Revising Commitments: Time Preference and Time-Inconsistency in the Field*

Xavier Giné, Jessica Goldberg, Dan Silverman, and Dean Yang†

July 18, 2011

Abstract

A new interest in intertemporal choice is fueled by evidence of non-constant time discounting and a much better understanding of its theoretical consequences. This paper reports the results of a field experiment in rural Malawi designed to evaluate, directly, the relationship between intertemporal choice under commitment and time-inconsistency. The experiment used real choices over money to elicit time preferences from approximately 2,200 subjects. The stakes of the intertemporal choices were high, representing a total of about one month's wages, on average. A randomly selected subset of respondents was later approached in the two weeks prior to their first disbursement of money. These subjects were reminded of their initial choices and asked if they would like to revise them and thus exhibit time-inconsistency. Choices under commitment reveal substantial heterogeneity in the consistency of choice with the law of demand, time preference, and dynamic consistency. When the commitment is broken, revisions are common, and they are only somewhat more likely to reallocate income backward in time. Following the theory, revisions are correlated with preferences under commitment.

---

*Niall Keleher, of Innovations for Poverty Action, was instrumental to the design and implementation of this study. We thank him for his important contributions to this project. We also thank James Andreoni, Stefano DellaVigna, Yoram Halevy, Glenn Harrison, Pam Jakiela, Damon Jones, David I. Levine, Stephan Meier, Ted Miguel, Matthew Rabin, and participants in several seminars for their many helpful comments.

†Giné: Development Economics Research Group, World Bank (xgine@worldbank.org); Goldberg: Ford School and Department of Economics, University of Michigan (jegoldbe@umich.edu); Silverman: Department of Economics, University of Michigan and NBER (dansilv@umich.edu); Yang: Ford School and Department of Economics, University of Michigan and NBER (deanyang@umich.edu).

1 Introduction

The long-standing interest in intertemporal choice has, in recent years, been further fueled by evidence of non-constant time discounting and a better understanding of its theoretical consequences. Several studies, drawing mostly on experimental data, can be interpreted to show that time discount rates decline as tradeoffs are pushed into the temporal distance.[1] Specifically, many of these studies document dynamic preference reversals: under commitment, subjects choose the larger and later of two rewards when both are distant in time, but prefer the smaller and earlier one as both rewards draw nearer to the present.

Interpreted as non-constant time discounting, these preference reversals have important implications. When utility is time-separable, non-constant time discounting implies time-inconsistency; the choices (plans) that a person makes now about consumption at later date are different from the choices he would make when that date arrives. Self-control problems and a demand for commitment thus emerge. If plans set at some earlier point will not be followed, then sophisticated decision-makers will want to limit their own ability to revise decisions about the future. The important consequences of this time-inconsistency and its associated self-control problems have generated a great deal of interest. They have now been studied, with both theoretical and empirical methods, in many different contexts.[2]

While time-inconsistency and self-control problems drive much of the interest in new models of non-constant time discounting, to our knowledge there has been no attempt to evaluate, directly, the link between time discounting and time-inconsistency. Do those who appear to have non-constant time discounting also exhibit time-inconsistency? In this paper, we present the results of a field experiment set in rural Malawi that we designed to address this question.

Why wouldn't those who exhibit preference reversals under commitment also always exhibit time-inconsistency? Several reasons are possible. First, some preference reversals may reflect predictable changes in the marginal utility of consumption. For example, even in a very standard model with exponential time discounting, if income rises (falls) with time and consumption smoothing is incomplete, then a consumer may appear more (less) patient regarding tradeoffs later when the marginal utility of consumption is lower (higher).[3] Such a consumer would not, however, exhibit time-inconsistency. Second, some preference refersals may reflect inattention, confusion about

[1] See Ainslie (1992), Thaler (1991) and several papers in Loewenstein and Elster (1992), for reviews of this evidence.

[2] Early contributions include Phelps and Pollak (1968), Laibson (1997) and O'Donoghue and Rabin (1999). See DellaVigna (2009) for a recent review of empirical applications.

[3] This observation is a special case of the more general point made by Andersen et al. (2008), Andreoni and Sprenger (2010) and Noor (2011), that proper inference about time discounting requires information about the curvature of the utility function.

tradeoffs, or responses to perceived experimenter demands.[4] It is unclear whether consumers who make such plans under commitment would also tend to revise them later, one way or the other, if the commitment is relaxed. Third, even if preferences under commitment were well-described by changing time discount rates, the simple act of making a plan may importantly limit what would otherwise be important problems of self-control.[5] Concerns about self-image and the importance of "following through" might trump the effects of changing time discount rates. In this case, again, dynamic preference refersals under commitment would not predict time-inconsistency.

These potential wedges between the evidence for non-constant time discounting and time-inconsistency motivate us to test whether, indeed, dynamic preference reversals under commitment predict the revision of choices when the commitment is unexpectedly relaxed.[6] This test is especially relevent in light of the increasing interest in the design of commitment savings products in developing countries. Growing evidence indicates that commitment savings instruments increase average saving and investment in low-income countries. See, e.g., Ashraf et al. (2006) and Brune et al., 2011. These commitments presumably work by alleviating self- and other-control problems.[7] However, such commitments also reduce flexibility and may, in some cases, restrain too much. By evaluating the link between intertemporal choices under commitment and time-inconsistency, and other motives for revising prior choices, our experiment can inform the design and targeting of future savings products, and further illuminate the mechanisms behind the demand for and consequences of commitment.

1.1 Overview of the Experiment

Our experiment involved approximately 1,100 households (2,200 individuals) in rural Malawi and proceeded in two stages. In stage one, we adapted the methods of Andreoni and Sprenger (2010) to elicit time preferences under commitment. The head of household and his spouse each

[4] The correlation between test scores, cogntive load and short-term patience in Benjamin et al. (2006) lends some support to this conjecture.

[5] Academic psychology points to the important effects of making plans or setting goals on self-control and "self-efficacy." Bandura (1997). Evidence of the importance of planning is also seen in Ameriks et al. (2003). This idea is also broadly consistent with economic models of costly self-control such as Gul and Pesendorfer (2001), Ozdendoren et al. (2009), and Fudenberg and Levine (2010), in which consumers may both seek commitment and, yet, not always exhibit time-inconsistency.

[6] Correlating dynamic preference reversals with the demand for commitment, as in Ashraf et al. (2006), serves a similar purpose. It also tests a key, but ancillary, prediction of changing time discount rates. Relative to this test, however, ours has the advantage of being robust to naivete. The qualitative link between changing rates of time discount and time-inconsistency does not depend on the sophistication.

[7] Citations about a demand for commitment deriving from a desire to protect assets from others.

made several independent choices about the allocation of a substantial amount of money over time. Each choice was an allocation of an endowment between two periods, one sooner and one later. Participants were paid a return on the part of the endowment that was saved for later. Each participant made 10 allocation choices; one choice facing each of five different rates of return and regarding two different time frames, one in the near future and one in the more distant future. Choices regarding the near future allocated money between tomorrow and 30 days from tomorrow. Choices regarding the more distant future allocated money between 60 days and 90 days from tomorrow. All choices were incentivized; participants knew that either the household head or his spouse would be randomly selected to have one of his or her, randomly selected, choices implemented. The stakes were high. On average, the total allocation amounted to approximately a month's wages for a worker in rural Malawi.

Stage two of the experiment applied only to those households whose implemented choice concerned the more distant future. At a randomly selected day in the two weeks prior to the arrival of the first disbursement of their money, the household was *unexpectedly* revisited. At that visit, the participant whose choice was implemented was reminded of the decision he or she made 45-60 days earlier under commitment. That previous choice was made clear and salient. Then, facing the same rate of return, the participant was allowed to revise the original allocation decision. That (un)revised decision was then implemented with certainty. Surveys at both stages measured household wealth, income, expenditures as well as the participants' expectations for each of these variables.

1.1.1 The Setting

As a setting for experimental study of intertemporal choice, rural Malawi has a number of advantages. Most important is that financial markets in the area are thin and participants thus lack effective methods for smoothing the relatively large amounts of new consumption that the experiment makes possible.

At the time of year when the experiment was implemented (the rainy season), borrowing substantial amounts of cash in rural Malawi is not merely expensive, it often appears effectively impossible. Similarly, short-term saving is difficult due to limited banking institutions, and familial or social demands for what appears like excess cash.[8] The lack of borrowing and saving opportunities is important because it sharply reduces the smoothing opportunities that should confound efforts to elicit time preferences from standard experiments in developed economies. In economies

[8] Longer-run saving instruments with positive rates of return are more common. They typically take the form of agricultural inputs.

with thick financial markets and low transaction costs, answers to the questions asked in typical experiments should, in theory, bracket only the market rates of return that participants face, and reveal very little about their "true preferences" (Fuchs, 1982).[9] To the extent that rural Malawi in the rainy season approximates autarchy, smoothing is made much more difficult. This is especially true for the participants in our experiment because the stakes are so large and involve cash rather than grain or other consumption goods. Subjects will find it very difficulty to smooth such a large amount experimental income by consuming from their own stores. In this way, studying intertemporal allocations in rural Malawi provides data that, in principle, is much closer to the theoretical concepts of interest: preferences over the time-allocation of consumption.

Rural Malawi also has disadvantages as a location for experiments on time preference. The region's low population density and relatively poor infrastructure make some experimental logistics difficult. One important consequence is that payments to participants can arrive no sooner than one day after they make their choices. Thus we cannot study preferences regarding consumption in the present; and we cannot observe the consequences of changes in time discounting that occur as just as intertemporal tradeoffs are made strictly later than the present. To the extent that changes in time discounting are largest then, we would expect any relationships between choice under commitment and revision behavior to be attenuated.

Another potential disadvantage of the setting is that participants have low levels of formal education and may therefore find the somewhat abstract experiment especially difficult to grasp, or view it much differently than we would expect. Participants make a living from seasonal crops, so they are uncommonly familiar with the problem of smoothing consumption over time. Nevertheless, because we asked them to make choices in an unfamiliar context, it is natural to worry about participants' abilites to quickly understand the experiment as an economic decision. For this reason, our analysis takes special care to evaluate the consistency of participants' choices with a basic feature of rational economic decision-making – the law of demand.[10] We interpret the degree of consistency with the law of demand as a measure of participants' understanding the trade-offs involved in their decisions.[11]

[9] To illustrate, suppose that outside of the lab a participant can borrow or save at market rate r. And suppose the participant faces no financial transaction costs. A typical experiment asks the participant to choose between $\$x$ sooner or $\$(1+r_e)\,x$ later, where r_e denotes the rate of return implied by the later option. The participant may view this as a choice between Option A, $\$x$ sooner and access to the interest rate r, and Option B, $\$(1+r_e)\,x$ later and access to the interest rate r. If $r_e > r$, then the set of allocations of consumption between sooner and later is strictly larger under option B than under option A.Thus, for any monotonic time preferences, option B is preferred. Analogously, if $r_e < r$, then A is preferred for any monotonic time preferences.

[10] Cf. Choi et al., (2007, 2011).

[11] Alternatively, one could interpret consistency with the law of demand as measure of the appropriateness of

1.2 Summary of Findings

Judging by the consistency of their decisions with the law of demand, we find that most participants appear to understand the trade-offs involved in the experimental choices. The decisions made under commitment (stage one) indicate that a large majority of participants increase the level of later consumption as the price of doing so goes down. Some, however, often deviated from the law of demand and thus indicate either that they did not clearly understand the trade-offs they faced, or that economic models are a poor description of their motives. Our analysis allows for the possibility that violations of the law of demand may be related to other aspects of choice behavior.

The choices under commitment also reveal considerable heterogeneity in intertemporal allocations. A large majority selected an interior solution to the allocation problem, but some were much more willing than others to wait in order to receive a larger amount of cash. This willingness to wait is, to some extent, predicted by observable characteristics of the participants. Those with more wealth at baseline and more food in store, those with greater cognitive skills and financial numeracy, and younger farmers were all more willing to wait.

The results of stage one of the experiment also show that the aggregate distribution of allocations is quite stable with respect to time frame. The overall distribution of choices when the trade-offs are between tomorrow and 30 days from tomorrow looks quite similar to the one when the trade-offs are between 60 and 90 days from tomorrow. Looking just at aggregate distribution would, thus, suggest that dynamic preference reversals under commitment are uncommon.

An examination of individual heterogeneity tells a different story, however. We find that intertemporal preference reversals are common, but that "present"-biased preference reversals are only somewhat more common than those in which participants appear less patient as intertemporal trade-offs get pushed out into the temporal distance. In addition, we find evidence that the time preference inconsistencies are not merely noise – perhaps trembles from an otherwise time-consistent preference ordering. Among the participants who exhibit preference reversals, a substantial fraction (18%) is present-biased in at least 4 of 5 decisions. The tendency to be consistent or present-biased is also predictable with observable characteristics of the participants. Finally, simulations of purely random choice and of time-consistent choice with error are poor fits with critical aspects of the observed data.

Turning to stage two of the experiment, we observe several important features of behavior upon revisiting. First, revision behavior is common.The experiment made the original allocation choice clear and salient, and yet 65% of participants made some adjustment to their allocation decision. Most of these adjustments are fairly large; two thirds of those who revised made a reallocation

interpreting these choices with the tools of economics.

involving a shift of at least 4 tokens. A second important feature of behavior upon revisiting is that revision decisions shift the allocation of income forward and backward in time with nearly equal frequency. Of participants who made some adjustment to their allocation decision, 53% shifted income backward in time and 47% shifted income forward in time.

Most important for the purposes of this study, we find that revision behavior is predicted, in anticipated ways, by choices under commitment. As the theory predicts, those who were more dynamically consistent under commitment are substantially less likely to make a reallocation upon revisiting. We also find that those who, under commitment, better adhere to the law of demand are also substantially less likely to make a reallocation upon revisiting. Last, we find that, consistent with the central premise of hyperbolic discounting models, those who tend to be "present"-biased under commitment are more likely to make present-biased reallocations upon revision. Future-biased choices under commitment are also correlated the future-biased revisions, but this result is not as robust. Adding controls for observable characteristics of the subjects, makes the correlation economically small and insignificant. In addition, unlike the results for "present"-biased choices under commitment, the relationship between future-biased choice under commitment and future-biased revision is not monotonic.

Overall, we interpret these findings to indicate that time preferences are importantly heterogeneous. Subjects commonly exhibit dynamic consistencies, but "present"-biased preference reversals are not much more common than reversals in the other direction. As the new models of changing rates of time-discount predict, "present"-biased inconsistencies are linked to time-inconsistency. Our findings regarding future-biased preference reversals suggest, however, that these may more likely be driven by predictable changes in the marginal utility of consumption or by mis-understandings of the choice envirnoment.

The rest of the paper is organized as follows. Section 2 presents details of the experimental design, the sample of participants and the setting of the experiment. Section 3 presents the results of stage one of the experiment and choices under commitment. Section 4 discusses the results of stage two and choices upon revisiting. In section 5 we discuss some prior studies of intertemporal choice and explain our contribution to the literature. Section 6 offers some conclusions and suggests avenues for work.

2 The Experiment

2.1 Sample and Setting

Participants in the experiment were drawn from a population of rural households in central Malawi who grow tobacco as their main cash crop. As of the 2008-2009 growing season, these farmers were under contract with (the subsidiaries of) two large tobacco companies. The companies organized the farmers into clubs that range in size from 3 to 43 members. To facilitate timely revisiting, we limited our sample to those farmers located near a main trading center in the town of Mponela (population 13,670), and who lived in six traditional authorities (TAs) in the Dowa and Ntchisi districts. Experimental payments were delivered in the form of vouchers (described below) redeemable for cash from an office set up for this purpose in Mponela. To allow relatively easy access to participants and to faciliate their access to the cash disbursements, we included all farmers in these TAs that were 2008-09 members of clubs in which the median club member lives 25 kilometers or less from the disbursement office in Mponela.[12] Finally, to facilitate study of interactions within the household, we further restricted our sample to farmers who were part of a married couple.[13]

These sample restrictions left us with 1,271 targeted farmer households. A total of 1139 households (89.6%) were successfully interviewed. Table 1 provides some summary statistics of the sample based on their responses to a baseline survey (Appendix A). The median respondent is 46 years old, has 4 years of formal education, lives in a village with 100 others, including 2 relatives other than his spouse. There is, however, considerable variation in these demographic variables. Despite being drawn from the same population of tobacco farmers in central Malawi, participants are heterogenous.

[Table 1 here]

These summary statistics reveal important features of participants' wealth and income. On a global scale, the households in the sample are poor. Because the households all farm tobacco, however, we would expect them to have higher incomes and wealth than most in rural Malawi. At the time of the baseline survey, the median participant's household has no cash in any bank account,

[12] Scheduling for the stage one visit of this study was done by stratifying scheduling across agricultural zones. Within a zone, the order by which clubs were visited was randomly assigned. Scheduling was done on a club-by-club basis in order to facilitate field work since members of the same club often live within the same village or in neighboring villages.

[13] Each household had also been involved in a previous commitment savings experiments. See Brunne et al. (2011) for details. Our regression analysis controls for any effects of the experimental treatments in the previous experiment.

and the 90th percentile of the cash distribution has just 7,000 MK (approximately $47 US). The survey also inventoried the household's non-cash assets including livestock and agricultural and household durables; and participants were asked to estimate how much they could earn by selling those assets. When we include the self-reported value of these assets, we find that the median household held just 4,350 MK of wealth and the 90th percentile held 255,700 MK ($1,700 US).

In the calendar year before the baseline survey, the median household reported cash income of 40,200 MK ($270). However, the baseline interview was conducted in the middle of the rainy season; a time of planting and cultivating, not harvesting. At this time of year, cash income is very low, and will be low for the next few months. Neither the cash crop or the primary staple (maize) is ready to be harvested until mid April or early May. The summary statistics bear this out. The median household expects virtually no cash income between the interview date and the April 2010. These households do, however, expect substantially more income in 2010 than they received in 2009. After selling the tobacco harvest, the median household anticipates annual cash income of 82,500 MK (approximately $553 US). We will return to the issue of income expectations below.

2.2 Stage One: Choices Under Commitment

To gather data on intertemporal choices under commitment, we adapted the methods of Andreoni and Sprenger (2010) for the Malawian field environment. Upon initiating the interview, the household head and his spouse were physically separated. Then, after a few basic questions regarding demographics, each made 5 independent choices regarding the allocation of 2000MK between tomorrow and 30 days from tomorrow.[14]

More precisely, each participant was presented with a small bowl containing 20 beans (tokens) and two empty dishes, dish A and dish B. A token allocated to dish A corresponded to 100MK tomorrow. A token allocated to dish B corresponded to $100MK * (1 + r)$ 30 days from tomorrow, where r is the rate of return for waiting 30 days. The rate of return took on five different values: 0.10, 0.25, 0.50, 0,75, and 1.00. The rates of return rose, in order, with each of the five allocation choices; and participants knew that. For each rate of return, once the participant set out an allocation of tokens to the dishes, the tokens were translated into Malawian kwacha and the total was written above each dish on a whiteboard. Having seen the allocation in kwacha, the participant had the opportunity to adjust the allocation. This process was repeated until the participant indicated that he is ready to move on to the next allocation choice.

[14] The income was presented as compensation for participation in the survey and the prior research project (Brunne et al. 2011).

After completing the first five choices, the participant answered several series of questions from the baseline survey. (See Appendix A for details.) Then, using the same elicitation method with cup, beans, and dishes, the participant again made 5 independent choices regarding 2000MK, while facing different rates of return for waiting. This time, each of the 5 choices concerned the allocation of money between 60 and 90 days from tomorrow. Figure 1 presents a schematic of the method used to elicit intertemporal choices under commitment.

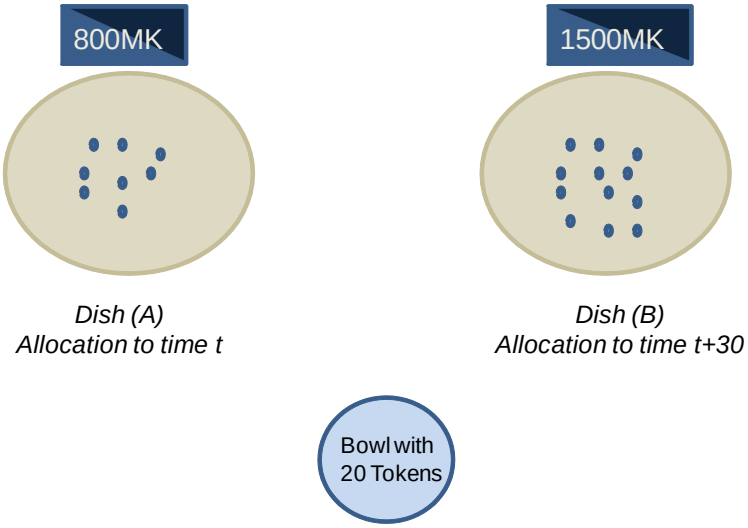

Figure 1: Schematic of the Preference Elicitation Method; Example with $r = 0.25$.

The interruption between the first and second set of five choices was intentional. We sought to avoid participants choosing the same allocations in both time frames simply for the sake of being (or appearing) consistent. In addition, the order in which the time preference sections of the questionnaire were administered was, in fact, randomly assigned between households within tobacco clubs. Approximately 50 percent, randomly determined, were first asked to make decisions about the allocation of money between tomorrow and 30 days from tomorrow. The remaining 50 percent were first asked to make decisions about allocations between 60 and 90 days from tomorrow. Our regression analysis controls for any effects of these alternative orders.

Before making their choices, each participant knew that one member of the couple would be randomly chosen to have one of his or her randomly selected choices implemented. The randomization was performed on site, by rolling dice. Implementation took the form of a voucher, redeemable at the disbursement office in Mponela. The voucher indicated the allocation (some amount at time t and another amount at time $t + 30$) and was issued to the member of the couple who was randomly chosen. Identity of the recipient was established with a name and a fingerprint placed

on the voucher.[15]

2.2.1 Advantages of the Elicitation Method

Andreoni and Sprenger's (2010) elicitation method has important advantages. The first feature is that participants choose a single allocation from convex budget set. A traditional elicitation method would ask participants a sequence of unfolding binary choice questions (do you prefer 2000MK tomorrow or 2500MK 30 days from tomorrow?), stopping the sequence when the participant flips his choice from sooner to later (or vice versa). Relative to that method, Andreoni and Sprenger (2010) permits much more rapid identification of marginal rates of substitution between sooner and later. By allowing rapid data collection, possibilities for studying individual heterogeneity in preferences increase. By avoiding a single stopping rule, the method also permits better analysis of whether participants appear to understand the choice environment and the trade-offs they face.

The rapid identification of marginal rates of substitution across time is especially important if the analyst wants to allow for the possibility of changing marginal utility from consumption. Traditional analysis has assumed that per period (flow) utility is linear. The Andreoni and Sprenger (2010) method encourages experimental manipulation of the return to waiting, and thus allows the analyst to relax the assumption of linear utility. In this way, it becomes possible to distinguish intertemporal tradeoffs driven by time discounting from tradeoffs driven by diminishing marginal flow utility from consumption. Relaxing the linear utility assumption seems especially relevant in the context of our experiment where the magnitude of the trade-offs is large relative to, for example, annual consumption.

2.3 Stage Two: Revisiting the Decision

In stage one of the experiment, one of each household's 20 decisions was randomly selected to be implemented. If the selected decision concerned an allocation between tomorrow and 30 days from tomorrow, then the experimental intervention was finished. The household redeemed its allocation and was not interviewed again. Stage two of the experiment applied only to those households whose randomly selected decision concerned an allocation between 60 and 90 days from tomorrow.

In stage two, this group of households was unexpectedly revisited. The targeted date for revisiting was randomly selected from the interval between 16 and 2 days prior to the first possible disbursement of funds. Thus, these households were targeted to be revisited between 45 and 59 days

[15] See Giné, Goldberg, and Yang (2010) for analysis of fingerprinting as a means of economic identification in rural Malawi.

from their baseline interview.[16] Successful revisiting required that the participant whose choice was implemented be located and interviewed. For this reason, actual dates for revisiting spanned the interval between 16 and 1 days prior to the first possible disbursement of funds.

At the revisit, the household head and his spouse were physically separated and the survey of wealth, income and expenditure was performed again. After the interview portion of the revisit, the participant whose choice from the baseline interview was set to be implemented was again presented with a cup containing 20 tokens. This time, however, four dishes were placed in front of the participant: dishes $A, B, A'$and B'. Dishes A and B contained a total of 20 tokens allocated to reflect the participant's original decision at baseline. Dishes A' and B' were empty. The participant was told that the first set of dishes showed his or her baseline choice; an allocation between what was effectively 1 to 16 days from the revisit and 30 days from then. This information was verified with the voucher that was issued at baseline. The participant was also reminded of the rate of return for waiting that applied at baseline, and the tokens on dishes A and B were translated into kwacha using whiteboards.

The participant was then asked to allocate the 20 tokens in the cup between the empty dishes A' and B', where the same rate of return for waiting applied. The allocation to the second set of dishes was again translated into kwacha and the participant was asked if he or she wanted to adjust the allocation. This process was repeated until the participant indicated he or she was finished. Then a new voucher was issued (regardless of whether the allocation was revised), and the interview was concluded. Figure 2 presents a schematic of this revising procedure.

The revising procedure is intended to measure the presence and magnitude of time-inconsistency. We therefore made the original allocation decision salient and unambiguous. The procedure is also designed to balance the consequences of implicit experimenter demands or connotations of a "right" answer. The participant must actively choose an allocation by placing tokens on the mats, and the status quo is thus discouraged. However, because the original allocation is set out just next to new allocation, there should be no difficulty replicating the original allocation and perhaps some mild, but implicit encouragement to do so. Especially given the rural setting, and the difficulty of double blind protocols, we cannot hope to eliminate the consequences of implicit experimenter demands. Instead we designed the experiment to limit the biases they might generate.

The conceptual key to the revising procedure is that participants recall, with perfect accuracy, the plan they chose at baseline. We then seek to quantify the extent to which they deviate from that plan, and learn more about why they deviate. In this way, the experiment is importantly

[16] Revisits occured even if the household chose a "corner solution" involving no disbursement of funds at 61 days from their interview.

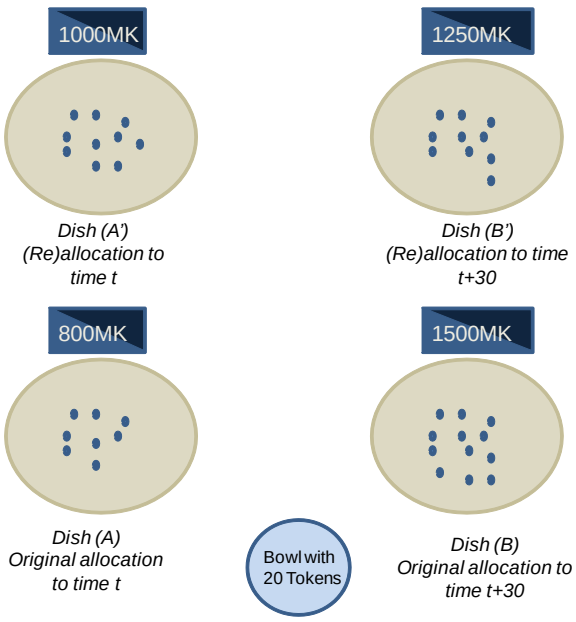

Figure 2: Schematic of the Revising Procedure; Example with $r = 0.25$.

different from an experiment designed to evaluate the stability of preferences regarding a delay of fixed length Δ. If that were the goal, we would not have reminded participants of their original choice and we would have repeated the elicitation method for a fixed delay. Instead, we make the original choice unambiguous and randomly choose the delay from a two week interval.

2.4 Summary and Timeline

To summarize, in January and February 2010, we asked the head of household and his spouse each to make 10 independent choices about the allocation of a substantial amount of money either between tomorrow and 30 days from tomorrow or between 60 and 90 days from tomorrow. We then revisited the households who were randomly selected to have one of their choices regarding the more distant future implemented. At that visit, which occurred in March and April of 2010, the participant whose choice was implemented is reminded of the decision he or she made under commitment. Then, facing the same rate of return for waiting, the participant is allowed to revise the original allocation decision. That (un)revised decision is then implemented with certainty. Figure 3 depicts the timeline of the experiment and Figure 4 presents the distribution of baseline interview and revisiting dates.

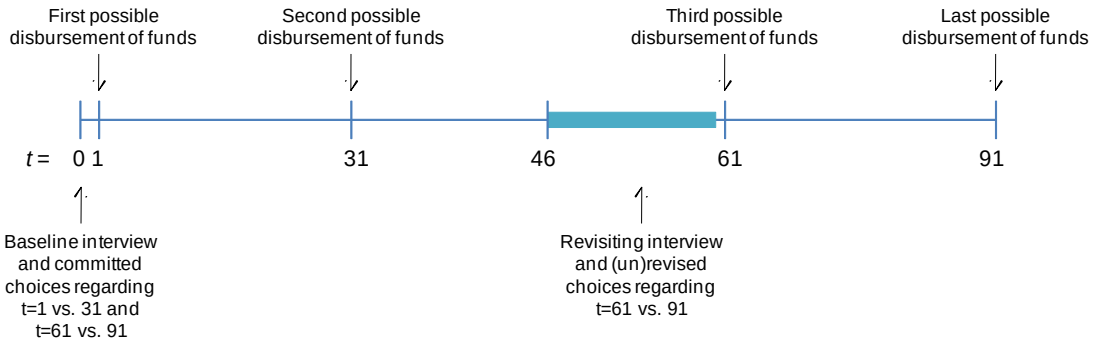

Figure 3: Timeline of Interviews, Choices and Disbursement of Funds

3 Stage One: Choices Under Commitment

3.1 Theoretical Framework

In this section we describe the results from stage one of the experiment, the choices made under commitment. To guide interpretation and fix ideas, consider first a canonical, deterministic problem of choice over time. In that problem the decision-maker solves

$$\max_{\mathbf{c}=(c_1,c_2,..,c_T)\in\mathbb{R}_+^T} U(\mathbf{c})$$

$$s.t. \quad k_{t+1} = (k_t - c_t)(1+r) \tag{1}$$

$$k_0 = \overline{k}, \; k_T \geq 0. \tag{2}$$

The consumer thus chooses a bundle of consumption, $\mathbf{c}$, the elements of which are indexed by time, to maximize a utility function U. The choice of $\mathbf{c}$ is restricted to a feasible set defined by the intertemporal budget constraints (1) and the boundary conditions (2). The usual assumptions are: (1) Monotonicity: U is increasing in each element of $\mathbf{c}$; other things equal more is preferred to less. (2) Diminishing marginal utility: U is concave in $\mathbf{c}$ and so smoother allocations tend to be preferred. (3) Impatience: other things equal, consumption is preferred sooner rather than later (impatience). Given a feasible set, choices about the allocation of consumption over time, i.e., time preferences, are driven by all three of these usual assumptions.

We will interpret participants' choices under commitment in stage one, as solving a special version the canonical model, where U is time separable. Abstracting from the discrete nature of

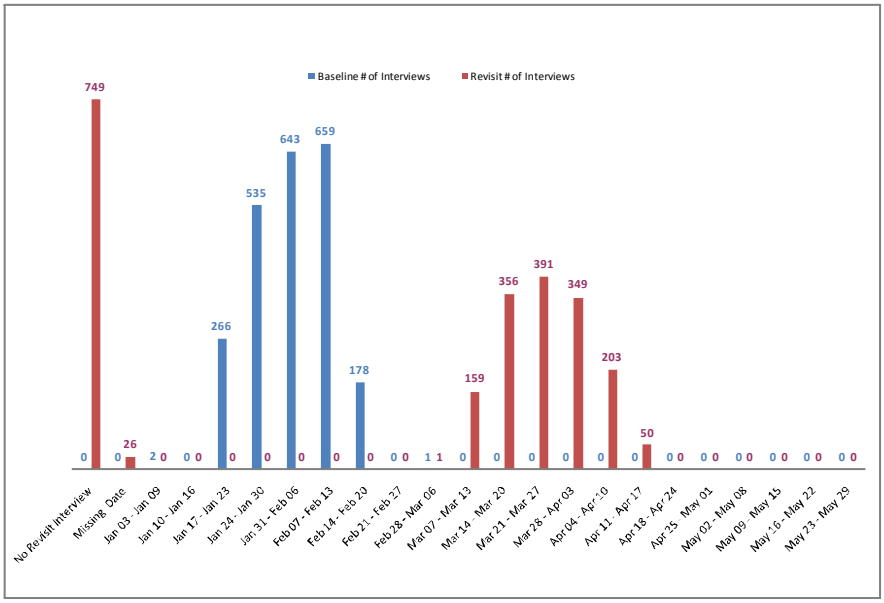

Figure 4: Number of Interviews and Revisits, By Date

the choice set, we will interpret stage 1 decisions as solving for each rate of return r:

$$\max_{(c_t, c_{t+30}) \in \mathbb{R}^2_+} u_t(c_t) + \delta_t u_{t+30}(c_{t+30}) \tag{P}$$
$$s.t. \quad c_t + \frac{c_{t+30}}{1+r} = 2000MK$$

Note that flow utilities and the time discount factors need not be stationary. Non stationary flow utilities allow for the possibility of predictable changes in the marginal utility of consumption. Non-stationary time discount factors allow for the possibility that, as has often been observed, time discount rates depend on the intertemporal distance to the trade-off. While this framework is deterministic, our analysis makes use of survey evidence regarding the real-life uncertainty that participants face.

Note, too, that the specification in problem P is, absent additional assumptions, generically unidentified from choices $(c_t, c_{t+30}) \in \mathbb{R}^2_+$ facing a exogeneous rates of return r. It is trivial to show that for any collection $(u_t, u_{t+30}, \delta_t)$ that reconcile the choice data, there is another collection $\left(u'_t, u'_{t+30}, \delta'_t\right)$ that can also reconcile the data. Restrictions, such as stationarity of the flow utility functions, or restricted form of non-stationarity (such as $u_t(c_t) = \frac{c^{1-\gamma_t}}{1-\gamma_t}$), are required in order to identify the preference parameters.

3.2 Adherence to the Law of Demand (Monotonicity)

While the preference parameters of (P) are not generically identified, time separability and montonicity of the flow utilities makes a strong prediction. If participants solve problem (P) then the allocation to the later period, measured in kwacha, should increase with the rate of return to waiting r. To see why this is true, it is useful to think of $\left(\frac{1}{1+r}\right)$ as the price of consumption later in terms of consumption sooner. When r goes up, the price of later consumption goes down. The result is an income effect creating incentives to increase consumption in both periods, and a substitution effect that is positive for consumption in the later period. Thus both income and substitution effects lead to increased consumption (kwacha) in the later period.[17]

As a first step in the analysis of stage 1 decisions, we evaluate the extent to which participants choices are consistent with this basic prediction of rational choice with time-separable, monotonic utilities. We view this as a logical pre-condition for interpreting the choices under commitment as revealing preferences in the context of a canonical model. To the extent that choices are inconsistent with the law of demand, it suggests either that participants did not understand the trade-offs involved in their decisions very well, or that the canonical model is very poorly suited for interpreting and making predictions about their behavior.

Each participant made ten intertemporal allocation decisions. To evaluate adherence with the law of demand, we partition these decisions into pairs, where each element of the pair is a choice regarding an allocation of kwacha over the same two dates. The first element of the pair is the allocation chosen when facing rate of return r. The other element is the allocation chosen when facing rate of return r', the next lowest rate of return. For each participant there are eight such pairs, four for each of the two time frames.[18] A total of 2,285 participants completed stage 1 of the experiment. The data thus contain 18,280 pairs of decisions where r increases by one increment; of these, 14,749 (81%) were such that the allocation to the later period increased with r. Thus, approximately 81% of pairs were consistent with this basic prediction of rational choice with monotonic, time-separable utility. In addition, the typical deviation from consistency is fairly modest in size. The median violation could be made consistent with a reallocation of less than two tokens. We interpret these results to indicate that, on average, participants understood the trade-offs they were facing and that the time-separable version of the canonical model is a reasonable starting point for describing average behavior.

[17] The allocation to the earlier period can go up or down depending on whether the income or substitution effect dominates.

[18] Note that a subject who did not change his or her allocation of tokens within a pair would still appear consistent with this prediciton; the allocation to the later period would increase with r.

The average rates of adherence with the law of demand mask some important heterogeneity, however. Table 2 presents the distribution of participants by the number of times they increased their later allocation of consumption with a single increase in the rate of return r. Recall, there are eight such pairs for each participant and thus the number of times a participant can be consistent by this measure ranges from zero to eight. Table 2 shows that, measured in this way, 31% of participants are always consistent and 75% are consistent at least three quarters of the time. At the other end of the spectrum, 11% of the sample violated this form of consistency at least half the time. The violations of this second group are also larger; their violations would require the reallocation of somewhat more tokens, on average, to be made consistent. We will return to consider the consequences of this heterogeneity below.

Number of Consistent Pairs	Frequency	Percent	Cumulative
0	3	0.13%	0.13%
1	6	0.26%	0.39%
2	16	0.70%	1.09%
3	48	2.10%	3.19%
4	168	7.35%	10.55%
5	321	14.05%	24.60%
6	482	21.09%	45.69%
7	530	23.19%	68.88%
8	711	31.12%	100.00%
Total	2285	100.00%	

Table 2: Number (of 8) Positive Changes in Later Consumption with Increase in r

3.3 Intertemporal Trade-offs Under Commitment

With greater confidence that most participants understood the trade-offs involved in their choices, and that the canonical model is a decent description of their behavior, we further examine decisions under commitment. Table 3 presents some summary statistics of these choices, displayed separately by time frame – tomorrow vs. 30 days from tomorrow (near) and 60 and 90 days from tomorrow (far). More precisely, Table 3 describes the distribution of allocations to the

later period, in kwacha.

Variable	Mean	Std. Dev	Percentiles					Percent at a corner
			10th	25th	50th	75th	90th	
Kwacha Allocated to Later Period								
t+30 at r=10%	1293.0	522.4	660.0	1100.0	1210.0	1650.0	2090.0	13%
t+30 at r=25%	1534.1	599.0	750.0	1250.0	1500.0	1875.0	2500.0	14%
t+30 at r=50%	1923.8	733.2	1050.0	1500.0	1950.0	2400.0	3000.0	16%
t+30 at r=75%	2250.6	883.4	1050.0	1750.0	2275.0	2975.0	3500.0	17%
t+30 at r=100%	2705.3	1067.9	1200.0	2000.0	2800.0	3600.0	4000.0	22%
t+90 at r=10%	1306.0	516.8	660.0	1100.0	1320.0	1650.0	2090.0	12%
t+90 at r=25%	1565.4	587.4	875.0	1250.0	1500.0	2000.0	2500.0	14%
t+90 at r=50%	1916.9	733.1	900.0	1500.0	1950.0	2400.0	3000.0	16%
t+90 at r=75%	2301.3	869.2	1225.0	1750.0	2275.0	2975.0	3500.0	17%
t+90 at r=100%	2746.1	1034.0	1400.0	2000.0	2800.0	3800.0	4000.0	23%

Table 3: Allocations to Later, in Malawian Kwacha, by Time Frame and Rate of Return

Several features of this distribution are worth noting. First, participants typically reveal a willingness to balance consumption between the two periods that is consistent with the canonical model described in Section 3.1. For example, when facing a rate of return to waiting of 50%, the median choice allocates 1,950MK to later and, thus, 700MK to sooner. A modest percentage (12%-13%) of allocations are "corner solutions." This willingness to locate at an interior allocation is consistent with participants not having, or not realizing they have, meaningful smoothing opportunities. The tendency toward interior solutions also points, in the absence of very high rates of time discounting, to the importance of diminishing marginal utilities of income.

A second important feature of this distribution is the heterogeneity in preferences that it reveals. In the nearer time frame, the tenth percentile allocates just 750MK (6 of 20 tokens) to the later period when the rate of return is 25%. The 90th percentile allocates all of its endowment to the later period. This heterogeneity is to some extent predictable with observable characteristics of the participants. Table 4 shows the results of a regression of the difference between the natural log of the allocation to sooner and later on the rate of return r and observable characteristics of the participants. These results provide some evidence that, conditional on the rate of return, those with more wealth at baseline allocate more consumption to later, as do those with more relatives who live in the village. There is also some weak evidence that those who scored higher on the word recall test and the financial literacy questions allocate more of their endowment to later, but that those who score higher on the Raven's test allocate less of their endowment to later. Measured in this way, we find no evidence that education has significant relationship with patience in this domain.

	Dependent Variable: Change in log(c) From Sooner to Later			
	delay of 1 vs. 31 days		delay of 61 vs. 91 days	
Rate of return (r)	0.948***	0.951***	0.932***	0.935***
	(0.028)	(0.028)	(0.028)	(0.028)
Male		0.018		0.018
		(0.040)		(0.041)
Less than 35 yrs old		0.066		0.069
		(0.055)		(0.057)
35-57 yrs old		0.058		0.035
		(0.043)		(0.046)
Some primary school		0.021		0.001
		(0.046)		(0.048)
Primary school		-0.031		-0.071
		(0.071)		(0.071)
More than primary school		0.037		0.014
		(0.096)		(0.089)
Have adequate maize		0.020		0.042
		(0.046)		(0.045)
log(Baseline wealth)		0.022		0.033**
		(0.014)		(0.015)
Words recalled		0.020		0.020
		(0.014)		(0.015)
Raven's Tests Correct		-0.027		-0.040*
		(0.020)		(0.021)
Finanical Literacy Questions Correct		0.026		0.033
		(0.025)		(0.025)
Number of relatives in the village		0.004*		0.007***
		(0.002)		(0.002)
Constant	0.382***	0.163	0.406***	0.132
	(0.029)	(0.125)	(0.030)	(0.133)
N	9264	9264	9257	9257
Adjusted R-squared	0.0870	0.0937	0.0837	0.0931

Table 4: OLS Estimates Change in ln(c) From Sooner to Later

The estimates in Table 4 have the advantage of being easily interpreted in terms of a simple economic model of intertemporal choice. If we adopt the canonical model in problem (P) and assume stationary, isoelastic utilities ($u(c) = \frac{c^{1-\rho}}{1-\rho}$), then the coefficient on r is an estimate of $\frac{1}{\rho}$. The disadvantage of this specification is that it excludes corner allocations, where the log of consumption (0) at one time or the other is undefined. Analysis of a levels specification gives qualitatively similar results (available upon request) with more evidence of a positive correlation between word recall and the willingness to postpone consumption.

3.4 Dynamic Consistency

A third important feature of the distribution choices displayed in Table 3 is its apparent stationarity. Comparing the top and bottom halves of Table 3 shows that the distribution of allocations to later is not dramatically altered by the change in time frame. For example, the mean allocation to later when facing a 25% rate of return is 1,534MK when the trade-off is between tomorrow and 30

days from tomorrow. The analogous number is 1,565MK when the trade-off is pushed 60 days out into the future. This average stationarity is, however, somewhat misleading. It both overstates the stability of individual choices across time frames and masks heterogeneity in individual tendencies to shift consumption forward or back, depending on the time frame.

Each participant makes five pairs of decisions where each element of a pair differs only in the time frame. Of the total of 12,660 such pairs, just 4,362 (34%) are identical and just 6,458 (51%) differ by a token or less. Thus, in nearly half of all such pairs their elements are substantially different from each other. There is a relatively modest tendency for this dynamic inconsistency to be "present"-biased. Of the 6,202 pairs that differ by strictly more than a token, 3,264 (53%) allocate more consumption to the earlier period when the trade-off is between 1 and 30 days. The remaining 47% of these preference reversals allocate more consumption to the later period when the trade-off is between 1 and 30 days from tomorrow. Consistent with the "magnitude effect," preference reversals are more common at lower rates of return for waiting.[19] When the rate of return for waiting is 10%, just 47% of pairs differ by a token or less.

In this way, the results of stage one of the experiment indicate that intertemporal preference reversals are common, but that "present"-biased preference reversals are only somewhat more common than those in which participants appear less patient as intertemporal trade-offs get pushed out into the temporal distance. There is evidence, however, that these dynamic inconsistencies are not merely noise. Among those participants who exhibit preference reversals, a substantial fraction (18%) is present-biased in at least 4 of 5 decisions.[20] In addition, the tendency to be consistent or present-biased is somewhat predictable with observable characteristics of the participants.

Table 5 presents the results of regressions that relate a participants tendency to be consistent or "present"-biased to certain observable characteristics. In each column the dependent variable is either the fraction of pairs of decisions in which the participant was dynamically consistent or the fraction the participant was present-biased. The first column of Table 5 indicates that males, those with greater maize stores and those who score better on financial literacy test all tend to be more dynamically consistent. Column 3 of Table 5 indicates that these variables have similar relationships (with opposite signs) with fraction present-biased, though these relationships are not statistically significant.

Columns 2 and 4 reveal an important relationship. There is a strong association between basic consistency as described in section 3.2 and dynamic consistency as measured here. Those with

[19] The magnitude effect refers to the finding that individuals appear more present-oriented when the stakes are smaller. See, e.g., Thaler, (1981) Loewenstein (1987) and Benzion et al. (1989).

[20] We more thoroughly examine the consistency of the data with random choice in Section XX.

higher basic consistency scores are much more likely to make dynamically consistent choices.[21] This link suggests that the tendency to exhibit preference reversals may be attributable, in part, to a poor understanding of the choice environment and the trade-offs involved. We pursue this hypothesis further as we examine the revising behavior in stage 2 of the experiment.

	Dependent Variable			
Variables	Fraction Consistent		Fraction Present-biased	
Male	0.049	0.048	-0.004	-0.004
	(0.015)	(0.014)	(0.014)	(0.014)
Less than 35 yrs old	-0.021	0.004	0.011	0.004
	(0.020)	(0.019)	(0.020)	(0.020)
35-57 yrs old	-0.015	-0.009	0.002	0.001
	(0.017)	(0.016)	(0.016)	(0.016)
Some primary school	-0.014	-0.016	0.022	0.023
	(0.019)	(0.017)	(0.018)	(0.018)
Primary school	-0.024	-0.014	0.021	0.017
	(0.027)	(0.025)	(0.025)	(0.025)
More than primary school	-0.034	-0.039	0.027	0.03
	(0.032)	(0.029)	(0.032)	(0.031)
Have adequate maize	0.025	0.021	-0.018	-0.016
	(0.017)	(0.015)	(0.015)	(0.015)
Words recalled	0.003	0.002	0.001	0.001
	(0.005)	(0.005)	(0.005)	(0.005)
Raven's Tests Correct	-0.003	0.001	-0.01	-0.012
	(0.008)	(0.007)	(0.008)	(0.007)
Finanical Literacy Questions Correct	0.02	0.009	-0.009	-0.005
	(0.009)	(0.008)	(0.008)	(0.008)
Basic Consistency Score (0-8)		0.081		-0.029
		(0.004)		(0.004)
Constant	0.243	-0.293	0.381	0.565
	(0.030)	(0.039)	(0.029)	(0.042)
R-squared	0.0104	0.1523	0.0009	0.018
N	2220	2220	2220	2220

Table 5: OLS Estimates of Fraction Consistent or Fraction Present-biased

4 Stage Two: Undoing Commitment

Having described some basic features of preferences under commitment, we now turn to examine the relationship between preferences under commitment, and other participant characteristics, and revision behavior in stage two of the experiment. Our goal is to evaluate, quantitatively, the importance of time preferences relative to other motives for revision.

[21] Note there is no mechanical reason why these two measures should be linked. The first regards the response of allocations to changes in r, within time frame. The second regards consistency of allocations, given r, across time frames. For example, a subject who always violated basic consistency could be perfectly dynamically consistent, simply by replicating his non-monontonic allocations in both time frames.

4.1 Qualitative Features of Revision Behavior

Because this is, to our knowledge, a first attempt to study revision behavior in an experimental study of time preference, we begin with a simple description of the choices upon revisiting. Recall that stage two of the experiment applies only to those households whose randomly selected choice was an allocation between 61 and 91 days from the baseline interview. The randomization was designed to slightly favor the later time frame and so we had 1,521 households to target and 761 individuals to present with a revision opportunity. Of these we were successful in collecting revision choice data from 701 (92%). Figure 5 presents a histogram of the changes in the participants allocation to later upon revisiting.

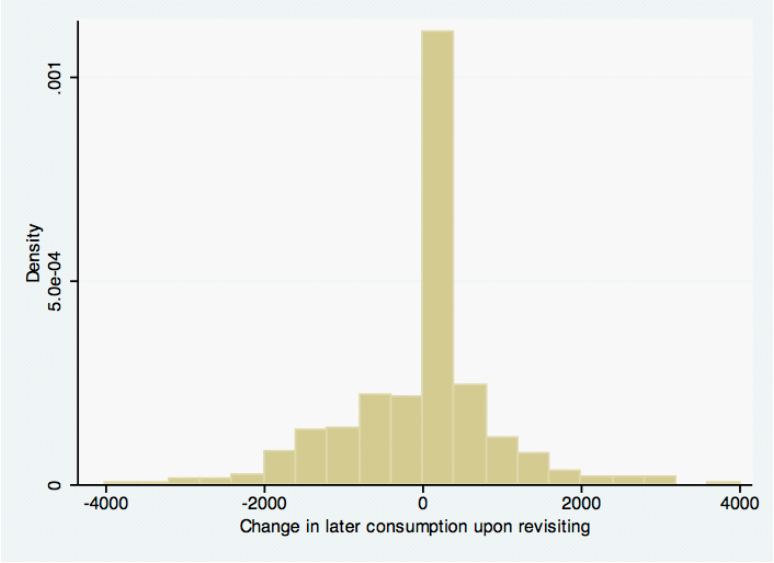

Figure 5: Distribution of Changes to Later Consumpation Allocation, Upon Revisiting

The first important feature of this distribution is the frequency of revision behavior. Despite an experimental design that made the original allocation choice clear and salient, 65% of participants (458) made some adjustment to their allocation decision. It could be that implicit experimenter demands caused some participants to feel as though some change was expected of them. However, among those who revised their original decision, a large majority (86%) made a reallocation involving a shift of at least two tokens, and 67% made a reallocation involving a shift of at least 4 tokens.

A second important feature of this distribution is that revision decisions shift the allocation of income forward and backward in time with nearly equal frequency. Of the 458 participants who made some adjustment to their allocation decision, 53% shifted income backward in time and 47%

shifted income forward in time. As the histogram also indicates, however, the adjustments that moved consumption forward in time tended to be more modest in size. Of these, approximately 55% involve the shifting of at least 4 tokens, and just 16% involve shifting 10 tokens or more. The comparable numbers for adjustments that moved consumption backward in time are 70% and 27%. Table 6 presents summary statistics of the revisits that resulted in reallocations, distinguishing between reallocations that moved time forward versus backward in time.

					Percentiles		
Variable	N	Mean	Std. Dev	Min	10th	50th	90th	Max
Revisiting decisions								
Increases in later allocation	217	806.82	696.32	110	175	600	1750	4000
Decreases in later allocation	241	-979.00	726.93	-4000	-1925	-770	-220	-110

Table 6: Distribution of Reallocations at Revisiting

4.2 Who Revises?

4.2.1 Any Revision

Having described some basic features of revision behavior, we move on to consider the correlates of these choices. Of primary interest is the ability of preferences under commitment to predict revision behavior. To begin this analysis, Table 7 presents the estimates of linear probability models relating revision behavior to observable characteristics and basic features of preferences under commitment. In each column, the dependent variable is an indicator equal to one if, upon revisiting, the participant reallocated his or her allocation by at least two tokens.[22]

[22] Results for for smaller reallocations are very similar.

	Dependent Variable					
Variables	I(Any Revision)					
Preferences under commitment						
Consistency ratio [0,1]		-0.259***	-0.202***	-0.200***	-0.178**	-0.178**
		(0.060)	(0.066)	(0.066)	(0.066)	(0.066)
Basic Consistency Score [0,8]			-0.032**	-0.033**	-0.017	-0.018
			(0.015)	(0.015)	(0.016)	(0.016)
Time from payment						
Days to first disbursement				0.005	0.004	0.004
				(0.004)	(0.004)	(0.004)
Spousal conflict						
Diff from Spouse's Allocation To Sooner (>=0)					0.030***	0.030***
					(0.007)	(0.007)
Diff from Spouse's Allocation To Sooner (<0)					-0.013**	-0.013**
					(0.005)	(0.005)
Shocks						
Death in the Family						0.021
						(0.106)
Shock to expected hh expenditure						0.009
						(0.013)
Shock to expected hh income						0.028
						(0.019)
Demographics						
Male	0.011	0.010	0.011	0.009	0.030	0.031
	(0.044)	(0.043)	(0.043)	(0.044)	(0.043)	(0.043)
Less than 35 yrs old	0.105*	0.097*	0.088	0.088	0.104*	0.097*
	(0.060)	(0.059)	(0.059)	(0.059)	(0.058)	(0.058)
35-57 yrs old	0.116**	0.106**	0.103**	0.096*	0.105**	0.101**
	(0.051)	(0.051)	(0.050)	(0.050)	(0.049)	(0.049)
Some primary school	0.004	-0.006	-0.009	-0.01	-0.013	-0.01
	(0.057)	(0.057)	(0.057)	(0.056)	(0.055)	(0.055)
Primary school	-0.079	-0.079	-0.089	-0.097	-0.083	-0.084
	(0.077)	(0.076)	(0.075)	(0.075)	(0.074)	(0.075)
More than primary school	-0.045	-0.07	-0.065	-0.072	-0.078	-0.076
	(0.096)	(0.094)	(0.093)	(0.093)	(0.091)	(0.091)
Have adequate maize	-0.026	-0.028	-0.032	-0.034	-0.028	-0.034
	(0.046)	(0.046)	(0.046)	(0.046)	(0.045)	(0.046)
Words recalled	-0.016	-0.018	-0.016	-0.014	-0.017	-0.017
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Raven's Tests Correct	-0.037*	-0.038*	-0.041*	-0.038*	-0.040*	-0.041*
	(0.023)	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)
Finanical Literacy Questions Correct	-0.005	-0.001	0.005	0.006	0.005	0.003
	(0.023)	(0.022)	(0.023)	(0.023)	(0.023)	(0.023)
Number of relatives in the village	-0.002	-0.002	-0.002	-0.002	-0.003	-0.003
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Constant	0.658***	0.790***	0.971***	0.926***	0.809***	0.822***
	(0.088)	(0.091)	(0.124)	(0.130)	(0.135)	(0.136)
Adjusted R-squared	0.0075	0.0314	0.0365	0.0372	0.0645	0.0625
N	675	675	675	675	675	675

Table 7: The Probability of Revising

The estimates presented in the first column of Table 7 indicate that younger farmers (or their spouses) are substantially more likely than older ones to make any revision to their original alloca-

tion upon revisiting. The probability of making a revision is approximately 0.1 higher among these younger groups. Recall that the average probability of making any revision of this magnitude is 0.56. These results also indicate that those who scored higher on the Raven's IQ test are less likely to make any revision upon revisiting, though this relationship is only just significant at the 10% confidence level. Other observable characteristics of the participants have little power to predict revision behavior.

Preferences under commitment are, however, strongly linked to revision behavior.[23] Column two of Table 7 adds to the demographic controls a "consistency ratio" measure. This measure takes, again, the five pairs of decisions where each element of a pair differs only in the time frame and gives the fraction of those five pairs in which the participant exhibited *no* preference reversals. Here, a pair of allocations are considered the same if they differ by a token or less. The coefficient on this consistency measure indicates that, as would be expected from Strotz (1956), the tendency to exhibit preference reversals under commitment predicts revision when that commitment is revisited. The coefficient on the consistency ratio indicates that moving from the 25th percentile of the ratio to the 75th percentile decreases the probability of revising upon revisiting by 0.16. This relationship is both economically substantial and statistically significant.

The results of column three, where we add the measure of basic consistency from Section (3.2) indicate that the relationship between preferences under commitment and revision behavior is not driven by those participants who appeared to have difficulty understanding the trade-offs involved in the decision problem. Those who better satisfied the basic predictions of rationality are significantly less like to make any revision. However, the coefficient on the consistency ratio is attenuated only modestly when we add our measure of basic consistency; the relationship between preferences under commitment and revision behavior remains large and statistically significant. Column 4 of Table 7 adds a control for the (randomly assigned) lag between the date of revisit and the date at which the first disbursement funds will occur. We find little evidence of an effect of the lag on the probability of any revision. The point estimate is positive, but an economically substantial negative relationship is also in the 95% confidence interval.

These results suggest that time preferences under commitment are, as indicated by theory, significant predictors of revision behavior when commitment is relaxed. In other words, these results provide evidence that the relationship between choices under commitment and revision behavior is of the sign predicted by theory. Next we turn to evaluate the relative magnitude of the relationship as we consider alternative motives for revision.

A natural motive for revising an allocation is intra-household disagreement regarding the pre-

[23] Results on the magnitude of the revision (not shown) reinforce these and subsequent findings.

ferred allocation. Spouses may have different objectives for the use of household income, or they may have different information or beliefs about household income, expenditure or wealth. If so, we would expect spouses to have different preferences regarding the experimental allocation; and, if they discussed these differences after the baseline experiment, we might expect the disagreement regarding the best allocation to influence revision behavior. In column 5 of Table 7 we add controls for the *magnitude* (measured in token) of the difference between the participant's preferred allocation under commitment and that of his or her spouse. We allow the effect to differ depending on whether the participant's preferred allocation was more front-loaded (the top row) or more back-loaded (the bottom row). The results indicate a moderate relationship that is asymmetric around zero. Participants who chose to allocate more tokens to soon than their spouses did are more likely to revise. Adding the two coefficients, a standard deviation increase in the conflict measure (amounting to nearly 4 tokens) is associated with an increase of approximately 0.067 in the probability of making a revision. Participants who chose to allocate fewer tokens to soon than their spouses did are, however, a bit less likely to revise. Taking only the bottom coefficient we see that a standard deviation increase in the conflict measure is associated with an increase of approximately 0.052 in the probability of making a revision.

Another natural motive for revising an allocation are shocks to expectations regarding health, expenditure or income. In column 6 of Table 7 we include measures of some of these shocks. The first is an indicator for whether, between baseline and revisiting, the household experienced a death in the (extended) family.[24] The second measures the absolute magnitude of the difference between expected and realized expenditures between baseline interview and revisit. The expectations were collected as part of the baseline interview. The realizations were collected upon revisiting. The third is the analogous measure for household income. The point estimates indicate that each of these shocks are, as predicted by theory, positively related to the probability of revision. The estimates are quite imprecise, however. In none of the cases, can we reject a null hypothesis of no relationship. Importantly, adding these measures of shocks does not much influence the coefficients on the consistency scores. In this way, we find no evidence that the relationship between choice under commitment and revision behavior is driven by a correlation between shocks and preferences.

4.2.2 "Present-biased" Revision

The results in Table 7 concerned the probability of making any substantial revision, shifting income either backward or forward in time. Next we look at a subset of these revisions. Table 8

[24] Unfortunately, this event was relatively common. Among the revisited households, 2.7% indicated that a member of their extended family had died since January.

presents analogous results for "present-biased" revisions; i.e., revisions that shift income backward in time.

	Dependent Variable					
Variables	I(Present-Biased Revision)					
Preferences under commitment						
Present-biased ratio [0,1]		0.145**	0.129**	0.129**	0.132**	0.130**
		(0.063)	(0.065)	(0.065)	(0.065)	(0.065)
Basic Consistency Score [0,8]			-0.017	-0.017	-0.015	-0.016
			(0.013)	(0.013)	(0.014)	(0.014)
Time from payment						
Days to first disbursement				-0.001	-0.001	-0.001
				(0.004)	(0.004)	(0.004)
Spousal conflict						
Diff from Spouse's Allocation To Sooner (>=0)					0.003	0.003
					(0.006)	(0.006)
Diff from Spouse's Allocation To Sooner (<0)					-0.002	-0.002
					(0.005)	(0.005)
Shocks						
Death in the Family						-0.033
						(0.102)
Shock to expected hh expenditure						0.016
						(0.013)
Shock to expected hh income						-0.019
						(0.029)
Demographics						
Male	0.075*	0.072*	0.072*	0.071*	0.072*	0.072*
	(0.042)	(0.042)	(0.042)	(0.042)	(0.043)	(0.043)
Less than 35 yrs old	0.138**	0.132**	0.127**	0.124**	0.126**	0.118**
	(0.054)	(0.054)	(0.054)	(0.054)	(0.055)	(0.055)
35-57 yrs old	0.095**	0.092**	0.089*	0.088*	0.089*	0.086*
	(0.046)	(0.046)	(0.046)	(0.046)	(0.047)	(0.047)
Some primary school	-0.003	-0.012	-0.014	-0.014	-0.015	-0.014
	(0.055)	(0.055)	(0.055)	(0.055)	(0.055)	(0.056)
Primary school	-0.081	-0.087	-0.092	-0.091	-0.09	-0.091
	(0.072)	(0.073)	(0.072)	(0.073)	(0.073)	(0.073)
More than primary school	-0.132	-0.140*	-0.139*	-0.139	-0.141*	-0.139
	(0.085)	(0.085)	(0.085)	(0.085)	(0.085)	(0.085)
Have adequate maize	0.01	0.013	0.01	0.01	0.011	0.002
	(0.043)	(0.043)	(0.043)	(0.043)	(0.043)	(0.043)
Words recalled	-0.023	-0.024*	-0.023	-0.023	-0.023	-0.023
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Raven's Tests Correct	-0.01	-0.01	-0.011	-0.012	-0.012	-0.012
	(0.023)	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)
Finanical Literacy Questions Correct	-0.003	-0.003	0.001	0.002	0.002	0
	(0.022)	(0.021)	(0.022)	(0.022)	(0.022)	(0.022)
Number of relatives in the village	-0.003	-0.003	-0.003	-0.003	-0.003	-0.004*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Constant	0.380***	0.353***	0.471***	0.479***	0.471***	0.485***
	(0.081)	(0.081)	(0.122)	(0.127)	(0.132)	(0.134)
Adjusted R-squared	0.0093	0.0155	0.0165	0.0146	0.0119	0.0112
N	675	675	675	675	675	675

Table 8: The Probability of Present-biased Revision

The results in column 1 of Table 8 show that both males and younger participants are more likely to make present-biased revisions, though the coefficient on male is significant only at the 10% level ($p\text{-}value = 0.076$). In column 2 we add a measure of dynamic consistency analogous to that in Table 7. In this case, the "present-biased ratio" takes the five pairs of decisions where each element of a pair differs only in the time frame and gives the fraction of those five pairs in which the participant exhibited "present-biased" preference reversals.[25]

The results of column 2 indicate that the tendency to make present-biased preference reversals under commitment is strongly related to present-biased revisions when that commitment is undone. In this case, the coefficient on the present-bias ratio indicates that moving from the 25th to the 75th percentile increases the probability of a present-biased revision by 0.06. Recall that the average probability of making a present-biased revision is 0.31. Again, conditioning on our measure of basic consistency (in column 3) has relatively modest effects on the estimates. In column 4, the point estimate on the lag between revisit and first disbursement of funds has the sign one would expect from a model of non-constant discounting, but it is very imprecisely estimated.

Turning again to alternative motives for revision, in column 5 here we allow the measure spousal conflict to enter in levels (rather than absolute magnitudes). We find no evidence that conflict influences the probability of present-biased revision. In column 6 we add controls for shocks, again in levels rather than absolute magnitudes. The point estimate on a death in the family is negative, but imprecisely estimated. The point estimates on the measures of shocks to expenditure and income each have the expected sign, but neither is precisely estimated. We cannot reject a null hypothesis of no relationship. Again, it is important that adding these measures of shocks does not much influence the coefficients on the consistency scores. Again, we find no evidence that the relationship between choice under commitment and revision behavior is driven by a correlation between shocks and preferences.

5 Related Literature (Incomplete)

There is a long tradition of estimating time preferences from observational data. Hausman (1979), Lawrance (1991) and Warner and Pleeter (2001) are prominent examples. In this tradition, the analyst observes the (implicit) price consumers are willing to pay in order to move consumption forward in time. In Hausman (1979), for example, a time discount rate is inferred from the price elasticity of demand for long-run energy efficiency in household appliances. The early contributions to this literature assumed that time discount rates were constant with respect to time.

[25] Again, we consider it a reversal only if the allocations differ by at least two tokens. Results are quantitatively similar if we reduce the tolerance to just one token.

More recently, observational data has been used to estimate potentially non-constant time-discount functions. This literature, which restricts itself to estimating quasi-hyperbolic discount functions, includes Paserman (2008), Fang and Silverman (2009) and Laibson et al. (2007). We depart from this literature by adopting experimental methods for eliciting time preferences.

The experimental literature on time preference is large. Frederick et al (2002) offer a review. More recent examples include, Andersen et al. (2008), Benhabib, et al. (2010), and Andreoni and Sprenger (2010). As noted above, we adapt to our field environment the methods that Andreoni and Sprenger (2010) developed for use in the lab. A key advantage of the Andreoni and Sprenger (2010) design is that, with a convex choice set, it permits the relatively rapid collection of information on time preferences regarding several different rates of return. The method thus allows analysis of the basic consistency of choice as well as quantitative measures of the degree of dynamic inconsistency.

Most recently, experimental studies of time preference have been taken into the field, often in developing countries. Prominent examples include Harrison et al. (2002), Ashraf et al. (2006), and Tanaka, et al. (2009). Two such field experiments are closely related to ours and thus merit special consideration. The first, Ashraf et al (2006), asked a sample of 1,777 in the Philippines hypothetical time preference questions on a survey. Later, a subset of this sample was offered a commitment saving product. Ashraf et al (2006) found that women who exhibited present-biased preference reversals on the survey questions were, as predicted by theory, more likely to take up the commitment saving product. Our paper differs from this prior study, most importantly, by studying directly the link between incentivized time preference decisions and time-inconsistency. We measure the extent of preference reversals, as well as the basic consistency of choice with rational economic models thus provide a quantitative assessment of the mechanisms behind time inconsistency and the demand for commitment. The second closely related paper, Harrison et al. (2005), elicited time preferences from 253 participants in Denmark. Of this sample, 97 were later revisited and, importantly, asked to perform the same time preference experiment again. Our experiment is distinguished from Harrision et al. (2005) by, among other things, making a participant's original choice clear and salient. Our goal is not to evaluate the stability of time preference, but rather to measure time-inconsistency and estimate its relationship to preferences under commitment.

6 Conclusion

To be added.

References

[1] Ainslie, G., *Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person* (Cambridge [England]; New York: Cambridge University Press, 1992).

[2] Ameriks, J., A. Caplin and J. Leahy (2003) "Wealth Accumulation and the Propensity to Plan." *Quarterly Journal of Economics*, *118*, pp. 1007-1047.

[3] Ashraf, Nava, Dean Karlan and Wesley Yin (2006). "Tying Odysseus to the Mast: Evidence form a Commitment Savings Product in the Philippines," *Quarterly Journal of Economics* 121(2), 635-672.

[4] Benhabib, Jess, Alberto Bisin and Andrew Schotter (2010). "Present-bias, Quasi-Hyperbolic Discounting, and Fixed Costs." *Games and Economic Behavior* 69(2), 205-223.

[5] Andersen, Steffen, Glenn Harrison, Morten Lau and Elisabet Rustrom (2008). "Eliciting Risk and Time Preferences," *Econometrica* 76(3), 583-618.

[6] Andreoni, James and Charles Sprenger (2010). "Estimating Time Preferences from Convex Budgets." NBER Working Paper Number 16347. NBER, Cambridge MA.

[7] Choi, Syngjoo, Ray Fisman, Douglas Gale and Shachar Kariv (2007). "Consistency and Heterogeneity of Individual Behavior Under Uncertainty," *American Economic Review* 97(5), 1921-1938.

[8] Choi, Syngjoo, Shachar Kariv, Wieland Muller, and Dan Silverman (2011). "Who Is (More) Rational?" Mimeo, University of Michigan.

[9] DellaVigna, Stefano (2009). "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47(2), 315-372.

[10] Fang, Hanming and Dan Silverman (2009). "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY," *International Economic Review*, 50(4), 1043-1077

[11] Frederick, Shane, George Loewenstein and Ted O'Donoghue (2002). "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature,* 40(2), 351-401.

[12] Fuchs, Victor R. (1982). "Time Preference and Health: An Exploratory Study," in Victor R. Fuchs, editor, *Economic Aspects of Health*. University of Chicago Press, Chicago, 93-120.

[13] Fudenberg, Drew and David K. Levine (2006). "A Dual Self Model of Impusle Control," *American Economic Review* 96 (5), 1449-1476.

[14] ____, (2010) "Timing and Self-Control," mimeo, Harvard University Department of Economics.

[15] Gul, Faruk, and Wolfgang Pesendorfer (2001). "Temptation and Self-Control," *Econometrica*, 69 (6), 1403-1435.

[16] Giné, Xavier, Jessica Goldberg and Dean Yang (2010). "Identification Strategy: A Field Experiment on Dynamic Incentives in Rural Credit Markets." Mimeo, University of Michigan.

[17] Harrison, Glenn W., Morten Lau and Elisabet Rustrom (2005). "Dynamic Consistency in Denmark: A Longitudinal Field Experiment," Mimeo, Georgia State University.

[18] Hausman, Jerry A. (1979) "Individual Discount Rates and the Purchase and Utilization of Energy-using Durables." *Bell Journal of Economics* **10(1):**33-54.

[19] Laibson, David (1997). "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 112(2), pp.443-77.

[20] Laibson, David, Andrea Repetto and Jeremey Tobacman (2007). "Estimating Discount Functions with Consumption Choices over the Lifecycle." Mimeo, University of Pennsylvania.

[21] Loewenstein, George (1987). "Anticipation and the Valuation of Delayed Consumption," *Economic Journal* 97, pp. 666-684.

[22] Loewenstein, George, and Jon Elster, *Choice Over Time (*Russell Sage: New York, 1992).

[23] Noor, Jawwad (2011). "Time Preference: Experiments and Foundations," mimeo, Boston University.

[24] O'Donoghue, Ted, and Matthew Rabin, (1999) "Doing It Now or Later," *American Economic Review* 89 (March), 103-124.

[25] Ozdenoren, Emre, Stephen Salant, and Dan Silverman, (2009) "Willpower and the Optimal Control of Visceral Urges," forthcoming in the *Journal of the European Economic Association*.

[26] Paserman, Daniele (2008). "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Implications." *Economic Journal*, 118(531), 1418-1452.

[27] Phelps, Edmund S., and Robert A. Pollak (1968). "On Second-best National Saving and Game-equilibrium Growth." *Review of Economic Studies* 35 (April), 185-199.

[28] Strotz, Robert H. (1956). "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies* 23, 165-180.

[29] Tanaka, Tomomi, Colin Camerer and Quang Nguyen (2010). "Risk and Time Preferences: Experimental and Household Survey Data from Vietnam." *American Economic Review* 100(1) 557-71.

[30] Thaler, Richard T., (1991). "Some Empirical Evidence on Dynamic Inconsistency," in Richard H. Thaler, ed.. *Quasi Rational Economics.* New York: Russell Sage Foundation, 127-33.

[31] Warner, John T. and Saul Pleeter (2001). "The Personal Discount Rate: Evidence from Military Downsizing." *American Economic Review* **91(1):** 33-53**.**

Variable	# of Non-missing Values	Mean	Std. Dev	Min	Percentiles 50th	Percentiles 90th	Max
Demographics							
Respondent's Own Age	2234	46.84	14.18	18	46	67	95
Respondent's Spouse's Age	2118	46.04	13.81	18	46	64	95
Years of schooling	2281	4.29	3.21	0	4	8	14
Number of Relatives in Village	2279	4.57	8.46	0	2	10	132
Total Number of People in Village	2048	176.63	267.06	0	100	350	4000
Wealth and Income (in 100 MK)							
Total Money in all HH Bank Accounts	2284	43.4	229.2	0.0	0.0	70.0	5400.0
Value of all money and material assets owned by HH	2284	1118.4	2664.5	0.0	435.8	2557.0	69502.5
Realized income (incl. gifts) for 2009	2284	986.1	2423.5	0.0	402.0	2044.0	60810.0
Realized revenue (excl. gifts) for 2009	2284	916.8	2369.3	0.0	360.0	1940.0	60810.0
Expected income (incl. gifts) for 2010	2284	1743.0	3698.0	0.0	825.0	3500.0	80030.0
Expected revenue (excl. gifts) for 2010	2284	1692.6	3660.6	0.0	800.0	3315.0	80030.0
Expected income (incl. gifts) by April 2010	2284	172.0	620.2	0.0	3.6	400.0	13770.0
Expected revenue (excl. gifts) by April 2010	2284	152.6	587.5	0.0	0.0	370.0	13350.0
Expected Expenditure in (100 MK)							
Expected Expenditures in February 2010 (Baseline)	2284	148.1	179.1	0.0	97.0	301.0	2315.0
Expected Expenditures in March 2010 (Baseline)	2284	145.6	174.2	0.0	96.5	303.0	2290.0
Expected Expenditures in April 2010 (Baseline)	2284	152.7	194.7	0.0	100.0	320.0	2595.0
Expected Expenditures in May 2010 (Baseline)	2284	182.7	251.1	0.0	112.8	385.0	4535.0
Knowledge and Ability							
Words Recalled - First Time	2284	4.78	1.33	0	5	6	10
Words Recalled - Second Time	2284	3.90	1.51	0	4	6	9
Words Recalled - Total	2284	8.68	2.60	0	9	12	19
Number Correct on Raven's Matrices	2284	1.51	0.93	0	1	3	3
Numeracy							
Number Correct	2284	0.71	0.98	0	0	2	3

Analysis sample is 2284 people from rural Malawi who began (either complete or incomplete) a survey at baseline

Table 1: Summary Statistics of Baseline (stage one) Sample