. For small shock frequencies λ , and in the region where savings are skewed towards period 1, $s_1 \geq \bar{s} \equiv \frac{p-1}{2}$, adopting a regular-instalment product increases the maximum the agent is willing to save, i.e., $s_{max}^B > s_{max}$. A sufficient constraint on the shock frequency is $\lambda < \frac{u'(1)}{u'(0.5)}$. In the region $s_1 < \bar{s}$, adopting the regular-instalment product unambiguously decreases s_{max} .

Proof. In the region $s_1 \geq \bar{s}$: From inequality 2.5, the introduction of a penalty D will increase s_{max} whenever $\beta\lambda[u(s_1 - \tilde{s}_2^{No}) - u(s_1 - D - \tilde{s}_2^{No})] < u(1 - s_1^{No}) - u(1 - D - s_1^{No})$. To a first-order approximation, this is equivalent to $\beta\lambda u'(s_1) \cdot D < u'(1) \cdot D$, which holds whenever $\lambda < u'(1)/u'(s_1)$. Given $s_1 \geq \bar{s} > 0.5$, it is sufficient that $\lambda < u'(1)/u'(0.5)$. Therefore, inequality 2.5 always holds using the original $s_{max}(\beta)$, and it still holds for $s_{max}(\beta) + \epsilon$. For the special case where $D > s_1$, limited liability applies: The left-hand side stays constant as D increases, while the right-hand side decreases in D . The positive effect of D on s_{max} is reinforced. The resulting $s_{max}^B(\beta)$ will be strictly higher than $s_{max}(\beta)$ for $s_1 \geq \bar{s}$.

In the region $s_1 < \bar{s}$: The agent compares

$$\begin{aligned} & u(1 - s_1 - D) + \beta(1 - \lambda)^2(u(2 + s_1 - p) + u(b)) \\ & + \beta(1 - \lambda)\lambda(u(2 + s_1 - p) + u(p - 1)) \\ & + \beta\lambda(u(s_1 - \tilde{s}_2^{No}) + E[u(y_3 + \tilde{s}_2^{No})]) \\ & \geq u(1 - D - s_1^{No}) + \beta E[u(y_2 + s_1^{No} - \tilde{s}_2^{No}) + u(y_3 + \tilde{s}_2^{No})] \end{aligned}$$

With a strictly concave utility function, the utility loss from D when starting at consumption level $1 - s_1$ is bigger than the utility loss from D when starting at consumption level 1: $u(1 - s_1) - u(1 - s_1 - D) > u(1) - u(1 - D)$ for $s_1 > 0$. In other words, the penalty D hurts the agent more when he is saving for the nondivisible than when he is not. With the left-hand side decreasing more than the right-hand side, willingness to save will decrease, shifting the $s_{max}^B(\beta)$ -curve below the original $s_{max}(\beta)$ -curve for $s_1 < \bar{s}$. Further note \tilde{s}_2^{No} is affected by D , but only through s_1^{No} . For s_1^{No} , the envelope condition applies. \square

Lemma 4. The threshold $\hat{s}(\beta)$ weakly decreases in β . Equivalently, as β increases, a larger range $s_1 \in (\hat{s}(\beta), \bar{s})$ is strictly dominated by \bar{s} .

Proof. The threshold $\hat{s}(\beta)$ is the lowest value of s_1 which satisfies

$$\begin{aligned} & u(1 - \hat{s} - D) + \beta(1 - \lambda)(u(2 + \hat{s} - p)) + \beta\lambda(u(\hat{s} - \tilde{s}_2^{No}(\hat{s})) + E[u(y_3 + \tilde{s}_2^{No}(\hat{s}))]) \\ & \leq u(1 - \bar{s}) + \beta(1 - \lambda)(u(2 + \bar{s} - p)) + \beta\lambda(u(\bar{s} - D - \tilde{s}_2^{No}(\bar{s})) + E[u(y_3 + \tilde{s}_2^{No}(\bar{s}))]). \end{aligned}$$

By construction, $\hat{s}(\beta) < \bar{s} - D$ for all $\beta > 0$. Given $u(1 - \hat{s} - D) > u(1 - \bar{s})$, it must be that

$$\begin{aligned} & \beta(1 - \lambda)(u(2 + \hat{s} - p)) + \beta\lambda(u(\hat{s} - \tilde{s}_2^{No}(\hat{s})) + E[u(y_3 + \tilde{s}_2^{No}(\hat{s}))]) \\ & < \beta(1 - \lambda)(u(2 + \bar{s} - p)) + \beta\lambda(u(\bar{s} - D - \tilde{s}_2^{No}(\bar{s})) + E[u(y_3 + \tilde{s}_2^{No}(\bar{s}))]). \end{aligned}$$

This inequality will still hold for $\beta' > \beta$, and become more slack. All values of s_1 which were strictly dominated at β are also strictly dominated at β' . The dominated region $(\hat{s}(\beta), \bar{s})$ becomes weakly larger. \square

Proposition 2. *For a given shock frequency λ , the minimum effective penalty D_{eff} that will enforce saving weakly decreases in the time-inconsistency parameter β .*

Proof. For a given β , $D_{eff}(\beta)$ is defined as the minimum penalty such that $s_{min}^B(\beta) \leq s_{max}^B(\beta)$. Holding the penalty D fixed, and increasing β to $\beta' > \beta$, Lemma 1 and 2 assert that $s_{min}^B(\beta') < s_{min}^B(\beta) \leq s_{max}^B(\beta) < s_{max}^B(\beta')$. Thus, penalty $D_{eff}(\beta)$ is effective for all $\beta' \geq \beta$. \square

Proposition 3. *The optimal Regular Saver contract for a fully sophisticated agent with $\beta < \hat{\beta}$ depends on the effect of the minimum effective penalty, $D = D_{eff}$: Where D_{eff} results in $s_{min}^B(\beta) \leq s_{max}^B(\beta) \leq \bar{s}$ (illustrated in Figure 2.3), equilibrium savings contributions are perfectly smooth, i.e., $s_1 = \bar{s} \equiv \frac{p-1}{2}$ and $s_2 = 2\bar{s} = p - 1$. Consequently, the optimal contract is to choose D_{eff} . Where D_{eff} results in $s_{max}^B(\beta) \geq s_{min}^B(\beta) > \bar{s}$ (illustrated in Figure 2.2), the optimal contract involves $D \geq D_{eff}$, with equilibrium savings weakly skewed towards period 1 ($s_1 \geq \bar{s}$).*

Proof. First, note that a fully sophisticated agent will never adopt a contract with $D < D_{eff}$: This results in $s_{max}^B(\beta) < s_{min}^B(\beta)$, and thus in certain default in period 1, which is dominated by not adopting the product. It then trivially follows that when D_{eff} results in $s_{max}^B(\beta) \geq s_{min}^B(\beta) > \bar{s}$, the optimal contract involves $D \geq D_{eff}$, and achieves equilibrium savings $s_1 \geq \bar{s}$.

Second, when D_{eff} results in $s_{min}^B(\beta) \leq s_{max}^B(\beta) \leq \bar{s}$, choosing D_{eff} necessarily results in equilibrium savings \bar{s} . To see this, recall that $s_{min}^B(\beta) = s_{min}(\beta)$ in the region $s_1 < \bar{s}$: Period 1 has already defaulted on the contract, implying the contract is no longer active in period 2. Further, by Proposition 1, $s_{max}^B(\beta) < s_{max}(\beta)$ in the region $s_1 < \bar{s}$. Therefore, starting from $\beta < \hat{\beta}$ and thus $s_{max}(\beta) < s_{min}(\beta)$, introducing a penalty will never lead to an intersection $s_{max}^B = s_{min}^B$ in the region $s_1 < \bar{s}$. The only possibility for $s_{min}^B(\beta) \leq s_{max}^B(\beta) \leq \bar{s}$ to occur is an intersection of the curves on the vertical (dominated) part of the s_{max}^B -curve, where $s_{min}^B \in [\hat{s}(\beta), \bar{s})$, and $s_{max}^B = \bar{s}$. This happens when the penalty is sufficiently high to make

the agent willing to jump to \bar{s} . From the equilibrium savings schedule, $s_1 = \bar{s}$ if $\max(s_{min}^B, s_{opt}^B) \in [\hat{s}, \bar{s})$. \square

Appendix C

Appendix to Chapter 3

C.1 Supplementary Figures and Tables

TABLE C.1: PERSONAL SAVINGS GOALS

	All	All (%)	ASA clients	Gihandom clients
Education	163	21.79	18	21
General Savings/Not specified	148	19.79	37	21
House/Lot purchase/construction/repair	106	14.17	20	12
Christmas/Birthday/Fiesta/Baptism	91	12.17	12	16
Capital for Business	69	9.22	9	5
Household Item (Appliance/Furniture)	41	5.48	5	4
TV/DVD Player/Laptop/Cellphone	33	4.41	3	2
Emergency Buffer	31	4.14	1	0
Health/Medical	26	3.48	3	2
Agricultural/Livestock	19	2.54	2	6
Motorbike/Car/Boat	17	2.27	4	2
Travel/Vacation	4	0.53	0	1
Total	748	100	114	92
Median Goal Amount (pesos)	2400		2400	2400
Median Time until Goal Date (days)	137		138	133
Median Termination Fee (pesos, if ASA)	–		150	–
Date-Based Goal (if Gihandom)	–		–	53
Amount-Based Goal (if Gihandom)	–		–	39

TABLE C.2: TAKE-UP RATES

	Assigned	Reached	Take-Up	Take-Up (% assigned)	Take-Up (% reached)
Regular Saver (ASA)	457	423	114	25%	27%
Withdrawal Restriction (Gihandom)	228	219	92	40%	42%
Standard Account (OSA) with P100	913	852	788	86%	92%

C.1. SUPPLEMENTARY FIGURES AND TABLES

TABLE C.3: QUANTILE REGRESSIONS

		(1) Change in Bank Savings	(2) Change in Other Savings	(3) Change in Outstanding Loans
10th	Regular Saver	0.00	252.00	-4,000.00*
Percentile	(ASA)	(0.00)	(2,353.66)	(2,282.56)
	Withdrawal Restr.	0.00	-148.00	-345.00
	(Gihandom)	(0.00)	(2,670.35)	(2,598.30)
20th	Regular Saver	0.00	-271.00	-2,000.00*
Percentile	(ASA)	(0.00)	(630.63)	(1,021.07)
	Withdrawal Restr.	0.00	-1,071.00	-1,000.00
	(Gihandom)	(0.00)	(715.48)	(1,162.30)
30th	Regular Saver	0.00	-150.00	-800.01**
Percentile	(ASA)	(0.00)	(261.67)	(394.72)
	Withdrawal Restr.	0.00	-240.00	-700.00
	(Gihandom)	(0.00)	(296.88)	(449.32)
40th	Regular Saver	0.00	0.00	0.00
Percentile	(ASA)	(5.45)	(53.89)	(129.39)
	Withdrawal Restr.	0.00	0.00	0.00
	(Gihandom)	(6.29)	(61.15)	(147.28)
50th	Regular Saver	0.00	0.00	0.00
Percentile	(ASA)	(5.23)	(97.89)	(41.80)
	Withdrawal Restr.	100.00***	56.67	0.00
	(Gihandom)	(6.03)	(111.06)	(47.58)
60th	Regular Saver	0.00	85.00	50.00
Percentile	(ASA)	(0.00)	(229.72)	(261.24)
	Withdrawal Restr.	100.00***	-135.00	-100.00
	(Gihandom)	(0.00)	(260.62)	(297.38)
70th	Regular Saver	0.00	110.00	-234.00
Percentile	(ASA)	(17.91)	(389.19)	(711.40)
	Withdrawal Restr.	100.00***	-343.44	-800.00
	(Gihandom)	(20.64)	(441.56)	(809.80)
80th	Regular Saver	200.00	-208.00	840.00
Percentile	(ASA)	(181.42)	(587.84)	(1,226.00)
	Withdrawal Restr.	150.00	-865.96	340.00
	(Gihandom)	(209.10)	(666.93)	(1,395.59)
90th	Regular Saver	2,051.87***	-635.00	925.00
Percentile	(ASA)	(329.68)	(1,290.76)	(3,737.72)
	Withdrawal Restr.	280.00	-1,050.00	-489.00
	(Gihandom)	(379.97)	(1,464.43)	(4,254.74)
Observations		748	603	720

Estimated standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Survey-based data (columns (2) and (3)) is truncated at 1 percent. All reported coefficients are Intent-to-Treat effects.

C.1. SUPPLEMENTARY FIGURES AND TABLES

Notes for Table C.3: The effect of offering the Regular Saver (ASA) product on total bank savings (ordinary plus commitment savings accounts) is not apparent until the 90th percentile. This is consistent with a large effect on the 51 ASA clients who successfully completed their contract, and a limited effect on non-adopters. The ASA product was offered to 423 individuals, of whom 114 adopted the product. The 63 ASA clients who defaulted largely achieved a zero change in savings - a majority of defaulters stopped depositing soon after opening their account (see Figure 3.1), and their opening balance was consumed by the default penalty.

The effect of offering the Withdrawal Restriction (Gihandom) product on bank savings is 100 pesos at the median - this is likely the mechanical result of a 42 percent take-up rate and a 100 pesos minimum opening balance. In contrast to ASA clients, those Gihandom clients who stopped depositing after their opening balance (79 percent) did not lose their savings to a default penalty, but their savings remain frozen in their account (up to a goal date or amount, see Section 3.5.2).

The regressions in columns (2) and (3) are based on survey responses on individuals' outstanding loan balance, as well as on savings at home and at other banks. While there is a large amount of noise in the survey data, there is no systematic evidence of a substitution from other sources of savings into savings at the partner bank. However, offering the Regular Saver product may have facilitated the biggest reductions in loan demand (at 10th, 20th, and 30th percentile).

C.1. SUPPLEMENTARY FIGURES AND TABLES




Figure C.1: Distributional Effect of Treatment on Savings, Loans and Expenditures




Figure C.2: Termination Fees (Chosen & Charged)

C.2 Robustness Checks

This section tests robustness along several dimensions. Table C.4 verifies that the estimation of average treatment effects is robust to the inclusion of unbalanced covariates (see Table 3.1). Table C.5 tests robustness of the take-up and default regressions of Section 3.5.3 with respect to the measurement of sophistication. As outlined in Section 3.3, sophistication is measured by interacting observed time-inconsistency (in MPLs) with a measure of *perceived* time-inconsistency. Instead of the previously used “Perceived Temptation” variable, Columns (4) and (8) of Table C.5 use “Perceived Self-Control” to capture perception of time-inconsistency (both measures are discussed in Section 3.3). Note that 316 out of 402 (79%) individuals in the R-sample report zero (or in 13 cases, negative) values of Perceived Self-Control. Interacted with the observed measure of present bias, this implies that only 21 out of 402 values of *Pres.Bias*Self-Control* are non-zero. While the relationship with take-up is not significant (likely due to a lack of variation), the coefficient on *Pres.Bias*Self-Control* is roughly comparable in magnitude and sign to the coefficient on *Pres.Bias*Temptation*.

Table C.6 looks at the effect of using real incentives instead of hypothetical questions in the measurement of time-inconsistency. Section 3.3 outlines the multiple price list method which was used to elicit individuals’ time preferences. The elicitation was first conducted with the entire sample using hypothetical questions. Towards the end of the survey (approximately 30min later), the elicitation was repeated for a randomly chosen half of the sample with real monetary rewards (Appendix C.3 describes the randomisation). During the hypothetical round, individuals were not informed about the existence of the real-rewards round.

The regressions in the main text use the incentivised measures where obtained (468 of 913 individuals, equivalent to 230 of 457 in group R), and rely on measures from the hypothetical round otherwise. Columns (2) and (5) of Table C.6 exploit the fact that ‘hypothetical measures’ are available for the whole sample, and re-run the ASA take-up and default estimations (treatment group R) from Section 3.5.3 using only unincentivised measures of present bias and impatience. In contrast, Columns (3) and (6) restrict the sample to those who received real rewards, and rely only on incentivised measures.

Table C.6 indicates that the main results of this paper (sophistication negatively predicts commitment take-up, (naive) present bias predicts default) appear to be driven by the incentivised measures of time-inconsistency: The estimated effects in the real-rewards sample are highly significant despite the much smaller sample size, while the coefficients for unincentivised measures of present bias (Columns

TABLE C.4: AVERAGE TREATMENT EFFECTS - ROBUSTNESS TO UNBALANCED COVARIATES

Dependent Variable	Change in Bank Savings (1)	(2)	Change in Other Savings (3)	(4)	Change in Outstanding Loans (5)	(6)
Regular Saver Treatment (ASA)	585.465*** (129.251)	588.830*** (132.821)	426.811 (671.844)	247.250 (688.523)	-840.258 (1,180.168)	-871.795 (1,197.380)
Withdrawal Restr.	148.243***	163.561***	-328.159	-613.247	-308.549	-678.111
Treatment (Gihandom)	(40.927)	(53.661)	(705.461)	(704.891)	(1,258.139)	(1,274.365)
Impatience		-121.374 (106.121)		843.994 (536.597)		1,635.189 (1,039.716)
Risk Aversion		24.084 (39.676)		-93.407 (133.289)		329.584 (230.476)
HH Income		24.678 (21.932)		-377.650* (192.628)		-88.329 (244.735)
Constant	27.160*** (9.399)	-113.780 (135.998)	63.451 (531.028)	1,377.945 (930.966)	1,882.729** (920.828)	286.813 (1,391.783)
R ²	0.02	0.03	0.00	0.03	0.00	0.01
Observations	748	745	603	600	720	715

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. All coefficients measure Intent-to-Treat (ITT) effects, i.e., the effect of being offered the commitment product. Columns (1)-(2) use administrative data from the partner bank. Columns (3)-(6) use survey-based data. The dependent variable in columns (3)-(4) is the change in reported savings at home and at other banks (other than the partner bank) between the baseline survey and the personal savings goal date (3-6 months later). The dependent variable in columns (5)-(6) is the change in the reported outstanding loan balance between the baseline survey and the endline survey, 6 months later. Both loan and survey-based savings data are truncated at 1 percent.

TABLE C.5: ASA TAKE-UP & DEFAULT: ROBUSTNESS (SOPHISTICATION MEASURE)

Dependent Variable	ASA Take-Up			ASA Default (R-Sample)				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Present Bias	0.0827 (0.0864)	-0.0405 (0.0613)	0.0623 (0.0825)	0.0000 (0.0685)	0.1119* (0.0654)	0.0281 (0.0451)	0.0933 (0.0616)	0.0749 (0.0510)
Soph. Present Bias (Pres.Bias*Temptation)	-0.0631** (0.0292)		-0.0570** (0.0288)		-0.0453** (0.0230)		-0.0371 (0.0237)	
Perceived Temptation	-0.0046 (0.0125)	-0.0137 (0.0115)	-0.0095 (0.0123)		-0.0202* (0.0105)	-0.0274*** (0.0097)	-0.0179* (0.0105)	
Pres.Bias*Self-Control				-0.0394 (0.0391)				-0.0651* (0.0375)
Perceived Self-Control				-0.0273 (0.0216)				-0.0211 (0.0182)
Full Controls	YES	YES	NO	YES	YES	YES	NO	YES
Marketer FE	YES	YES	YES	YES	YES	YES	YES	YES
Observations	402	402	408	402	402	402	408	402

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Entries in the table represent the marginal coefficients of the corresponding probit regression. In columns (3) and (7), the set of control variables has been limited to age and gender (all other control variables may directly or indirectly represent choice variables). Columns (2) and (6) omit the interaction term *tempted* – *ideal* * *presentbias* (“Pres.Bias*Temptation”), which is used as a measure of sophisticated time-inconsistency (see Section 3.3). Columns (4) and (8) instead use *expected* – *ideal* * *presentbias* (“Pres.Bias*Self-Control”) as an alternative measure of sophistication.

TABLE C.6: ASA TAKE-UP & DEFAULT: ROBUSTNESS (REAL VS. HYPOTHETICAL INCENTIVES)

Dependent Variable	ASA Take-Up			ASA Default (R-Sample)		
	(1) Aggregate	(2) Hypothetical Questions	(3) Real Incentives	(4) Aggregate	(5) Hypothetical Questions	(6) Real Incentives
Present Bias	0.0827 (0.0864)	-0.0508 (0.0820)	0.3203** (0.1417)	0.1119* (0.0654)	-0.0191 (0.0595)	0.3383*** (0.1095)
Soph. Present Bias (Pres.Bias*Temptation)	-0.0631** (0.0292)	0.0005 (0.0276)	-0.1937*** (0.0558)	-0.0453** (0.0230)	0.0025 (0.0175)	-0.1517*** (0.0511)
Perceived Temptation	-0.0046 (0.0125)	-0.0149 (0.0130)	0.0058 (0.0178)	-0.0202* (0.0105)	-0.0294*** (0.0112)	-0.0266* (0.0158)
Impatience	-0.0008 (0.0464)	-0.0213 (0.0474)	0.0487 (0.0681)	-0.0030 (0.0372)	-0.0247 (0.0387)	0.0142 (0.0554)
Full Controls	YES	YES	YES	YES	YES	YES
Marketer FE	YES	YES	YES	YES	YES	YES
Observations	402	401	199	402	401	199

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Entries in the table represent the marginal coefficients of the corresponding probit regression. Columns (2) and (5) regress the dependent variable on unincentivised measures of present bias and impatience for the entire sample. Columns (3) and (6) restrict the analysis to the real-rewards sample, and use only incentivised measures of present bias and impatience. Columns (1) and (4) use incentivised measures where available, and hypothetical measures otherwise. They are identical to Columns (1) and (2) of Table 3.6.

C.2. ROBUSTNESS CHECKS

TABLE C.7: REAL VS. HYPOTHETICAL INCENTIVES

A. BETWEEN-INDIVIDUAL COMPARISON (CROSS SECTION)			
Dependent Variable	Present Bias	Future Bias	Impatience
Real Incentives	-0.0264 (0.0250)	-0.0117 (0.0264)	0.0253 (0.0323)
Mean Dep. Variable	0.166	0.189	0.357
Observations	882	882	882
B. WITHIN-INDIVIDUAL COMPARISON (PANEL DATA)			
Dependent Variable	Present Bias	Future Bias	Impatience
Real Incentives	-0.0825 (0.1127)	-0.3049*** (0.1064)	0.5086*** (0.1452)
Mean Dep. Variable	0.161	0.219	0.337
Individuals	462	462	462
Observations	903	903	903

Standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Entries in the tables represent the marginal coefficients of probit regressions. Table B restricts the sample to those individuals who received real incentives, and uses a panel structure with 'real vs. hypothetical incentives' as the time dimension (thus, $T = 2$).

(2) and (5)) are close to zero. This result raises a series of new questions, most notably: What is the effect of real monetary incentives in the measurement of time preferences? Clearly a research topic of great interest in its own right, this question is beyond the scope of this dissertation, and will be the subject of a separate working paper.

Table C.7 provides some preliminary evidence on the effect of real rewards: The between-individual analysis is a simple cross-section regression of time-preference outcomes (incentivised where obtained, otherwise hypothetical) on whether or not the individual received monetary incentives. The within-individual analysis is restricted to the real-rewards sample, and uses two observations per individual: One to capture her time preferences using hypothetical questions, and one under monetary rewards. To illustrate, the estimated equation for present bias is $presentbias_{it} = \alpha + \beta * real_{it} + \mu_i + \epsilon_{it}$, where μ_i is assumed to be random. The results suggest that monetary incentives may decrease the occurrence of time-inconsistency: Individuals were less likely to exhibit either present bias or future bias (although only the latter effect is significant), but developed more general impatience. The between-individual analysis confirms the sign of this effect (less time-inconsistency,

C.2. ROBUSTNESS CHECKS

more impatience), but remains statistically insignificant. In combination with the strong predictive power for commitment take-up and default observed in Table C.6, these results are consistent with the idea that incentivising survey questions reduces noise and improves the quality of the answers.

C.3 Survey Measurement and Marketing Material

The ad-hoc randomization to determine who would receive real rewards for the time-preference questions was implemented as follows: At the start of the survey, enumerators verified respondents' ID as a part of the screening process. Enumerators then performed a calculation based on an individual's birth day, month and year. If the calculated number was odd, the respondent received a survey containing questions with real rewards. If the calculated number was even, the survey was administered with hypothetical questions.¹ Individuals were not informed about this randomisation when starting the survey, but the nature of rewards was transparent at the time of asking the questions. Serious consideration was given to the possibility of an uncertainty bias: In the presence of uncertainty about whether they would receive a promised future payment, even time-consistent agents would have an incentive to always pick the immediate reward. Choices in the future time frame would be unaffected, resulting in an upward bias on the present bias measure. To assure individuals that all payments were guaranteed, both cash and official post-dated bank cheques were presented during the game.

In addition to the measures for present bias and sophistication, the baseline survey obtained measures of other covariates of interest: A measure of the strength of financial claims from others is obtained using a methodology similar to that in Johnson et al. (2002): Individuals were presented with a hypothetical scenario in which they keep 3000 pesos in their house, set aside for a particular expenditure that is due in one month. If the people around them knew about this money, how many would ask for assistance, and how much would they ask? This hypothetical framing avoids the endogeneity inherent in asking respondents directly about actual transfers made to others (actual transfers were also observed, but not used in the analysis). The 'Financial Claims' variable used in this paper is an indicator for individuals who reported to face above-median claims from others (the median was 500 pesos, which was also the mode). Risk aversion is a score in $[1, 6]$, and represents the individual's choice when faced with a set of lottery options with increasing expected value and increasing variance (see Figure C.3). Choosing the 'no-risk' lottery A yielded a score of 6, for extreme risk aversion (this option was chosen by 48 percent of the sample). Cognitive ability is proxied by the number of correct answers (out of five possible) from a culture-free intelligence test (see Fig-


¹The calculation was designed to give an odd number if the individual's birth year was odd, and even otherwise. The survey team was unaware of this connection. Given the availability of verified IDs which included birthdays, it was possible to check ex-post that the correct type of survey had been administered.

C.3. SURVEY MEASUREMENT AND MARKETING MATERIAL

H10: Coin Flip Game

Suppose we play a game where you flip a coin and win a prize of money depending on if it is heads or tails. Example: Barangay lottery.

Which game would you prefer to play?




	 100-Peso	 500-Peso
Game A:	P100	P100
Game B:	P90	P190
Game C:	P80	P240
Game D:	P60	P300
Game E:	P20	P380
Game F:	P0	P400

Figure C.3: Test of Risk Aversion (Methodology: Binswanger (1980))





Figure C.4: Illustration: Test of Cognitive Ability

C.3. SURVEY MEASUREMENT AND MARKETING MATERIAL



Personal Savings Plan



Name: Sample

Address: Gingoog City, Mindanao

Purpose of Savings: Tuition Fees

Goal Date: 4 July

Goal Amount: 5000Php

wk	Date Due	Deposit Due	Date of Deposit	Deposit made? (tick!)
1	22 Feb	250		
2	29 Feb	250		
3	07 Mar	250		
4	14 Mar	250		
5	21 Mar	250		
6	28 Mar	250		
7	04 Apr	250		
8	11 Apr	250		
9	18 Apr	250		
10	25 Apr	250		

wk	Date Due	Deposit Due	Date of Deposit	Deposit made? (tick!)
11	2 May	250		
12	9 May	250		
13	16 May	250		
14	23 May	250		
15	30 May	250		
16	6 Jun	250		
17	13 Jun	250		
18	20 Jun	250		
19	27 Jun	250		
20	4 Jul	250		

Figure C.5: Personal Savings Plan (All Treatment Groups)

C.3. SURVEY MEASUREMENT AND MARKETING MATERIAL

ure C.4 for a sample question). A financial literacy score is given by the number of correct answers (again, out of five possible) to basic numeracy questions. Household bargaining power is measured as follows: Individuals were asked who was the main decisionmaker for five types of household expenses (market purchases, durable goods, transfers to others, personal recreation, and schooling of children). For each type of expense at their discretion, their bargaining score increased by one, resulting in a measure with a range $[0, 5]$. 95 percent of respondents were female; thus the variable measures predominantly female bargaining power. Distance to the bank branch is measured as the linear geographic distance to the partner bank, obtained using GPS coordinates. An existing savings account indicates that the individual reported to have an existing savings or checking account at any bank (not necessarily the partner bank) at the time of the baseline survey. ‘Donates to charity’ is a dummy that switches on if the individual reported to have given any positive amount of money to charity in the past 12 months. It is a proxy for the individual’s attitude towards charitable giving, motivated by the fact that the ASA default penalty was framed as a charitable contribution. ‘Charity buckets’ are common even in low-income areas of the Philippines, especially for disaster relief and the Red Cross. While charitable giving is unsurprisingly related to income, 40 percent of the population reported positive contributions, many as small as five pesos (the median was 100 pesos, conditional on giving). Finally, the shock arrival rate is proxied by the number of unexpected emergencies (such as death or illness of a household member, redundancy, natural disasters, damage to house and crops, theft, and a range of others) that a household suffered in the last 12 months before the start of the treatment (‘Emergencies last yr’).