When Commitment Fails – Evidence from a Field Experiment*

ANETT JOHN†

September 2016

Commitment products can remedy self-control problems. However, imperfect knowledge about their preferences may (discontinuously) lead individuals to select into incentive-incompatible commitments, which reduce their welfare. I conduct a field experiment in the Philippines ($N = 913$), where low-income individuals were randomly offered a regular-instalment commitment savings product. Individuals chose a personalised savings plan and a default penalty themselves. A majority appears to choose a harmful contract: While the average effect on bank savings is large, 55 percent of clients default, and incur monetary losses. A possible explanation is that the chosen penalties 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, hyperbolic discounting, partial sophistication, microsavings
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. 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 to Nageeb Ali, Douglas Bernheim, Andrea Canidio, Kristina Czura, Stefano DellaVigna, Jonathan de Quidt, Paul Heidhues, Alex Imas, Supreet Kaur, David Laibson, Matthew Levy, Lance Lochner, George Loewenstein, Johannes Spinnewijn, Betty Wilkinson, 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, 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, or not keeping chocolate in the house. In developing countries, documented demand for commitment devices goes back to the literature on rotating savings and credit organisations (ROSCAs), the wandering deposit collectors of South Asia and Africa, and more recent studies on newly introduced commitment savings products.[3]

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. Among the most frequently cited models to rationalise the observed demand for commitment are those of quasi-hyperbolic discounting (Laibson (1997) and O'Donoghue and Rabin (1999)).[4] They 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, agents 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 a long-run perspective. Empirically, commitment devices have been shown to increase savings levels (Ashraf et al. (2006)), 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)), as well as to reduce heavy daytime drinking (Schilbach (2015)).

But are people good at choosing the 'right' commitment contract? By construction, correctly choosing a welfare-improving contract requires knowledge about one's future preferences: To determine whether a contract will enable her to follow through with a plan, the agent needs to anticipate

[1] This paper focuses purely on self-commitment. It does not address commitment contracts adopted with strategic motives with respect to others, or 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.

[2] See DellaVigna and Malmendier (2006) for gym memberships, Benartzi and Thaler (2004) for 401(k) pension savings, and Ariely and Wertenbroch (2002) for academic assignments. For an overview of commitment devices, see Bryan et al. (2010).

[3] On ROSCAs, see Anderson and Baland (2002), or Gugerty (2007). See e.g. Besley (1995) on West Africa's susu collectors. Ashraf et al. (2006), Brune et al. (2011) and Dupas and Robinson (2013) study withdrawal-restriction savings accounts. Duflo et al. (2011) study commitments to use fertilizer.

[4] Models of temptation and self-control (Gul and Pesendorfer (2001)) and dual-self models (Fudenberg and Levine (2006)) also predict a demand for commitment, with similar implications for observed behaviour.

how her future selves will behave under the contract. Telling her social network that she plans on running a marathon, for instance, requires the agent to assess whether the threat of embarrassment upon withdrawal is sufficient to overcome her laziness. If the agent is overconfident, the contract may result in undesirable behaviour (she does not run the marathon), and she may be harmed rather than helped (she gains no health benefits, but pays the cost of embarrassment). The very nature of most commitment contracts is to impose monetary or social penalties for undesirable behaviour.[5] Thus, adopting a commitment device that is ill-suited to one's preferences may backfire and become a threat to welfare.

This paper argues that commitment can be harmful if agents select into the wrong commitment contract - and that they frequently do. In a framework where a penalty is conditional on a binary savings decision, I first outline theoretically why offering commitment improves the welfare of fully sophisticated agents, but is likely to harm partially sophisticated agents. Given a continuous penalty, this effect is 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,[6] as naive agents underestimate the cost of effective commitment. Second, I conduct a randomised experiment where individuals could sign up for a new commitment savings account with fixed regular instalments. Adopters 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 three 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. This is 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 (described in Section 4), risk aversion, financial claims from others, cognitive ability, financial literacy, intrahousehold bargaining power, household demographics, and measures of saving, borrowing, and household

[5] Examples of monetary penalties include any type of commitment contract with front-loaded fees, such as retirement savings products with acquisition or management costs. Such fees generate a 'J curve' in the asset value, resulting in high negative returns if the contract is cancelled during the initial years. A similar argument can be made for front-loaded gym membership costs.

[6] Conditional on an awareness that one cannot save without commitment. The latter rules out perfect naiveté.

expenditures. After the baseline survey, all individuals were provided with a marketing treatment, which included a personalised savings plan for an upcoming expenditure and a free non-commitment savings account with 100 pesos (U.S. $2.50) opening balance.[7] 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.[8] The interest rate was equal to the standard market interest rate.[9]

As a control treatment, 25 percent of the sample (the 'Withdrawal Restriction' group) were offered the commitment savings account studied in Ashraf et al. (2006), Giné et al. (2010), and Brune et al. (2011): The 'Gihandom' account (Visayan for 'dream') allowed individuals to restrict withdrawals before either the goal date or the goal amount from their savings plan 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 the control group, none of the savings plan features were binding. Since individuals' expenditures were due at different times, the outcome of interest are individuals' savings at the time of their goal date. The study concluded with a comprehensive endline survey, as well as a repeat marketing stage where ASA clients could opt to 'pre-order' the product for a second round.

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, despite the prior universal provision of free standard savings accounts. 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 429 pesos (U.S. $10.20, ITT) relative to the control group. The corresponding effect for the Withdrawal Restriction group was 148 pesos (U.S. $3.50).[10] 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, these averages hide 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 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.

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

[8] The concept is similar to the Stickk initiative (www.stickk.com), applied to a development context.

[9] As of September 2012, the interest rate was 1.5 percent p.a. for all offered accounts. It decreased to 1 percent in November 2012. The inflation rate for 2012 was 3.1 percent.

[10] The ITT measures the effect of being *offered* the product, and already accounts for any penalty charges. An increase of 429 pesos (148 pesos) corresponds to 20 percent (7 percent) of median weekly household income.

For those who had chosen to make their goal amount binding (45 percent), this meant their initial savings were tied up indefinitely, or until the bank would exhaust their account with dormancy fees.[11]

I show that ASA completion is strongly bi-modal, in that most 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. Using measures of time-inconsistency and sophistication from the baseline survey, 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: For a given level of time-inconsistency, more sophisticated agents are *less* likely to adopt commitment, but also less likely to default, conditional on take-up. This is consistent with the theory: Sophistication increases both the perceived cost of commitment and the agent's ability to choose an incentive-compatible contract. The conjecture that partial sophisticates chose 'weak commitments' 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).

This paper builds and expands on the literature in several ways. To the author's knowledge, it is among the first papers to explicitly discuss heterogeneous effects (and possible welfare risks) of commitment contracts, and the first study to link such effects 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. The study further relates to Heidhues and Kőszegi (2009), who show theoretically that commitment is likely to reduce welfare when agents are partially sophisticated. Their model differs from the one presented in this paper in some key assumptions and welfare implications: Heidhues and Kőszegi (2009) assume that commitment is desirable at all levels of sophistication, and derive that welfare is lowest just below full sophistication, as agents commit to high penalties which remain insufficient to deter consumption. In my model, commitment adoption is endogenous, and decreases in the amount of sophistication. The cost of commitment endogenously results from the presence of uncertainty, as unexpected shocks force the agent to default. The negative welfare effect of large but insufficient penalties is compensated by the probability that these large penalties are prohibitive for adopting commitment. Consequently, the group most at risk of negative welfare effects are near-naifs, who widely adopt weak commitments and subsequently default. Other related theoretical contributions are DellaVigna and Malmendier (2004), Heidhues and Kőszegi (2010) and Eliaz and Spiegler (2006), who address supply-side responses to partial sophistication.

In the realm of commitment savings, the literature has largely focused on positive average effects. I suggest that these effects may be very heterogeneous, including the possibility of a majority being hurt by the product. My results are consistent with previous findings: Ashraf et al. (2006) find that a withdrawal-restriction product increased savings by 81 percent on average, but 50 percent of the clients made no further deposits after the opening balance. Out of 62 clients who selected an amount goal, only six reached this goal within a year. Giné et al. (2010) offered smokers in the Philippines a commitment

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

savings contract for smoking cessation, in which participants forfeit their savings if they fail a nicotine test after 6 months. Offering the contract increased the likelihood of smoking cessation by 3 percentage points. However, 66 percent of smokers who took up the product failed the nicotine test, forfeiting an average of 277 pesos in savings. In summary, a closer look at the heterogeneity behind average treatment effects may reveal that adverse effects of commitment products are 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. 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 a one-off investment (Ananth et al. (2007), Rutherford (2000)). 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.[12] 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 (2015).[13] If a significant share of the demand for loans is a demand for commitment to fixed instalments, then we should 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. The Gihandom account used in this study is identical to Ashraf et al. (2006)'s SEED, and their effect on savings replicates within a 4 percent margin.[14] The estimated effect of the regular-instalment product on bank savings is almost three times higher, consistent with the theoretical work of Amador et al. (2006): When individuals value both commitment and flexibility, the optimal contract involves a minimum (per-period) savings requirement.

Third, the paper proposes a simple measure of sophistication that is practicable in a field setting in a developing country. It consists of an interaction of two established measures – one of observed time-inconsistency, and one of self-perceived temptation. 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. Such a mapping does not allow for the possibility that individuals may take up commitment products by mistake. Recent advances with prediction tasks (Augenblick and Rabin (2015), Acland and Levy (2015)) are promising, but challenging to implement in a field context.

[12] ROSCAs are inflexible to an individual's needs, and not widely available in the study region. Deposit collectors usually do not commit individuals to fixed instalments, and social sanctions can be avoided using small deposits or convenient excuses.

[13] See also Afzal et al. (2015), who give participants a direct savings vs. loan choice in a lab-in-the-field setting.

[14] Ashraf et al. (2006)'s SEED increases savings by 411 pesos (ITT) after 12 months, while Gihandom increases savings by 148 pesos after 4.5 months. Monthly equivalents are 34.25 and 32.88 pesos, respectively.

The paper proceeds as follows. Section 2 outlines a model of commitment under partial sophistication. Section 3 describes the experimental design. Section 4 explains the survey instrument, focusing on the measurement of time-inconsistency and sophistication. Section 5 outlines the empirical strategy. Section 6 presents empirical results. Section 7 discusses alternative explanations. Section 8 concludes and discusses policy implications.

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.

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, 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-instalment product, at the expense of tractability. A companion paper, John (2015), does this for the case of full sophistication.

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 planning period 0, when the agent may choose to adopt commitment (Section 2.3). The agent receives a per period income of $y_t = 1$, which she can either consume or save. She cannot borrow. Throughout, assume the interest rate is $R = 1$ and $\delta = 1$ for simplicity. Define s_1 as the amount of savings that she 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*: She 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 her present bias is captured in the parameter $\tilde{\beta} \in [\beta, 1]$, which she believes she will use in all future periods. In particular, the agent believes in period t that her 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 she 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 her 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.[15] 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 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.[16]

Following 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.

2.2. No-Commitment Equilibrium

The model is easily solved using backward induction. In period 2, the agent will buy the nondivisible whenever she can afford it, i.e., whenever $y_2+s_1\geq p$. 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 insufficient savings $s_1<p-1$. The consumption profile is $c_2=y_2+s_1+(b-p)\cdot 1(y_2+s_1\geq p)$.

In period 1, the agent realizes the good will be bought if and only if she sends $s_1\geq p-1$, and absent shocks. She 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.*[17](All proofs are in Appendix II.)

The threshold $\hat{\beta}\equiv\frac{p-1}{\lambda(p-1)+(1-\lambda)(b-1)}$ has an intuitive interpretation: It is the ratio of the cost of saving today, $p-1$, to its expected benefit tomorrow, $\lambda(p-1)+(1-\lambda)(b-1)$. Because $b>p$, a time-consistent agent will always save (i.e., $\hat{\beta}<1$). Note that the no-commitment equilibrium is invariant to the degree of sophistication (see footnote 20).

2.3. Introducing Commitment

For agents with high levels of time-inconsistency, the no-commitment equilibrium is inefficient: Saving for the nondivisible good is welfare-improving regardless of β, since the period 0 planner

[15]Suppose the sudden illness of a family member changes preferences such that utility stays unchanged if a hospital visit (at cost 1) is consumed and paid for, and drops to $u(c)=-\infty$ without a hospital visit.

[16]Time-consistent explanations to abandon savings plans include state-dependent preferences.

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

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. Suppose the individual is now given the possibility to commit to save, enforced via a penalty for non-compliance which the individual chooses herself ex-ante. While not the only form of commitment (see the discussion in Section 8), 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 5).

In the current setting, individuals commit to a self-chosen default penalty, imposed if they fail to follow a regular-instalment savings plan. Applied to a simple two-period model, the period 0 agent can choose a penalty D that is imposed on her in period 1 if she fails to meet the savings requirement, $s_1 \geq p-1$. Given D, the period 1 agent now prefers to save iff

$$1-(p-1)+\beta[\lambda(p-1)+(1-\lambda)b] \geq 1-D+\beta(1-\lambda). \tag{1}$$

Equivalently, she 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.

$$D \geq D_{min}(\beta) \equiv \underbrace{(p-1)}_{cost\,today} - \beta \underbrace{[\lambda(p-1)+(1-\lambda)(b-1)]}_{benefit\,tomorrow}. \tag{2}$$

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.*

A note on enforceability: The model assumes that the penalty D is fully enforceable, even if the agent loses her income to a shock, and is consequently unable to save. In a two-period model, full enforceability is required to prevent the agent from adopting infinitely large penalties, reasoning that she 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. Second, D may represent a discounted stream of future losses, such as those from losing access to microfinance. Third, in settings with multiple periods, 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.[18] This creates a cost of commitment, and a demand for flexibility.[19]

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

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

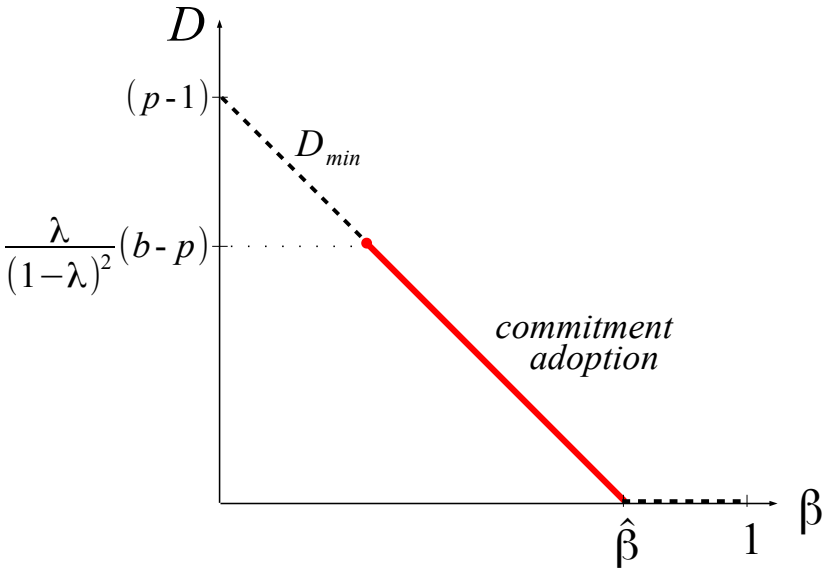

FIGURE 1. COMMITMENT TAKE-UP DECISION (FULL SOPHISTICATION)

Figure 1 illustrates the relationship between time-inconsistency β and commitment adoption D assuming full sophistication, $\tilde{\beta} = \beta$. The dotted line is the minimum default penalty that is effective in enforcing the savings plan, D_{min}. The solid line indicates the range of β for which D_{min} is adopted. Generalizing to partial sophistication ($\beta \leq \tilde{\beta} \leq 1$), β and D_{min} are replaced by $\tilde{\beta}$ and $\tilde{D}_{min}$: Perceived values rather than actual values determine the choice of commitment. Thus, for those unable to save in autarky ($\beta \leq \tilde{\beta} < \hat{\beta}$), commitment is attractive for high $\tilde{\beta}$, and prohibitively expensive for low $\tilde{\beta}$.

2.4. Commitment Choice in Equilibrium

The commitment adoption decision has two steps: The period 0 planner first identifies which commitment contract (which penalty) will enable her to save while maximising flexibility. In a second step, she decides whether to adopt this contract. Commitment is costly: If a shock hits in period 1, the agent loses her 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-inconsistency. The equilibrium behaviour with full sophistication is summarized in Proposition 3. Figure 1 illustrates the relationship between time-inconsistency and commitment adoption.

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 expected cost of commitment due to 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).*

A key intuition is that period 1's incentive constraint 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 $D = 0$. Further note that the ex-ante benefit of commitment, $(1-\lambda)^2(b-p)$, does not depend on the time-consistency parameter β. In contrast, β determines the cost of commitment, λD_{min}. Perhaps counter-intuitively, commitment is most attractive to those with the *lowest* degree of time-inconsistency (conditional on $\beta < \hat{\beta}$), as the penalty required to enforce the savings plan is small, and poses little risk in the presence of shocks. In consequence, agents adopt commitment for sufficiently high β.

2.5. The Effect of Partial Sophistication

A partially sophisticated agent believes that her future selves will discount the future at rate $\tilde{\beta} > \beta$ rather than at β. This corresponds to the classic pattern of procrastination, where an agent believes she 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 her period 1 self to save, resulting in default and welfare losses. The results are robust to more stochastic types of partial sophistication (see Section 2.6).

Proposition 4. *Equilibrium with Partial Sophistication: (a) Conditional on adopting commitment, partially sophisticated individuals 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.*

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.[20] Holding true β fixed,[21] a higher degree of naiveté $\tilde{\beta} - \beta$ implies that a lower penalty is regarded as effective, which decreases the perceived cost of commitment. Conditional on β, as well as on the agent's perceived inability to save without commitment ($\tilde{\beta} < \hat{\beta}$), adoption increases with naiveté.

[20] In 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 belief $\tilde{\beta}$ 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.

[21] Is is plausible to hold β fixed while $\tilde{\beta}$ increases? While this study is agnostic about the functional relationship between β and $\tilde{\beta}$, measures of both are uncorrelated in the current dataset (see Appendix Figure A3).

2.6. Discussion and Extensions

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. Appendix I.1 discusses the case of stochastic sophistication, and argues that the findings observed in the data are hard to reconcile with rational expectations about β: Comparing stochastic with full sophistication, commitment becomes less attractive due to the downside risk that a given penalty will not be effective. This leads to low commitment take-up, high conditional penalties, and low default rates. To reconcile stochastic beliefs about time-inconsistency with the observed low penalties and high default rates, one needs to allow for belief distributions which are systematically biased towards naiveté.

Is it plausible that individuals persistently hold incorrect beliefs about their time preferences, despite being able to observe their own past behaviour? Appendix I.2 discusses why Bayesian learning may fail (see also Ali (2011)). For instance, learning may be specific to context: An individual may realize from past observation whether she is able to save for the nondivisible good by herself (i.e., the inequality $\beta \geq \hat{\beta}$ is observed). However, she may be unfamiliar with her savings behaviour under commitment. Other impediments to Bayesian learning include self-serving beliefs and neuroscientific explanations (such as stress-induced time-inconsistency). Appendix I.3 discusses pessimistic beliefs ($\tilde{\beta} < \beta$) and overcommitment.

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. The bank agreed to offer both the regular-instalment product ASA and the withdrawal-restriction product Gihandom in two of their 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 branches, ASA and Gihandom constituted new product additions.

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 (usually the mother of the family). The baseline survey was completed with all such individuals who (i) had some form of identification,[22] (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

[22] Birth certificates, tax certificates, voter's ID, and several other substitutes were accepted.

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.[23] 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 Appendix Figure A6). Median savings goals were 2400 pesos across all groups (close to the 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 Appendix Table A1). The duration of savings plans was limited to 3–6 months (median: 137 days). In addition, everyone was offered a non-commitment savings account (henceforth called 'ordinary savings account'), containing a free 100 pesos opening balance, as a 'welcome gift' from the bank, along with an encouragement to use this account to save for the expenditure.[24]

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 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 Gihandom product. Up to the point of offering the commitment products, the marketing script was identical across groups R, W, and C.[25] Out of 852 marketing visits, 788 individuals 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.[26] In group W, 92 (out of 219 offered) accepted the Gihandom product. Appendix Table A2 summarizes the take-up results.

The regular-instalment product ASA committed clients to a fixed instalment plan with weekly (84 percent) or bi-weekly (16 percent) due dates. An account was considered in default from the day the client fell three instalments behind. In case of 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. The amount of the fee was chosen by the client upon signing the ASA contract, and donated to charity in case of default. Clients were

[23] A 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 groups R and C were reached slightly less often than individuals in group W.

[24] The 100 pesos constituted the minimum maintaining balance. Individuals were permitted to close the account, but incurred a 50 peso closing fee. During the period of observation, no client closed their account. This closing fee did not apply to successful ASA or Gihandom accounts. 18 clients had previously existing 1st Valley Bank savings accounts. For these clients, 100 pesos was deposited in their existing account instead of opening a new account.

[25] 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.

[26] One 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.

given a choice of three national (but not local) Philippine charities.[27] A variety of flexibility features allowed 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 A6). The possibility of making future deposits early effectively provided a form of insurance against uneven income streams. Withdrawals during the savings period were only possible by allowing default to occur.

Enforceability of the termination fee was facilitated through the account opening balance: To complete the opening of an ASA, clients had to deposit an opening balance equal to their first weekly instalment, but at least 150 pesos (250 pesos) for savings goals below (above) 2500 pesos. Using the same threshold, clients could choose a termination fee as high as they liked, but no lower than a minimum of 150 pesos (250 pesos). Consequently, the minimum termination fee could always be enforced. Higher termination fees could be enforced only if the client continued to save, or if their opening balance exceeded the minimum. By construction, all ASAs were either successfully completed or in default by the goal date,[28] and any remaining savings were transferred to clients' ordinary savings accounts.

The withdrawal-restriction account Gihandom was simpler in structure: Clients chose to restrict withdrawals before a specified goal was reached. Out of 92 Gihandom clients, 39 chose a goal amount, and 53 chose a goal date. The goal amount can be interpreted as the stronger restriction, since additional deposits need to be made in order to receive savings back.[29] There was no time limit for reaching the goal amount. However, as is common for Philippine banks, significant dormancy fees were applied after two years of inactivity. The minimum opening balance for Gihandom was 100 pesos.[30] For both ASA and Gihandom, opening balances were collected one week *after* contract signing. The practical motivation 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.[31] Finally, 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, or a natural disaster, clients could close their account and access their savings without any penalties. Within the six months of observation, no client exercised this option.

[27] Attitudes towards charities were measured in the baseline, and are controlled for in the analysis.

[28] There was a one-week grace period to make any outstanding deposits, but no client made use of this.

[29] Neither Ashraf et al. (2006)'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.

[30] Different minimum opening balances for ASA and Gihandom may drive the observed difference in take-up.

[31] See e.g. Benartzi and Thaler (2004). Financially, the late collection just delayed when individuals entered the commitment contract. 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. Out of 159 (116) individuals who initially signed the ASA (Gihandom) contract, 45 (24) failed to deposit an opening balance.

Individuals were left to themselves during the savings period, without help from deposit collectors or reminders. After all goal dates had been reached, a comprehensive endline survey asked about savings, outstanding loans, expenditures, changes in income, and various types of shocks experienced. In addition, existing ASA clients were offered to 'pre-order' ASA for a second round, and told that the bank would continue the product conditional on sufficient demand. While the Pre-Order did not involve a financial commitment, it required the completion of substantial paperwork and a new savings plan (to deter cheap talk).

4. The Survey Instrument

In the context of a comprehensive baseline survey, I measured time-inconsistent preferences using multiple price lists (MPLs):[32] 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, for the others the questions were hypothetical (see Appendix V for details). After six questions using a near time frame (now versus one month), the same six questions were asked 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. 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. Exponential discounters will be time-consistent – i.e., the amount necessary to make them switch from the earlier to the later reward will be the same in both time frames. I identify as hyperbolic discounter (or 'present biased') those who put a higher premium on waiting for one month in the present than in the future. Individuals who exhibit more patience now than in the future are 'future biased'. An individual who always prefers the earlier reward in all questions and 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 anchoring. For those with real rewards, one randomly selected question was paid out, using draws of numbered ping pong balls from a 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, 18.9 percent future-biased, and 35.7 percent impatient. No

[32] In a developing country field setting, MPLs were more practicable than the convex budget method recently introduced by Andreoni and Sprenger (2012). While MPLs fail to identify point values for β and δ, they do identify time-inconsistency, i.e., $\beta < 1$. Furthermore, the use of monetary payments in measuring consumption preferences has been questioned by Augenblick et al. (2015). While their criticism is valid here, severe liquidity constraints are widespread in the low-income study population, creating a tight link between cash inflows and consumption.

significant difference is apparent between those offered real and those offered hypothetical rewards.[33] These estimates are roughly comparable to others in the literature.[34]

In addition to a measure of preference reversals, the analysis requires a measure of sophistication which is not in itself derived from commitment demand. To the author's knowledge, the only study which identifies $\tilde{\beta}$ on an individual level using experimental data is Augenblick and Rabin (2015), using incentivised predictions of future behaviour.[35] Their prediction task requires a sophisticated laboratory setting due to the various possible biases: Individuals may adjust their behaviour to comply with previous predictions (due to anchoring, prediction bonuses, or a desire to be consistent), and they may adjust their predictions to commit their future selves to desirable behaviour. Augenblick and Rabin (2015) address these challenges using extensive variations in the size of the bonuses for accurate predictions, and in whether a prediction is elicited at all. The required extent of task repetitions is not practicable in a field experiment with over 900 participants, calling for a survey-based proxy measure.

The approach pursued in this paper relies on a simple idea: In that sophistication is a function of both β and $\tilde{\beta}$, it can be proxied by interacting a measure of *observed* time-inconsistency with a measure of *perceived* time-inconsistency. MPL questions are used for the former. For perceived time-inconsistency, I use the self-control measure proposed by Ameriks et al. (2007), henceforth referred to as ACLT. Like Augenblick and Rabin (2015), ACLT infer sophistication from predictions about future temptations and behaviour. Simple hypothetical survey questions elicit individuals' ideal, tempted, and expected allocation of a fixed resource over time. While designed to identify the parameters of the Gul and Pesendorfer (2001) model, the resulting measure reflects an individual's *perceived* (rather than actual) self-control problems. Finally, the consumption framing of the ACLT questions addresses recent concerns by Augenblick et al. (2015) and Frederick et al. (2002) that time preferences over consumption cannot be inferred from choices over fungible monetary payments.

The setup is as follows: Respondents were presented with a hypothetical scenario of winning ten certificates for "dream restaurant nights". Each certificate entitled the holder and a companion to an evening at any local restaurant of their choice, including an unlimited budget for food and drink, and all gratuities. The certificates were valid for two years starting immediately, and expired thereafter. The scenario included an example list of local middle-class restaurants, chosen to be above what respondents could usually afford, and regarded as highly desirable. This was intended to prevent substitution of certificates into everyday consumption. Piloting of the scenario revealed not only that it was easily understood – the low income levels of the study population also avoided several confounds which might be expected in richer countries: Respondents were used to eat at home or in simple street eateries, *carinderias*. The restaurant widely regarded as the most desirable was *Jollibee*, a large Philip-

[33] A bias from uncertainty would predict more present bias under real rewards. See Appendix IV for a comparison of real and hypothetical incentives.

[34] In their MPL benchmark, Andreoni and Sprenger (2012) find 16.7 percent present-biased and 10.7 percent future-biased. Kaur et al. (2015) find 17 percent present-biased. Using hypothetical questions, Ashraf et al. (2006) find 27.5 percent present-biased and 19.8 percent future-biased. Since future-biased and time-consistent preferences generate the same predictions for commitment demand, they are not distinguished in the regression analysis.

[35] Related structural approaches include Acland and Levy (2015) and Skiba and Tobacman (2008).

pine family fast-food chain. Restaurant vouchers were thus understood as a carefree temptation good, without the need for advance bookings or a babysitter. Using the wording of ACLT,[36] I then asked for

1. the *ideal* allocation of the ten certificates to year 1, as opposed to year 2,
2. the allocation individuals would be *tempted* to consume in year 1, and
3. the allocation individuals *expected* to consume in year 1, taking into account both the ideal and the temptation.

These questions provide two important measures: *Perceived lack of self-control* (from $(3)-(1)$, $expected-ideal$) and *perceived temptation* (from $(2)-(1)$, $tempted-ideal$). They were designed for the Gul-Pesendorfer model. Using the $\beta\delta$-model (where costly self-control does not exist), the two measures should be the same: The difference between the ex-ante optimal allocation, and the allocation the agent expects to play in a subgame perfect equilibrium. In this sense, both measures are equally suitable to capture perceived time-inconsistency.

I focus on $tempted-ideal$, and interact it with observed time-inconsistency (in MPLs) to obtain a measure of sophisticated hyperbolic discounting: $tempted-ideal*presentbias$.[37] This choice is motivated by the prediction that temptation, not self-control, drives commitment demand: Suppose costly self-control does exist, and thus the two measures diverge. An individual who achieves the ideal allocation through the use of costly self-control (implying $expected-ideal=0$) will still benefit from commitment. While a commitment device will not affect the allocation consumed, it increases utility by removing temptation.[38] Therefore, $tempted-ideal$ closely relates to the concept of sophistication, in that it captures the perceived benefit from commitment. A second difference between $tempted-ideal$ and $expected-ideal$ emerges in the presence of image concerns: Admitting to an interviewer that one is subject to temptation may be easier than admitting that one gives in to these temptations. Consistent with these conjectures, observed levels of *perceived temptation* are much higher (median 2, mean 2.35) than levels of *perceived lack of self-control* (median 0, mean 0.61).

Surprisingly, neither measure predicts being present biased ($p=0.9$ and 0.17, Figure A3 shows the bivariate distribution). The orthogonality of temptation and present bias facilitates the later interpretation of marginal effects, as the data suggests that one variable can be increased while the other is held constant. Several factors may explain the lack of correlation: First, individuals may use costly self-control to appear time-consistent. Second, all variables may be measured with error. While the ACLT questions were well-understood, the MPL questions were more susceptible to anchoring, and also required higher numeracy skills. Third, individuals may be pessimistic about their degree of time-inconsistency (discussed in Appendix I.3). The analysis proceeds with $tempted-ideal*presentbias$

[36] Questions 1, 2 and 3 correspond to questions (a), (c) and (d) in ACLT, respectively.

[37] I 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. Negative values occurred in 4 (42) out of 910 cases for temptation (self-control).

[38] In a lab experiment, Toussaert (2015) finds between 23 and 36 percent 'self-control types', who are willing to pay for commitment even though they expect to resist temptation, in order to reduce the need for costly self-control.

TABLE 1. SUMMARY STATISTICS BY TREATMENT ASSIGNMENT

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

as a measure of sophistication. Robustness checks with respect to the sophistication measure are reported in Appendix IV.

In addition to measuring time-inconsistency, the baseline survey obtained measures of financial claims from others, risk aversion, cognitive ability, financial literacy, intra-household bargaining power, 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 V.

Table 1 presents summary statistics. Randomisation into treatment groups occurred shortly after the baseline survey. Means were statistically different across treatment groups for 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 IV.

5. Empirical Strategy

The main outcomes of econometric interest are a range of treatment effects (on savings, loan demand, and expenditures), as well as predictors of take-up, contract outcome (successful or default), and the pre-order decision. Denote by R_i an indicator variable for assignment to the 'Regular Saver' (ASA) group, and by W_i an indicator variable for assignment to the 'Withdrawal Restriction' (Gihandom) group. Treatment effects can be estimated using OLS on

$$\Delta Y_i = \alpha_0 + \alpha_R R_i + \alpha_W W_i + \varepsilon_i \tag{3}$$

where ΔY_i denotes the change in the outcome of interest. The obtained coefficients $\hat{\alpha}_R$ and $\hat{\alpha}_W$ estimate the Intent-to-Treat effect (ITT) – the causal effect of having been *offered* the corresponding commitment product. 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, and a zero effect on those who did not.[39] The TOT can be estimated by dividing the ITT by the fraction of take-ups, or using an instrumental variables approach with assignment to treatment as orthogonal instruments.

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

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

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 ASA, or to pre-order ASA for a second round. The vector X_i contains

[39] As outlined in Section 3, the marketing treatment was identical for all groups up until the commitment offer. See Imbens and Angrist (1994) for a discussion on ITTs and local average treatment effects.

TABLE 2. SAVINGS OUTCOMES (OLS, PROBIT)

	(1) Change in Bank Savings	(2) Purchased Savings Goal	(3) Borrowed to Purchase Goal (given purchase)	(4) Change in Other Savings (survey-based)
Regular Saver Treatment (ASA)	428.633***	0.1156**	0.0509	426.811
	(65.587)	(0.0486)	(0.0621)	(671.844)
Withdrawal Restr. Treatment (Gihandom)	148.243***	0.1322**	0.2109***	-328.159
	(40.927)	(0.0545)	(0.0808)	(705.461)
Constant	27.160***			63.451
	(9.399)			(531.028)
Mean Dep. Variable		0.4992	0.1922	
R^2	0.04			0.00
Observations	746	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* is the change in a client's total savings balance at the partner bank, summed across ordinary savings accounts and commitment savings products (ASA or Gihandom). Two outliers are excluded, see footnote 40. 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.

demographics and survey-based preference measures. All binary choice regressions contain marketer fixed effects to filter noise from differences in marketer ability.

6. Results

6.1. Average Treatment Effects

Effect on Total Bank Savings This section presents estimated average effects of the two commitment treatments. 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 savings plan. The reason for focusing on the goal date is that all savings are expected to be spent on the planned expenditure after this date. The cost of this choice is that it diminishes the sample to those 748 individuals who were willing to make a savings plan with the marketer. This form of attrition is orthogonal to assignment to treatment.

Column (1) in Table 2 estimates that assignment to the Regular Saver treatment group increased average bank balances by 429 pesos (U.S.$10.20) relative to the control group.[40] This estimate is net of default penalties, and excludes the 100 peso gift contained in the ordinary savings account. Individuals assigned to the Withdrawal-Restriction group saved on average 148 pesos more than the control group. When normalised by the savings duration, the latter effect replicates the result of Ashraf et al. (2006): Their SEED account increased savings by 34.25 pesos per month (411 pesos over 12 months), while the Gihandom account increased savings by 32.88 pesos per month (148 pesos over 4.5 months). The terms and conditions of SEED and Gihandom were identical.[41] A small but significant savings increase of 27 pesos in the control group may reflect the effect of marketing, specifically the universal provision of savings plans and ordinary savings accounts. Alternatively, it may reflect income gains from the monetary rewards provided in the baseline survey.[42]

Treatment-on-the-Treated effects (TOT) can be obtained by instrumenting take-up of ASA and Gihandom with assignment to the treatment groups (R and W). The identifying assumption is discussed in Section 5. The TOT estimates suggest that taking up the regular-instalment product ASA increased savings by 1392 pesos, while taking up the withdrawal-restriction product Gihandom increased savings by 303 pesos. The increased gap in the TOT effects relative to the ITT effects is a result of the higher take-up rate for Gihandom.

Six months after the baseline survey, when all goal dates had been reached, an endline survey asked whether individuals had purchased the good (respectively, paid for the expenditure) they had been saving for. Out of the 615 individuals who had a) made a savings plan and were b) reached by the endline survey,[43] 307 reported to have bought the desired good. When asked how they paid for this expenditure, slightly below 20 percent of respondents declared to have used loans from formal or informal sources. Columns (2) and (3) in Table 2 present probit estimations of the effect of treatment on the likelihood of purchasing the good, and on borrowing for the purchase. Because the sample of those who purchase the good is self-selected, the coefficients in column (3) do not have a causal interpretation, and merely represent correlations.

Table 2 confirms that both the Regular Saver treatment and the Withdrawal Restriction treatment increased an individual's chances of purchasing their savings goal. The coefficients for the two treatments are not significantly different from each other. However, the Withdrawal Restriction group was 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) increases the likelihood of obtaining one's savings goal, but does not significantly affect the probability of

[40] The ITT estimate excludes two outliers, both in the Regular Saver group, whose savings increased by 15 and 18 standard deviations, respectively. Including these outliers changes the ITT estimate to 585 pesos.

[41] The study locations are 70km (2 hours by local bus) apart. The study populations differed: Ashraf et al. (2006) studied a sample of previous clients of the partner bank, while this paper studies a general low-income population with little previous bank exposure.

[42] The randomisation was stratified on the receipt of monetary rewards in the time preference elicitation.

[43] Attrition in the endline survey and completion of a savings plan are both orthogonal to treatment status.

borrowing for it. This suggests that the ASA product indeed helped individuals to purchase a savings goal using their own funds.

Appendix Figure A1 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 (also see Table A8). Table A6 presents quantile treatment effects. Table A5 examines treatment effect heterogeneity across a number of covariates, and finds that such heterogeneity is most pronounced for existing savings account holders – consistent with trust and basic familiarity with the banking system.

Testing for Crowd-Out of Savings Column (4) of Table 2 investigates 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. At baseline, 46 percent of the sample reported to have an existing savings or checking account with another institution. This number is partly driven by a formal requirement to open a savings account when obtaining microloans. 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 outcome variable in column (4) is the change in an individual's total savings balance outside of the partner bank, as measured in the baseline and endline 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 endline survey asked about savings kept around at the goal date, as opposed to the survey date. Unfortunately, the survey data is noisy, and coefficients are estimated with substantial imprecision.[44] The available evidence does not suggest a substitution between increased savings at the partner bank, and reduced savings at home or at other institutions. All coefficients are insignificant, and the sign on being offered ASA is positive.

6.2. Heterogeneity: Descriptive Results

The ASA results were strongly bi-modal: 51 ASA clients (45 percent) successfully completed their savings contract, with goal dates between December 2012 and April 2013. They completed all scheduled instalments with a median of 12 transactions, and reached savings goals between 950 and 7150 pesos (U.S.$170).[45] 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 this was not an immediate possibility, the repeat marketing stage included the option to 'pre-order' the product for a second round. The pre-order contract was not financially binding, but included substantial official paperwork. Two

[44] The savings data has been truncated at 1 percent, reducing the sample from 615 to 603 observations.

[45] Savings contracts had a median of 16 scheduled instalments, with a range from 8 to 25.

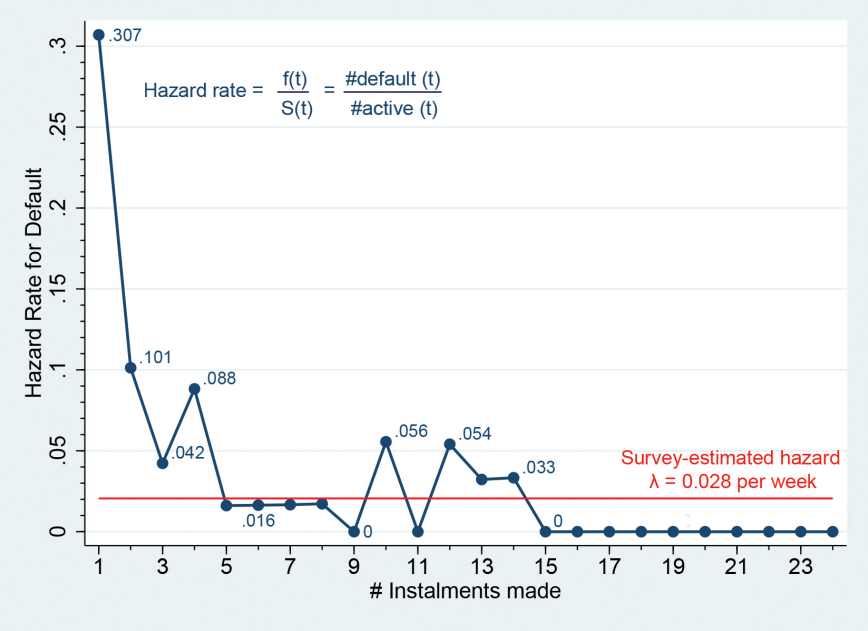

FIGURE 2. HAZARD RATE FOR ASA DEFAULT BY NUMBER OF INSTALMENTS MADE

thirds of the successful clients took up this offer (see Table A4), devised a new savings plan, and chose a new termination fee. The bank has since offered new ASA contracts to those enquiring about them.

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. What happened? Two possibilities emerge: (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 defaults should depend on the shock distribution: Assume that, as modeled in Section 2, savers get hit by large shocks with a per-period probability of λ. If shocks are i.i.d., a fraction λ of the surviving population should default each period. In other words, the default hazard rate should be constant over time. In sharp contrast, Figure 2 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, 8 clients deposited one more instalment, and only 10 clients defaulted after depositing more than five instalments. Figure 2 shows the default hazard rate by instalments, where the number of defaults at instalment t is given by those who discontinued depositing after making t instalments. The active population at t includes those who have neither defaulted nor successfully completed their contract at this point. Far from being constant, the default hazard rate spikes after the first instalment, and trails off afterwards. The null hypothesis of a constant hazard rate is rejected at $p<0.01$ despite using only 24 instalment-observations.[46]

Approximating instalments with weeks (84 percent of clients chose weekly instalments), Figure 2 further illustrates the predicted default timing given a hazard rate of 0.028 per week. This estimate

[46]Using instalments like a time axis is a simplification: Individuals were allowed to fall two instalments behind at any time. Consequently, time lags occur, specifically so between the last completed instalment and the official time of account closure. Furthermore, individuals could deposit multiple instalments in a single transaction.

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.

The observed default timing is difficult to reconcile with i.i.d. shocks, but individuals may have had heterogeneous shock arrival rates, or the shocks may have been correlated. Specifically, heterogeneous shock frequency rates λ_i would generate a declining rate of defaults, as those with the highest λ_i drop out first. However, a default rate of 31 percent after the first instalment requires a substantial fraction of the population to have λs well in excess of 0.31 shocks per week. The maximum observed number of shocks for an individual was 0.19 per week (0.12 among adopters). Even if this survey measure is flawed, and true λs are much higher, theory predicts that individuals with high λs are likely to stay away from commitment (see Section 2, Proposition 3). Admittedly, the data do not confirm this prediction: An ex-ante measure of individuals' shock frequency from the baseline survey is negatively, but not significantly, related to commitment adoption (Table 3). It is possible that individuals were naive about their latent proneness to shocks. An example of biased beliefs is that one's λ_i corresponds to the average shock frequency $\bar{\lambda}$ in the population. While this would explain a bulk of defaults soon after opening, it does require $\bar{\lambda} = 0.31$, far beyond the elicited frequencies. Section 7.1 discusses the evidence for aggregate shocks (specifically, tropical storms). While a minor tropical storm hit the Philippines during the period of observation, it occurred several weeks after the main wave of defaults.

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 her time-inconsistency is likely to select into a commitment contract that is too "weak" to overcome her 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 or P250), 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, they would realise which penalty will be effective for them on average. Building on the discussions in Section 2.6 and Appendix I.1, uncertainty in β makes commitment less attractive. Theory predicts low commitment take-up, high conditional penalties, and low default rates – in sharp contrast to the patterns observed.

Appendix Figure A2 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 account opening, losing the opening balance (via the termination fee) was the only economic consequence of the contract, thus leaving them worse off. Clients who defaulted later (potentially following shocks) may have derived benefits from the commitment contract through higher savings – at a premium of 150 or 250 pesos, which may be cheaper than interest payments to local moneylenders. A cautious estimate of the frequency of 'mistakes' is provided by the pre-order results: 55 percent of all clients (71 percent of defaulting clients and 35 percent of successful clients) chose not to order ASA again (see Table A4).

For the Gihandom accounts, both benefits and risks were less pronounced: Out of 92 accounts, only five reached the specified goal amount (3 out of 53 date-based accounts, and 2 out of 39 amount-based). Median savings were 100 pesos, equivalent to the minimum opening balance. 79 percent of Gihandom clients (85 percent of amount-based accounts) made no further deposits after the opening balance. Similarly to ASA clients, 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 ASA default penalty is fixed.

Out of 582 clients who exclusively had an ordinary savings account, one reached their specified goal amount. Summary statistics on account usage can be found in Appendix Table A3.

6.3. Heterogeneity: Regressions

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

Predicting Take-Up of the Commitment Savings Products Columns (1) to (3) in Table 3 present probit regressions of the ASA take-up decision. The first notable fact is the lack of any demographic predictors for the adoption of the regular-instalment product: Age, gender, income, assets, marital status, education and household size are all uncorrelated with take-up.

Instead, ASA take-up is predicted by the proposed measure of sophisticated hyperbolic discounting (see Section 4) and a measure of cognitive ability (based on Raven's matrices, see Figure A5). The positive association with cognitive ability is reassuring: The ASA product is more complex than traditional savings accounts (but no more complex than a loan contract). The significance of cognitive skills suggests that clients who were more likely to understand the rules were also more likely to

TABLE 3. PREDICTING DEMAND FOR COMMITMENT (PROBIT)

Commitment	Regular Saver (ASA)			Withdrawal Restriction (Gihandom)		
Take-Up	(1)	(2)	(3)	(4)	(5)	(6)
Age (yrs)	-0.0015	-0.0004	-0.0003	0.0017	0.0009	0.0011
	(0.0019)	(0.0020)	(0.0020)	(0.0029)	(0.0028)	(0.0028)
Female	0.0328	0.0592	0.0536	0.2418	0.2347	0.2300
	(0.0914)	(0.0936)	(0.0868)	(0.1687)	(0.1501)	(0.1516)
Married	0.0076	0.0095	0.0165	-0.0932	-0.0848	-0.0952
	(0.0650)	(0.0640)	(0.0630)	(0.0939)	(0.0872)	(0.0887)
Weekly HH income	0.0018	0.0013	-0.0028	0.0003	-0.0054	-0.0052
(in thousand pesos)	(0.0081)	(0.0080)	(0.0079)	(0.0118)	(0.0108)	(0.0109)
No. of appliances owned	0.0007	-0.0048	-0.0084	0.0220	0.0253	0.0258
	(0.0145)	(0.0143)	(0.0143)	(0.0224)	(0.0213)	(0.0216)
No. of HH members	0.0125	0.0105	0.0130	0.0229	0.0202	0.0217
	(0.0100)	(0.0097)	(0.0095)	(0.0157)	(0.0154)	(0.0158)
Education (yrs)	-0.0055	-0.0075	-0.0094	0.0276***	0.0316***	0.0315***
	(0.0066)	(0.0064)	(0.0064)	(0.0094)	(0.0093)	(0.0093)
Present Bias	0.0757	0.0636	0.0827	0.0809	0.1020	0.1046
	(0.0866)	(0.0870)	(0.0864)	(0.1260)	(0.1256)	(0.1301)
Soph. Present Bias	-0.0622**	-0.0579**	-0.0631**	-0.0363	-0.0524	-0.0530
(Pres.Bias*Temptation)	(0.0291)	(0.0290)	(0.0292)	(0.0491)	(0.0510)	(0.0532)
Perceived Temptation (0-10)	-0.0114	-0.0067	-0.0046	-0.0058	0.0005	0.0006
	(0.0123)	(0.0124)	(0.0125)	(0.0212)	(0.0205)	(0.0209)
Impatience	-0.0047	0.0053	-0.0008	-0.0124	-0.0224	-0.0219
	(0.0476)	(0.0467)	(0.0464)	(0.0717)	(0.0684)	(0.0688)
Faces strong financial	-0.0022	-0.0076	-0.0038	0.1079	0.1166*	0.1185*
claims from others	(0.0426)	(0.0418)	(0.0414)	(0.0663)	(0.0638)	(0.0646)
Risk Aversion (0-6)		-0.0049	-0.0059		0.0497***	0.0500***
		(0.0106)	(0.0105)		(0.0167)	(0.0168)
Cognitive Ability (0-5)		0.0353*	0.0363*		0.0157	0.0174
		(0.0188)	(0.0187)		(0.0236)	(0.0239)
Financial Literacy (0-5)		0.0425*	0.0328		-0.0115	-0.0120
		(0.0246)	(0.0250)		(0.0322)	(0.0321)
HH Bargaining Power (0-5)		0.0063	0.0053		0.0444***	0.0456***
		(0.0113)	(0.0113)		(0.0164)	(0.0167)
Distance to Bank (km)			-0.0265			0.0084
			(0.0207)			(0.0260)
Existing Savings Account			0.0998**			-0.0022
			(0.0444)			(0.0667)
Donated to Charity			0.0221			-0.0100
			(0.0424)			(0.0625)
#Emergencies last yr			-0.0161			-0.0215
			(0.0277)			(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 marginal coefficients. Variables are as described in Table 1. The sample is restricted to clients who could be located for the marketing visit (see footnote 23).

take up the product. This may be interpreted as evidence against possible manipulation by the bank marketers (discussed in Section 7.5).

Present bias on its own is not a predictor of take-up, consistent with the intuition that *awareness* of time-inconsistency, rather than *actual* time-inconsistency, determines demand for commitment. 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 (observed using MPLs) 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. To interpret interaction coefficients, note that present bias is a binary variable, whereas temptation is in the interval [0,10], with a median of 2.

The model in Section 2 outlines a possible explanation for the negative link between sophistication and commitment adoption: Commitment is attractive for partially sophisticated agents, who anticipate that a low default penalty will be sufficient to make them save. In contrast, agents who perceive themselves as strongly tempted have two options: Either they choose a sufficiently large penalty, or they stay away from commitment. Non-adoption may be optimal if the required effective penalty is prohibitively high. An additional channel is that penalties were only enforceable up to the current savings level. An agent who anticipates that her first instalment is not a sufficient stake to get her to make the second one, should also stay away.

Column (3) investigates correlations with other choices. While neither distance to the bank branch, attitude towards charities (proxied by having donated any positive amount 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 probit regressions of the Gihandom take-up decision, limiting the sample to group W. If Gihandom and ASA were perceived as close substitutes, and individuals in need of commitment merely took up whichever commitment product was offered to them, then the factors predicting ASA take-up should also predict Gihandom adoption. Far from this, their sets of predictors are mutually exclusive. Gihandom adoption is predicted by high education (in years of schooling), high risk aversion (choosing a safe lottery in Figure A4), 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 Anderson and Baland (2002), who argue that Kenyan women use commitment devices to protect their savings against claims from their husbands.[47] The estimates in Table 3 are consistent with the hypothesis that women took up Gihandom to safeguard their savings from intra-household

[47]The estimated linear relationship of commitment take-up with household bargaining power is unable to capture Anderson and Baland (2002)'s 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.

conflicts. The strong predictive power of risk aversion is consistent with a precautionary savings motive. Compared to ASA, Gihandom was indeed better suited as a safeguarding device: Its withdrawal restriction prevented other household members from accessing savings, while preserving the woman's flexibility in when to make deposits. In contrast, ASA allowed withdrawals at any time, albeit at the cost of defaulting on the account. A reservation must be made with respect to statistical power: Group W is half the size of group R, reducing the precision of estimates.

Predicting Default and Repeat Take-Up Table 4 presents probit regressions of the ASA default and ASA pre-order decisions, in direct comparison with the predictors of take-up. In addition to the previous regressors, Table 4 also includes the number of emergencies which the household suffered since the baseline survey (discussed in Section 6.2).

Column (2) predicts ASA default among those assigned to the Regular Saver group (R). Individuals who do not commit cannot default. Thus, the coefficients are best understood as predictors of who took up the commitment product 'by mistake', proxied by take-up and subsequent default. This interpretation notably abstracts from the possibility of rational default following a shock. The results provide further support to the partial sophistication hypothesis: Present-biased individuals are significantly more likely to take up the ASA product and then default. This effect is particularly strong for agents who report low levels of temptation, representing naive and partially sophisticated hyperbolics. In contrast, more sophisticated hyperbolics are *less* likely to default: Aggregating the coefficients for present bias (0.11*), sophistication (-0.045**) and temptation (-0.02*) yields a lower likelihood of default for all present-biased agents with temptation values higher than the median of 2. The estimates are consistent with the idea that sophisticated hyperbolics either do not adopt commitment, or make sure to choose a contract which is incentive-compatible (adjusted through the weekly instalment or the penalty, 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 insignificant.

Columns (3) and (4) restrict the analysis to clients who took up the ASA product, and should be interpreted as correlational evidence: The regressions condition on an endogenous variable, and are likely subject to sample selection bias. In predicting default occurrence, the marginal coefficient on present bias has quadrupled, and kept its significance. The link between present bias and default seems to be stronger than the link between present bias and take-up, consistent with the intuition that *awareness* of time-inconsistency drives commitment adoption, while actual time-inconsistency determines the success of the contract. The temptation measure now has strong predictive power on its own, even when not interacted with present bias. Individuals who report feeling tempted but who do not exhibit hyperbolic preference reversals could be (i) time-inconsistent, but incorrectly classified by the MPLs, for instance because their switching point is beyond the measured range (35.7 percent of respondents chose the earlier reward in all MPL questions). Or (ii), individuals have Gul-Pesendorfer preferences (see the discussion in Section 4). Heidhues and Kőszegi (2009) allow individuals with Gul-Pesendorfer preferences to be partially naive about their level of temptation. Put into the current

TABLE 4. ASA DEFAULTS & REPEAT TAKE-UPS (PROBIT)

Dependent Variable	(1) ASA Take-Up	(2) Default (R-Sample)	(3) Default (takeup-Sample)	(4) Pre-Order (takeup-Sample)
Age (yrs)	-0.0003	-0.0024	-0.0057	-0.0080*
	(0.0020)	(0.0016)	(0.0041)	(0.0043)
Female	0.0536	0.1189	0.3663*	0.0745
	(0.0868)	(0.0911)	(0.1871)	(0.1942)
Married	0.0165	0.0064	0.0363	-0.2433
	(0.0630)	(0.0536)	(0.1369)	(0.1490)
Weekly HH income (1000 pesos)	-0.0028	0.0034	0.0206	0.0124
	(0.0079)	(0.0057)	(0.0159)	(0.0177)
No. of appliances owned	-0.0084	-0.0093	-0.0384	0.0007
	(0.0143)	(0.0115)	(0.0238)	(0.0298)
No. of HH members	0.0130	0.0123	0.0159	0.0015
	(0.0095)	(0.0077)	(0.0143)	(0.0179)
Education (yrs)	-0.0094	-0.0030	0.0016	-0.0133
	(0.0064)	(0.0053)	(0.0121)	(0.0142)
Present Bias	0.0827	0.1119*	0.4837*	-0.4860**
	(0.0864)	(0.0654)	(0.2550)	(0.2353)
Soph. Present Bias	-0.0631**	-0.0453**	-0.1718	0.2375**
(Pres.Bias*Temptation)	(0.0292)	(0.0230)	(0.1299)	(0.1170)
Perceived Temptation (0-10)	-0.0046	-0.0202*	-0.0655***	0.0109
	(0.0125)	(0.0105)	(0.0244)	(0.0301)
Impatience	-0.0008	-0.0030	0.0293	0.0319
	(0.0464)	(0.0372)	(0.0882)	(0.1001)
Faces strong financial	-0.0038	-0.0113	0.0095	0.0097
claims from others	(0.0414)	(0.0330)	(0.0848)	(0.0905)
Risk Aversion (0-6)	-0.0059	-0.0181**	-0.0673***	0.0199
	(0.0105)	(0.0084)	(0.0192)	(0.0240)
Cognitive Ability (0-5)	0.0363*	0.0365**	0.0658*	-0.0385
	(0.0187)	(0.0143)	(0.0384)	(0.0428)
Financial Literacy (0-5)	0.0328	-0.0168	-0.1403***	0.0582
	(0.0250)	(0.0204)	(0.0401)	(0.0465)
HH Bargaining Power (0-5)	0.0053	-0.0116	-0.0778***	0.0792***
	(0.0113)	(0.0090)	(0.0236)	(0.0257)
Distance to Bank (km)	-0.0265	-0.0115	-0.0365	0.0643
	(0.0207)	(0.0165)	(0.0536)	(0.0615)
Existing Savings Account	0.0998**	0.0296	-0.0641	0.1694*
	(0.0444)	(0.0363)	(0.0860)	(0.0890)
Donated to Charity	0.0221	0.0422	0.0802	0.2122**
	(0.0424)	(0.0332)	(0.0888)	(0.0896)
#Emergencies last yr	-0.0161	0.0005	0.0558	-0.0855
	(0.0277)	(0.0213)	(0.0601)	(0.0677)
#Emergencies since baseline		-0.0033	0.1156*	-0.0414
		(0.0182)	(0.0687)	(0.0658)
Marketer FE	YES	YES	YES	YES
Mean Dep. Variable	0.2687	0.1468	0.5463	0.4630
Observations	402	402	108	108

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Entries in the table represent the marginal coefficients of the probit regressions. Variables are as described in Table 1.

context, such individuals would likely be classed as time-consistent in MPLs, where commitment is absent. However, faced with a commitment choice, they would underestimate the default penalty required to make them save. Concluding, both cases (i) and (ii) imply that higher awareness of temptation will prompt individuals to choose more incentive-compatible contracts, reducing default.

Moving on to the pre-order (repeat take-up) decision, the coefficients on present bias (-0.49**) and sophistication (0.24**) are large and significant. The aggregate coefficient for a present-biased individual with the median value of perceived temptation is approximately zero. This has a convenient interpretation: Relatively naive hyperbolic discounters (those with below-median reported temptation) are unlikely to take up the ASA product again. The same group is most likely to have defaulted on their previous contract. The result is encouraging, suggesting that individuals who 'burnt their fingers' learned about their true preferences. The reverse holds true for present-biased individuals with above-median reported temptation (sophisticated hyperbolics): They were more likely to pre-order ASA for a second round, potentially following a positive experience with their first contract.

A number of other factors can help in explaining the observed default rates. The most obvious candidate - the occurrence of shocks during the savings period - finds some support in the take-up sample (*emergencies since baseline*, column (3)). The positive correlation of defaults with shocks, in combination with the fact that 45 percent of clients successfully completed their ASA contract, suggests that a significant portion of clients likely did choose a contract which was optimal for them in expectation.[48] The theoretical prediction that shock realisation should be irrelevant to the pre-order decision (as it does not affect contract optimality in expectation) is supported by the data (see column (4) of Table 4). Other factors predicting default include financial literacy (-), household bargaining power (-), risk aversion (-) and cognitive ability (+). Financial literacy is perhaps the least surprising: Individuals with poor numeracy skills tend to do worse at managing their household finances, and may fail to allocate a portion of the household budget to regular ASA deposits. The positive significance of cognitive ability is 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. Explaining the correlation between cognitive ability and default within the take-up sample is more difficult. Possible channels include rationalisability of behaviour and a sunk-cost fallacy: An individual with high cognitive skills is more likely to realise when a contract is no longer optimal, even for time-inconsistent reasons.

A similar puzzle arises from the negative correlation of risk aversion with default, but not with take-up. Analogue to the role of financial literacy, risk averse agents may be better managers of household finances, and set aside 'buffers' which can be used to make the ASA deposits. An alternative channel is loss aversion regarding the default penalty: Risk aversion is measured from choices over lotteries (Figure A4). If the no-risk lottery was perceived as a reference point, then lottery choice would also capture loss aversion.

[48] Clients who defaulted may have had a stronger incentive to report shocks, in order to preserve their self-image or reputation. This would bias the shock coefficient upwards. However, the endline survey was framed as coming from a research organisation, with no direct link to the bank. It made no reference to ASA. Note that attrition in the endline survey was compensated by imputing the median shock value for those who did not participate.

Finally, within the take-up sample, ASA default is strongly related to household bargaining power. Individuals may have learned soon after opening their account that it causes household conflicts to put aside a portion of the household budget every week. Clients with low bargaining power are likely to have yielded to these disagreements, and defaulted on their contracts. This can be interpreted as costly experimentation with a new savings technology, and is discussed further in Section 7. The large positive association of household bargaining power with the pre-order decision further supports a learning explanation: Once individuals had learned about the difficulties of regularly diverting a share of the household budget, only those with sufficient autonomy chose to take up the product again.

7. Alternative Explanations

So far, the paper has focused on time-inconsistency and partial sophistication in attempting to explain why individuals may adopt and subsequently default on commitment contracts. Section 6.2 further argued that the timing of defaults is hard to reconcile with shocks to income or expenditures, at least provided that individuals had correct beliefs about the probability of such shocks. However, numerous other factors could have triggered a wave of defaults shortly after adopting commitment. Aggregate shocks may have affected the study region, or individuals may have been too optimistic about their future income. Another range of explanations can be summarised as 'costly experimentation': Individuals were learning what it involves to save each week. How hard would it be to extract 150 pesos from their household every week? How much time will it take to walk to the bank, and how difficult will it be to remember to make instalments on time? Finally, how much discipline would it take to curb consumption? While aspects of these questions may be product-specific, they all relate to more general preference and household parameters, as well as prior experiences. How prone is the individual to experience conflict within the household? How far away do they live from the bank, and how constrained is their attention? How much did they struggle to curb consumption, for instance when paying for previous loan instalments?

This section will address some main concerns relating to costly product experimentation, specifically, household conflict and transaction costs. It further discusses the evidence for aggregate shocks, income optimism, and persuasion by the bank marketers. As a rule, it is not difficult to think of reasons why individuals may have defaulted once they had entered commitment contracts. However, few of these reasons parsimoniously explain why individuals would demand commitment in the first place. Further, unless individuals were fully naive about potential challenges of saving, measures associated with these challenges should predict selection into (or out of) the contracts. For instance, those with very high shock arrival rates are unlikely to be among the adopters. Table 5 reviews the consistency of the study's findings with a number of alternative explanations.

7.1. Aggregate Shocks

An aggregate shock around the time of account opening may explain the large wave of defaults. The Philippines is a well-known area for earthquakes and tropical storms, and had recently been

TABLE 5. FINDINGS AND ALTERNATIVE EXPLANATIONS

				Costly Experimentation				
	Aggregate Shocks	**Hetero-geneous shock rates**	**Income Optimism**	**Limited At-tention/ Trans-action Costs**	**Household Conflicts**	**Time-Inconsistency + Partial Sophistication**	**Time-Inconsistency + Full Sophistication**	**Marketer Persuasion**
27% Commitment Take-Up (ASA)					Yes, if used to safeguard	Yes	Yes	Yes, but limited by 2-step account opening process
High Average Treatment Effect on Bank Savings					No, predicts higher savings under W-account	Needs full sophistication for part of the population	Yes	
55% Default, incl. 31% imme-diately after opening	Yes, given shock during account opening time	Ambiguous: High λs discourage take-up	Yes, upon realizing the true income	Ambiguous: High t or scarce attention discourage take-up	Yes, if conflict unexpected	Yes	No, 'rational defaults' only	Yes
Associated Measure Predicts Take-Up?	No evidence of an aggregate shock during account opening.	No, using ex-ante measured λ	Yes, contrary to prediction	No, using distance to bank	No, using HH bargaining power	Yes, sophistication negatively predicts take-up, conditional on β		No, cognitive ability *positively* predicts take-up
Associated Measure Predicts Default?		Yes, using ex-post measured λ. But λs quantitatively too small.	No, using average (predicted–real) income	No, using distance to bank	Yes, using HH bargaining power, but only within take-up sample	Yes, sophistication negatively predicts default, conditional on β. Observed time-inconsistency positively predicts default.		No, cognitive ability *positively* predicts default

Cells in this table summarise whether a given finding would be expected under a given explanation. The design of this table is inspired by DellaVigna and Malmendier (2006).

hit by tropical storm Washi in December 2011, causing 1,268 casualties.[49] 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 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 (20.4 percent of defaulting ASA clients) reported some damage to their house or crops. Conditional on non-zero damage, the median damage value was 1400 pesos (U.S. $33). Only 3 out of 732 respondents reported to have lost income due to the power outages.

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, usually three weeks later. By the time of the storm in early December, most of the contract defaults had already occurred.

7.2. Income Optimism

Following Browning and Tobacman (2007), the link between time-inconsistency in MPL questions and ASA default incidence could have been caused by overoptimistic beliefs about future income: If individuals expect their future income to be higher than their current income, they may select the smaller, sooner reward in the 'now vs. 1 month' frame, but choose the larger, later reward in the '1 month vs. 2 months' frame (see Section 4). As a result, they would be falsely classified as present-biased. Income optimism could also explain default incidence if it caused individuals to commit to overly ambitious savings plans.

Using data on predicted and realised incomes, I construct a measure which plausibly captures income optimism for groups. Details are described in Appendix Table A7. 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, or she may have systematically biased beliefs. However, given correct beliefs, the law of large numbers implies that individuals should correctly predict their income *on average*. If the utilised MPL questions capture income optimism rather than time-inconsistency, then individuals classified as 'present-biased' should have higher predicted-minus-realised income gaps. Furthermore, if defaults were caused by systematic mispredictions of future income, then defaulting clients should have higher prediction gaps than successful clients.

The average prediction gap for six-months income growth across the sample is 3.36 income brackets (see Table A7), suggesting that moderate income optimism may be common. However, the average prediction gap is not higher for individuals classified as present-biased. In contrast,

[49] Statistics from the Philippine National Disaster Risk Reduction and Management Council (NDRRMC).

the average prediction gap is significantly higher for ASA adopters than for non-adopters. This is counterintuitive, as income optimism does not predict a demand for commitment. Finally, individuals who defaulted on ASA did not mispredict their income significantly more than did clients who successfully completed their contract.

Summing up, the available evidence does not suggest a connection between income optimism and the observed measure of present bias. Commitment adopters were more optimistic than those who rejected commitment, but there is no significant difference between successful and defaulting clients. In addition, income optimism alone does not explain why individuals demand commitment. Less parsimonious explanations may combine different factors, such as fully sophisticated hyperbolic preferences in combination with income optimism. This combination may predict both a demand for commitment and subsequent default. However, it fails to explain why measures of sophistication (which are robust to income optimism) are negatively associated with take-up and default. In this sense, partial sophistication provides a parsimonious explanation that is consistent with the evidence.

7.3. Limited Attention and Transaction Costs

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 (such as a child's sickness, or running out of stock for one's business) before they cause damage. Given the relatively low stakes of the Regular Saver account (with default penalties roughly equivalent to a day's wage), it would be understandable if individuals prioritised their attention on their home and work lives, rather than on their bank accounts. However, if the returns to investing attention were higher at work or at home, individuals should not have adopted the Regular Saver product in the first place: During the marketing stage, ASA was clearly presented as attention-intensive. Clients were presented with an explicit savings plan including due dates for each week, and given the instruction to physically deposit their instalments at the bank. Most respondents received their income in cash, and bank transfers were uncommon. In the Banerjee and Mullainathan (2008) model, not investing attention in one aspect of one's life incurred a risk that a costly problem would go unnoticed. In contrast, not investing attention in the weekly ASA schedule resulted in *certain* default. Consistent with the prediction, attention concerns reduced take-up: Being 'too busy to go to the bank' was a common reported reason among those who chose not to adopt the Regular Saver product. The logic for transaction costs is analogue to that for limited attention. For those who did adopt the product, 'distance to the bank branch' does not predict default (as measured by GPS coordinates, see Table 4).

7.4. Household Conflicts

A desire to safeguard savings from other household members is the leading alternative hypothesis in explaining a demand for commitment savings devices. The Regular Saver account ASA did not explicitly restrict withdrawals, but it imposed a cost on them: Individuals could choose to terminate their

account and withdraw their savings (less the default penalty) at any time. Arguably, the withdrawal-restriction account Gihandom was more suitable as a pure safeguarding device. Consistent with this conjecture, Gihandom take-up is predicted by high education, high risk aversion, high household bargaining power, and strong claims from others on liquid assets (see Table 3). None of these factors are associated with the take-up of ASA. Moreover, if safeguarding is the driving motive for commitment, the treatment effect on bank savings is likely to be larger under Gihandom than under ASA, due to the account's stronger withdrawal restrictions. The opposite holds true in the data (Table 2).

Independent of safeguarding motives, household conflicts may also be responsible for the observed defaults: Other household members may not agree with the weekly extraction of household resources. Indeed, Section 6.3 shows that defaults are higher among those with low household bargaining power (defined in Appendix V). As previously argued for other default risks, the anticipation of household conflict should determine selection into commitment. It may also predict a higher optimal default penalty, which can be used to justify the need to honour the commitment to one's family. Neither relationship finds empirical support: There is no correlation of household bargaining power with ASA take-up (Table 3), and a near-zero correlation with the penalty choice ($\rho = -0.056$). While household conflicts may have caused some of the defaults, the data suggests that these conflicts were unexpected.

7.5. Marketer Persuasion

The bank marketers received a fixed daily wage (roughly three times the local average) in addition to a small commission. This raises the possibility of persuasion by the marketers - which may explain both commitment take-up and subsequent default. While persuasion cannot be fully ruled out, three factors worked to contain it: First, marketers were employed and trained exclusively by the research team. A script detailed every aspect of the conversation with a client, and auditors ensured that this script was followed. Strong emphasis was placed on clear explanations of the product features, and clients were encouraged to make sensible, conservative savings plans which were suitable for their income and usual expenditures. Second, opening an ASA or Gihandom account was a two-step process: After signing the contract, individuals had a one-week 'cool-down' period before the marketer returned to collect the opening balance. Those who signed the contract but failed to deposit the opening balance (see footnote 31) are considered non-adopters, thus filtering those most likely to have been momentarily swayed. Third, Table 3 shows that taking up ASA is positively related to cognitive ability. If commitment adoption was a consequence of gullibility, the opposite would be expected. Instead, a plausible interpretation for a positive relationship with cognitive ability is that those who struggled to understand the rules of an unfamiliar product were more likely to stay away.

8. Discussion

Commitment devices are receiving substantial attention both in the academic literature and in the public eye, and are generally portrayed as a promising way to overcome intrapersonal conflict. In the context of a regular-instalment 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 her degree of sophistication about her own time-inconsistency. I observe that a majority of individuals who take up a regular-instalment commitment product choose very low stakes for this commitment, and then default on it. Both take-up and default decisions are systematically linked to low measures of sophistication, suggesting that individuals may underestimate the amount of commitment required to make them save. Choosing unsuitable commitment devices can lead to real welfare losses – in the current setting, many clients incurred monetary penalties without building up any savings. This mechanism potentially extends to rich country applications such as gym contracts (as suggested by DellaVigna and Malmendier (2006)), diet clubs, and long-term pension savings plans.

From a policy perspective, there is no simple solution. One could offer commitment contracts exclusively with high penalties to ensure incentive-compatibility. However, this would deter commitment adoption – both by individuals who need higher penalties but fail to realise it, and by individuals who genuinely require only small penalties. Alternatively, one could legally mandate commitment savings, as prominently done in state pension schemes in Germany and Australia. In most contexts, this is unlikely to be optimal, as it requires the social planner to have a large amount of information on individual preferences.

The arguments in this paper focus on commitment contracts which impose penalties on binary consumption decisions. Other forms of commitment exist, with potentially different welfare implications. First, it may be possible to condition penalties on continuous events. For instance, choosing a social network full of health-conscious friends may expose an agent to increasing amounts of shame (or unsolicited health advice) the more weight she gains, or the more cigarettes she smokes. When penalties depend on continuous events, partial sophisticates would still choose commitments with insufficient rates of punishment. However, the marginal nature of the involved trade-offs imply that she would move closer to the optimal behaviour, and likely be better off than without commitment.

Second, some commitments work entirely without penalties, by directly eliminating undesirable options from the agent's choice menu. An example is the date-based Gihandom account featured in this study: Early withdrawals were simply disallowed. All savings were returned on the goal date, regardless of how much the agent had saved.[50] A different example is the decision to go shopping with a limited amount of cash, and no credit cards. At first sight, such commitments appear like a safe choice for partial sophisticates. Two qualifications are necessary: First, it can be hard to actually *eliminate* a bad option from the choice set, as opposed to just increasing its cost. Heidhues and Kőszegi (2009) discuss how individuals find costly ways to overcome the barriers they have set up for themselves. For instance, choosing not to buy alcohol during the weekly grocery shop may result in a last-minute trip to the store in the middle of one's favorite TV show, following a craving for beer. Second, the choices that *can* be eliminated may not be sufficient to enforce desirable behaviour. In this study, the Gihandom account restricted withdrawals, but it imposed no pressure to make any savings deposits after opening the account – indeed, 79 percent of clients did not.

[50] Note that amount-based Gihandoms imposed a penalty on failing to save by freezing savings indefinitely.

The comparison between ASA and Gihandom may point to a more general trade-off between hard and soft commitments: Offering stronger commitments with more pressure may provide greater benefits on *average* – as observed by a threefold treatment effect of ASA on average bank savings. However, stronger commitments may imply an increased risk of adverse welfare effects. Softer commitments may be limited in both their risks and their efficacy. Further research is needed to identify ways of overcoming time-inconsistency, which are both effective and at the same time 'safe' for 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.

Afzal, Uzma, Giovanna d'Adda, Marcel Fafchamps, Simon Quinn, and Farah Said, "Two Sides of the Same Rupee? Comparing Demand for Microcredit and Microsaving in a Framed Field Experiment in Rural Pakistan," Working Paper, Centre for the Study of African Economies 2015.

Ali, S Nageeb, "Learning Self-Control," *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.

Ameriks, John, Andrew Caplin, John Leahy, and Tom Tyler, "Measuring Self-Control Problems," *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.

Andreoni, James and Charles Sprenger, "Estimating Time Preferences from Convex Budgets," *American Economic Review*, 2012, *102(7)*, 3333–3356.

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

Augenblick, Ned and Matthew Rabin, "An Experiment on Time Preference and Misprediction in Unpleasant Tasks," Working Paper 2015.

_ , Muriel Niederle, and Charles Sprenger, "Working Over Time: Dynamic Inconsistency in Real Effort Tasks," *The Quarterly Journal of Economics*, 2015, *130(3)*, 1067–1115.

Banerjee, Abhijit V. and Sendhil Mullainathan, "Limited Attention and Income Distribution," *American Economic Review*, 2008, *98(2)*, 489–493.

Basu, Karna, "Commitment Savings in Informal Banking Markets," *Journal of Development Economics*, 2014, *107*, 97–111.

Bauer, Michal, Julie Chytilova, and Jonathan Morduch, "Behavioral Foundations of Microcredit: Experimental and Survey Evidence from Rural India," *American Economic Review*, 2012, *102(2)*, 1118–1139.

Benartzi, Shlomo and Richard H. Thaler, "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, 2004, *112(1)*, 164–187.

Besley, Timothy, "Savings, Credit and Insurance," in Jere Behrman and T.N. Srinivasan, eds., *Handbook of Development Economics*, Vol. 3A, Elsevier, 1995.

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.

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.

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.

Eliaz, Kfir and Ran Spiegler, "Contracting with Diversely Naive Agents," *The Review of Economic Studies*, 2006, *73(3)*, 689–714.

Fischer, Greg and Maitreesh Ghatak, "Repayment Frequency in Microfinance Contracts with Present-Biased Borrowers," EOPP Discussion Papers 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.

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.

Heidhues, Paul and Botond Kőszegi, "Futile Attempts at Self-Control," *Journal of the European Economic Association*, 2009, *7(2-3)*, 423–434.

__ **and** __ , "Exploiting Naivete About Self-Control in the Credit Market," *American Economic Review*, 2010, *100(5)*, 2279–2303.

Imbens, Guido W and Joshua D Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 1994, *62(2)*, 467–475.

John, Anett, "Can Fixed Regular Deposits Overcome Savings Constraints?," Working Paper 2015.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan, "Self-Control at Work: Evidence from a Field Experiment," *Journal of Political Economy*, 2015, *123(6)*, 1227–1277.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112(2)*, 443–478.

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 1999, *89(1)*, 103–124.

Rutherford, Stuart, *The Poor and Their Money*, Oxford: Oxford University Press, 2000.

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

Skiba, Paige Marta and Jeremy Tobacman, "Payday Loans, Uncertainty and Discounting: Explaining Patterns of Borrowing, Repayment, and Default," Technical Report 08-33, Vanderbilt Law and Economics Research Paper 2008.

Toussaert, Séverine, "Eliciting Temptation and Self-Control Through Menu Choices: A Lab Experiment on Curiosity," Working Paper 2015.