

When Commitment Fails

- Evidence from a Regular Saver Product in the Philippines*

Anett John[†]

Version 17 May 2015

(a current draft can be found [here](#))

Abstract

Commitment products are widely regarded as a remedy for self-control problems. However, imperfect knowledge about one's preferences implies that individuals may fail to anticipate their behaviour under commitment, and consequently choose ill-suited commitment contracts. I conduct a randomised experiment in the Philippines, where low-income individuals were randomly offered a regular-installment commitment savings product. Individuals chose the stakes of the contract (a default penalty) themselves. A majority appears to choose a harmful contract: While the intent-to-treat effect on individuals' bank savings is large, 55 percent of clients default on their savings contract. A possible explanation that is supported by the data is that the chosen stakes were too low (the commitment was too weak) to overcome clients' self-control problems. Both take-up and default are *negatively* predicted by measures of sophisticated hyperbolic discounting - suggesting that partial sophisticates adopt weak commitments and then default, while full sophisticates are more cautious about committing, but better able to choose incentive-compatible contracts.

Keywords: commitment savings, hyperbolic discounting, partial sophistication

JEL classification: C93, D03, D14, O12

*I would like to extend my gratitude to Oriana Bandiera, Maitreesh Ghatak, and Gharad Bryan, for their invaluable support and advice throughout this research project. I am deeply indebted to Dean Karlan, Ann Mayuga, Faith McCollister, Megan McGuire, Yoeri Suykerbuyk and Eva Ghirmai of IPA Philippines, without whom this project would not have been possible. I am grateful for advice from Stefano DellaVigna, Nageeb Ali, Matthew Levy, George Loewenstein, Lance Lochner, Francesco Nava, Johannes Spinnewijn, as well as from Betty Wilkinson and Michiko Katagami at the Asian Development Bank, and John Owens at the Rural Bankers' Association of the Philippines. Thanks to Jonathan de Quidt, Erina Ytsma, and various seminar audiences for helpful comments and discussions. I also thank 1st Valley Bank of Cagayan de Oro, Philippines, for a productive collaboration. I gratefully acknowledge the financial support of the Yale Savings and Payments Research Fund at Innovations for Poverty Action, sponsored by a grant from the Bill & Melinda Gates Foundation. This research further received generous support from the Royal Economic Society. All errors and omissions are my own.

[†]CREST-ENSAE Paris, email: anett.john@ensae.fr. This paper was previously circulated under the name Anett Hofmann.

1 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 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)).

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. Also see Duflo et al. (2011) for commitment to fertilizer use via advance purchase.

But are people good at choosing the ‘right’ commitment contract? 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, the contract may result in undesirable behaviour, and the agent may be harmed rather than helped. The very nature of most commitment contracts is to impose monetary or social penalties for undesirable behaviour. Thus, adopting a commitment device that is ill-suited to one’s preferences may ‘backfire’ and become a threat to welfare.⁷

This paper argues that commitment can be harmful if agents select into the wrong commitment contract - and that they frequently do. I show theoretically why offering commitment improves the welfare of fully sophisticated agents, but is likely to harm partially sophisticated agents. The effect is strongly discontinuous at full sophistication, as any amount of naiveté leads agents to select into contracts which violate their incentive constraints. Furthermore, commitment adoption *increases* in naiveté,⁸ as naive agents underestimate the cost of effective commitment. I conduct a randomised experiment where individuals could sign up for a new commitment savings account with fixed regular instalments, and where they are 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, 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. A possible explanation that is 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 repeat marketing stage with the offer of ‘pre-ordering’ the product for a second round support the hypothesis that a significant share of clients took up the commitment contract by mistake. Alternative explanations for default that find some empirical support are income optimism and household conflict. A pure stochastic shock explanation appears unlikely.

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

⁷Consider 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.

⁸Conditional on a given level of time-inconsistency, and an awareness that one cannot save without commitment. The latter rules out the case of perfect naiveté.

non-commitment savings account with 100 pesos (U.S. \$2.50) opening balance.⁹ Personal savings plans featured a self-chosen goal date, goal amount, and a weekly or bi-weekly instalment plan. 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 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.¹⁰ There was no compensation for the restrictions, no added help (such as deposit collectors or text message reminders), and a standard market interest rate.¹¹

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.¹²

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.¹³ 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 control.¹⁴ 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

⁹At the time of marketing (October 2012), the exchange rate was 42 Philippine pesos per U.S. dollar.

¹⁰The 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.

¹¹The 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.

¹²Pre-orders were not legally binding, but involved a cost through substantial paperwork.

¹³The 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.

¹⁴The 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.

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.¹⁵

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 more cautious about adopting commitment, but also less likely to default, conditional on take-up (consistent with an increased ability to choose a suitable contract). 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. I interpret this as evidence against a shock explanation, where individuals rationally default following large random shocks to their income or expenditures.

This paper builds and expands on the literature in three ways. First, 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 paper 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. Dupas and Robinson (2013) document that offering Kenyan women savings accounts with withdrawal fees 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. Finally, Giné et al. (2010) offered smokers in the Philippines a commitment savings contract for smoking cessation, in which participants forfeit their savings to charity if they fail a nicotine test after 6 months. 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.¹⁶ 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.

Second, the paper provides the first analysis of a commitment savings product with fixed regular instalments in a randomised setting. 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. It 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

¹⁵Dormancy fees are very common with Philippine banks, and commonly start after two years of inactivity.

¹⁶Interestingly, 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.

a one-off investment.¹⁷ 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 rarely available in savings products.¹⁸ The idea that time-inconsistent agents benefit from commitment to regular fixed instalments has been discussed by Fischer and Ghatak (2010), Bauer et al. (2012), and John (2014). 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 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 paper provides the first direct comparison of a regular-instalment savings product with a pure withdrawal-restriction product.¹⁹ 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.

Third, the paper 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.²⁰ 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.

The paper proceeds as follows. Section 2 outlines a model of commitment under partial sophistication. Section 3 describes the experimental design employed. Section 4 explains the survey instrument, with particular view to the measurement of time-inconsistency and sophisticated hyperbolic discounting. Section 5 outlines the empirical strategy. Section 6 presents empirical results. Section 7 discusses other explanations. Section 8 concludes.

2 Theory: Commitment under Partial Sophistication

The following section presents a simple model of commitment under partial sophistication. It sheds light on (i) why sophisticated hyperbolic discounters can benefit from commitment, (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.

¹⁷See e.g., Ananth et al. (2007).

¹⁸Informal 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.

¹⁹The 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.

²⁰With the notable exception of Tarozi and Mahajan (2011), who follow a structural approach.

For expositional ease, I use a two-period linear-utility model of consumption and savings. A linear model is sufficient to highlight the basic mechanism of undercommitment, and provides a reasonable approximation for small stakes. All main results hold in a model with concave utility. Concave utility and three or more periods are needed to illustrate the consumption smoothing benefits of a regular-installment product, at the expense of tractability. A non-stochastic version is described in John (2014).

2.1 The Model

Consider an agent who chooses whether to save for a nondivisible good which costs the lump-sum $1 < p < 2$ and yields a benefit $b > 2$. The agent lives for 2 periods. In addition, there is a period 0 ‘planner’s perspective,’ when the agent may choose to adopt commitment (Section 2.3). The agent receives a per period income of $y_t = 1$, which he can either consume or save. He cannot borrow. Throughout, assume the interest rate is $R = 1$ and $\delta = 1$ for simplicity. Define s_1 as the amount of savings that he sends from period 1 to period 2, so that $c_1 = y_1 - s_1 \geq 0$. Lifetime utility as evaluated in each period is given by the discounted stream of future consumption:

$$U_t = c_t + \beta E(c_{t+1}).$$

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. $t + s + 1$, $s > 0$). Up to this point, the savings model resembles the autarky savings framework in Basu (2014) (where ‘autarky’ refers to the absence of banking). I now generalize the model to allow for partial sophistication and stochastic income, thus creating a need for flexibility. 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

$$U_{t+s} = c_{t+s} + \tilde{\beta} E(c_{t+s+1}).$$

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, such that $y_t = 0$. 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.²¹ The implication of a shock is that the agent’s lifetime income drops to (at most) 1, which means the nondivisible good can no longer be purchased. When a shock hits, any plans to save for the nondivisible are abandoned, and any existing savings are consumed. 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 good.²²

While there is no consensus on how to define welfare for time-inconsistent agents, the paper follows O’Donoghue and Rabin (1999): An agent’s welfare is understood to be the lifetime utility of the period 0 agent from an ex-ante perspective: $W = U_0 = E[c_1 + c_2]$. The advantage of this convention is that no particular period is favoured (since no consumption takes place in period 0).

²¹Suppose 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.

²²Time-consistent explanations to abandon savings plans include state-dependent preferences.

2.2 No-Commitment Equilibrium

The model is easily solved using backward induction. In period 2, the agent will buy the nondivisible whenever he can afford it, i.e., whenever $y_2 + s_1 > p$. Given $p > 1$, this requires $y_2 = 1$ (there is no shock) and $s_1 \geq p - 1$. Additional savings $s_1 > p - 1$ are simply consumed, as are savings that are not sufficient to buy the good. The consumption profile is

$$c_2 = \begin{cases} y_2 + s_1 - p + b & \text{if } y_2 + s_1 \geq p \\ y_2 + s_1 & \text{if } y_2 + s_1 < p. \end{cases}$$

In period 1, the agent realizes the good will be bought if and only if he sends $s_1 \geq p - 1$, and absent shocks. He responds by either sending $s_1 = p - 1$, or sending zero: In a linear model with $\beta < 1$, it is never optimal to shift excess consumption to the future.

Proposition 1. *In the No-Commitment Equilibrium, the nondivisible good is bought by sufficiently time-consistent agents, i.e., those with a time-consistency parameter β above a threshold $\hat{\beta}$. The threshold $\hat{\beta}$ increases in the shock frequency λ and the price p , and decreases in the benefit b .²³*

Proof. The period 1 agent prefers to save $s_1 = p - 1$ rather than zero iff

$$\underbrace{1 - (p - 1)}_{c_1} + \beta \underbrace{[\lambda(p - 1) + (1 - \lambda)b]}_{c_2} \geq \underbrace{1}_{c'_1} + \beta \underbrace{(1 - \lambda)}_{c'_2} \quad (1)$$

Rearranging yields that, absent shocks in periods 1 and 2, the nondivisible good is purchased for

$$\beta \geq \hat{\beta} \equiv \frac{p - 1}{\lambda(p - 1) + (1 - \lambda)(b - 1)}.$$

It is easy to see that $\frac{\partial \hat{\beta}}{\partial \lambda} > 0$, $\frac{\partial \hat{\beta}}{\partial p} > 0$, and $\frac{\partial \hat{\beta}}{\partial b} < 0$. □

The threshold condition has an intuitive interpretation: The time-consistency parameter β has to be larger than the ratio of the cost of saving today, $p - 1$, and its expected benefit tomorrow, $\lambda(p - 1) + (1 - \lambda)(b - 1)$. Because $b > p$, a time-consistent agent with $\beta = 1$ will always save for the nondivisible good. What is the effect of (partial) sophistication? In any given period, the agent believes that his future selves will be more patient than his current self, applying $\tilde{\beta}$ rather than β to their own future. Since the game ends after period 2, period 1's savings decision is unaffected by $\tilde{\beta}$. The no-commitment equilibrium is thus invariant to the degree of sophistication.²⁴

2.3 Introducing Commitment

For agents with high levels of time-inconsistency, the no-commitment equilibrium is inefficient: From a period 0 'planner's perspective', saving for the nondivisible good is welfare-improving regardless of β , since the planner values consumption in periods 1 and 2 equally, and the benefit of the good exceeds its cost. However, agents with $\beta < \hat{\beta}$ cannot save on their own, creating potential gains from commitment.

²³Apart from the added shocks λ , the no-commitment equilibrium result is shared with Basu (2014). The models diverge from here onwards.

²⁴This result is specific to two-period models. For longer savings horizons, coordination failure presents an additional welfare risk: Partially sophisticated agents will underestimate the savings required to induce their future selves to save.

Suppose the individual is now given the possibility to commit to saving for the nondivisible good, enforced via a penalty for non-compliance which the individual chooses himself ex-ante. Self-imposed conditional penalties have many applications: People tell their friends (or worse, their enemies) about a plan to lose weight, and then suffer an embarrassment cost when they are seen eating fast food. They join ROSCAs, knowing there will be social sanctions when they fail to contribute. Monetary penalties appear in most commitment contracts with front-loaded fees (see footnote 7).

In the current setting, individuals commit to a self-chosen default penalty, imposed if they fail to follow a regular-installment savings plan (see Section 3 for full details on the product design). Applied to a simple two-period model, the period 0 agent can choose a penalty D that is imposed on him in period 1 if he fails to meet the savings requirement, $s_1 \geq p - 1$. For a given penalty D , the period 1 agent now prefers to save $s_1 = p - 1$ rather than zero iff

$$\underbrace{1 - (p - 1)}_{c_1} + \beta \underbrace{[\lambda(p - 1) + (1 - \lambda)b]}_{c_2} \geq \underbrace{1 - D}_{c'_1} + \beta \underbrace{(1 - \lambda)}_{c'_2}. \quad (2)$$

Equivalently, he is willing to save if the penalty is higher than a minimum effective threshold, denoted D_{min} , which bridges the gap between the current costs and the future benefits of saving. Note that $D_{min} \leq p - 1 < 1$ for all β .

$$D \geq D_{min}(\beta) \equiv \underbrace{(p - 1)}_{\text{cost today}} - \beta \underbrace{[\lambda(p - 1) + (1 - \lambda)(b - 1)]}_{\text{benefit tomorrow}}. \quad (3)$$

Proposition 2. *The minimum penalty that is effective in enforcing the savings plan, denoted D_{min} , strictly decreases in the time-consistency parameter β . Further, D_{min} strictly increases in the shock frequency rate λ .*

Proof. The result directly follows from equation 3, noting that $\frac{\partial D_{min}}{\partial \beta} < 0$ and $\frac{\partial D_{min}}{\partial \lambda} > 0$. \square

A note on enforceability: The current setting assumes that the penalty D is fully enforceable, even if the agent loses his income to a shock, and is consequently unable to save (this situation is termed ‘rational default,’ and discussed in Section 2.4). In a two-period model, full enforceability is required to prevent the agent from adopting infinitely large penalties, reasoning that he will never incur them on the equilibrium path. Empirical arguments support that commitment penalties can often be enforced irrespective of current financial resources: First, D may be a non-monetary cost, such as embarrassment towards friends upon failing to achieve a stated goal. Second, D may represent a discounted stream of future losses, such as those from being barred from microfinance products in the future. Third, if the savings setting is extended to multiple periods (as it is in this study), penalties can be enforced using the stock of past savings. What matters for the theoretical results is that the agent incurs the penalty with positive probability on the equilibrium path, due to future uncertainty which cannot be contracted upon.²⁵ This creates a cost of commitment, and a demand for flexibility.²⁶

²⁵Inability to contract on the arrival of shocks may be caused by unobservability or moral hazard.

²⁶Note that the commitment versus flexibility trade-off discussed here differs from that in Amador et al. (2006). In the latter setting, agents can enforce the desired savings behaviour perfectly, i.e., the implied penalty is infinite, but is never incurred. The demand for flexibility comes from taste shocks, rather than from the risk of being financially unable to save.

2.4 Commitment Choice in Equilibrium

The commitment adoption decision has two steps: The period 0 ‘planner’ self first has to identify which commitment contract (which penalty size) will enable him to save for the nondivisible good while maximising flexibility in the event of a shock. In a second step, he decides whether to adopt this contract or not. Commitment is costly: If a shock hits in period 1, the agent loses his income, and thus his ability to save. Default becomes unavoidable, and the agent incurs a loss of D . This risk is referred to as ‘rational default’, as it occurs independently of time-inconsistent behaviour.

The equilibrium behaviour with full sophistication is summarized in Proposition 3.

Proposition 3. *Equilibrium with Full Sophistication: (a) Conditional on adopting commitment, individuals will adopt the minimum effective penalty, D_{min} . (b) Individuals who are sufficiently time-consistent to save in autarky (those with $\beta \geq \hat{\beta}$, see Proposition 1) never adopt commitment. (c) Individuals who cannot save in autarky (those with $\beta < \hat{\beta}$) adopt commitment if i) β is sufficiently high, and ii) the shock frequency rate λ is sufficiently low. The adoption decision is summarized in the condition $\lambda D_{min} \leq (1 - \lambda)^2(b - p)$, where λD_{min} represents the cost of commitment in case of a rational default, and $(1 - \lambda)^2(b - p)$ captures the expected benefit of a successful savings plan. (d) With full sophistication, offering commitment weakly increases welfare (it strictly increases the expected welfare of adopters).*

Proof. (a) For any commitment contract with an effective penalty, $D \geq D_{min}(\beta)$, expected utility from period 0’s perspective is

$$E(U_0^D) = (1 - \lambda) \underbrace{[1 - (p - 1)]}_{c_1} + \underbrace{\lambda(p - 1) + (1 - \lambda)b}_{c_2} + \lambda \underbrace{[-D]}_{c'_1} + \underbrace{(1 - \lambda)}_{c'_2}, \quad (4)$$

where $\lambda[-D + (1 - \lambda)]$ captures the risk of ‘rational default’ due to a shock in period 1. Choosing $D > D_{min}$ is costly. Period 1’s incentive constraint (equation 2) only depends on whether $D \geq D_{min}(\beta)$, thus choosing the minimum penalty always dominates choosing larger penalties. Choosing $D < D_{min}$ is strictly dominated by choosing no penalty at all, since period 1’s incentive constraint is violated, and default occurs with certainty. As a result, the period 0 agent chooses either $D = D_{min}$ or $D = 0$.

(b) For those who are sufficiently time-consistent to save in autarky ($\beta \geq \hat{\beta}$), $D_{min} = 0$. Part (a) then implies that it is never optimal to adopt positive amounts of commitment.

(c) Without commitment, individuals who cannot save in autarky ($\beta < \hat{\beta}$) face an expected utility of $E(U_0^A) = 2(1 - \lambda)$. Comparing $E(U_0^A)$ to the expected utility with commitment (equation 4), the period 0 agent prefers to adopt commitment if

$$\lambda D_{min} \leq (1 - \lambda)^2(b - p) \quad (5)$$

From period 0’s perspective, the benefit of commitment is the ability to purchase the nondivisible good (absent shocks), $(1 - \lambda)^2(b - p)$. A key result is that this benefit does not depend on the time-consistency parameter β . In contrast, β determines the cost of commitment: The expected loss due to ‘rational default’ is λD_{min} , which decreases in β by Proposition 2. Perhaps counter-intuitively, commitment is most attractive to those with the *lowest* degree of time-inconsistency (or the highest β), as the penalty required to enforce the savings plan is small, and poses little risk in the presence of shocks. In consequence, conditional on $\beta < \hat{\beta}$, agents adopt commitment for sufficiently high β . The connection between commitment adoption and the shock frequency rate $\lambda \in [0, 1]$ is straightforward: By Proposi-

tion 2, the cost of commitment λD_{min} increases in λ , while the benefit from commitment $(1 - \lambda)^2(b - p)$ decreases in λ . Commitment is adopted if shocks are sufficiently rare.

(d) Welfare has been defined as the ex-ante utility of the period 0 planner, $W = U_0 = E[c_1 + c_2]$. With full sophistication, the period 0 planner correctly identifies which amount of commitment will be effective for him (D_{min}), and whether commitment is optimal ($E(U_0^D) > E(U_0^A)$, which simplifies to equation 5). By construction, the welfare of adopters is increased by offering commitment, while the welfare of non-adopters is unaffected. \square

2.5 The Effect of Partial Sophistication

Suppose the agent is partially sophisticated about his degree of time-inconsistency, believing that his future selves will discount their own future at rate $\tilde{\beta} > \beta$ rather than at β . This corresponds to the classic pattern of procrastination, where an agent believes he will be patient enough to complete an unpleasant task (dieting, exercising, doing housework) tomorrow, but not today. The current setting examines the simplest form of partial sophistication, where beliefs are deterministic and incorrect. The result is that the period 0 planner underestimates the size of the penalty that will be required to induce his period 1 self to save, resulting in default and welfare losses.²⁷ The results are robust to allowing for more stochastic types of partial sophistication (see Section 2.6, which also addresses learning).

Proposition 4. *Equilibrium with Partial Sophistication: (a) Conditional on adopting commitment, individuals who are partially sophisticated about their time-inconsistency will adopt penalties strictly below the required effective minimum, $\tilde{D}_{min} < D_{min}$. As a result, adopters' incentive constraints in period 1 are systematically violated, triggering contract default. (b) Individuals who believe themselves to be sufficiently time-consistent to save in autarky (those with $\tilde{\beta} \geq \hat{\beta}$) never adopt commitment. (c) For those who realize they cannot save in autarky ($\tilde{\beta} < \hat{\beta}$), sophistication negatively predicts commitment adoption. For a given β , commitment is most attractive to those with the largest amount of naiveté, $\tilde{\beta} - \beta$. (d) With partial sophistication, offering commitment weakly decreases welfare (it strictly decreases the expected welfare of adopters).*

Proof. All arguments are analogue to the case of full sophistication, except that the period 0 agent believes the period 1 agent will apply $\tilde{\beta} > \beta$ in making intertemporal choices.

(a) From period 0's perspective, a penalty is perceived to be effective to ensure compliance when it satisfies $D \geq D_{min}(\tilde{\beta})$. For ease of notation, denote the perceived minimum effective penalty as $\tilde{D}_{min} \equiv D_{min}(\tilde{\beta})$. D_{min} strictly decreases in β , thus $\tilde{\beta} > \beta$ implies $\tilde{D}_{min} < D_{min}$. By Proposition 3, conditional on adopting commitment, agents will adopt \tilde{D}_{min} . Upon reaching period 1, and realizing one's true value of β , the incentive constraint (equation 2) is violated. Period 1 prefers to abandon the savings plan and incur the penalty.

(b) Directly follows from Proposition 3(b), noting that $\tilde{D}_{min} = 0$ for $\tilde{\beta} \geq \hat{\beta}$.²⁸

(c) Agents compare the *perceived* cost of commitment with the *perceived* benefit, and adopt commitment if $\lambda \tilde{D}_{min} \leq (1 - \lambda)^2(b - p)$ (this is a corollary of Proposition 3c). The cost of commitment decreases in $\tilde{\beta}$, while the benefit is invariant to it. Holding factual time-inconsistency β fixed,²⁹ a higher

²⁷In a two-period model, the degree of sophistication only affects period 0 decisions: Period 2 makes no decisions about the future, and thus period 1's $\tilde{\beta}$ is irrelevant. In a multi-period model, additional coordination problems arise, as the agent may consume savings made by past selves rather than to accumulate them over time.

²⁸An interesting empirical question is whether perceived ability to save without commitment ($\tilde{\beta} \geq \hat{\beta}$) necessarily results from factual ability to save without commitment ($\beta \geq \hat{\beta}$). See Section 2.6 on learning.

²⁹Is it plausible to hold β fixed while $\tilde{\beta}$ increases? This study is agnostic about the functional relationship between β and $\tilde{\beta}$, but the data suggest indeed that measures of time-inconsistency and perceived time-inconsistency are uncorrelated.

degree of naiveté $\tilde{\beta} - \beta$ implies that a lower penalty is regarded as effective, and thus less is at stake in case of a ‘rational default.’ Conditional on β , as well as on the agent’s perceived inability to save without commitment ($\tilde{\beta} < \hat{\beta}$), adoption increases with naiveté.

(d) Faced with an ineffective penalty $\tilde{D}_{min} < D_{min}$, the period 1 agent chooses to default, and pays D . Thus, offering commitment contracts decreases the ex-ante utility of adopters from $E(U_0^A) = 2(1 - \lambda)$ to $E(U_0^D) = 2(1 - \lambda) - \tilde{D}_{min}$. Note that naiveté has an ambiguous effect on this welfare loss: A more naive agent believes a lower penalty \tilde{D}_{min} to be effective, and is therefore more likely to adopt commitment, but he also suffers a smaller welfare loss after default. \square

2.6 Discussion and Extensions

Stochastic Sophistication and Rational Expectations The model studies a simple form of partial sophistication. The assumption of a deterministic and incorrect $\tilde{\beta}$ was introduced by O’Donoghue and Rabin (1999), and suffices to highlight the main mechanism of undercommitment. In reality, individuals may have more complex belief distributions about their future preferences, or preferences themselves may be stochastic. Consider the special case of rational expectations, and let future β be distributed according to some known $F_\beta(x)$, with an implied distribution for the minimum effective penalty, $G_{D_{min}}(x)$. The agent now faces an additional risk: The required D_{min} realized in period 1 may exceed a given penalty chosen in period 0, leading to costly default. The probability of default can be reduced by choosing a higher penalty, but this increases his vulnerability to shocks. The period 0 agent responds by maximising expected utility under commitment with respect to the optimal penalty D^* : He weighs equation 4 by the probability that a chosen penalty will be effective, $P[D_{min} \leq D^*] = G_{D_{min}}(D^*)$, otherwise he receives $2(1 - \lambda) - D^*$. This generalized optimization accounts for the possibility of no commitment, $D^* = 0$ (the agent does not know whether he is able to save in autarky, thus $G_{D_{min}}(0) > 0$).

Comparing a situation with stochastic versus full sophistication, commitment becomes less attractive: In both cases, the agent correctly assesses the effective penalty level in expectation, $E(D_{min})$. Even with risk-neutral preferences, expected utility from adopting $E(D_{min})$ is lower in the stochastic case: The default risk from a higher realized $D_{min} > E(D_{min})$ is not compensated by any benefit in the case of $D_{min} < E(D_{min})$. The agent’s response depends on the shape of $G_{D_{min}}(x)$, as well as on how much he gains from effective commitment, $(1 - \lambda)^2(b - p)$. If this gain is large, he will move towards higher penalties, $D^* > E(D_{min})$. If it is not, he will move away from commitment altogether. Uncertainty in agents’ beliefs thus results in low commitment take-up and high conditional penalties. This prediction is amplified if the agent is risk-averse, as the variance of consumption increases under commitment.

The predictions of stochastic sophistication and rational expectations are in contrast with the patterns observed in the data - they cannot explain why individuals choose very low default penalties, and why default is so frequent. Furthermore, the theory is inconsistent with high observed commitment adoption (relative to the observed frequency of time-inconsistent preferences). To reconcile stochastic beliefs about time-inconsistency with the empirical evidence, one needs to allow for belief distributions which are systematically biased towards naiveté.³⁰ With beliefs that are skewed towards $\beta = 1$, predictions closely resemble those of Section 2.5.

³⁰For instance, Duflo et al. (2011) consider farmers who are stochastically present-biased, and who systematically overestimate the probability that they will be patient in the next period.

Learning Empirical evidence suggesting partial sophistication about time-inconsistent preferences has been widely discussed.³¹ But is it plausible that individuals permanently hold incorrect beliefs about their time preferences, despite being able to observe their own past behaviour? Ali (2011) shows that rational learning about time-inconsistency may be perpetually partial: The agent learns about his β only if he exposes himself to temptation. If he is (or stochastically becomes) sufficiently pessimistic about being present-biased, he may use commitment to restrict his choice set to a point where he no longer learns. However, this argument cannot rationalise perpetual overoptimism, captured in the assumption that $\tilde{\beta} > \beta$: Optimistic planners undercommit and expose themselves to temptation, which should allow them to learn and update their beliefs over time. Yet, there are numerous reasons why Bayesian learning may be slow, or fail altogether: First, learning may be context-specific. The nature and degree of an individual's time-inconsistency may vary across dimensions (saving money, gym visits, food choices, work effort) or even across settings (saving on a regular-installment plan versus unscheduled savings). Context-specificity is likely to impede and slow down learning. In addition to learning about the temptations he faces, a new decisionmaking environment may also present the agent with a need to learn about the benefits and costs of his actions: How difficult is it to take a set amount of money out of the budget each period?³² Ali (2011) discusses how such multidimensional learning problems may create challenges of identification, and further slow the updating of beliefs.

Applying context-specificity to the model, a plausible case is that the period 0 agent is familiar with his savings behaviour absent commitment, but unfamiliar with his savings behaviour under a formal commitment savings product. In particular, he may realize from past observation whether he is able to save for the nondivisible good by himself, i.e., the inequality $\beta \geq \hat{\beta}$ is observed. Nevertheless, not having first-hand experience with commitment, knowing that $\beta < \hat{\beta}$ does not help him to assess how much commitment it will take to make him save. This argument relates to Bénabou and Tirole (2004)'s model of willpower and limited recall, in which the agent fails to remember past motives and feelings, but forms beliefs about his willpower based on past actions.

A second potential impediment to Bayesian updating are self-serving beliefs: Despite observing his own behaviour, the agent may prefer to attribute past failures to save to taste shocks ("I didn't want the good anymore") or to income shocks ("business was bad last month"), rather than admit to himself that he is present-biased. It is intuitive that agents may like to think of themselves as disciplined savers, healthy eaters, or frequent gym-goers. To the author's knowledge, no study directly considers the possibility that individuals derive utility from believing they are time-consistent. However, concerns about self-image and 'ego utility' have recently attracted research interest in the domain of skills and ability (Bénabou and Tirole (2002); Köszegi (2006); Möbius et al. (2011)). Bénabou and Tirole (2002) argue theoretically that rational individuals may prefer optimistic views of themselves to accurate ones, due to consumption, signaling and motivational reasons. Möbius et al. (2011) find empirical evidence that subjects systematically overweigh positive feedback relative to negative, and update their beliefs too little in response to either type of signal. They reconcile the evidence with a model of 'optimally biased Bayesians,' who misinterpret the informativeness of signals, but then correctly apply Bayes' rule. The information processing bias is 'optimal' in that it weighs the utility from enjoying a favorable

³¹ Examples include, but are not limited to, DellaVigna and Malmendier (2006); Duflo et al. (2011); Acland and Levy (2015), and Shui and Ausubel (2005).

³² Admittedly, almost all individuals in this study had previously experienced regular instalment structures through their loan repayments. However, the different labels 'savings' versus 'loan repayment' could imply a difference in cost, for instance in justifying them to one's family.

self-image with the cost of being more likely to make bad decisions. Similar arguments may apply to beliefs about time-inconsistency, in that individuals ‘optimally’ choose their degree of naiveté.

A potential third impediment comes from neuroscience: Recent evidence suggests that time-inconsistent behaviour may itself be a result of stress (Cornelisse et al. (2013); McClure et al. (2004)), or a cognitive consequence of scarcity (Mani et al. (2013); Shah et al. (2012)). It is an open question to what extent the same factors affect learning. For instance, if present-biased behaviour results from a lack of mental bandwidth allocated to specific tasks, and is thus situational and temporary in nature, then people may struggle to anticipate when, and for which tasks, they will adopt this behaviour. Similarly, at calm and reflective moments, they may underestimate how much their tastes will change when they are stressed, as evidenced by the literature on projection bias (Loewenstein et al. (2003)). This argument suggests an interesting contrast: In the standard model, partial sophistication is caused when the agent *underestimates* how much his future taste for immediate gratification resembles his current taste for immediate gratification. With contextual time-inconsistency and projection bias, partial sophistication is caused when the agent *exaggerates* how much his future tastes will resemble his current tastes, in that his current self is patient but his future ‘stressed’ self is not.

Pessimism and Overcommitment While the model focuses on agents who are optimistic about their degree of time-inconsistency ($\tilde{\beta} \geq \beta$, following O’Donoghue and Rabin (1999)), it easily accomodates pessimistic beliefs ($\tilde{\beta} < \beta$). Pessimistic agents overestimate the penalty required to make them save. Where pessimism results from agents believing they are strongly present biased in absolute terms (low $\tilde{\beta}$), commitment becomes very unattractive: As outlined in Propositions 3 and 4, the cost of commitment $\lambda \tilde{D}_{min}$ decreases in $\tilde{\beta}$ while the benefit $(1 - \lambda)^2(b - p)$ is invariant to it. The perceived minimum effective penalty is likely to be prohibitive. Where it is not, agents will overcommit, with excessively large penalties $\tilde{D}_{min} > D_{min}$. Excessive penalties are costless absent shocks, but they increase the damage in the case of ‘rational default.’ Summarising, pessimistic agents do not adopt commitment when it would indeed be optimal;³³ and when they do, they choose penalties which harm them more than necessary in the case of shocks. Yet, they correctly assess that their chosen contracts are ex-ante improvements relative to autarky. Thus, offering commitment contracts remains weakly welfare-improving.

It is difficult to derive testable predictions for pessimistic agents ($\tilde{\beta} < \beta$): They will be pooled with time-consistent and sophisticated time-inconsistent agents at all times. The following empirical analysis should be viewed in this light: Theory predicts that the pool of non-adopters will consist of time-consistent agents, time-inconsistent agents who believe they can save by themselves, and time-inconsistent agents with a prohibitively high perceived minimum effective penalty.³⁴ Similarly, the group of successful commitment adopters pools fully sophisticated and pessimistic agents, since both adopt effective penalties. Finally, the group of unsuccessful (defaulting) commitment adopters pools partial sophisticates and all adopters who suffered a shock. While the latter can be distinguished in an empirical setting by exploiting the randomness in their default timing, distinguishing pessimistic from fully sophisticated agents is not possible empirically.

³³This occurs when $\lambda D_{min}(\beta) \leq (1 - \lambda)^2(b - p) < \lambda \tilde{D}_{min}(\tilde{\beta})$: The benefit of commitment outweighs its cost using the true effective penalty, but not using the higher penalty \tilde{D}_{min} the pessimist believes he requires.

³⁴Allowing for multiple periods, and liability limited to current assets, there is a fourth group of non-adopters: Sophisticates who realize their penalty is not enforceable until a stock of savings has been accumulated. However, accumulating this stock is not incentive-compatible without an enforceable penalty. This is a relevant consideration for the empirical analysis in the remainder of this paper.

3 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.³⁵

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,³⁶ (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.³⁷ 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 6). 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 pesos. Common savings goals were school tuition fees, house repairs, and Christmas gifts (see Table IX 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.³⁸

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.³⁹ It is to be

³⁵Gihandom had been previously offered at other 1st Valley Bank branches, but not at the two branches in question.

³⁶Accepted forms of identification included a birth certificate, tax certificate, voter's ID, and several other substitutes.

³⁷A test for equality of means in the probability of being reached by the marketer yields an *F*-statistic with a *p*-value of 0.16. Individuals in group R were as likely to be reached as those in group C, but slightly less likely than those in group W.

³⁸Individuals 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.

³⁹Neither Ashraf et al. (2006b)'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.

expected that the marketing treatment itself influenced individuals' savings behaviour, as evidenced in the literature on mental accounting.⁴⁰ 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.⁴¹ In group W, 92 (out of 219 offered) accepted the Gihandom product. Table X 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 savings 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.⁴² 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 6). 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 account 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. Note that by nature of the contract, all ASAs were either successfully completed or in default by the goal date,⁴³ and any remaining savings were transferred to clients' ordinary savings accounts.

⁴⁰Most prominently, see Thaler (1985) and Shefrin and Thaler (1988).

⁴¹One member of the control group was mistakenly offered the ASA product and accepted, thus a total of 115 ASA accounts were opened. The estimated ASA treatment effect therefore represents a lower bound on the true effect.

⁴²Attitudes towards charities were measured in the baseline, and are controlled for in the analysis.

⁴³After the goal date, there was a one-week grace period to make any outstanding deposits, but no client made use of this.

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 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 Gihandom was 100 pesos. The difference of 50–150 pesos (depending on the savings goal) in the mandatory opening balances between ASA and Gihandom is a possible explanatory factor in the difference between take-up rates. Finally, two features were common to both ASA and Gihandom: 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 1.⁴⁴ 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),⁴⁵ 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 of a new savings plan, a new ‘Voluntary Donation Form’, and a decision on a new termination fee (to deter cheap talk).

4 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 6.1).

⁴⁴This 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 Gihandom is 24 out of 116 initially signed contracts.

⁴⁵Provided appropriate documentation, i.e. a hospital bill, death certificate, or proof of relocation.

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 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 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. No systematic difference is apparent between those offered real and those offered hypothetical rewards.⁴⁶ 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.

⁴⁶At 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.

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 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 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*.⁴⁷ 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.

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.

Table I 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 I 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 B.

5 Empirical Strategy

The primary objective of the study was to analyse the demand for and the effects of a regular-instalment commitment savings product, and to compare its performance to traditional withdrawal-restriction

⁴⁷I 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).

TABLE I: SUMMARY STATISTICS BY TREATMENT ASSIGNMENT

	Regular Saver Treatment	Withdrawal Restr. Treatment	Control	F-stat P-value
Age (yrs)*	43.834 (0.603)	43.449 (0.821)	44.250 (0.841)	0.804
Female*	0.941 (0.011)	0.943 (0.015)	0.943 (0.015)	0.991
Married	0.862 0.016	0.873 0.022	0.851 0.024	0.795
Weekly HH income (pesos)	2890.89 (124.26)	2485.78 (165.13)	3194.43 (272.45)	0.048
No. of appliances owned	2.276 0.084	2.180 0.110	2.250 0.124	0.802
No. of HH members	5.072 (0.091)	5.180 (0.140)	5.430 (0.140)	0.108
Education (yrs)	10.556 (0.166)	10.392 (0.242)	10.564 (0.251)	0.840
Received real rewards*	0.503 (0.023)	0.522 (0.033)	0.526 (0.033)	0.837
Present Bias*	0.172 (0.018)	0.161 (0.025)	0.156 (0.025)	0.839
Impatience	0.322 (0.022)	0.404 (0.033)	0.333 (0.031)	0.096
Perceived Temptation (range 0-10)	2.384 (0.089)	2.185 (0.112)	2.471 (0.121)	0.225
Faces strong financial claims from others*	0.393 (0.023)	0.388 (0.032)	0.386 (0.032)	0.987
Risk aversion (range 0-6)	4.225 (0.093)	4.636 (0.122)	4.132 (0.129)	0.010
Cognitive ability (range 0-5)	2.937 (0.059)	2.886 (0.089)	2.934 (0.096)	0.887
Financial literacy (range 0-5)	1.856 (0.047)	1.838 (0.068)	1.851 (0.069)	0.977
Existing savings account	0.468 (0.023)	0.465 (0.033)	0.425 (0.033)	0.518
Donated to charity in the last 12 months	0.396 (0.023)	0.386 (0.032)	0.452 (0.033)	0.284
Observations	457	228	228	913

Note: A starred variable indicates that the randomisation was stratified on this variable. Individuals were classified as facing *strong financial claims from others* if they reported strictly above-median values for the financial requests they would face from relatives, friends and neighbours in a hypothetical scenario where they kept cash at home. *Existing savings account* is an indicator for holding a bank account with any bank at the time of marketing. *Real rewards*, *impatience*, and *present bias* refer to the time-preference elicitation via multiple price lists: *Real rewards* is an indicator for receiving monetary incentives. Individuals are defined as *present biased* if the reward needed to make them wait for one month is larger in the present than in the future. They are classified as *impatient* if they always chose the earlier reward in all time preference questions. *Perceived temptation* is the difference between individuals' 'tempted' and their 'ideal' allocation of 10 restaurant nights across 2 years, described in Section 4. Risk aversion represents the individual's choice from a set of six lotteries with increasing expected value and variance, where the 'no risk' option yields a score of 6. *Cognitive ability* is the number of correct answers (0-5) from a culture-free intelligence test. *Financial literacy* is the number of correct answers (0-5) to basic numeracy questions.

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 Giandom product. Treatment effects can be estimated using the equation

$$\Delta Y_i = \alpha_0 + \alpha_R R_i + \alpha_W W_i + \epsilon_i \quad (6)$$

where ΔY_i denotes the change in the outcome variable of interest. In Section 6.1, I focus on bank savings, but also provide estimates for total savings, loan demand, and expenditures (see Figure 2). 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 Giandom clients, it is the sum of their Giandom 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 survey-based measure of other savings, 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 6 provides $\hat{\alpha}_R$ and $\hat{\alpha}_W$ – estimates of the Intent-to-Treat effects of the regular-installment product ASA and the withdrawal-restriction product Giandom. 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, 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.⁴⁸ 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. Alternatively, equation 6 can be estimated using an instrumental variables approach, with take-up (ASA_i and $Giandom_i$) as the regressors and assignment to treatment (R_i and W_i) as orthogonal

⁴⁸See Imbens and Angrist (1994) and Duflo et al. (2007) for a discussion on ITTs and local average treatment effects.

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 4. In addition, all binary choice regressions contain marketer fixed effects. This is to filter any noise from differences in marketer ability.

6 Results

6.1 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.⁴⁹ 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 II 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 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

⁴⁹All 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 II: 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 R ²	0.02	0.4992	0.1922	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. *Change in Bank Savings* captures 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 period of observation starts with the baseline survey and ends with the goal date specified in an individual's personal savings plan (median duration: 137 days). *Purchased Savings Goal* is an indicator for whether the individual reported having purchased or paid for the savings purpose specified in their savings plan at the time of the endline survey. *Change in Other Savings* is the change in self-reported savings at home and at other banks. Survey-based savings data are truncated at 1 percent.

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 II, 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 5). The TOT regression suggests that taking up the regular-installment 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 II presents the results from a probit estimation of whether individuals purchased the savings goal (see Table IX) 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.⁵⁰ Note that this is a distinct 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

⁵⁰The survey team was informed about this savings purpose, in case individuals had forgotten.

made a savings plan during the marketing stage and were b) reached by the endline survey.⁵¹ 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 effect of treatment on the likelihood of borrowing for the purchase (conditional on purchase). Borrowing was not uncommon: Slightly below 20 percent of individuals chose loans or family borrowing as a means of affording the expenditure.

Table II 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: Converting the probit coefficients into marginal effects, assignment to group W increases the likelihood of borrowing by 19.6 percentage points (from 11.4 to 31 percent). In comparison, assignment to group R (being offered ASA) increased the likelihood of obtaining one's savings goal, but did not significantly affect the probability of borrowing for it. 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 2 (Appendix A) 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.⁵²

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 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 II 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.⁵³ 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

⁵¹Both 'having a savings plan' and 'being reached by the endline survey' are orthogonal to assignment to treatment.

⁵²The 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.

⁵³To 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.

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.

6.2 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,⁵⁴ 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 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 IV), 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:⁵⁵ (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 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 1). Approximating transactions with weeks (85 percent of clients chose weekly instalments), Figure 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 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

⁵⁴ One transaction can cover several weeks’ instalments.

⁵⁵Strictly speaking, this assumes that an individual would not take up a contract if he knows that default is certain.

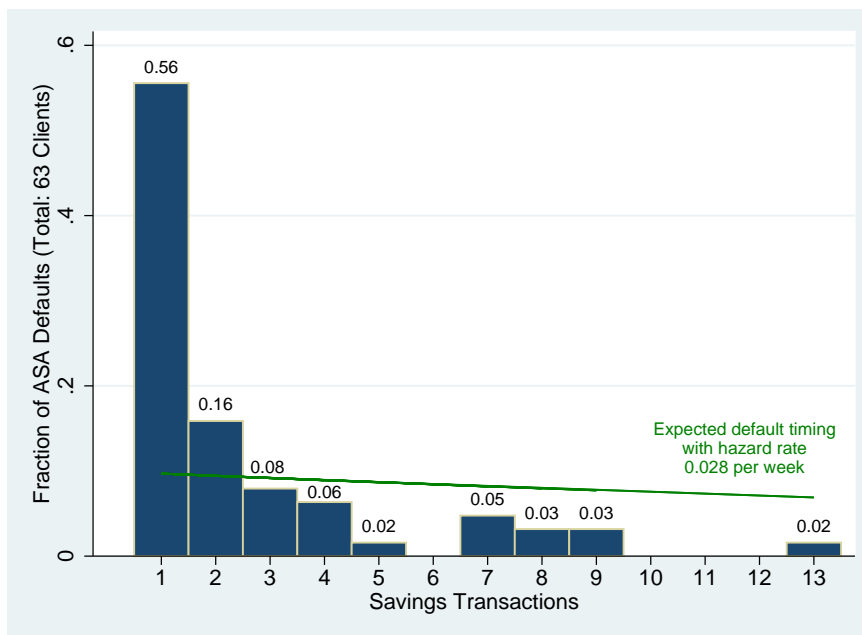


Figure 1: Savings Transactions: Defaulted ASA Clients

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 7.2. Heterogeneous hazard rates are discussed in Section 7.1.

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. Section 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 penalties 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 3 (Appendix A) 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 IV).

TABLE III: 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 IV: ASA PRE-ORDER

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

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 savings 141 pesos).⁵⁶ 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

⁵⁶Two accounts were closed after reaching the goal amount, and another two were closed after the bank mistakenly treated them as date-based accounts.

these clients, one reached the goal amount specified within their savings plan. Summary statistics for transactions in all accounts can be found in Table III.

6.3 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.

6.3.1 Predicting Take-Up of the Commitment Savings Products

Columns (1) to (3) in Table V 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.⁵⁷ 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 4) and a measure of cognitive ability (see Figure 5 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 4 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 $\beta < \tilde{\beta} < 1$) 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 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.

⁵⁷The 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 V: PREDICTING DEMAND FOR COMMITMENT (PROBIT)

Commitment Take-Up	Regular Saver (ASA)			Withdrawal Restriction (Gihandom)		
	(1)	(2)	(3)	(4)	(5)	(6)
Age (yrs)	-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)
Weekly HH income (in thousand pesos)	0.0018 (0.0081)	0.0013 (0.0080)	-0.0028 (0.0079)	0.0003 (0.0118)	-0.0054 (0.0108)	-0.0052 (0.0109)
No. of appliances owned	0.0007 (0.0145)	-0.0048 (0.0143)	-0.0084 (0.0143)	0.0220 (0.0224)	0.0253 (0.0213)	0.0258 (0.0216)
No. of 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 (0-10)	-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)
Faces strong financial claims from others	-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 (0-6)		-0.0049 (0.0106)	-0.0059 (0.0105)		0.0497*** (0.0167)	0.0500*** (0.0168)
Cognitive Ability (0-5)		0.0353* (0.0188)	0.0363* (0.0187)		0.0157 (0.0236)	0.0174 (0.0239)
Financial Literacy (0-5)		0.0425* (0.0246)	0.0328 (0.0250)		-0.0115 (0.0322)	-0.0120 (0.0321)
HH Bargaining Power (0-5)		0.0063 (0.0113)	0.0053 (0.0113)		0.0444*** (0.0164)	0.0456*** (0.0167)
Distance to Bank (km)			-0.0265 (0.0207)			0.0084 (0.0260)
Existing Savings Account			0.0998** (0.0444)			-0.0022 (0.0667)
Donated 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. Variables are as described in Table I.

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 4), 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 V 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 accessing 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.

6.3.2 Predicting Default and Repeat Take-Up

Table VI 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 V) has been added for comparison. In addition to the regressors from the take-up regressions, Table VI also

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

Dependent Variable	ASA Take-Up	Default (R-Sample)	Default (takeup-Sample) Heckit Estimation	Pre-Order (takeup-Sample) Heckit Estimation
Age (yrs)	-0.0003 (0.0020)	-0.0024 (0.0016)	-0.0024 (0.0019)	-0.0007 (0.0017)
Female	0.0536 (0.0868)	0.1189 (0.0911)	0.1423 (0.1078)	-0.0028 (0.0552)
Married	0.0165 (0.0630)	0.0064 (0.0536)		
Weekly HH income (1000 pesos)	-0.0028 (0.0079)	0.0034 (0.0057)	0.0035 (0.0077)	0.0075 (0.0078)
No. of appliances owned	-0.0084 (0.0143)	-0.0093 (0.0115)		
No. of HH members	0.0130 (0.0095)	0.0123 (0.0077)	0.0135 (0.0094)	-0.0121* (0.0062)
Education (yrs)	-0.0094 (0.0064)	-0.0030 (0.0053)	-0.0051 (0.0069)	-0.0005 (0.0045)
Present Bias	0.0827 (0.0864)	0.1119* (0.0654)	0.2082** (0.0839)	-0.1999*** (0.0633)
Soph. Present Bias (Pres.Bias*Temptation)	-0.0631** (0.0292)	-0.0453** (0.0230)	-0.1026** (0.0441)	0.1074** (0.0543)
Perceived Temptation (0-10)	-0.0046 (0.0125)	-0.0202* (0.0105)	-0.0227 (0.0143)	0.0022 (0.0126)
Impatience	-0.0008 (0.0464)	-0.0030 (0.0372)		
Faces strong financial claims from others	-0.0038 (0.0414)	-0.0113 (0.0330)		
Risk Aversion (0-6)	-0.0059 (0.0105)	-0.0181** (0.0084)	-0.0267** (0.0136)	0.0074 (0.0110)
Cognitive Ability (0-5)	0.0363* (0.0187)	0.0365** (0.0143)	0.0344* (0.0183)	-0.0197 (0.0175)
Financial Literacy (0-5)	0.0328 (0.0250)	-0.0168 (0.0204)		
HH Bargaining Power (0-5)	0.0053 (0.0113)	-0.0116 (0.0090)	-0.0194 (0.0123)	0.0176* (0.0091)
Distance to Bank (km)	-0.0265 (0.0207)	-0.0115 (0.0165)	-0.0098 (0.0262)	0.0349 (0.0226)
Existing Savings Account	0.0998** (0.0444)	0.0296 (0.0363)		
Donated to Charity	0.0221 (0.0424)	0.0422 (0.0332)	0.0414 (0.0460)	0.0357 (0.0370)
#Emergencies last yr	-0.0161 (0.0277)	0.0005 (0.0213)		
#Emergencies since baseline		-0.0033 (0.0182)	0.0333 (0.0358)	0.0035 (0.0204)
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. Variables are as described in Table I.

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) in Table VI can be understood as an analysis of which type of individuals took up the commitment product ‘by mistake’, proxied by take-up and subsequent default (this interpretation abstracts from the possibility of rational default). The results provide further support to the partial sophistication hypothesis: Among those randomly assigned to group R, 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 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 don’t select into the product (if the effective penalty would be prohibitively high), or they make sure to choose a contract which is incentive-compatible for their preferences (which could be through the size of the weekly instalment, or the size of the penalty). Furthermore, the temptation measure has some predictive power on its own, even when not interacted with present bias. Expanding on the discussion from Section 4, 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.

To gain a more direct interpretation of the factors predicting default, column (3) analyses default occurrence in the take-up sample, and uses the Heckit method to correct for selection bias. The selection equation is given by the take-up regression in column (1). The outcome equation needs an exclusion restriction: Collinearity between the regressors and the inverse Mills ratio prevents the maximum likelihood estimation from converging, and calls for a reduced set of regressors. This is achieved by restricting the set of regressors to those which are fundamental to the analysis (such as income), or which test an important hypothesis (such as distance to the bank as a proxy for transaction costs). Regressors may be dropped if they predict selection into the sample (take-up), but not default. While their effect on take-up varies, none of the seven dropped regressors predict default in the R-group sample, with all z-statistics close to zero. An illustrative example is ‘Emergencies last year’: The amount of shocks to income and expenditures that an individual suffered in the year before the observation period proxies the shock hazard rate, and may be expected to predict whether an individual is willing to take up a commitment product. However, defaulting on said commitment product is more likely to depend on the realisation of shocks during the savings period, which is directly captured in ‘Emergencies since baseline’ (elicited during an independent endline survey).

Looking at the results in column (3), the marginal coefficient on present bias has doubled, and become more significant. Interestingly, the link between present bias (as proxied by time-inconsistency in

MPLs) and default seems to be 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). Moving on to the pre-order (repeat take-up) decision, the effects of present bias (-0.20^{***}) and sophistication (0.11^{**}) are now 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. In other words, they 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 imperfect mapping is supported by the data, see Table IV).

Several other factors can help in explaining the observed default rates. The most obvious candidate - the occurrence of shocks during the savings period - finds weak support in the take-up sample ('Emergencies since baseline' has the expected positive sign, albeit being insignificant). Shock occurrence was estimated by asking for common income or consumption emergencies during the endline survey.⁵⁸ A positive relationship between defaults and shocks, in combination with the fact that 45 percent of clients completed their ASA contract successfully, would indicate 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 a near-zero coefficient on emergencies in column (4) of Table VI. Other factors predicting default include risk aversion (-) and cognitive ability (+). 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. More generally, cognitive ability may reflect predictability and rationalisability of behaviour: An individual with high cognitive skills is more likely to realise that 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 4). If pref-

⁵⁸There 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.

erences 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 worth mentioning that household bargaining power is borderline significant ($p = 0.115$) for default, and significantly predicts repeat take-up. A possible explanation is that intra-household conflicts initially 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 positive association of household bargaining power with the pre-order decision provides 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.

6.3.3 Heterogeneous Treatment Effects

Table VII examines treatment effect heterogeneity across a number of dimensions of interest. The regression set-up is identical to that in column (1) of Table II: 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.

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.⁵⁹ 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 (see Section 2). After taking into account the default penalty,

⁵⁹It 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.

TABLE VII: HETEROGENEOUS TREATMENT EFFECTS: CHANGE IN BANK SAVINGS

	(1)	(2)	(3)	(4)	(5)
Regular Saver (R)	707.175*** (257.145)	481.493*** (134.955)	287.405*** (63.639)	455.476*** (133.291)	351.944*** (114.122)
Withdrawal Rest. (W)	129.494*** (40.626)	146.343*** (39.698)	147.299*** (40.732)	154.503*** (42.628)	148.877*** (41.430)
R * Present bias	-58.840 (543.374)				
Present Bias	57.671 (83.915)				
R * Soph. Present Bias (R*Pres.Bias*Temptation)	20.760 (81.973)				
Soph. Present Bias	-6.042 (18.188)				
R * Perceived Temptation	-48.836 (62.291)				
Perceived Temptation	-4.94 (9.286)				
R * High Cognitive Ability		261.295 (291.074)			
High Cognitive Ability		-30.709 (36.453)			
R * Ex. Savings Account			621.5350** (269.7281)		
Existing Savings Account			74.9955* (41.3577)		
R * Weekly HH income				44.6249 (56.0615)	
Weekly HH income				7.6109 (5.4829)	
R * HH bargaining power					88.1604* (50.5904)
HH bargaining power					6.9664 (10.3868)
Constant	34.948 (26.710)	41.148** (17.922)	-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. Variables are as described in Table I.

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 VI. 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 VII 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 V). Thus, a lower percentage of sophisticated agents were 'treated'. Theoretical arguments in Section 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 cognitive 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 6.3.2). 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).

7 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.

7.1 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.

TABLE VIII: INCOME OPTIMISM

	Not Present-Biased	Present-biased	All	T-stat P-value
Prediction Gap (growth)	3.290378 (0.6976)	3.677686 (1.2298)	3.357041 (0.6146)	0.81
Prediction Gap (level)	1.269759 (1.1704)	-2.22314 (2.6292)	0.6685633 (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.

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.

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

come 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.⁶⁰ 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 VIII 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 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 (discussed in Section 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

⁶⁰For 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).

sophistication are *negatively* associated with take-up and default. In this sense, partial sophistication provides a parsimonious explanation that is consistent with the evidence.

7.2 Aggregate Shocks

Idiosyncratic and independent shocks are unlikely to cause the default timing pattern apparent in Figure 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 in December 2011, causing 1,268 casualties (more than half of them in Cagayan de Oro, a city 126km west of the study location).⁶¹ 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 hit the Mindanao region between December 2 and December 9, 2012. Fortunately, 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 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.⁶² An additional 15 clients made one or two deposits after opening (see Figure 1 for the distribution of transactions). By the time of the storm in early December, most of the contract defaults had already occurred.

7.3 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 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

⁶¹Statistics from the Philippine National Disaster Risk Reduction and Management Council (NDRRMC).

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

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 V).

8 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. 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 takes up a commitment product chooses very low stakes for this commitment, and then defaults 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 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 to purchase 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 Gihandom 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.

References

- Acland, Dan and Matthew R. Levy**, "Naiveté, Projection Bias, and Habit Formation in Gym Attendance," *Management Science*, 2015, 61(1), 146–160.
- 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 Intrahousehold Resource Allocation," *Quarterly Journal of Economics*, 2002, 117(3), 963–995.
- 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.
- Banerjee, Abhijit V. and Sendhil Mullainathan**, "Limited Attention and Income Distribution," *American Economic Review*, 2008, 98(2), 489–493.
- and —, "The Shape of Temptation: Implications for the Economic Lives of the Poor," Working Paper, National Bureau of Economic Research 2010.
- 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.
- Bénabou, Roland and Jean Tirole**, "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 2002, 117(3), 871–915.
- and —, "Willpower and Personal Rules," *Journal of Political Economy*, 2004, 112(4), 848–886.
- 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.

- Besley, Timothy**, "Savings, Credit and Insurance," in Jere Behrman and T.N. Srinivasan, eds., *Handbook of Development Economics*, Vol. 3A, Elsevier, 1995.
- , **Stephen Coate, and Glenn Loury**, "The Economics of Rotating Savings and Credit Associations," *American Economic Review*, 1993, 83(4), 792–810.
- 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," Working Paper, The Wharton School, University of Pennsylvania 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.
- Cornelisse, Sandra, Vanessa van Ast, Johannes Haushofer, Maayke Seinstra, and Marian Joels**, "Time-Dependent Effect of Hydrocortisone Administration on Intertemporal Choice," Technical Report, Available at SSRN: <http://ssrn.com/abstract=2294189> 2013.
- 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, 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**, "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya," *American Economic Journal: Applied Economics*, 2013, 5(1), 163–192.
- , **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.
- 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 18065, 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.
- Imbens, Guido W and Joshua D Angrist**, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 1994, 62(2), 467–475.
- John, Anett**, "Just A Few Cents Each Day: Can Fixed Regular Deposits Overcome Savings Constraints?," EOPP Discussion Papers 51, London School of Economics 2014.
- 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.
- Köszegi, Botond**, "Ego Utility, Overconfidence, and Task Choice," *Journal of the European Economic Association*, 2006, 4(4), 673–707.
- 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.
- , **Ted O'Donoghue, and Matthew Rabin**, "Projection Bias in Predicting Future Utility," *The Quarterly Journal of Economics*, 2003, 118(4), 1209–1248.
- Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao**, "Poverty Impedes Cognitive Function," *Science*, 2013, 341(6149), 976–980.
- McClure, Samuel M, David I Laibson, George Loewenstein, and Jonathan D Cohen**, "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science*, 2004, 306(5695), 503–507.
- Möbius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat**, "Managing Self-Confidence: Theory and Experimental Evidence," Working Paper 17014, National Bureau of Economic Research 2011.
- O'Donoghue, Ted and Matthew Rabin**, "Doing It Now or Later," *American Economic Review*, 1999, 89(1), 103–124.

- Sayman, Serdar and Ayse Öncüler**, “An Investigation of Time Inconsistency,” *Management Science*, 2009, 55(3), 470–482.
- Shah, Anuj K., Sendhil Mullainathan, and Eldar Shafir**, “Some Consequences of Having Too Little,” *Science*, 2012, 338(6107), 682–685.
- Shefrin, Hersch M and Richard H Thaler**, “The Behavioral Life-Cycle Hypothesis,” *Economic Inquiry*, 1988, 26(4), 609–643.
- Shui, Haiyan and Lawrence M. Ausubel**, “Time Inconsistency in the Credit Card Market,” Working Paper, University of Maryland 2005.
- Sinn, Miriam**, “Risk and Time-Inconsistency: Evidence from a Field Experiment in West Bengal,” Working Paper, London School of Economics 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.
- Thaler, Richard**, “Mental Accounting and Consumer Choice,” *Marketing Science*, 1985, 4(3), 199–214.

A Appendix: Supplementary Figures and Tables

TABLE IX: 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 X: 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%



Figure 2: Distributional Effect of Treatment on Savings, Loans and Expenditures

TABLE XI: 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.

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 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 6.2).

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).




Figure 3: Termination Fees (Chosen & Charged)

B Appendix: Robustness Checks

This section tests robustness along several dimensions. Table XII verifies that the estimation of average treatment effects is robust to the inclusion of unbalanced covariates (see Table I). Table XIII tests robustness of the take-up and default regressions of Section 6.3 with respect to the measurement of sophistication. As outlined in Section 4, 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 XIII use “Perceived Self-Control” to capture perception of time-inconsistency (both measures are discussed in Section 4). 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 XII: AVERAGE TREATMENT EFFECTS - ROBUSTNESS TO UNBALANCED COVARIATES

Dependent Variable	Change in Bank Savings		Change in Other Savings		Change in Outstanding Loans	
	(1)	(2)	(3)	(4)	(5)	(6)
Regular Saver	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)
Treatment (ASA)						
Withdrawal	148.243***	163.561***	-328.159	-613.247	-308.549	-678.111
Restr.						
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. Control variables are defined as in Table I. 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 XIII: REGULAR SAVER (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. Columns (4) and (8) instead use *expected* — *ideal* * *presentbias* (“Pres.Bias*Self-Control”) as an alternative measure of sophistication. See Section 4 for a detailed description of the sophistication measures.

TABLE XIV: 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. Variables are as defined in Table I. Columns (2) and (5) regress the dependent variable on uncentrified 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 VI.

TABLE XV: 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, $T = 2$).

Table XIV looks at the effect of using real incentives instead of hypothetical questions in the measurement of time-inconsistency. Section 4 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 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 XIV 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 6.3 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 XIV 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 (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 XV 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, more impatience), but remains statistically insignificant. In combination with the strong predictive power for commitment take-up and default observed in Table XIV, these results are consistent with the idea that incentivising survey questions reduces noise and improves the quality of the answers.

C Appendix: 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.⁶³ 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 4). 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 Figure 5 for a sample question). A financial literacy score is given by the number of correct answers (again, out of five possible) to basic numeracy

⁶³The 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.

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?




		
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 4: Test of Risk Aversion (Methodology: Binswanger (1980))





Figure 5: Illustration: Test of Cognitive Ability



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 6: Personal Savings Plan (All Treatment Groups)

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').