Just A Few Cents Each Day: Can Fixed Regular Deposits Overcome Savings Constraints?

Evidence from a Commitment Savings Product in Bangladesh*

Anett John†

Abstract

Empirical evidence suggests that there is a high demand for informal savings mechanisms even though these often feature negative returns - such as deposit collectors, ROSCAs, microloans, and informal borrowing. This paper argues that individuals may face even higher negative returns to saving at home due to hyperbolic discounting and claims on savings by relatives. I outline a model that shows why hyperbolic discounters cannot reach their welfare-maximising level of savings, and why a commitment savings product with fixed period contributions can increase their achievable savings level. Using a novel dataset obtained from a small microfinance institution in Bangladesh, the paper then presents some first empirical evidence on the effects of a commitment savings product with fixed regular instalments. I find that the introduction of the regular saver product was associated with an increase in individuals' savings contributions of 180 percent after a periods of five months. The paper concludes that the provision of commitment savings products with fixed contributions may reduce savings constraints and increase individuals' welfare, providing a substitute for costly informal mechanisms. However, since the data originates from a field study with self-selection problems rather than a randomized controlled experiment, further studies are needed to confirm this effect.

*I would like to thank Stuart Rutherford, Billal Hossain and M. Maniruzzaman of *SafeSave* Bangladesh for their continued help, support, and patience. I would also like to thank Greg Fischer and Oriana Bandiera for advice and guidance. I am indebted to Maitreesh Ghatak and Daniel Stein for helpful comments and suggestions. All errors and omissions are the sole responsibility of the author.

†Department of Economics and STICERD, London School of Economics. Email: a.hofmann1@lse.ac.uk. This paper was previously circulated under the name Anett Hofmann.

1 Introduction

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

A similar puzzle is the popularity of costly or inflexible savings mechanisms: Consider the wandering deposit collectors found in South Asia and Africa.[3] Rutherford (1999) illustrates their work with the example of a woman in India who collects a fixed amount of 5 or 10 Rupees per day from a client. She returns their savings after 220 days, charging a fee of 20 days' deposits for her service. This is equivalent to an annualized interest rate of *minus* 30 percent. Demand for such services is high – but there are few individuals who have build up the trust and reputation needed to provide this service, severely limiting supply.

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

The obvious question to ask is: Why do people not just save at home, instead of relying on such costly devices? An answer suggested in the existing literature are hyperbolic (or quasi-hyperbolic) preferences.[4] Hyperbolic discounters are impatient over current trade-offs (now vs. tomorrow) and patient over future trade-offs (one year vs. one year plus one day). As a result, they procrastinate saving. This leads either to failure to reach their savings target, or to a failure to smooth consumption over time by saving too much in the last minute. If hyperbolic individuals realize their own time inconsistency, they will be willing to pay for commitment savings products which enable them to overcome their present bias by forcing them to save.

A second factor which receives increasing attention are financial claims from relatives, neighbours and spouses. Slum communities and rural villages are often close-knit social networks. If an individual has managed to save money at home, she may be informally obligated to help out a neighbour who has fallen on hard times, a relative who needs medical treatment, or simply a husband who

[1] See e.g., Collins et al. (2009), the Indian vegetable vendors studied in Ananth et al. (2007), or Rutherford (1999).
[2] See e.g., Rutherford (1999) or Ananth et al. (2007).
[3] See e.g. Besley (1995) on West Africa's susu collectors.
[4] Most prominently, Laibson (1997) and O'Donoghue and Rabin (1999).

claims the remains of the household budget for personal expenses.[5] Awareness of claims on liquid assets may sharply reduce incentives to save in the first place. As in the previous case, such individuals may develop a preference for commitment to "safeguard" their savings.

In recent years, a burgeoning literature on commitment savings has confirmed the empirical link between hyperbolic discounting, intra-household bargaining and a preference for commitment: Bauer et al. (2012) find that individuals in their sample who have been identified as hyperbolic keep a lower proportion of their savings at home and are more likely to borrow from microcredit organizations.[6] Ashraf et al. (2006b) find that hyperbolic preferences predict take-up of a commitment savings product which features withdrawal restrictions. Gugerty (2007) finds ROSCA participation patterns to be consistent with both time-inconsistent preferences and low intra-household bargaining power (only husband has formal sector income) – i.e. individuals value public pressure to save. Finally, Anderson and Baland (2002) find an inverted U-shaped relationship between women's autonomy from their husbands (as proxied by female share of household income) and participation in ROSCAs: Participation in ROSCAs is high when female autonomy is high enough to persuade their husband to let them join, but not so high that the protection which the ROSCA provides is no longer needed.

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

While this view is neither new nor surprising, the existing literature on formal commitment savings has focused almost exclusively on savings products featuring withdrawal restrictions, but without an obligation to deposit in the first place. Notable examples include Ashraf et al. (2006b), who find a positive effect on savings levels using a withdrawal-restriction product in the Philippines. Dupas and Robinson (2013) study a similar (though not as strict) withdrawal-restriction product in Kenya and find increased investment in clients' small businesses. Brune et al. (2011) randomly of-

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

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

fered withdrawal-restriction accounts to smallholder cash crop farmers, and found a positive effect on agricultural input use, as well as on subsequent crop sales and household expenditures. Other notable savings interventions included text message reminders (Karlan et al. (2010)) and deposit collectors (Ashraf et al. (2006a), Giné et al. (2010)).[7] However, to the author's knowledge, none of the formal commitment savings accounts which have been studied so far commit clients to make future deposits.

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

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

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

While it would ultimately be interesting to compare the performance of a savings product with restricted withdrawals to a savings product with fixed regular contributions, this paper will focus on the question whether regular saver products can increase savings levels (by overcoming savings

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

constraints) relative to autarky. In an attempt to answer this question and complement the empirical analysis, the paper takes a theoretical approach to show that with hyperbolic discounting, a regular saver product increases the range of parameter values where a savings goal (the price of a nondivisible) can be reached, as well as smoothing savings contributions.

This paper proceeds as follows: Section 2 outlines a model of savings with hyperbolic discounting and looks at the effect of a regular saver product. Section 3 describes the data. Section 4 describes the empirical strategy, discusses possible sources and remedies of the endogeneity issue, and presents estimation results. Section 5 concludes.

2 A Conceptual Framework

2.1 A Savings Model with Hyperbolic Discounters

The introductory section has painted the picture of a large unmet demand for commitment savings products with regular contributions (from now on: "regular saver products"). It has been suggested that this demand may result from present-biased preferences, as well as from the desire to protect one's savings for necessary household expenditures from claims by relatives and other community members. There is, however, no theoretical foundation which explicitly justifies a need for regular saver products, as opposed to commitment products which simply carry withdrawal restrictions.

The following section attempts to provide a stylised model that is simple, yet capable of explaining how regular saver products may help hyperbolic discounters achieve higher savings levels. Higher savings levels, in turn, will enable them to afford welfare-increasing lump-sum expenditures - such as commonly needed for education, a long-awaited medical treatment, a new roof, a bride's dowry, or a business investment. The benchmark model without banking services follows Basu (2012)'s theory of saving under hyperbolic preferences, but departs from it in the design of the commitment savings product in the second part of the model.

To model achievable savings levels, consider an agent who decides whether to save up for a non-divisible good which costs the lump-sum $1 < p < 3$ and yields a monetary benefit $b > 3$. The agent lives for 3 periods and receives a non-stochastic per period income of 1, which he can either consume immediately or save. He cannot borrow. His instantaneous utility function is assumed to be twice differentiable and strictly concave, i.e. $u'(c) > 0$ and $u''(c) < 0$. His lifetime utility as evaluated in each period is given by the discounted sum of the instantaneous utilities:

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

When $\beta < 1$, the discount rate between the current period t and t +1 is lower than the discount rate between subsequent periods – i.e. the agent exhibits a present bias. The model will assume that he is a *sophisticated* hyperbolic discounter – i.e. he realizes that his future selves will display the same present bias that he has, and apply the discount function $1, \beta, \beta$, resulting in an inconsistency between the preferences of his current self and the preferences of his t+k-selves.[8] The model further assumes that there is no other form of discounting (i.e. $\delta = 1$). While the literature disagrees on how to evaluate the welfare of hyperbolic discounters, the most commonly applied criterion which will be used for the purposes of this paper is the one proposed by O'Donoghue and Rabin (1999): The welfare of a hyperbolic discounter is what he would like to maximise in a hypothetical period 0, i.e. just before the start of his life. His welfare is then given by

$$W = u(c_1) + u(c_2) + u(c_3).$$

Finally, I will assume a gross return on savings of R=1 in both the autarky and the banking scenario. This will later make it possible to isolate the effect of commitment on savings levels, since it is not profitable for individuals to use banking services simply to gain higher returns.

We can infer the welfare-maximising solution to the savings problem: Given the agent's desire to smooth consumption over time, it is optimal for him to split the necessary savings burden of $p - 1$ evenly over periods 1 and 2, and then spend his period 3 income plus accumulated savings on the good (the implied consumption profile is $c_1 = c_2 = (3 - p)/2$, $c = b$).[9] This is preferable to not buying the good (in which case he should not save at all) as long as

$$2u(\frac{3-p}{2}) + u(b) \geq 3u(1).$$

The remaining part of the section will assume that this equation holds – i.e. it is optimal for the agent to buy the good.

2.1.1 The Autarky Equilibrium

Since the agent is aware that he will follow different time preferences every period, his equilibrium behaviour represents the Subgame Perfect Nash Equilibrium of a game that he plays with his consecutive selves. Basu (2012) shows that without banking, the nondivisible will be purchased for any β above a critical $\hat{\beta}$. Since the focus of this paper is on achievable savings levels for given time preferences, we will show how this translates into a critical price $p_{max} = f(\beta)$. If the nondivisible costs less than p_{max}, it will be purchased.

Finding an equilibrium of this game is straightforward if $\beta = 1$: The agent's objectives are time-consistent, and his maximisation problem will coincide exactly with the welfare maximisation described above – in other words, he will always buy the nondivisible given the above condition, and save $\frac{p-1}{2}$ in periods 1 and 2.

In the hyperbolic discounting case of $\beta < 1$, the individual has to be treated as three separate optimising agents, one in each period. In period 1 and 2, the agent decides how much money s_t to send to the next period, and the period 3 self decides whether to buy the good (if possible). The game can be solved by backward induction.

[8]The results of the model will hold to some extent if he is partially sophisticated, but not if he is perfectly naïve, i.e. if his current self believes that future selves will share his preferences.

[9]If $p < 2$, he could save $p - 1$ in period 1 and buy the good in period 2. However, this yields no additional benefit given our assumption of $\delta = 1$ and implies a more uneven consumption profile.

Period 3

Given $b > p$, the agent will buy the nondivisible whenever he can afford it, i.e. whenever $s_2 \geq p - 1$. Any additional savings $s_2 > p - 1$ are simply consumed. The same happens with savings that are not sufficient to buy the good. Period 3 consumption is thus

$$c_3 = \begin{cases} 1 + s_2 - p + b & if\ s_2 \geq p - 1 \\ 1 + s_2 & if\ s_2 < p - 1. \end{cases}$$

Period 2

The period 2 self knows the good will be bought if and only if he sends $s_2 \geq p - 1$. This reduces the analysis to two relevant cases: If he wants the good to be bought, he will send $s_2 = p - 1$ exactly: Given that utility is concave and period 3 will get the utility associated with b, the marginal utility from sending an additional unit to period 3 will be much lower than that of consuming it in period 2. The utility he receives is $U_2 = u(1 + s_1 - (p - 1)) + \beta u(b)$. If he does not want the good to be bought, he will decide to optimally allocate the savings s_1 that he received from the first period between period 2 and 3, i.e. he finds $s_2 * (s_1) = argmax[u(1 + s_1 - s_2) + \beta u(1 + s_2)]$ subject to $0 \leq s_2 < p - 1$. In consequence, he will prefer that the good is bought if

$$u(1 + s_1 - (p - 1)) + \beta u(b) \geq u(1 + s_1 - s_2*) + \beta u(1 + s_2*).$$

Proposition 1. *The above equation holds if s_1 is bigger than some threshold value, $s_1 \geq s_{min}$.*

(Proofs of all propositions are in Appendix B.)

Proposition 2. *s_{min} is weakly decreasing in* β*.*

Hence, once s_1 crosses some threshold s_{min}, saving for the good will be optimal for all $s_1 \geq s_{min}$. For $\beta = 1$, we know that $s_{min} \leq \frac{p-1}{2}$, since the agent is willing to follow the welfare-maximising savings plan. For very low β, s_{min} is likely to be higher than 1, i.e. the good cannot be purchased as the agent would rather consume his savings.

Period 1

Analogous to finding the minimum s_1 that the period 1 agent *needs* to send to period 2 to make him save for the nondivisible, we can find the maximum s_1 that he is *willing* to save in order to receive the good. If this maximum is bigger than the minimum needed, the good will be purchased. Note that, conditional on the good not being purchased, the period 1 agent prefers to save $s_1 = 0$ to smooth his consumption path (if anything, he would like to borrow from the future, but is not able to do so by assumption). He knows this will lead to $s_2 = 0$ and thus to $U_1 = (1 + 2\beta)u(1)$. The maximum that he is willing to pay for the purchase of the nondivisible (i.e. for $s_2 = p - 1$) can be found by comparing

$$u(1 - s_1) + \beta u(2 + s_1 - p) + \beta u(b) \geq (1 + 2\beta)u(1).$$

Define s_{max} as the maximum value of s_1 such that this inequality holds (if there is no such value, let $s_{max} = 0$).

Proposition 3. *s_{max} is weakly increasing in β.*

We further know $s_{max}(0) = 0$ and $s_{max}(1) > \frac{p-1}{2}$.

Equilibrium

With s_{min} weakly decreasing and s_{max} weakly increasing in β, there will be a threshold level $\hat{\beta}$ such that $s_{min}(\beta) \leq s_{max}(\beta)$ for any $\beta \geq \hat{\beta}$. How do we know $\hat{\beta}$ is in the relevant interval $[0,1]$? We know that $s_{min}(0) > s_{max}(0)$ since $s_{min}(0) > 1$, $s_{max}(0) = 0$ and that $s_{min}(1) \leq s_{max}(1)$ as the non-hyperbolic agent always purchases the good. This means that the intersection $s_{min} = s_{max}$ occurs for some $\hat{\beta} \in (0,1]$. Therefore, the nondivisible is purchased in equilibrium for the range $\beta \in [\hat{\beta}, 1]$.

Since we have assumed $1 < p < 3$, it is possible that the agent may be willing to save $p-1$ in period 1, enjoy the good in period 2, and consume his income in period 3. This is symmetric to the case in which a period 2 agent is willing to save $p-1$ in order to receive the good in the third period, i.e. the case that $s_{min}(\beta) = 0$. Hence, this option is nested in our model. This case will generally not be of practical importance, since the implied restrictions on β are stronger than those when the necessary savings $p-1$ can be split over two periods.

How does the optimal savings schedule look in equilibrium? For $\beta \in [0, \hat{\beta})$ we know $s_1 = s_2 = 0$. Above $\hat{\beta}$, if β is large and $s_{min}(\beta)$ low, period 1 may want to save more than s_{min} voluntarily in order to smooth consumption. Ideally, he would like to save $s_1 = s_{opt}$ such that $u'(1 - s_{opt}) = \beta$. He will follow this optimal plan whenever $s_{opt} \geq s_{min}$.[10] This leads us to an equilibrium optimal savings schedule:

$$s_1 = \begin{cases} max(s_{min},\, s_{opt}) & if\ \beta \in [\hat{\beta},\, 1] \\ 0 & if\ \beta \in [0,\, \hat{\beta}) \end{cases},\quad s_2 = \begin{cases} p-1 & if\ \beta \in [\hat{\beta},\, 1] \\ 0 & if\ \beta \in [0,\, \hat{\beta}) \end{cases}.$$

Achievable Savings Levels in the Autarky Equilibrium

Using the savings model of Basu (2012), it has been derived that for a given price p of the nondivisible, the agent will choose to save up for the good if β is above a threshold level $\hat{\beta}$. However, in empirical field work, researchers are unlikely to know the exact time preference parameters of the agents concerned, unless extensive surveys or lab experiments have been conducted. In consequence, a natural question to ask is: For *given* time preferences, what is the maximum savings level that agents can reach? In our context, this can be translated to the question which price p the agent is willing to pay for a good yielding a constant benefit b. Having set up Basu's benchmark model, this question is easily answered: The threshold $\hat{\beta}$ will be an increasing function of p, so the maximum achievable price p_{max} must be an increasing function of β. To see this formally, it can be shown that

Proposition 4. *s_{min} is weakly increasing in p*

[10] Basu shows that s_{opt} is weakly increasing in β, and always smaller than s_{max}, yielding an intersection $s_{min} = s_{opt}$ at $\beta_{max} > \hat{\beta}$ above which period 1 will save s_{opt} rather than s_{min}.

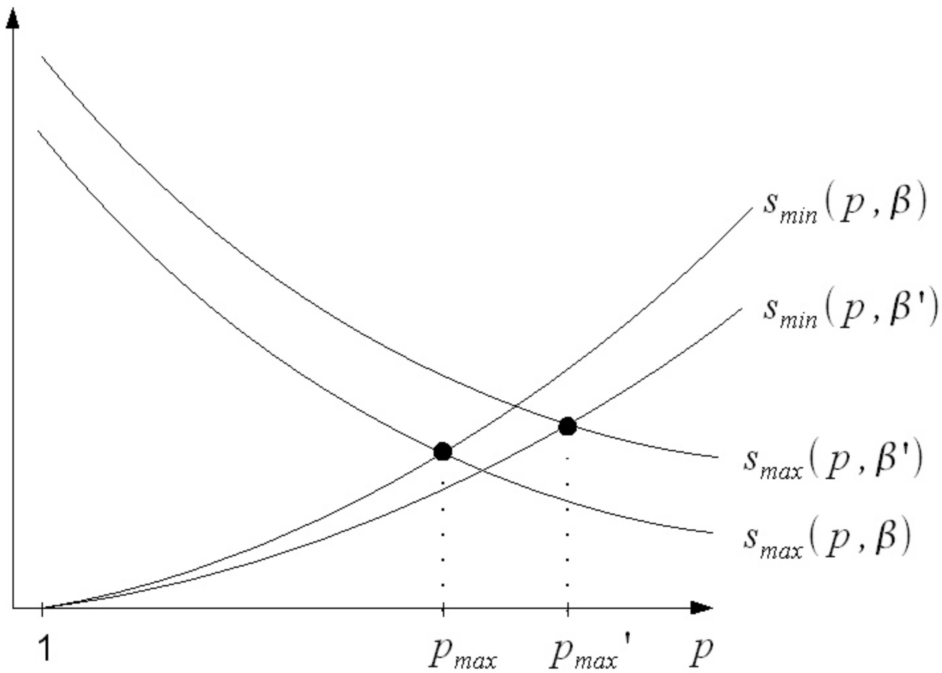

Figure 1: Achievable Savings Levels in Autarky

Proposition 5. *s_{max} is weakly decreasing in p.*

An agent will choose to save for the nondivisible if $s_{min}(p, \beta) \leq s_{max}(p, \beta)$, yielding a unique threshold p_{max} (at given β) and a price range $p \in (1, p_{max}]$ for which the good is bought. To see how p and β relate, consider p_{max} for a given β. Note that $s_{min}(p, \beta)$ is weakly decreasing in β and $s_{max}(p, \beta)$ is weakly increasing in β. Increasing β to $\beta' > \beta$ will thus shift down $s_{min}(p)$ and shift up $s_{max}(p)$. As can be seen from Figure 11, the intersection $s_{min} = s_{max}$ shifts to the right, increasing the threshold p_{max} to p'_{max}. It can be concluded that $p_{max}(\beta)$ is weakly increasing in β.

2.1.2 Equilibrium with a Regular Saver Commitment Product

We are now in a position to look at the effect of a commitment savings product with fixed contributions on an individual who cannot reach his welfare-maximising savings level due to hyperbolic discounting. The analysis will first look at the effect on the threshold $\hat{\beta}$, and then derive the effect on the achievable savings level p_{max}. The commitment savings product is defined in the following way:

1. The agent can choose a fixed per period savings contribution before contributions start.

2. In each period, the agent either pays the contracted-upon contribution, or he "defaults", in which case he pays a penalty and receives his accumulated savings back.

The first assumption will simplify things tremendously, as the agent is not subject to temptation in the period when he signs the contract. Studies like the SMarT project of Benartzi and Thaler (2004) show that this may be a plausible approach in practice. The second assumption requires the bank to have contract enforcing power in the case that the agent defaults in the first period, when there are no accumulated savings and he needs to make a net payment upon default. For the purposes of our model, we will assume that the bank has sufficient claims on the agent's income to ensure that the penalty is paid. In practice, this may be implemented via a small item of collateral which the bank holds. A different approach with the same implications is to think of the penalty as a psychological cost which the agent pays upon default: After bothering to set up a formal savings contract, he may face some cost of "renegotiation" with himself. Finally, the penalty could be paid via an account opening fee at the time of signing the contract, which is reimbursed with the first savings contribution (this case differs slightly in modelling strategy, see footnote X).

To see how regular contributions can help agents to save more, consider an agent who can commit in period 0 to save a fixed amount $\bar{s}$ in a designated bank account in both period 1 and 2. Once he chooses not to contribute $\bar{s}$ in a period, he pays a penalty D for cancelling the contract but receives back any accumulated savings. In addition, he is free to save at home independently of his bank contributions, so we can capture his total savings (in the bank plus at home) transferred from period t to $t+1$ as s_t. The penalty D is imposed in period 1 if $s_1 < \bar{s}$, and in period 2 if $s_1 \geq \bar{s}$, $s_2 < 2\bar{s}$. Assume $D < \bar{s}$.

Given an exogenous price p of the nondivisible, if the period 0 agent does want the commitment product, he will choose the welfare-maximising contribution $\bar{s} = \frac{p-1}{2}$. With $\bar{s}$ known, backward induction can be used to solve for the equilibrium.

Period 3

As before, the period 3 agent will buy the good whenever he can afford it, i.e. whenever $s_2 \geq p-1$, and consume any remaining savings.

Period 2

The savings decision of the period 2 self will depend on s_1, with a discontinuity at $s_1 = \bar{s}$. Suppose $s_1 \geq \bar{s}$: Period 1 has made the required contribution, so period 2 faces a penalty if he does not. If $s_1 > \bar{s}$, period 1 has left additional savings at home which can be used to make the period 2 contribution. Period 2 will prefer to stick to the contract, thus transferring $2\bar{s} = p-1$ and purchasing the good if

$$u(1+s_1-(p-1)) + \beta u(b) \geq u(1+s_1-D-s_2*) + \beta u(1+s_2*)$$

where s_2* is the optimal way to split remaining savings ($s_1 - D$) between period 2 and period 3, conditional on *not* buying the good:

$$s_2*(s_1) = argmax[u(1+s_1-D-s_2) + \beta u(1+s_2)]$$

subject to $0 \leq s_2 < p-1$. Since the inequality differs from the non-banking case only in the penalty D, the same proof can be used to show that the nondivisible is bought for any $s_1 \geq s^B_{min}(\beta)$, where $s^B_{min}(\beta)$ is the critical threshold in the banking case. For a given β, this threshold will be strictly lower than $s_{min}(\beta)$ in the autarky case: The right-hand side of the inequality decreases when D is introduced, while the left-hand side stays unchanged. Thus, the inequality will always hold using the original $s_{min}(\beta)$, and it will still hold for $s_{min}(\beta) - \epsilon$. In other words, a lower amount of savings is needed to incentivise period 2 to purchase the good because the option to just consume savings is less attractive.

The effect of the penalty disappears for $s_1 < \bar{s}$: Period 1 has already cancelled the contract and incurred the penalty, so period 2 is no longer committed to contribute $\bar{s}$. As in autarky, he compares whether

$$u(1+s_1-(p-1)) + \beta u(b) \geq u(1+s_1-s_2*) + \beta u(1+s_2*)$$

and purchases the nondivisible if and only if $s_1 \geq s_{min}(\beta)$. As a result, $s^B_{min}(\beta) = s_{min}(\beta)$ for $s_1 < \bar{s}$.

Figure 2 shows that the combined $s^B_{min}(\beta)$-function is horizontal between β'' and β''', where $s^B_{min}(\beta'') = s^B_{min}(\beta''') = \bar{s} = \frac{p-1}{2}$. In this interval, the period 2 agent needs to be given less than $\bar{s}$ if he faces the penalty and more than $\bar{s}$ if he does not. To ensure that he faces the penalty, the period 1 agent needs to save $s_1 \geq \bar{s}$ to keep the contract active. Between β'' and β''', the minimum s_1 needed to ensure the good is purchased is thus $\bar{s}$. [11]

It can be concluded that, with the regular saver product, the period 2 agent will prefer to save for the good if $s_1 \geq s^B_{min}(\beta)$, and accordingly save $s_2 = p-1$. For $s_1 < s^B_{min}(\beta)$, he will save $s_2 * (s_1) < p-1$. The threshold s^B_{min} is weakly lower than the autarky threshold s_{min}. Note that for $s^B_{min}(\beta) > \bar{s}$, the period 2 agent is not willing to make the contracted-upon contribution unless period 1 makes additional savings at home.

Period 1

As in autarky, period 1's behaviour is most easily captured by looking at the maximum s_1 that she is willing to save in order to receive the good. In contrast to autarky, she now incurs the penalty D if $s_1 < \bar{s}$. In the region $s_1 \geq \bar{s}$, the agent prefers to save for the nondivisible if

$$u(1-s_1) + \beta u(2+s_1-p) + \beta u(b) \geq u(1-D) + 2\beta u(1).$$

This condition is identical to the autarky case apart from D on the right-hand side. Note that if the agent does not want to save for the good (i.e. is not willing to pay the necessary s_1), she will always prefer to save $s_1 = 0$ and incur the penalty.[12] In consequence, the agent will transfer the necessary s_1 as long as $s_1 \leq s^B_{max}(\beta)$, where $s^B_{max}(\beta)$ is the critical threshold in the banking case. Since the left-hand side of the inequality is identical to the autarky case and the right-hand side decreases in D,

[11] Note that period 1 will never save in the region $\bar{s} - D \leq s_1 < \bar{s}$, as the contribution made plus the penalty incurred exceeds $\bar{s}$. This does not affect s_{min}, but it is captured in s_{max}.

[12] The alternative is to transfer $\bar{s}$ just to delay the penalty, but given $D < \bar{s}$ and a desire to smooth consumption, we know $u(1-D) + \beta u(1) > u(1-\bar{s}) + \beta u(1+\bar{s}-D)$.

the inequality always holds using the original $s_{max}(\beta)$, and it still holds for $s_{max}(\beta) + \epsilon$. Thus, the resulting $s^B_{max}(\beta)$ will be strictly higher than $s_{max}(\beta)$ for $s_1 \geq \bar{s}$.

Finally, consider the case where necessary savings are $s_1 < \bar{s}$, i.e. period 1 could ensure the good is bought even if he does not contribute $\bar{s}$. In this case, he faces a penalty whether or not he saves for the good. He compares

$$u(1 - s_1 - D) + \beta u(2 + s_1 - p) + \beta u(b) \geq u(1 - D) + 2\beta u(1),$$

which differs from the original inequality in the penalty D on both sides. The resulting threshold $s^B_{max}(\beta)$ now turns out to be *lower* than the original threshold $s_{max}(\beta)$. To see this, note that with a strictly concave utility function, the utility loss from D when starting at consumption level $1 - s_1$ is bigger than the utility loss from D when starting at consumption level 1: $u(1 - s_1) - u(1 - s_1 - D) > u(1) - u(1 - D)$ for $s_1 > 0$. In other words, the penalty D hurts the agent more when he is saving than when he is not. With the left-hand side decreasing more than the right-hand side, willingness to save will decrease, shifting the $s^B_{max}(\beta)$-curve below the original $s_{max}(\beta)$-curve for $s_1 < \bar{s}$.

Figure 2 shows that the two parts of the $s^B_{max}(\beta)$-curve combine with a vertical rather than a horizontal line. To see why, consider extending the lower part of $s^B_{max}(\beta)$ to the $\bar{s}$-line. For any β in this range, e.g. for $\hat{\beta}_B$, s_{max} is below $\bar{s}$ if the agent has to pay the penalty along with her savings, and above $\bar{s}$ if she does not. Since she does not have the pay the penalty for $s_1 \geq \bar{s}$, the maximum that she is willing to pay is given by the upper part of the $s^B_{max}(\beta)$-curve. The vertical part of $s^B_{max}(\beta)$ results from period 1 never choosing s_1 in the region $\hat{s} < s_1 < \bar{s}$, where $\hat{s}$ is defined such that the upper part of $s^B_{max}(\beta)$ reaches $\bar{s}$ at the same β' where the lower part of $s^B_{max}(\beta)$ reaches $\hat{s}$ (see Figure 2). Intuitively, if the necessary s_1 is $\bar{s} - \epsilon$, period 1 is better off to save $\bar{s}$ voluntarily instead of paying $s_1 + D > \bar{s}$. Furthermore, it can be shown that $\hat{s} < \bar{s} - D$: Consider a situation where the necessary $s_1$is $s_1 = \bar{s} - D - \epsilon$. Rather than incurring the penalty and paying a total of $s_1 + D = \bar{s} - \epsilon$, he will prefer to save $s_1 = \bar{s}$, which costs ϵ more in period 1 but yields an added benefit of $D + \epsilon$ in period 2. Summarizing period 1's behaviour, he incurs the penalty D when $\beta < \beta'$, with s_{max} given by the lower part of $s^B_{max}(\beta)$. For $\beta \geq \beta'$, he is willing to save $s_1 \geq \bar{s}$, with s_{max} by the upper part of $s^B_{max}(\beta)$.

Equilibrium with a Regular Saver Product

The nondivisible is purchased whenever $s^B_{max}(\beta) \geq s^B_{min}(\beta)$, which occurs for any $\beta \in [\hat{\beta}_B, 1]$. In Figure 2.1.2, the critical $\hat{\beta}_B$ is lower than the threshold $\hat{\beta}$ obtained from s_{min} and s_{max} in autarky: The nondivisible good can be purchased for a larger range of β when agents use the commitment product. For $\beta \in [\hat{\beta}_B, \hat{\beta})$, the regular saver product is welfare-enhancing as the nondivisible can be purchased with the product but not without it. The second relevant factor for welfare analysis is how evenly savings are spread across period 1 and 2. Note that in spite of the agent signing a contract to contribute the welfare-maximising amounts $s_1 = \bar{s}$, $s_2 = 2\bar{s}$, the equilibrium amount of savings need not adhere to this schedule. However, it can be shown that given a sufficiently large penalty, the

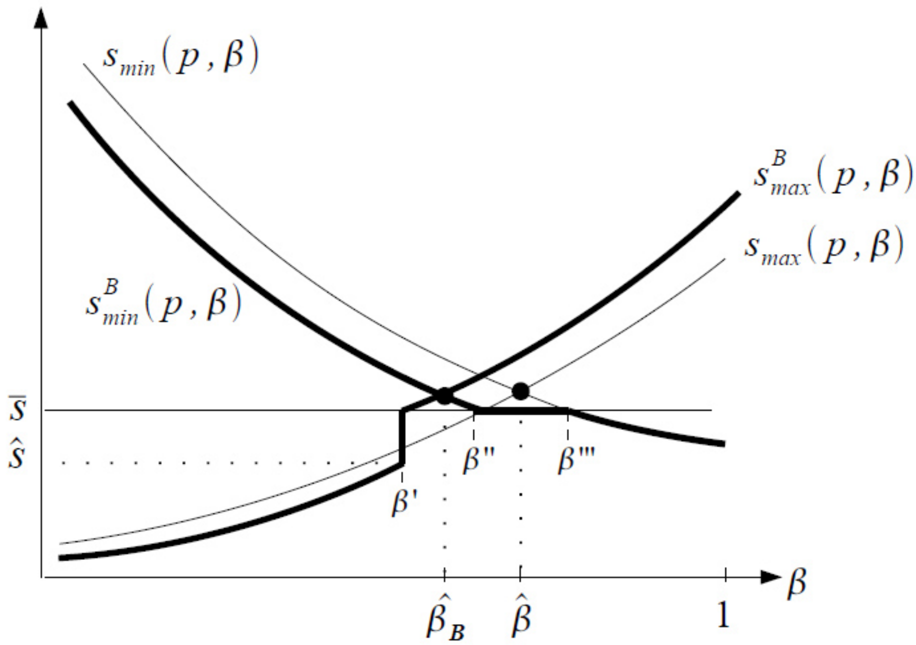

Figure 2: Equilibrium with a Regular Saver Product

regular saver product weakly decreases the difference between equilibrium s_1 and $\bar{s}$, thus increasing welfare by smoothing consumption.

Proposition 6. *Given a sufficiently large penalty $D \geq \bar{s} - s_{min}(1)$, introduction of the regular saver product increases the range of parameters where the nondivisible good can be purchased, i.e. $\hat{\beta}_B < \hat{\beta}$ holds.*

Proposition 7. *Given $D \geq \bar{s} - s_{min}(1)$, the regular saver product weakly decreases $|\bar{s} - s_1|$ for all $\beta \in [\hat{\beta}_B, 1]$.*

The situation looks different for low values of β, where the commitment product may even hurt agents: First, for β such that the nondivisible is bought neither with nor without commitment, agents are worse off as they pay a penalty for cancelling the contract. Second, if the intersection $s^B_{max}(\beta) = s^B_{min}(\beta)$ occurs on the lower part of the $s^B_{max}(\beta)$-curve, then $\hat{\beta} < \hat{\beta}_B$, so commitment *decreases* the range of β where the nondivisible can be bought. Intuitively, this happens when the agent was in an autarky equilibrium where the good was purchased with a large share of the savings burden on period 2. With the commitment product, if he is not willing to jump to $s_1 = \bar{s}$, he incurs the penalty, is willing to save even less in period 1, and may thus no longer be able to buy the good. The method to determine whether $\hat{\beta}_B < \hat{\beta}$ for a specific combination of $u(c_t)$, D, p, b is therefore to check whether at the autarky threshold $\hat{\beta}$, the agent is willing to pay $s_1 \geq \bar{s}$ when subjected to the penalty. If yes, the upper part of the $s^B_{max}(\beta)$-curve applies, and we can conclude $\hat{\beta}_B < \hat{\beta}$. Another way of stating the same criterion is to check whether $\beta' < \hat{\beta}$.

It has been shown that the product will hurt the agent for $\beta < min\left\{\hat{\beta}, \hat{\beta}_B\right\}$, and for $\hat{\beta} \leq \beta < \hat{\beta}_B$. For $\hat{\beta}_B \leq \beta < \hat{\beta}$, the product will enable him to buy the nondivisible, and for $\beta \in [\hat{\beta}_B, 1]$, the product will bring savings contributions weakly closer to the welfare-maximising $s_1 = \bar{s}$. The implications for

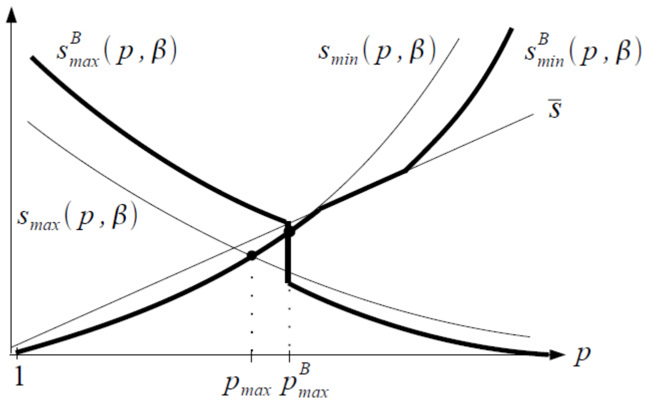

Figure 3: Achievable Savings Levels with a Regular Saver Product

take-up are straightforward: Since the contract is signed in period 0 rather than in period 1, the agent will choose the regular saver product whenever it increases his welfare, i.e. if and only if $\beta \in [\hat{\beta}_B, 1]$.[13]

Achievable Savings Levels with a Regular Saver Product

Returning to the original question which effect a regular saver product has on the maximum achievable savings level p_{max} for given time preferences β, it is analogous to show that p_{max} increases as long as the penalty is high enough. First, to ensure that $D < \bar{s} = \frac{p-1}{2}$ is true for all p, it is convenient to define D as a fraction of $\bar{s}$, $D = \lambda \frac{p-1}{2}$, $\lambda \in (0,1)$. Second, Section 2.1.1 has derived that s_{max} weakly decreases in p, while s_{min} weakly increases in p. The effect of the penalty on the two curves is analogous to the effect shown for variable preferences β:

$$s^B_{min} \equiv min\{s_1 \mid u(1+s_1-(p-1)) + \beta u(b) \geq u(1+s_1-D-s_2*) + \beta u(1+s_2*)\} \text{ for } s_1 \geq \bar{s}$$

Consider the above inequality at previous levels of s_{min} and s_{max}. When the penalty D is introduced, the right-hand side decreases, so the inequality will hold for $p_{max} + \epsilon$. Thus the previous s_{min} holds at higher levels of p, shifting the s_{min}-curve to the right (see Figure 3). For $s_1 < \bar{s}$, the contract has been cancelled in period 1, and period 2 no longer faces a penalty, thus the original s_{min} applies.

$$s^B_{max} \equiv max\{s_1 \mid u(1-s_1) + \beta u(2+s_1-p) + \beta u(b) \geq u(1-D) + 2\beta u(1)\} \text{ for } s_1 \geq \bar{s}$$

$$s^B_{max} \equiv max\{s_1 \mid u(1-s_1-D) + \beta u(2+s_1-p) + \beta u(b) \geq u(1-D) + 2\beta u(1)\} \text{ for } s_1 < \bar{s}$$

For s_{max}, the same logic shows that the upper part of the s_{max}-curve shifts right, while the lower part shifts left below the original s_{max} due to concavity of $u(c_t)$.

Analogously to the criterion for $\hat{\beta}_B < \hat{\beta}$, the new threshold p^B_{max} will be higher than p_{max} if and only if, at the original p_{max}, the agent is willing to pay $s_1 \geq \bar{s}$ when subjected to the penalty – i.e.

[13] A possible extension could make the optimal choice of $\bar{s}$ a function of β. A lower choice of $\bar{s}$ for low β may prevent the case of $\hat{\beta} < \hat{\beta}_B$. Purchasing the good would then require additional savings above $2\bar{s}$.

he jumps to the upper part of the s^B_{max}-curve. Note that this holds irrespective of the shape of the s^B_{min}-curve (as long as it increases), and of whether the original equilibrium s_1 is above or below $\bar{s}$. As with $\hat{\beta}_B$, whether $p^B_{max} > p_{max}$ depends critically on the size of the penalty. The bigger the penalty, the larger the shift in s_{max}, the higher the likelihood that the agent would rather pay $\bar{s}$ at the original p_{max} instead of incurring the penalty. Trivially, at $D = \bar{s}$, any $p < 3$ is achievable.

To summarize the take-up decision in the scenario with fixed β and variable p, the period 0 agent will adopt the product whenever it is welfare-maximising. This is the case whenever $2u(\frac{3-p}{2}) + u(b) \geq 3u(1)$, i.e. receiving the benefit b is worth paying p, and when $p \in [1, p^B_{max}]$, i.e. the product enables him to buy the nondivisible and smooth his savings contributions relative to buying the good in autarky.[14]

Extension: Adding Relatives' Claims

So far the model has assumed that the individual has full control over his assets. This may not be true in reality. A very crude way to capture this is to assume there is a negative return on savings kept at home, i.e. period 2 receives μs_1 and period 3 receives μs_2, where $\mu < 1$. In autarky, the s_{min}-curve will shift upwards, as the left-hand side of the inequality decreases more than the right-hand side due to concave utility. Further, s_{max} shifts down, thus increasing the threshold $\hat{\beta}$ (and decreasing p_{max}). The good can be bought less often with relatives' claims, which is intuitive given that saving is less attractive with negative returns. When a regular saver product is introduced, period 2 receives $\mu(s_1 - \bar{s}) + \bar{s}$ if $s_1 \geq \bar{s}$, and μs_1 otherwise. For $s_1 \geq \bar{s}$, the product now causes s_{min} to shift down for *two* reasons: Because the penalty D makes not saving less attractive, and because the product increases returns on $\bar{s}$ from $\mu$to 1. For the same two reasons, s_{max} shifts up, provided that $s_1 \geq \bar{s}$. For $s_1 < \bar{s}$, s_{min} is unchanged relative to autarky, while s_{max} shifts below its autarky level as the penalty appears on both sides of the inequality. The conclusion is the same as before: If, at the original $\hat{\beta}$ (or p_{max}), period 1 is willing to pay $\bar{s}$ when subjected to the penalty, then the commitment product leads to $\hat{\beta}_B < \hat{\beta}$, respectively $p^B_{max} > p_{max}$. In contrast to the situation without relatives' claims, the intervals $[\hat{\beta}_B, \hat{\beta}]$ and $[p_{max}, p^B_{max}]$ are likely to be larger, as s_{min} decreases more and s_{max} increases more in response to the product. In other words, relative to the original model, the product enables more agents to achieve higher savings levels due to its added benefit of protecting savings from relatives' claims. Finally, note that the banking equilibrium will occur at the point $s_1 = \bar{s}$ more often than without relatives' claims, i.e. horizontal and vertical parts of the curves are larger: While in the original model, the only pressure was not to save *less* than $\bar{s}$, now there is added pressure to not save *more* than $\bar{s}$ either in order to avoid negative returns to saving at home.

[14] As a possible extension, consider the case where the penalty is only enforceable in period 2 as a deduction from savings, but not in period 1. Instead, the agent pays D as an opening fee in period 0, which is reimbursed upon payment of $s_1 \geq \bar{s}$. For the period 0 agent, this will imply an additional condition $\beta \in [\beta_{min}, 1]$ for take-up as he now faces temptation. The condition for s^B_{min} is unchanged, while the condition for s^B_{max} becomes

$$\begin{aligned} u(1 - s_1 + D) + \beta u(2 + s_1 - p) + \beta u(b) &\geq (1 + 2\beta)u(1) \text{ for } s_1 \geq \bar{s} \\ u(1 - s_1) + \beta u(2 + s_1 - p) + \beta u(b) &\geq (1 + 2\beta)u(1) \text{ for } s_1 < \bar{s} \end{aligned}$$

The s_{max}-curve shifts up for $s_1 \geq \bar{s}$ relative to autarky and stays at its autarky level for $s_1 < \bar{s}$. Hence, the results of the model hold.

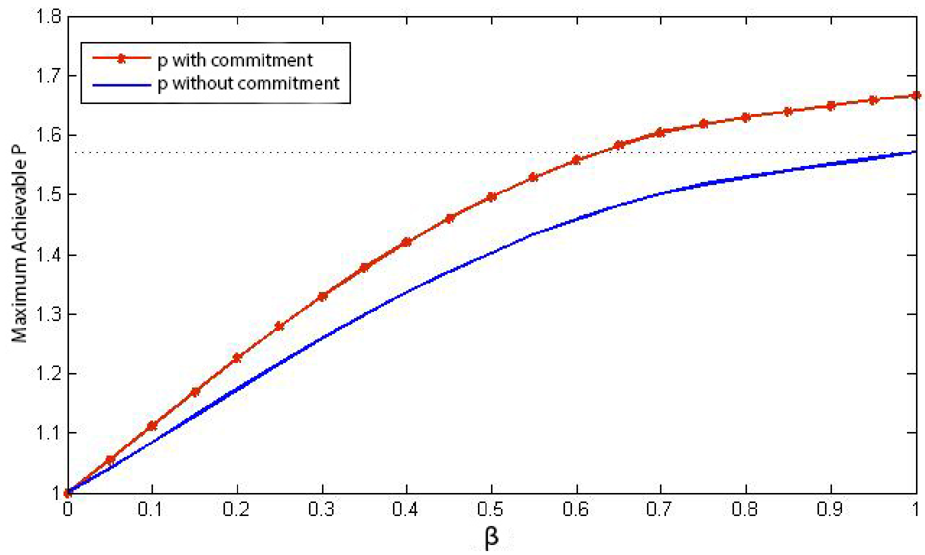

Figure 4: Numerical Illustration: Achievable Savings Levels

2.2 Numerical Illustration

For illustrative purposes, consider the following parameter specifications. Note that the penalty D has been set to $D = 0.5\bar{s}$ to ensure that $D < \bar{s}$ is true for all p.

$$
\begin{aligned}
u(c_t) &= -\frac{1}{c_t} \quad (i.e.\ CRRA\ with\ \gamma = 2) \\
b &= 5 \\
D &= \frac{\bar{s}}{2} = \frac{p-1}{4}
\end{aligned}
$$

Requiring that $s_{max} \geq s_{min}$ at given β and solving numerically for p_{max} yields maximum achievable savings levels, which are illustrated in Figure 2.2 for both autarky and when using the regular saver product. Note that by the condition $2u(\frac{3-p}{2}) + u(b) \geq 3u(1)$, the maximum price at which the good is welfare-improving is $\bar{p} = 11/7$ (the dotted line in Figure 2.2). Further, it can be shown that the optimal contribution $s_1 = \bar{s}$ is reached more frequently. Hence, the period 0 agent desires commitment for any $\beta \leq 1$ given an achievable price $p \leq min\,\{p_{max}, 11/7\}$.

3 The Data

The previous chapter has attempted to provide a theoretical foundation of why individuals subject to short-term consumption temptations and financial claims from relatives may benefit from a savings product which enables them to commit to fixed contributions every period. Section 4 will look at the empirical results of one such product: The Long Term Savings (LTS) product offered by *SafeSave*, an MFI located in Dhaka, Bangladesh. *SafeSave*, which was founded by microfinance researcher Stuart Rutherford in 1996, aims to provide day-to-day money management services such as basic savings and loan instruments to the poor.[15] In June 2009, *SafeSave* introduced the LTS product, enabling clients to save either 50Tk or a multiple of 100Tk (about U.S. $1.3) each month for a chosen contract length of 3, 5, 7 or 10 years. A particular feature of SafeSave is branchless banking - i.e. "field officers" visit each client's house every working day to collect general savings, and once a month LTS savings. The only time clients need to visit a bank branch is to open an account (general or LTS), or to take withdrawals or loans in excess of 500Tk. These field officers informed clients about the new LTS product in June 2009.

The analysis in this paper focuses on the 333 clients who opened an LTS account at a particular *SafeSave* branch (Bauniaband branch, the first to offer LTS) between June 2009 and February 2010, conditional on having had a general savings account since August 2008 or longer.[16] The variable of interest is the increase in clients' savings contributions after LTS adoption: Do clients save more when they use LTS? Data is available on clients' general savings level plus LTS savings as of 31 August and 31 January each year, between August 2006 and January 2010. This makes it possible to compare savings contributions made during the five month period from August to January, for the years 2006 to 2009.[17] During the last of these four periods, from August 2009 to January 2010, LTS was available. This paper explores the variation contained in the timing of LTS adoption: Those who adopted LTS soon after its introduction in June 2009 were "treated" with the product for the full five month period of interest. Those clients who adopted LTS later in 2009 received "partial treatment" in the sense that they used LTS for only a fraction of the period from August 2009 to January 2010. The analysis then compares whether the savings contributions of "fully treated" individuals increased more relative to previous years' contributions than the savings of "partially treated" or untreated individuals did. The obvious difficulty with approaching LTS as a treatment is that assignment into the treatment group does not occur randomly, as individuals select into the treatment group by adopting early. The effects that can arise from such endogeneity of treatment are discussed in detail in section 4.2 and Appendix A.

While the theoretical model described in section 2 provides a plausible foundation for the empirical analysis to follow, two areas of departure from the theory are worth mentioning: First, interest rates. Interest rates for LTS are 7, 8, 9, and 10 percent for a 3, 5, 7, and 10 year contract, respectively, where the average chosen contract length is 5.3 years. With inflation ranging around 9 percent in 2008 and 2009, this implies the real interest rate on LTS is currently negative or roughly zero. Hence, it seems unlikely that clients adopted LTS to increase the size of the pie, rather than for its commitment features. However, we cannot exclude the possibility that individuals simply needed to save for the future and chose LTS to evade inflation, not because they wanted or needed commitment. It

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

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

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

is helpful to note that the general savings account pays 6 percent interest without any commitment.[18] Given the uncertainty of most slum dwellers' incomes and consumption needs, it seems doubtful whether a client who does not desire commitment would sign a 5-year contract for an additional 2 percent of interest.

Secondly, the penalty upon cancelling the LTS contract differs from the simple structure assumed in the model: If a client is three months or less behind on payments, a "service fee" needs to be paid to catch up on payments. If he is more than three months behind, the LTS contract is cancelled, and the penalty imposed is that he loses all LTS interest payments above the general account interest rate, with retrospective effect since he opened his LTS. In terms of our theory, the three months delay can be captured by thinking of one period as three months. The penalty for cancelling LTS, however, is not an actual penalty, in the sense that agents receive the same returns which they would have received, had they never joined LTS. If agents are rational, this will reduce the efficacy of the product. For behavioural agents who exhibit reference dependence and loss aversion, the penalty would still be effective: While holding an LTS, individuals will adjust their reference point to the higher interest payments they believe to accumulate over time, and will be reluctant to give them up. The fact that only 8 out of 333 clients in our sample cancelled their LTS during the period of observation supports this explanation.

4 Estimation

4.1 The Differences-in-Differences Approach

An obvious first step to estimate the effect that LTS has on client's savings is to look at how much early adopters saved between August 2009 and January 2010, i.e. while using LTS, and compare this to their savings during earlier August to January periods. To isolate time effects, this change in savings can be compared to the change in savings of a control group during the same time period. The resulting differences-in-differences (DD) estimate can be interpreted as a causal treatment effect only under the strong assumption that the treatment group would have evolved identically to the control group, had they not been treated (i.e. had they not used LTS). We will discuss the plausibility of this assumption in Section 4.2. The first part of the analysis considers all clients who adopted LTS between June 2009 and September 2009 ("early adopters") as "treated", while those clients who adopted LTS between October 2009 and February 2010 ("late adopters") are used as a control group. Note that LTS deposits must be made on the 20th of each month, and "adopting LTS in September" means that clients chose to first contribute to LTS on 20th of September. Hence, clients adopting LTS in September can be considered as "fully treated" in the sense that they contributed on all LTS payment dates between 31 August 2009 and 31 January 2010.

Table I presents descriptive statistics of clients' characteristics (as far as available), the LTS contract options they chose, and their average savings contributions by group. Column (1) of Table II presents the regression estimates of a DD approach generalized to multiple time periods, which at this point

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

should be interpreted as correlations only. The regression equation is

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

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

The cut-off between early and late adopters may seem arbitrary, given that joining LTS was possible every month. As an alternative, we can generalize the concept of a treatment and a control group to allow for separate treatment groups depending on month of LTS entry. This allows us to test whether savings increase in the number of months that a client used LTS – in other words, whether the coefficient on treatment in the savings equation increases in the number of months that a client was "treated" with LTS.[19] The regression equation becomes

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

where m indexes months of treatment. The equation contains a vector of treatment group effects γ_m (dummy T_i is one if individual i was "treated" m months), a vector of time effects λ_t, and a vector of treatment effects β_m, where D_{it} is a vector of dummies that switches on for observations in 2009 for individuals treated for m months. Treatment varies at the individual level, and we explicitly assume that observations are independent at the individual level, i.e. there is no group structure to the time-individual specific errors ϵ_{it} (see Section 4.3 for concerns regarding the error structure).

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

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

TABLE I: DESCRIPTIVE STATISTICS

	Early Adopters	Late Adopters	All
Age	30.6	32.9	31.4
Female	89.7	89	89.5
Occupation (percent)			
Housewife	66.5	52.3	61.9
Small Business	6.7	15.6	9.6
Private Service	11.6	9.1	10.8
Handicrafts	6.7	12.8	8.7
Student	8.5	10.2	9.0
LTS Entry Month (percent)			
June 2009			27.9
July 2009			17.7
August 2009			11.4
September 2009			10.2
October 2009			7.2
November 2009			4.5
December 2009			3.0
January 2010			11.7
February 2010			6.3
LTS Contract Details			
Contract Length (years)	5.1	5.76	5.3
Monthly Deposit (Taka)	210	235	218
Purpose of LTS Contract			
Education	56	51.4	54.6
Ornaments	1.4	1.8	1.5
Land Purchase	5.8	9.2	6.9
House Construction	1.3	0.9	1.2
Medical Treatment	4.9	0.9	3.6
Marriage	10.7	12.8	11.4
Business	19.9	23	20.8
Mean Savings Contributions			
Aug 2006 - Jan 2007	318	197	277
Aug 2007 - Jan 2008	323	301	316
Aug 2008 - Jan 2009	292	285	290
Aug 2009 - Jan 2010	1064	499	879
Observations	224	109	333

Notes: 1) "Early Adopters" are all clients who adopted an LTS contract at Bauniaband branch between June and September 2009. "Late adopters" are all clients who adopted an LTS contract at Bauniaband branch between October 2009 and February 2010.
2) While income data was not available on client level, bank staff estimate average monthly household income between 5000 and 6000 Tk.

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

	(1)	(2)	(3)
Treatment group*2009	523.1***	654.6***	750.9***
	(195)	(209)	(264)
Treatment group	42.45	39.90	148.3
	(58.3)	(61.9)	(240)
Treatment group trend			-51.16
			(106)
Female		4.513	4.570
		(174)	(134)
Age		24.70*	24.70*
		(13.0)	(14.6)
Age²		-0.308**	-0.307*
		(0.15)	(0.18)
Interaction Area*2009	NO	YES	YES
Time effects	YES	YES	YES
Occupation FE	NO	YES	YES
Observations	1185	1185	1185
R-squared	0.06	0.08	0.08
Robust SEs	YES	YES	NO

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

4.2 Endogeneity Concerns

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

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

where z_i is unobserved, the classic two-period DD estimator $\hat{\beta} = \Delta\bar{S}_t^T - \Delta\bar{S}_t^C$ differences out any effects that are constant either across time or across individuals. As an illustration, suppose that higher income causes early take-up, and higher income also causes higher savings levels, but it does not cause different savings time trends. Then the effect of income is time-invariant, and thus income

TABLE III: IMPACT OF A REGULAR SAVER PRODUCT – MULTIPLE TREATMENT GROUPS

Dependent Variable: Change in Total Savings from August (t) to January (t+1)

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

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

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

will not bias the DD estimator. The real problem for the DD approach adopted are unobservables which cause both early adoption *and* a change in the savings time trend. The following decomposition for the two-period case illustrates this: The treatment effect that we are ultimately interested in is

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

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

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

which can be decomposed into

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

In other words, the identifying assumption is that

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

The estimator is biased if the (counterfactual) change in savings for the treatment group absent treatment is different from the change in savings for the control group.[20] Given the lack of a suitable instrumental variable, this paper will not be able to fully solve the endogeneity problem. Instead, the remainder of this section will take a closer look at the selection bias by suggesting possible differences between early and late adopters, controlling for them where possible, and discussing their implications for savings behaviour.

Shedding Light on LTS Adoption

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

Table V suggests that early adopters were more likely to choose shorter contracts, but committed to roughly similar monthly contributions. They also opened their general bank accounts at the same time. Table V reveals that while age, gender and occupation had little effect on adoption timing, the area where clients live strongly influenced when they took up the product. The importance of area may seem striking, considering that *SafeSave* only operates within a one kilometre radius of a bank branch (in this case, Bauniaband branch), so the areas considered are relatively small and

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

TABLE IV
PREDICTING LTS ADOPTION TIMING

	Dependent Variable: LTS Entry Month (1=June 2009)
Female	-0.0859
	(0.55)
Age	-0.00975
	(0.061)
Age²	0.000243
	(0.00077)
Occupation FE	
Housewife	0.151
	(0.65)
Handicrafts	0.972
	(0.77)
Small Business	0.462
	(0.80)
Private Service	0.0203
	(0.75)
Area of Residence	
Bauniaband Block B	-1.302***
	(0.43)
Bauniaband Block C	-1.396***
	(0.42)
Bauniaband Block D	-0.587
	(0.49)
Bauniaband Block E	-0.574
	(0.48)
Kurmitola	2.374***
	(0.90)
Mirpur	1.617*
	(0.92)
Constant	4.228***
	(0.97)
Observations	333
R-squared	0.12

*significant at 10 percent; **significant at 5 percent; *** significant at 1 percent.
Notes: The omitted category for occupation is student. The omitted category for area is Bauniaband Block A.

TABLE V
LTS CONTRACT CHOICES

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

*significant at 10 percent;
**significant at 5 percent;
***significant at 1 percent.
Notes: The omitted category for purpose is Business.

close together. At first sight, this may look like an omitted variable bias – for instance, area may be correlated with income, which may influence take-up. While we do not have individual income data, observational evidence from field officers who visit each client every day suggests that all areas are comparatively poor. Second, insights such as those from Bertrand et al. (2004b) suggest that situational factors such as walking distance to the branch (to sign an LTS contract) may be important for take-up. Walking distances are Block A (5min), Block B (10min), Block C (16min), Block D (12min), Block E (12min), Kurmitola and Mirpur (5-15min). With Block C being the farthest from the branch and at the same time the earliest to adopt LTS, the distance explanation does not seem very likely.

There are various other explanations for a "patchwork pattern" of adoption, such as herding behaviour and other mechanisms of technology adoption.[21] Appendix A discusses these explanations in detail and shows that, in particular with herding behaviour, random factors may lead to a profitable innovation spreading very quickly in one area, and not at all or with considerable delay in another area, even if individuals are completely homogeneous in their savings preferences (implying zero selection bias). Factors which may determine adoption timing independent of area include different degrees of risk aversion and the amount of benefit which individuals derive from the product: Suppose information about the manageability of LTS is a public good. The theory in this paper then suggests that early adopters may be those who can reach their desired savings goal p with LTS but not without it, i.e. $\beta \in [\hat{\beta}_B, \hat{\beta})$ – they derive the largest benefit from the product and have the highest stake in learning about it. Individuals who could have reached their p in autarky, but at less smooth savings contributions, may still adopt LTS, but wait for others to adopt first. Appendix A elaborates on this argument and shows that the implications for a possible selection bias are generally ambiguous, instead pointing towards heterogeneous treatment effects.

Unfortunately, the data do not allow us to disentangle further which factors were at work in adoption timing. Most plausibly, it was a mixture of them plus some factors which have not been discussed – like shocks that occurred simultaneous with treatment. In the end, what matters for our estimate of the treatment effect is whether the factors which determined adoption also determined savings trends absent treatment. The following subsection will attempt to control for two sources of selection bias.

4.3 Adjusting the Estimation Strategy

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

[21] Compare Besley and Case (1994) and Foster and Rosenzweig (1995).

A second general source of selection bias that we can control for are pre-existing trends. In other words, if early adopters are different from late adopters in a way that causes them to follow different savings time trends before treatment, we can control for this by including a time trend specifically for the treatment group. Column (3) of Tables II and III show that the treatment effects are generally robust to including treatment-specific time trends. Time trends are mostly insignificant, suggesting that savings patterns of the treatment and the control group were comparable before treatment. Nevertheless, the power of this test may be low given the small sample size and the availability of only four years of data. Clients treated for four months (i.e. October entrants) constitute a special case as their treatment effect is fully absorbed by a large negative group effect and a large positive time trend. This is almost certainly due to the fact that the regressors explain only four periods of data and would disappear with a longer time series.

Unfortunately, an important source of bias remains in time-individual specific shocks which occurred simultaneous with LTS introduction: For instance, if there was a change in circumstance for a particular group of people which caused them to both adopt early and increase their savings (absent treatment), then the estimate of the treatment effect is biased upwards. Unless these shocks target specific age, occupation or area groups, the scarce nature of the data will not allow us to control for the resulting bias. As a consequence, the treatment regressor is plausibly still correlated with the time-individual specific error. The logical next step is an instrumental variable analysis – but lacking a suitable instrument, this step has to be postponed to a point when more comprehensive data is available.

Concerns about Standard Errors

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

Another frequent concern in the DD literature is clustering of the error terms. Grouped error terms are a concern when the outcome of interest varies at individual level while treatment varies only at group level.[22] This is not the case here, as both savings and treatment vary for each individual.

[22] See Angrist and Pischke (2009), Chapter 8.2.

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

5 Conclusion and Outlook

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

In attempting to understand the nature of the endogeneity, the paper finds a significant effect of area on adoption timing, but no correlation of adoption timing with other savings choices such as monthly contribution or bank entry timing. I also estimate that early and late adopters follow the same savings trends before treatment. A story of herding behaviour or other technology adoption mechanisms would be consistent with this evidence, explaining differences in adoption timing without necessarily implying a selection bias. However, the endogeneity problem cannot be fully resolved due to possible unobservable time-individual specific effects.

The theoretical model presented in the first part further suggests that heterogeneity of treatment effects may be an issue: While the benefit of the product may consist in reaching a previously unattainable savings goal for some individuals, it may simply consist in smoothing savings contributions for others.

The objective of the paper is to provide a motivation for future research into regular saver products, and to tackle the plethora of questions it does not answer: First, a randomized controlled experiment would help to account for unobserved determinants of participation by giving the product only to a fraction of those who want it, and using the others as a control group. Second, the present analysis lacks data on individuals' time preferences. Surveys or lab experiments are needed to identify individuals' time preferences as well as the strength of the social claims on their savings, in order to link hyperbolic discounting and relatives claims' to a demand for regular saver products. Third, it is desirable to more explicitly compare the performance of different types of commitment savings products – e.g. high frequency versus low frequency deposit products (daily versus monthly contributions), versus the pure withdrawal-restriction products which have been a subject of academic interest in recent years.

An important aim of this strand of research has to be the advancement of microfinance products which are tailored specifically to the needs of the poor. Examples such as that of Bank Rakyat Indonesia, who has 16 million low-income depositors versus two million borrowers,[23] and the famous Grameen bank, who holds $1.6 in savings deposits for every dollar it lends, illustrates that microfinance reality has long realized the importance of savings.

References

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.

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

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

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

Banerjee, Abhijit V., "A Simple Model of Herding Behavior," *Quarterly Journal of Economics*, 1992, *107(3)*, 797–817.

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

Basu, Karna, "Commitment Savings in Informal Banking Markets," 2012.

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

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

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119(1)*, 249–275.

_ , Sendhil Mullainathan, and Eldar Shafir, "A Behavioral Economics View of Poverty," *American Economic Review*, 2004, *94(1)*, 419–423.

Beshears, John Leonard, James J. Choi, David Laibson, Brigitte C. Madrian, and Jung Sakong, "Self Control and Liquidity: How to Design a Commitment Contract," Working Paper Series WR-895-SSA, RAND November 2011.

[23]Morduch (1999).

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

— **and Anne Case**, "Diffusion as a Learning Process: Evidence from HYV Cotton," Papers 174, Princeton, Woodrow Wilson School - Development Studies 1994.

— **and** — , "Unnatural Experiments? Estimating the Incidence of Endogenous Policies," *Economic Journal*, 2000, *110(467)*, F672–F694.

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

Brune, Lasse, Xavier Giné, Jessica Goldberg, and Dean Yang, "Commitments to Save: A Field Experiment in Rural Malawi," Policy Research Working Paper Series 5748, World Bank 2011.

Bryan, Gharad, Dean Karlan, and Scott Nelson, "Commitment Devices," *Annual Review of Economics*, 2010, *2(1)*, 671–698.

Collins, Daryl, Jonathan Morduch, Stuart Rutherford, and Orlanda Ruthven, *Portfolios of the Poor - How the World's Poor Live on $2 a Day*, Princeton: Princeton University Press, 2009.

Duflo, Esther, "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment," *American Economic Review*, 2001, *91(4)*, 795–813.

Dupas, Pascaline and Jonathan Robinson, "Why Don't the Poor Save More? Evidence from Health Savings Experiments," Working Paper 17255, National Bureau of Economic Research 2011.

— **and** — , "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya," *American Economic Journal: Applied Economics*, 2013, *5(1)*, 163–192.

Foster, Andrew D. and Mark R. Rosenzweig, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, 1995, *103(6)*, 1176–1209.

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.

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman, "Getting to the Top of Mind: How Reminders Increase Saving," Working Paper 16205, National Bureau of Economic Research July 2010.

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

Morduch, Jonathan, "Between the State and the Market: Can Informal Insurance Patch the Safety Net?," *World Bank Research Observer*, 1999, *14(2)*, 187–207.

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

Rutherford, Stuart, "The Poor and Their Money," FDRP Working Paper 3/99, University of Manchester, Institute for Development Policy and Management 1999.

Appendix A: Explaining Adoption Timing

Section 4.2 has pointed out that area of residence strongly influenced when clients took up the product. However, early adoption areas do not seem richer than others and are not closer to the bank branch. One possible factor at work may be a herding behaviour story in the style of Banerjee (1992): Consider homogeneous individuals who decide in a sequential fashion whether or not to adopt LTS. The sequential decision process could be due to reasons exogenous to savings, e.g. because families need to find time to sit down at the kitchen table and solve the optimization problem whether or not they want to commit to a 5 year contract.[24] Assume there is a network structure within each block or area (but not between areas): Communication between neighbours is good enough that agents know whether someone in their block has chosen to i) adopt, ii) not adopt, iii) is yet undecided. Everyone has a private signal whether the LTS product is good or bad for them: "Good" means it will be manageable with their budget, and will improve their situation. "Bad" means it will prove to be unmanageable, taking away slack from their budget, causing transaction cost and withdrawal fees, and constrain them until they decide to cancel it without reaping any returns. Assume further that the true return to the product is the same for everyone and positive, and that a large majority of individuals (say, 90 percent) receive a good signal. Finally, add a tie-breaking rule that someone who is indifferent will not bother to go to the bank to adopt the product.

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

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

The consequence is that even with completely homogeneous individuals, a profitable innovation may spread very quickly in one area, and not at all or with considerable delay in another area, depending on the signals of the first one or two individuals who find the time to optimise. What

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

would this explanation imply for our results on the treatment effect? If random factors determine who chooses first, then early adopter areas need not be different at all from late adopter areas in their savings behaviour, and what we are observing is an actual treatment effect. This explanation is supported empirically by the fact that early adopters did not choose higher monthly contributions than late adopters did, as can be seen from the lack of correlation between entry month and monthly deposit in Table V.

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

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

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

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

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

Appendix B: Proofs

Proposition 1: *$u(1+s_1-(p-1)) + \beta u(b) \geq u(1+s_1-s_2*) + \beta u(1+s_2*)$ holds if s_1 is bigger than some threshold value, $s_1 \geq s_{min}$.*

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

$$u(1+s_1'-(p-1)) + \beta u(b) \geq u(1+s_1'-s_2) + \beta u(1+s_2) \quad \text{for } \forall s_2 < p-1.$$

Since the equation holds for all $s_2 < p-1$, it will also hold for $s_2 * (s_1'')$ - the s_2 that is optimal at a higher level $s_1'' > s_1'$, conditional on the nondivisible not being bought. Due to strict concavity of

$u(c_t)$, we further know that

$$u(1+s_1'-s_2*(s_1''))-u(1+s_1'-(p-1))\geq u(1+s_1''-s_2*(s_1''))-u(1+s_1''-(p-1)),$$

i.e. the consumption gain $(p-1)-s_2*$ from deciding not to save for the good in period 2 gives a higher utility gain when starting from consumption level $1+s_1'$ than when starting from consumption level $1+s_1''$. Since

$$\beta u(b)-\beta u(1+s_2*(s_1''))\geq u(1+s_1'-s_2*(s_1''))-u(1+s_1'-(p-1))$$

is true by the optimality of buying the good at s_1', substitution yields

$$u(1+s_1''-(p-1))+\beta u(b)\geq u(1+s_1''-s_2*(s_1''))+\beta u(1+s_2*(s_1''))\quad \text{for all } s_1''>s_1'.$$

Hence, once s_1 crosses some threshold s_{min}, saving for the good will be optimal for all $s_1\geq s_{min}$. □

Proposition 2: *s_{min} is weakly decreasing in β.*

Proof. Consider a $\beta<1$ and its corresponding $s_{min}(\beta)$. We know that

$$u(1+s_{min}-(p-1))+\beta u(b)\geq u(1+s_{min}-s_2)+\beta u(1+s_2)\quad \text{for } \forall s_2<p-1.$$

Increasing β to $\bar{\beta}>\beta$ at unchanged s_1 will always preserve the inequality, since $u(b)>u(1+s_2)$ for $\forall s_2<p-1$, and generally allow for a lower s_{min}. Thus, s_{min} is weakly decreasing in β. □

Proposition 3: *s_{max} is weakly increasing in β.*

Proof. For a given β,

$$u(1-s_1)+\beta u(2+s_1-p)+\beta u(b)\geq(1+2\beta)u(1)$$

will hold for $s_1=s_{max}(\beta)$. Given that $u(1-s_1)<u(1)$, it must be true that $u(2+s_1-p)+u(b)\geq 2u(1)$. This statement does not depend on β, so it will also hold for a higher $\bar{\beta}>\beta$, at unchanged s_1. Can the same argument be used to show that the inequality also holds for a lower $\tilde{\beta}<\beta$? No: As β decreases, the weight of the inequality $u(1-s_1)<u(1)$ will rise, which will eventually flip the inequality in lifetime utility into the opposite direction. Hence, $s_{max}(\beta)$ is weakly increasing in β. □

Proposition 4: *s_{min} is weakly increasing in p.*

Proof. Recall that for $\forall s_1\geq s_{min}$,

$$u(1+s_1-(p-1))+\beta u(b)\geq u(1+s_1-s_2*(s_1))+\beta u(1+s_2*(s_1)).$$

Consider a lower $p'<p$. At unchanged s_1, the left-hand side of the equation unambiguously increases. The optimal s_2* conditional on not buying the good only depends on s_1 and not on p, thus

the right-hand side is unchanged.[25] Hence, the inequality will always hold. So s_{min} never increases for lower p, hence $s_{min}(p, \beta)$ must be weakly increasing in p. $\square$

Proposition 5: *s_{max} is weakly decreasing in p.*

Proof. Recall that for $\forall s_1 \leq s_{max}$,

$$u(1 - s_1) + \beta u(2 + s_1 - p) + \beta u(b) \geq (1 + 2\beta)u(1).$$

Consider a lower $p' < p$. At unchanged s_1, the left-hand side of the equation unambiguously increases, while the right-hand side is unchanged. Therefore, the inequality always holds, so s_{max} never decreases for lower p, thus $s_{max}(p, \beta)$ must be weakly decreasing in p. $\square$

Proposition 6: *When $\hat{\beta}_B < \hat{\beta}$, introduction of the commitment product weakly decreases $|\bar{s} - s_1|$ for all $\beta \in [\hat{\beta}_B, 1]$.*

Proof. For $\beta < \hat{\beta}_B \leq \hat{\beta}$, savings are zero both in autarky and the banking case. For $\beta \geq \hat{\beta}_B$, if equilibrium savings in autarky are within $[\hat{s}, \bar{s}]$, introduction of the penalty leaves s_{min} unchanged and increases s_{max}. Period 1 will never save $s_1 \in (\hat{s}, \bar{s})$, but instead jump to $s_1 = \bar{s}$, leading to an even distribution of savings. If the equilibrium savings in autarky are above $\bar{s}$ (such as for $\hat{\beta}$ in Figure 2), the commitment product decreases s_{min} and increases s_{max}. Conditional on $\bar{s} \leq s_{min} \leq s_{max}$, equilibrium savings are $s_1 = max\{s_{min}, s_{opt}\}$, where $s_{opt} \leq \frac{p-1}{2} = \bar{s}$ ($s_{opt} = \frac{p-1}{2}$ for $\beta = 1$ and then decreases in β).[26] Thus, if $s_1 > \bar{s}$ in autarky, then it must be that $s_{min} > s_{opt}$, so a regular saver product will reduce equilibrium s_1 towards $\bar{s}$, improving consumption smoothing from a welfare perspective. $\square$

[25] An exception arises if $p' - 1 < s_2 * (s_1) < p - 1$, i.e. the optimal way to split period 1's savings is already sufficient to buy the good at the new lower price. In this case the agent will definitely buy the good, so the statement that s_{min} is sufficient at the lower p' still holds.

[26] Basu (2012).