THE ECONOMIC JOURNAL

The Economic Journal, **128** (*February*), 159–188. Doi: 10.1111/ecoj.12378 © 2016 Royal Economic Society. Published by John Wiley & Sons, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

REVISING COMMITMENTS: FIELD EVIDENCE ON THE ADJUSTMENT OF PRIOR CHOICES*

Xavier Giné, Jessica Goldberg, Dan Silverman and Dean Yang

We implement an artefactual field experiment in rural Malawi to study revisions of prior choices regarding future income receipts. This allows examination of intertemporal choice revision and its determinants. New tests provide evidence of self-control problems for some participants. Revisions of money allocations towards the present are positively associated with refined measures of present-bias from an earlier survey and with the randomly assigned closeness in time to the first possible date of money disbursement. We find little evidence that revisions of allocations towards the present are associated with spousal preferences for such revision, household shocks, or the financial sophistication of respondents.

The well-being of individuals, especially those who live close to subsistence, depends importantly on the ability to make and execute intertemporal plans. The world over, however, individuals close to subsistence appear to leave consumption unsmooth, save at a low rate, or fail to use inexpensive agricultural and health inputs.

While these observed choices may be optimal given the constraints that individuals face and the incompleteness of markets, researchers have suggested that they may be the result of self-control problems.[1]

In this article, we investigate several potential sources of failure to pursue intertemporal plans by studying why choices about future consumption are revised. The article makes two contributions. First, we test for the presence of self-control problems using a novel and robust method. Second, we provide a quantitative analysis of this and other motives for the adjustment of prior choices.

Applied research typically models self-control problems as the result of present-biased (quasi-hyperbolic) time discounting. This modelling strategy is founded, in part, on evidence of non-constant time discounting. Several studies can be interpreted to show that time discount rates decline as tradeoffs are pushed into the temporal distance.[2] In particular, many experimental studies document 'static' preference reversals: 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.

* Corresponding author: Dan Silverman, Department of Economics, W.P. Carey School of Business, Arizona State University, PO Box 879801 Tempe, AZ 85287-9801, USA Email: dsilver3@asu.edu.

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, Pascaline Dupas, Yoram Halevy, Vivian Hoffmann, 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. We thank Lasse Brune, Jason Kerwin and Prachi Jain for excellent research assistance. Data and codes that replicate the results are available on the Economic Journal website.

[1] Some of the seemingly puzzling evidence regarding intertemporal choices of the poor were first summarised by Theordore Schultz in his 1979 Nobel Prize lecture and more recently in Banerjee and Duflo (2011).

[2] Thaler (1991), Ainslie (1992) and Loewenstein and Elster (1992) provide reviews.

Interpreted as present-biased time discounting and assuming time-separable preferences, these static preference reversals imply time-inconsistency: the choices (plans) that a person makes now about consumption at a later date are different from the choices he would make when that date arrives.[3] Self-control problems and a demand for commitment may thus emerge.[4]

However, until recently there have been no studies in the literature of whether static preference reversals are associated with time-inconsistency. To our knowledge, Halevy (2015) is the sole experiment in which the revision of previous decisions is a variable of interest. Augenblick *et al.* (2015) study revision of prior choices, focusing on dynamic inconsistency in monetary *versus* real effort choices. Otherwise, existing work has either studied the static preference reversals themselves, the stability of time preferences, or the relationship between static preference reversals and the demand for commitment.

While demand for commitment is, like time-inconsistency, a signature prediction of (quasi-)hyperbolic discounting models, studies that focus on the demand for commitment may understate self-control problems either because commitment devices are poorly designed and thus not demanded (Beshears *et al.*, 2015) or because demand for commitment requires some sophistication on the part of respondents: individuals who are naïve about their self-control problems should not want to limit their future choices.

Testing the central mechanism linking static preference reversals to self-control problems – by investigating the correlation between them and the revision of prior choices – is important because the static reversals can be driven by different factors.[5] For example, static preference reversals may reflect predictable changes in the marginal utility of consumption.[6] Alternatively, static preference reversals may reflect inattention, confusion about tradeoffs, or responses to perceived experimenter demands.[7] Finally, even if preferences under commitment were well-described by changing time discount rates, simply making a plan may limit self-control problems.[8] Individuals making static preference reversals for any of these reasons need not exhibit time-inconsistency.

In addition, there may be other explanations for the revision of prior choices. For example, individuals from close-knit communities in developing countries are often obliged to share their income with relatives and friends and such social pressure may prevent individuals from pursuing privately optimal choices and the revision of previous decisions.[9] Unexpected events could also motivate revisions to otherwise

[3] Early contributions include Phelps and Pollak (1968), Laibson (1997) and O'Donoghue and Rabin (1999). See DellaVigna (2009) and Bryan *et al.* (2010) for recent reviews of empirical applications.

[4] Ashraf *et al.* (2006), Duflo *et al* (2011), Dupas and Robinson (2011) and Brune *et al.* (2016).

[5] Halevy (2015) distinguishes between time-consistency, time-invariance and stationarity, making clear that static preference reversals are identified with non-stationarity but need not imply time-inconsistency.

[6] This observation has been made by Andersen *et al.* (2008), Andreoni and Sprenger (2012) and Ericson and Noor (2015), who note that proper inference about time discounting requires information about the curvature of the utility function.

[7] Benjamin *et al.* (2013) document correlations among test scores, cognitive load and short-term patience.

[8] Making plans or setting goals can affect self-control and self-efficacy (Bandura, 1997; Ameriks *et al.*, 2003). This idea is also consistent with economic models of costly self-control such as Gul and Peendorfer (2001), Ozdenoren *et al.* (2012) and Fudenberg and Levine (2012), in which consumers may both seek commitment and, yet, not always exhibit time-inconsistency.

[9] See, e.g., Platteau (2000), Maranz (2001), Anderson and Baland (2002), Ligon *et al* (2002), Hoff and Sen (2006), Ashraf (2009), Baland *et al.* (2011), Jakiela and Ozier (2011) and Schaner (2015).

optimal consumption paths. Finally, individuals could simply make mistakes in their original decisions and seek to revise them later. Our analysis explores the role of these three alternative explanations.

From a policy standpoint, it is important to understand what drives revision behaviour because it will influence the design of commitment devices and their welfare impact. If social pressure, shocks, or mistakes affect revisions, then commitment devices could be designed either to shield resources from one's social network (while maintaining access for oneself), or to allow access in case of emergency or error. In contrast, if self-control problems are important then commitment devices should protect resources from one's future self.

To assess the drivers of revision behaviour, we implement an artefactual field experiment where the key dependent variable is revision of a previous decision under commitment. Our sample consists of several hundred wife-husband pairs in rural Malawi. We elicited intertemporal choices by adapting Andreoni and Sprenger's (2012) convex time budget method, with large real stakes (roughly a month's wages). Subjects made several choices regarding an allocation of money to be disbursed at two points, 61 and 91 days, in the future. A subset of these subjects was revisited some time prior to $t = 61$ and given the opportunity to revise the allocation between $t = 61$ and $t = 91$. A measure of this revision is our dependent variable. We examine correlates of this revision corresponding to each of the four potential determinants of revision outlined above.

The experiment also provides a complementary test of quasi-hyperbolic discounting models. In those models, average revisions towards sooner should be larger when the time lag between the revision decision and the first disbursement ($t = 61$) is sufficiently small. We randomised the number of days prior to $t = 61$ when each subject had to make the revision decision.

Analysis of initial allocations indicates that they usually, but not always, adhere to the law of demand; individuals typically allocated more income to later periods when offered higher rates of return to waiting. We interpret this to indicate that most subjects understood the choices made but that some preference reversals may simply reflect confusion. We also find that 'static' preference reversals are frequent, but only slightly more likely to be present-biased (as opposed to future-biased).[10]

Turning to revision behaviour, we find that revisions are common, often substantial in size and shift money both sooner and later. We find some evidence that time-inconsistency induces these revisions: subjects shift more money towards sooner when:

(*i*) their initial allocations are present-biased; and
(*ii*) the time lag to disbursement is shorter (when the revision decision is made six or fewer days prior to day $t = 61$).

Importantly, the relationship between present-biased and revisions towards sooner is concentrated among individuals that do not exhibit anticipated changes in the marginal utility of consumption. This finding is significant because it demonstrates, in a developing context, that predictable changes in the marginal utility of consumption

[10] This finding contrasts other studies using the multiple price list method but is consistent with Andreoni and Sprenger (2012).

may drive the observed static preference reversals. Put differently, we find evidence of a reason why not all present-biased preference reversals are the result of time-inconsistency.

We find no evidence that social pressure affects revision decisions in a meaningful way: respondents' revisions are not much higher when one's spouse's sooner allocations are larger than one's own, or when they have many other relatives in the village. We also find little evidence that shocks or financial sophistication (a proxy for mistakes) strongly predict revisions (although the impact is less precisely estimated).

The next Section presents details of the experimental design, the sample of participants and the experimental setting. Section 2 presents the theoretical framework and derives the testable implications. Then, Section 3 describes the choices under commitment and the drivers of revision behaviour. Section 4 clarifies our contribution to the related literature and Section 5 concludes.

1. The Experiment

The experiment proceeded in two stages. In stage one, we elicited intertemporal choices under commitment. Husbands and wives each separately 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'. In stage two of the experiment, some households were revisited on a randomly selected day in the two weeks prior to the arrival of the first disbursement of their money in the far period and given an opportunity to revise their original far-period allocation. Surveys at both stages measured household wealth, income, and expenditures as well as the participants' expectations for each of these variables.

1.1. *The Setting*

Rural Malawi has a number of advantages as a setting for experimental study of intertemporal choice. Most important, financial markets are thin especially during the rainy season when the experiment was conducted. During this lean period, study participants have virtually no cash and borrowing is not merely expensive but it is often impossible. Similarly, short-term saving can be difficult due to limited access to banking institutions and familial or social demands for what appears like excess cash.[11]

This financial market incompleteness is important because it reduces smoothing opportunities that confound efforts to elicit time preferences in developed economies.[12] When financial markets are thick and transaction costs low, answers to the questions asked in typical time-preference experiments should, in theory, reflect only the market rates of return participants face and reveal little about their

[11] In Malawian survey data collected by Finscope in 2008, only 26% of respondents use a formal financial product and around 60% had never heard of a savings account (Finscope, 2016).

[12] Grain and other consumption goods in store are used to smooth consumption but only partially. We rely on the fact that stakes are high and that they involve cash.

preferences (Fuchs, 1982; Chabris *et al.*, 2008).[13] Augenblick *et al.* (2015) address this issue by giving respondents in a US university campus choices over leisure that is hard to smooth instead of monetary prizes.

Our study location also has some disadvantages. Poor infrastructure makes the logistics of a large-scale experiment challenging. In addition, participants have low levels of formal education and may therefore find the experiment difficult to grasp. We therefore evaluate the consistency of participants' choices with a basic prediction of standard models of economic decision-making: the law of demand. The degree of consistency with the law of demand will provide a measure of participants' understanding of the trade-offs involved in their decisions.

1.2. *The Sample*

Participants in the experiment were recruited in January and February 2010 from a population of rural households in central Malawi who were growing tobacco as their main cash crop. Participants were a subset of respondents who were participating in another simultaneous experiment on savings.[14] To be eligible for inclusion in this experiment, respondents had to be located within 25 kilometres of the town of Mponela, to facilitate our cash disbursements. Due to our interest in interactions within the household, we further restricted our sample to farmers who were part of a married couple.

These sample restrictions left us with 1,268 targeted households. A total of 1,071 households (84.4%) and 2,142 respondents were successfully interviewed at baseline. A subset of 661 respondents (randomly selected from the full set of baseline respondents) make up the stage two sample to be revisited.

Table 1 provides summary statistics of baseline survey responses. In the full sample (panel (*a*)), the median respondent is 46 years old, has four years of formal education, lives in a village with 177 inhabitants, including four relatives other than his or her spouse. When compared to typical households from low-income countries, the households in the sample are poor and in the central Malawi region we study, tobacco farmers have similar poverty and income levels to those of non-tobacco-producing households.[15] At the time of the baseline survey, the median household in the household has a zero balance in formal bank accounts; the 90th percentile of the bank balance distribution is just 700 Malawi Kwacha (MK), or approximately US$4.67. Including the self-reported value of assets, the median household held just 4,446 MK of

[13] To illustrate, suppose that outside of the laboratory a participant can borrow or save at market rate r without transaction costs. A typical experiment asks the participant to choose between $\$x$ sooner or $\$(1 + r_e)x$ later, where r_e is 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 under option *B* contains the set under option *A* and more. Thus, for any monotonic preference ordering, option *B* is preferred. Analogously, if $r > r_e$ then is *A* preferred.

[14] See the online Appendix for further details on sampling and Brune *et al.* (2016) for details on the broader study from which our study participants were drawn. We note that the inclusion of a dummy indicating the treatment status in the savings experiment does not change the results significantly.

[15] Based on our calculations from the 2004 Malawi Integrated Household Survey (IHS), individuals in tobacco farming rural households in central Malawi live on PPP$1.48/day on average, while the average for central Malawian rural households overall is PPP$1.51/day.

Table 1

Summary Statistics

Variable	*N*	Mean	SD	Min	10th percentile	50th percentile	90th percentile	Max
Panel (*a*): baseline sample (stage 1)								
Change in allocation to sooner, avg. across all interest rates (MK)	2,142	15.70	302.66	−2,000	−280	0	340	2,000
Fraction present biased, all interest rates	2,142	0.29	0.27	0	0	0.2	0.6	1
Fraction present biased, non-implemented interest rates	2,142	0.28	0.28	0	0	0.25	0.75	1
Indicator: more elastic in the far time frame	2,142	0.33	0.47	0	0	0	1	1
Fraction of decisions consistent with law of demand	2,142	0.81	0.18	0.125	0.5	0.875	1	1
Fraction of all tokens allocated to 'sooner'	2,142	0.37	0.19	0	0.1	0.389	0.578	
Implemented interest rate	2,142	0.62	0.33	0.1	0.1	0.75	1	1
Demographics								
Male	2,142	0.50	0.50	0	0	0	1	1
Respondent's own age	2,142	46.47	14.02	18	28	46	65	95
Age 35 or under	2,142	0.24	0.43	0	0	0	1	1
36–56 years old	2,142	0.51	0.50	0	0	1	1	1
Respondent's spouse's age	2,142	46.48	14.04	18	28	46	65	95
Years of schooling	2,142	4.47	4.20	0	0	4	8	77
Some primary school	2,142	0.61	0.49	0	0	1	1	1
Primary school	2,142	0.15	0.36	0	0	0	1	1
More than primary school	2,142	0.07	0.26	0	0	0	0	1
Have adequate maize	2,142	0.22	0.41	0	0	0	1	1
Number of relatives in village	2,142	4.64	8.63	0	0	2	10	132
Total number of people in village	2,142	177.08	258.47	0	35	120	320	4,000
Aptitude questions								
Words recalled – first time	2,142	4.81	1.31	0	3	5	6	10
Number correct on Raven's matrices	2,142	1.53	0.92	0	0	2	3	3
Financial literacy questions correct	2,142	0.73	0.99	0	0	0	2	3
Wealth and income								
Total HH wealth	2,142	11,449	27,313	40	1,020	4,446	25,800	695,025
HH total in bank	2,142	447.86	2,358.96	0	0	0	700	54,000
HH total cash	2,142	156.26	1,353.26	0	0	0	100	34,000
HH items	2,142	6,218	19,737	0	600	2,346	11,625	588,290
HH animals	2,142	4,627	10,776	0	0	1,250	12,150	123,600
Expected income (in period between baseline and revisit)	2,142	1,758	6,307	0	0	50	4,470	137,700

Table 1
(*Continued*)

Variable	*N*	Mean	SD	Min	10th percentile	50th percentile	90th percentile	Max
Panel (*b*): revisit sample (stage 2)								
Change in sooner allocation upon revisiting (MK)	661	61.42	595.98	−2,000	−600	0	900	2,100
Indicator: change in sooner allocation upon revisiting is negative	661	0.31	0.46	0	0	0	1	1
Indicator: change in sooner allocation upon revisiting is positive	661	0.34	0.47	0	0	0	1	1
Fraction present biased, all interest rates	661	0.30	0.28	0	0	0.20	0.80	1
Fraction present biased, non-implemented interest rates	661	0.30	0.29	0	0	0.25	0.75	1
Indicator: more elastic in the far time frame	661	0.33	0.47	0	0	0	1	1
Fraction of decisions consistent with law of demand	661	0.81	0.17	0.25	0.63	0.88	1	1
Fraction of all tokens allocated to 'sooner'	661	0.36	0.18	0	0.11	0.38	0.57	
Days to first disbursement at revisit (targeted)	661	9.22	4.42	2	3	10	15	16
Days to first disbursement at revisit (actual)	661	8.98	4.45	1	3	9	15	16
Indicator: days to first disbursement (targeted) is six days or less	661	0.33	0.47	0	0	0	1	1
Implemented interest rate	661	0.58	0.32	0.1	0.1	0.75	1	1
Shocks								
Death in family	661	0.02	0.15	0	0	0	0	1
Shock to expected HH income	661	114.21	714.59	−2,985	−90	0	350	13,735
Spouse minus own allocation to sooner (MK)	661	712.25	488.89	0	0	700	1,400	2,000

Notes. Both baseline and revisit datasets are at individual level. Baseline dataset (panel (*a*)) composed of wife-husband pairs interviewed separately in Jan–Feb 2010. Revisit dataset (panel (*b*)) constructed by first randomly choosing 2/3 of households surveyed at baseline and then randomly choosing either husband or wife within household. Revisit interviews occurred in Mar–Apr 2010, with target revisit date randomly chosen to fall between 46 and 59 days after baseline interview (16 to 2 days prior to first 'far' period disbursement at day 61).

wealth and the 90th percentile held 25,800 MK. Because the baseline survey was conducted during the rainy season, several months would elapse before the cash crop or primary staple (maize) would be harvested in mid-April or early May. As a result, the median household expects virtually no income between the interview date and April 2010.

1.3. *Implementation of Stage One*

Figure 1 displays the timeline of the experiment. At the baseline interview, the household head and spouse were physically separated. After demographics questions, each made five independent choices regarding the allocation of 2,000 MK between tomorrow ('sooner') and 30 days from tomorrow ('later').

Each participant was given a bowl containing 20 beans (tokens) and two empty dishes, *A* and *B*. One token allocated to dish *A* corresponded to 100 MK tomorrow. One token allocated to dish *B* corresponded to 100 MK $\times$ $(1 + r)$ 30 days from tomorrow, where r is the rate of return for waiting. 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 the order before making any choices. For each rate of return, the participant made an allocation of tokens to dishes, the tokens were translated into Malawi Kwacha and the total was written above each dish on a whiteboard. The participant was then allowed to adjust the allocation. This process was repeated until the participant was ready to make the next allocation.

After completing the first five choices, the participant answered a series of questions from the baseline survey. Then, using the same elicitation method with cup, beans, and dishes, the participant again made five independent choices regarding 2,000 MK, while facing different rates of return for waiting. This time, each of the five choices concerned the allocation of money between 60 and 90 days from tomorrow (the 'far' time frame). Online Appendix Figure B1 presents a schematic of the allocation decision.

The interruption between the five choices in the near time frame and the five choices in the far was intentional. We sought to avoid having participants choose the same allocations in both 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 randomly assigned between households within clubs. With probability 1/2, a participant was first presented with the 'near' time frame allocations; otherwise, the 'far' allocations were presented first. Controlling for order effects does not affect the results, and the order in which time frames were presented does not predict choices.

Before making their choices, each participant was told that one member of the couple would be randomly chosen to have one of his or her choices implemented. The randomisation was performed on site by rolling dice, and it was designed to favour (with two-thirds probability) the far time frame to have a large enough sample of stage two revisits. Implementation took the form of a voucher, redeemable at a disbursement office set up for this purpose in the nearest town, Mponela. The voucher indicated the allocation and was issued to the member of the couple who was randomly chosen. The

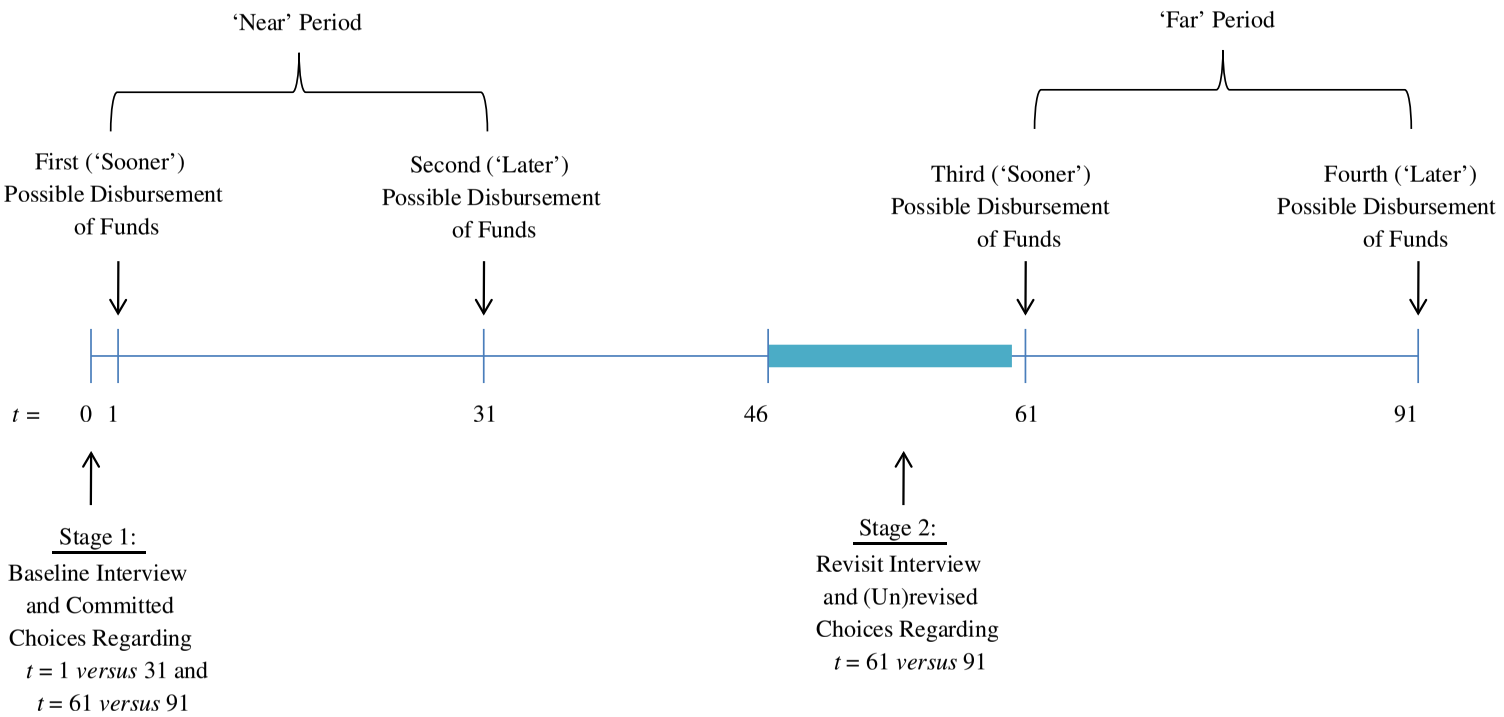

Fig. 1. *Timeline of Interviews, Choices, and Disbursement of Funds*
Note. Colour figure can be viewed at wileyonlinelibrary.com.

recipient's identity was established with a name and a fingerprint placed on the voucher.

We made key aspects of payment delivery symmetric between the 'near'and 'far' time frames. In particular, we provided two vouchers, one for the 'sooner' period (either the day after the visit or 60 days from then) and one for the 'later' period (30 days from the day of the visit or 90 days from then, depending on time frame) redeemable for cash at the disbursement office. This symmetry has advantages over a design where near payments are made in cash during the experiment. That design could favour allocations to the 'sooner' period in the 'near' time frame if participants mistrusted the experimenters or if the infrastructure in the area induced substantial transaction costs to redeeming the 'later' period voucher. A disadvantage of this symmetry is that payments were available no sooner than one day after the choices were made. Therefore, we cannot study preferences regarding consumption in the present. To the extent that changes in time discounting are largest when tradeoffs are pushed just beyond the present, any relationships between choice under commitment and revision behaviour should be attenuated.[16]

1.4. *Implementation of Stage Two*

Stage two of the experiment was only carried out with those households whose randomly selected decision concerned an allocation in the far time frame.[17]

In stage two, these households were unexpectedly revisited. The target revisit date was randomly selected from the interval between 16 and 2 days prior to day 61 (the first far-frame disbursement date). Revisits occurred even if the household chose an allocation involving no disbursement of funds at day 61.[18] Revisits occurred in March and April 2010.[19]

[16] This 'front end delay' payment method has been used in the literature by Pender (1996), Andersen *et al.* (2008) and Bauer *et al.* (2010), among others.

[17] Recall that in stage one of the experiment, one of each household's 20 decisions (10 of the husband's and 10 of the wife's) was randomly selected to be implemented. If the selected decision concerned an allocation in the near time frame (which happened with probability one-third by design), the experimental intervention was completed for that household. The chosen individual in the household redeemed the allocation and was not interviewed again.

[18] In all that follows, we focus on the randomly-assigned targeted lag (in days) to first disbursement, since it is exogenous to farmer actions. We made the first attempt to revisit each respondent on the date implied by the randomly-assigned target lag. In some cases, the actual lag was shorter than the targeted lag, because some farmers could not immediately be located. The actual lag is highly correlated with the target lag; the correlation coefficient is 0.99. 84.9% of respondents were revisited with exactly the targeted lag and 97.4% were revisited no more than two days after their target date. The maximum difference between target and actual lag is six days.

[19] In stage one, participants were told, 'We will give you one voucher for the money that you want sooner and one voucher for the money that you want later. Each voucher will have a date written on it, you will not be able to change these dates and will not be able to redeem the voucher before the date written on it'. Participants were not told that vouchers might be replaced or reissued. This framing, followed by the unannounced opportunity to revise the decision, may be perceived as deception. Inference in the experiment depends on respondents being unaware of the potential revision opportunity. The prohibition on deception in economic experiments derives in large part from circumstances where participants are drawn from a common pool and take part in multiple experiments (Jamison *et al.*, 2008). The concern is that deception in one experiment will induce scepticism about the experimenters' 'real' intent and affect behaviour in later experiments. The participants in this field experiment are not part of such a common pool.

At the revisit, the wife and husband were physically separated and a survey of wealth, income and expenditure was taken. Then, the participant whose choice had been selected to be implemented was presented with a bowl with 20 tokens. This time, 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 reflecting 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 one to 16 days from the revisit and 30 days thereafter. 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', with the same rate of return for waiting. 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 set of vouchers was issued (regardless of whether the allocation was revised) and the interview was concluded. Online Appendix Figure C1 presents a schematic of the revising procedure.

Because we sought to measure revisions of prior choices, we made the original allocation decision salient and unambiguous. This procedure is also designed to balance the consequences of implicit experimenter demands. The participant must actively choose an allocation by placing tokens in the dishes and the status quo is thus discouraged. The mere fact that we revisited the household and allowed a revision might also imply that some change is appropriate. However, because the original allocation is set out just next to the new allocation, there should be no difficulty replicating the original allocation and perhaps some mild, implicit encouragement to do so. Given the difficulty of double blind protocols in this field setting, we cannot hope to eliminate the consequences of implicit experimenter demands. Instead we designed the experiment to limit the biases they might generate.

A key element of the revisit is that participants recall the allocation they chose at baseline. The experiment therefore does not seek to study the stability of preferences after a fixed time delay (Harrison *et al.*, 2005). If that were the goal, we would not have reminded participants of their original choice and we would have repeated the elicitation method after a fixed delay. Our decision to make the allocation chosen at baseline salient also implies that the choice made at the revisiting stage is deterministic in a way that the baseline choices were not. The choice made at the revisiting stage will be implemented with certainty, while only one baseline choice (selected at random) was implemented. This difference in the choice setting may attenuate the underlying relationship between baseline choices and choices at revisiting.

The two randomisations carried out in stage one generated exogenous variation in two independent variables of interest in the regression analysis. First, the implemented choice generated exogenous variation in the interest rate that applied to the revision decision. Second the targeted revisit date, generated exogenous variation in the time to first disbursement. Consistent with the fact that these two variables were randomly assigned, both the implemented interest rate and targeted days to first disbursement are for the most part uncorrelated with key baseline respondent and household characteristics (see online Appendix C.1 for further details).

2. Theoretical Framework

In this Section we develop a theoretical framework to aid interpretation and the definition of measures used to analyse the revision behaviour.

We model participants' choices in stage one as solving a problem that is simple but sufficiently flexible to allow static preference reversals both due to changing time discount rates (quasi-hyperbolic discounting) and due to time-specific marginal utilities of consumption. We define $U_1(c)$, utility from consumption over four periods as follows:

$$U_1(c) = u_1(c_1) + \beta \times \sum_{\tau=2}^{4} \delta^{\tau-1} u_\tau(c_\tau).$$

The familiar '$\beta - \delta$' formulation of the utility function allows static preference reversals if $\beta \neq 1$. This formulation of utility also allows for a certain form of time-dependence. While utility is separable in consumption across periods, the marginal utilities of consumption may depend on time (thus the time subscript s on $u_s(\cdot)$). This captures the possibility that consumption has different marginal value at different times.

Abstracting from the discrete choice set of the experiment, we can interpret the stage one decisions about the 'near' time frame as solving:

$$\begin{aligned} &\max_{c_1, c_2 \in \mathbb{R}^+} u_1(c_1) + \beta\delta u_2(c_2) \\ &\text{subject to } c_2 \leq (2{,}000 - c_1)(1 + r) \end{aligned} \tag{Near}$$

for each rate of return r and assuming an endowment of 2,000 MK. Similarly, decisions about the 'far' time frame solve:

$$\begin{aligned} &\max_{c_3, c_4 \in \mathbb{R}^+} \beta\delta^2 u_3(c_3) + \beta\delta^3 u_4(c_4) \\ &\text{subject to } c_4 \leq (2{,}000 - c_3)(1 + r). \end{aligned} \tag{Far}$$

Interior solutions to these two problems satisfy the first-order conditions:

$$u_1'(c_1^*) = (1 + r)\beta\delta u_2'(c_2^*), \tag{FOC Near}$$

$$u_3'(c_3^*) = (1 + r)\delta u_4'(c_4^*). \tag{FOC Far}$$

This formulation is useful as it allows two distinct sources of static preference reversals but additional assumptions on the functional form of utility are necessary for choices to identify discount factors in problems (Near) and (Far).[20]

We now turn to the choices in stage two of the experiment. If the revisit is sufficiently close to period 3 then the respondent solves:

$$\begin{aligned} &\max_{c_3, c_4 \in \mathbb{R}^+} U_{revisit}(c_3, c_4) = u_3(c_3) + \beta\delta u_4(c_4) \\ &\text{subject to } c_4 \leq (2{,}000 - c_3)(1 + r). \end{aligned}$$

[20] More formally, for any u_1, u_2, $\beta\delta$ that can reconcile choices regarding the near term, there exists another $\tilde{u}_1$, $\tilde{u}_2$, $\tilde{\beta\delta}$ that can do so as well and therefore one needs additional assumptions on the functional forms to identify β, δ and the curvature parameters of the utility function.

Interior solutions here satisfy:

$$u_3'(\tilde{c}_3^*) = (1+r)\beta\delta u_4'(\tilde{c}_4^*). \quad (1)$$

Recall, the solution to the stage one problem (Far) satisfied:

$$u_3'(c_3^*) = (1+r)\delta u_4'(c_4^*).$$

Thus, abstracting from uncertainty, social pressure and mistakes, if time discounting is exponential ($\beta = 1$) then the respondent will not revise ($\tilde{c}_3^* = c_3^*$). If instead the respondent is present-biased ($\beta < 1$) then behaviour is time-inconsistent $\tilde{c}_3{}^* > c_3^*$. Analogously, if ($\beta > 1$) then $\tilde{c}_3{}^* < c_3^*$.

2.1. *The Tests*

This deterministic analysis suggests the following two tests of non-constant time discounting.

Test 1. If the respondent exhibits static, present-biased preference reversals in stage one and thus appears to have $\beta < 1$, he will shift more consumption towards sooner upon revisiting. Similarly, if the respondent exhibits static, future-biased preference reversals in stage one and thus appears to have $\beta > 1$, he would shift more consumption towards later upon revisiting.

Test 2. If the revisit occurs sufficiently close to the date of first disbursement (period 3 in the above framework) then first order condition (1) applies and present (or future) bias will be evident in a revision towards sooner (later). If instead the revisit falls far before the date of first disbursement, then first order condition (FOC Far) continues to apply and the model predicts no revision.

2.1.1. *Random choice*

Test 1 is appropriate if one assumes that choice data are dictated by the deterministic model above and so the difference between the choice and the model's prediction (or error) is interpreted as an unobserved determinant of preferences. If, however, we allow for error in the implementation of 'true' preferences, estimates of the empirical model may exaggerate the correlation between static preference reversals and time-inconsistency.

To see why, consider an extreme version of that error: a respondent that makes allocations completely at random both at stage one and at the revisit. Now consider choices exhibiting present-bias. By definition, the allocation to sooner in the far time frame is lower than for the near time frame. When choice is entirely random, therefore, the individual will, on average, allocate more tokens to sooner upon revision. In this way, participants appearing present-biased due to implementation error are mechanically more likely to revise towards sooner.[21] An

[21] Consider the following numerical example with interest rate $r = 10\%$. An individual that appears present-biased randomly allocates 1,000 to sooner and 1,100 to later in the near time frame and 600 to sooner and 1,540 to later in the far time frame. Note that since the individual appears present-biased, the allocation to sooner in the far time frame has to be smaller than the allocation to sooner in the near time frame. In our example, the allocation to sooner is 600. But because this allocation to sooner will tend to be small, the probability that more tokens will be randomly allocated to sooner upon revisit is high, and therefore individuals that appear present-biased mechanically will be more likely to allocate more tokens to sooner upon revision.

analogous effect applies to future-biased static preference reversals and revisions towards later.

We tackle this confounding effect due to implementation error in our analysis of Section 3 by constructing measures of present or future bias only from the stage one choices that were not implemented. If implementation errors are independent of each other, then measuring the tendency for static preference reversals from the non-implemented choices will break the mechanical relationship between reversals and time-inconsistency in the experiment.[22]

2.1.2. *Time-specific marginal utilities*

Alternatively, while Test 1 assumes that static preference reversals are only due to non-constant time discounting, they can also emerge from time-specific marginal utilities of consumption, which may be relevant in Malawi. For example, the marginal utility of consumption may be especially high at the time of tilling or harvest (when farmers need more calories to maintain work effort) or during the period immediately prior to harvest (when caloric consumption is low).

To illustrate, suppose time discounting is constant ($\beta = 1$) but 'flow' utility is a function of time. Suppose, in particular, that utility is iso-elastic and varies only across, but not within, time frame:

$$u_\tau(c_\tau) = \frac{c_\tau^{1-\sigma}}{1-\sigma} \quad \text{for } \tau = 1,2 \quad \text{and} \quad u_\tau(c_\tau) = \frac{c_\tau^{1-\rho}}{1-\rho} \quad \text{for } \tau = 3,4 \qquad (2)$$
$$\sigma, \rho \geq 0.$$

Interior solutions to stage one problems (FOC Near) and (FOC Far) imply:

$$\left(\frac{2,000 - c_1^*}{c_1^*}\right)^\sigma = \left(\frac{2,000 - c_3^*}{c_3^*}\right)^\rho.$$

If optimal consumption (weakly) rises within the time frame (i.e. $(1 + r) \geq \delta$), then respondents with a higher elasticity of intertemporal substitution in the 'far' time frame will exhibit a present-biased static preference reversal and thus appear less patient in the 'near'.[23] Similarly, if the participant has a higher elasticity of intertemporal substitution within the 'near' time frame ($\sigma < \rho$) then $c_1^* < c_3^*$. Such a participant would not revise his or her original allocation (and thus would not exhibit time inconsistency) because the first order condition for the stage one problem (FOC Far) is the same as that of the revisit problem (1).

While this example relies on special functional forms, the insight is general. Differences in the curvature of flow utility across time frames can induce static preference reversals that are not driven by time inconsistency.

We accommodate this in our empirical analysis of Section 3 by identifying respondents who show differences in curvature across time frames and by allowing them to have a different correlation between static preference reversals and revisions of prior choices.

[22] See online Appendix D for simulations that illustrate the consequences of using only non-implemented choices to measure a participant's tendency to make static preference reversals.

[23] More formally, if $(1 + r) \geq \delta$ and $\sigma > \rho$ then $c_1^* > c_3^*$.

3. Results

We begin with an analysis of whether intertemporal choices are consistent with the law of demand and the prevalence of static preference reversals in stage one choices. We thus use all the 2,142 observations available. We then turn to stage two choices only available for the 661 individuals that were revisited.

3.1. *Adherence to the Law of Demand*

The additive separability and monotonicity of the flow utilities assumed in Section 2 above makes the strong prediction that if participants solve problems (Near) and (Far), then the allocation to the later period, measured in Kwacha, should increase with the rate of return to waiting r.[24]

We use the degree of consistency with this prediction of standard theory as a metric for judging the appropriateness of simple economic models to interpreting choices in the experiment: if choices are inconsistent with the law of demand, either poor participants did not understand the trade-offs involved, or standard economic models have little validity in this setting.

We evaluate adherence with the law of demand by dividing each participant's ten decisions into pairs, where each element of the pair is an allocation over the same two dates. The first element of the pair is the allocation to later when facing rate of return r. The other element is the allocation to later when facing the next lowest rate of return, r'. For each participant there are eight such pairs, four for each of the two time frames. Out of 17,136 such pairs in the data, in 13,859 pairs the allocation to the later period increased with r. Thus, 81% of pairs were consistent with the law of demand. The median violation is moderate in size in the sense that it could be made consistent with monotonicity with a reallocation of less than two tokens.[25]

Becker (1962) indicates that adherence with the law of demand is not a particularly stringent test of rationality because even random choice will, on average, obey the law of demand. We therefore compare the share of consistent pairs we observe in the experiment with the share generated from a simulation where the same-sized sample makes choices purely at random (see online Appendix D for details). In the simulation

[24] To see why, think of $1/(1 + r)$ 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 in the later period. The near allocation, on the other hand, can go up or down depending on whether the income or substitution effect dominates.

[25] A comparison with existing studies in developed countries is informative as we are not aware of similar statistics being provided in studies based in developing countries. For example, in Andreoni and Sprenger (2012), the percentage of individuals that would have six or more consistent pairs of choices is 92% (using the later allocation). According to Table 2, the percentage in this experiment is somewhat lower at 76%. Similarly, using a multiple price list elicitation format Meier and Sprenger (2015) found that only 11% of a US based sample exhibited multiple switch points and thus violated monotonicity – though studies of risk preferences have exhibited much higher rates of violation (Jacobson and Petrie, 2009) than what we observe. Finally, while the published statistics are not directly comparable, the US based subjects in Augenblick *et al.* (2015) also appear to adhere to the law of demand at higher rates than those in our study.

Table 2

Number (of 8) Positive Changes in Later Allocation with Increase in r

Number of consistent pairs	(1) Real data (%)	(2) Simulated data (%)
0	0.00	0.00
1	0.28	0.04
2	0.61	1.56
3	1.96	12.08
4	7.38	33.92
5	14.05	35.49
6	21.48	14.77
7	22.97	2.05
8	31.28	0.09

Notes. The Table presents share of individuals whose allocations in eight pairs of choices (with adjacent interest rates) are consistent with law of demand. Data in column (1) are from baseline sample (for details, see Table 1). Data from column (2) are from random-choice simulations described in online Appendix D.

57% of pairs are consistent with the law of demand.[26] While substantially lower than the average rate of consistency in the experiment, this simulation suggests some caution in interpreting the choices as resulting from simple optimisation and motivates disaggregated analysis.

Indeed there is important heterogeneity in consistency with the law of demand. Table 2 presents the distribution of participants by the number of times (out of eight) they increased their later allocation with a single increase in the rate of return *r*. Column 1 shows that, measured this way, 31.3% of participants are always consistent and 75.7% are consistent at least in six out of eight allocations. At the other end of the spectrum, 10.2% of the sample violated this form of consistency in at least four allocations.[27]

In sum, these levels of consistency with the law of demand suggest that many, but not all, participants understood the trade-offs they were facing and that, for this majority, their violations of monotonicity might be attributed to occasional 'trembles' in the allocation process.

Further examination of decisions in stage one reported in Table 3 reveals that choices are usually in the interior of the budget set. For example, at a 50% rate of return to waiting, the median allocation to later is 1,950 MK and 700 MK to sooner. A minority of allocations (12% to 23%) are 'corner solutions'. The high frequency of interior allocations is consistent with participants not having adequate tools outside the experiment to facilitate consumption smoothing and also points (in the absence of very high time discount rates) to the importance of diminishing marginal utilities of consumption.

[26] In contrast to the actual data, the median violation in the simulation of random choice could be made consistent with an allocation of six tokens.

[27] Column 2 reports the simulated distribution of consistent choices if participants were to choose consumption randomly. Virtually no-one is always consistent under random choice and only 16.9% are consistent in at least six out of eight allocations.

Another important feature of this distribution of stage one allocations is the heterogeneity in the willingness to wait in exchange for a larger reward. For example, for 'later' allocations in the 'near' time frame, at a 25% rate of return, the 10th percentile is 750 MK, while at the 90th percentile it is the entire endowment. This heterogeneity is somewhat predictable with observable subject characteristics. Regression analysis in online Appendix C.2 reveals that those with more wealth at baseline allocate more to later, as do those with more relatives who live in the village.

3.2. *Static Preference Reversals*

Table 3 shows a remarkable stability across time frames. The distribution of allocations to later is not dramatically altered by the change from the 'near' to 'far' time frame. For example, the mean allocations to later at the 25% rate of return are 1,536 MK and

Table 3

Allocations to Later, in Malawi Kwacha, by Time Frame and Rate of Return

			Percentiles					Percent at a corner
	Mean	SD	10th	25th	50th	75th	90th	
Panel (*a*): allocations to later								
Near period								
$t + 30$ at $r = 10\%$	1,295.9	524.8	660	1,100	1,320	1,650	2,090	13
$t + 30$ at $r = 25\%$	1,535.8	602.1	750	1,250	1,500	1,875	2,500	14
$t + 30$ at $r = 50\%$	1,930.5	734.0	1,050	1,500	1,950	2,550	3,000	16
$t + 30$ at $r = 75\%$	2,256.8	885.1	1,050	1,750	2,275	2,975	3,500	17
$t + 30$ at $r = 100\%$	2,713.7	1,045.4	1,200	2,000	2,800	3,600	4,000	22
Far period								
$t + 90$ at $r = 10\%$	1,306.7	518.7	660	1,100	1,320	1,650	2,090	12
$t + 90$ at $r = 25\%$	1,565.4	590.0	875	1,250	1,500	2,000	2,500	14
$t + 90$ at $r = 50\%$	1,922.9	733.2	900	1,500	1,950	2,400	3,000	16
$t + 90$ at $r = 75\%$	2,306.5	872.0	1,225	1,750	2,275	2,975	3,500	18
$t + 90$ at $r = 100\%$	2,757.1	1,030.8	1,400	2,000	2,800	3,800	4,000	23
Panel (*b*): allocations to sooner								
Near period								
$t + 30$ at $r = 10\%$	821.8	477.2	100	500	800	1,000	1,400	10
$t + 30$ at $r = 25\%$	771.4	481.5	0	500	800	1,000	1,400	11
$t + 30$ at $r = 50\%$	712.9	489.3	0	300	700	1,000	1,300	14
$t + 30$ at $r = 75\%$	710.2	505.9	0	300	700	1,000	1,400	14
$t + 30$ at $r = 100\%$	643.0	522.6	0	200	600	1,000	1,400	20
Far period								
$t + 90$ at $r = 10\%$	812.3	471.7	100	500	800	1,000	1,400	9
$t + 90$ at $r = 25\%$	747.3	471.8	0	400	800	1,000	1,300	12
$t + 90$ at $r = 50\%$	718.1	488.8	0	400	700	1,000	1,400	14
$t + 90$ at $r = 75\%$	681.8	498.3	0	300	700	1,000	1,300	16
$t + 90$ at $r = 100\%$	621.4	515.4	0	100	600	1,000	1,300	21

Notes. Data are from baseline sample (for details, see Table 1). The Table presents allocations to 'later' date (either $t = 30$ or $t = 90$) for each of 10 choices presented to respondents. Baseline interview is at $t = 0$. First set of five choices is in 'near' period, when allocations are between $t = 1$ and $t = 31$. 2nd set of five choices is in 'far' period, when allocations are between $t = 61$ and $t = 91$. Rates of return to waiting until 'later' date (interest rates) take on values of 10%, 25%, 50%, 75% and 100%. Allocations between sooner and later date must be made in 100 MK increments, out of total budget of 2,000 MK.

1,565 MK in the 'near' and 'far' time frames, respectively. We find, however, that this average stability obscures substantial volatility of individual choices across time frames and masks heterogeneity in individual tendencies to shift allocations forward or back, depending on the frame.

Each participant makes five pairs of decisions where each element of a pair differs only in time frame. Of all 10,710 such pairs, just 2,927 (27%) are identical and just 4,895 (46%) differ by a token or less. Thus, in more than half of all such pairs the elements are substantially different from one another. There is a modest tendency for these static preference reversals to be present-biased. Of the 5,815 pairs that differ by strictly more than a token, 3,061 (53%) allocate more to the sooner date in the near time frame. The remaining 47% allocate more to the later date in the near time frame.[28]

These patterns in stage one indicate that static preference reversals are common and that present-biased reversals are only somewhat more common. While the distribution of these static reversals is roughly symmetric around consistency, there is evidence that they are not just the result of random trembles. Among those participants who exhibit static reversals, 18% is present-biased in at least four of five decisions. Simulations of purely random choice indicate that the percentage of individuals with at least four of five present-biased pairs would be about 8%. The tendency to be consistent or present-biased is also somewhat predictable with observable characteristics of the participants.

Table 4 presents regression results that relate a participant's tendency to be consistent or present-biased to 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. Column (1) indicates that males and those with greater maize stores tend to be more dynamically consistent. Column 3 reveals that these variables have similar relationships (with opposite signs) with the fraction present-biased, though these relationships are not statistically significant. Indeed, the reported p-value in the last row suggests that household characteristics are jointly insignificant except for column (1).

Columns (2) and (4) reveal however two important relationships. First, there is a strong association between adherence to the law of demand (subsection 3.1) and static preference reversals.[29] Greater adherence to the law of demand is associated with more dynamically consistent choices. This suggests that for many the tendency to exhibit static preference reversals may be due to a poor understanding of the choice environment. Second, there is a strong association between being more responsive to the interest rate in the far time frame and present-biased static preference reversals. As explained in Section 2.1.2, this is what we would expect if some respondents exhibit static preference reversals because their marginal utilities of consumption depend on

[28] In the simulation of random choice, 4.77% are equal, 13.85% differ by one token or less and preference reversals are equally split between present and future biased (43% each).

[29] There is no mechanical reason why these two measures must be linked. The first regards the response of allocations to changes in within time frame. The second regards consistency of allocations across time frames. For example, a subject who always violated the law of demand could be perfectly dynamically consistent, simply by replicating his non-monotonic allocations in both time frames.

Table 4

Determinants of Fraction Consistent or Fraction Present-biased

	Dependent variable			
	Fraction consistent		Fraction present-biased	
	(1)	(2)	(3)	(4)
Male	0.029*	0.024*	0.001	0.003
	(0.015)	(0.014)	(0.013)	(0.013)
Age 35 or under	−0.029	−0.013	0.017	0.011
	(0.021)	(0.019)	(0.018)	(0.018)
36–56 years old	−0.021	−0.020	0.009	0.008
	(0.017)	(0.015)	(0.015)	(0.014)
Some primary school	−0.032*	−0.030*	0.031*	0.030*
	(0.019)	(0.016)	(0.016)	(0.016)
Primary school	−0.036	−0.021	0.017	0.011
	(0.027)	(0.023)	(0.023)	(0.022)
More than primary school	−0.068**	−0.070**	0.046	0.046
	(0.034)	(0.030)	(0.031)	(0.030)
Have adequate maize	0.032*	0.022	−0.007	−0.003
	(0.018)	(0.016)	(0.015)	(0.015)
Baseline wealth (100s of MK)	−0.000	−0.000	−0.000	−0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Words recalled	0.004	0.001	0.002	0.003
	(0.006)	(0.005)	(0.005)	(0.005)
Raven's tests correct	−0.001	0.006	−0.006	−0.009
	(0.008)	(0.007)	(0.007)	(0.007)
Financial literacy questions correct	0.008	−0.001	−0.004	−0.001
	(0.009)	(0.008)	(0.007)	(0.007)
Number of relatives in the village	0.001	0.000	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Adherence to law of demand ratio [0, 1]		0.695***		−0.255***
		(0.038)		(0.034)
Indicator: more elastic in the far time frame		−0.100***		0.044***
		(0.012)		(0.012)
Constant	0.456***	−0.069*	0.252***	0.442***
	(0.032)	(0.041)	(0.028)	(0.040)
N	2,142	2,142	2,142	2,142
Adjusted R^2	0.00	0.20	−0.00	0.04
p-value that all HH characteristics = 0	0.09	0.18	0.75	0.69

Note. Dependent variable in columns 1 and 2 is the fraction of the five choices pairs that were dynamically consistent. Dependent variable in columns 3 and 4 is the fraction of choice pairs that exhibited present bias. Unit of observation is individuals included in the baseline sample. All allocations made in Jan–Feb 2010. *, **, *** denote significance at the 10%, 5% and 1% levels respectively.

time. We investigate this possibility, as well as the role of confusion about the experiment, in our analysis of stage two revision behaviour below.

3.3. *Revision Behaviour*

Before studying the determinants of revision behaviour, we first describe basic features 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. We aimed to revisit 722 respondents and we successfully collected revision choice data from 661 (91.6%).

Revisions are common. While their original choice was clear and salient, 65% of participants (432) made some adjustment to that decision. Implicit experimenter demands may have caused some participants to feel as though some change was expected of them. A large majority (87%) made a reallocation involving a shift of at least two tokens, and 64% made a reallocation involving a shift of at least four tokens. Online Appendix Figure C2 presents a histogram of changes in the participants' allocations to sooner ($t = 61$) upon revisiting, excluding those who made no change (35% of observations), illustrating the frequency of relatively large revisions.

Furthermore, revisions shift the allocation of income forward and backward in time with nearly equal frequency. Of the 432 participants who made some revision, 52% shifted income towards sooner and 48% shifted income towards later. As the histogram also indicates, the revisions towards later tended to be more modest in size. Of these, approximately 56.5% involve the shifting of at least four tokens, and just 15.5% involve shifting 10 tokens or more. The comparable figures for revisions towards sooner are 70.2% and 25.8%.

Table 5 presents the results of ordinary least-squares regressions relating revision behaviour to potential determinants of revision. The dependent variable is the change in sooner allocations upon revisiting (in MK).[30]

In column (1), independent variables are restricted to baseline characteristics and the implemented interest rate. Respondents appear to revise less towards sooner at higher rates of return: the coefficient on the interest rate is negative and statistically significant at the 10% level. Males and younger individuals (those aged 56 or below) revise more towards sooner, while more-educated individuals (primary and more than primary) revise less towards sooner. Characteristics of the respondent's spouse, and baseline maize stores and wealth add relatively little explanatory power. With evidence on these basic correlates of revisions, we now turn to Tests 1 and 2.

Test 1 evaluates present-bias as the source of static preference reversals.[31] We construct a non-parametric measure based on the number of times that a respondent made a present-biased preference reversal in stage one.[32] We account for the effects of implementation error (see subsection 2.1.1) by taking just four of the five pairs of decisions where each element of a pair differs only in the time frame (excluding the pair associated with the implemented interest rate) and calculating the fraction of those four pairs in which the participant exhibited present-biased static preference reversals.[33]

As discussed in subsection 2.1.2, static preference reversals can also be driven by changes in the marginal utility of consumption. We therefore construct a non-

[30] In online Appendix Figure C1's example, the dependent variable would take the value 200, as two tokens were added to the time t dish compared to the original allocation.

[31] In the interest of brevity, we focus here on the test for $\beta \leq 1$ and leave analysis of future bias to online Appendix C.4.

[32] An alternative approach would parameterise the utility functions in problems (Near) and (Far) and estimate individual-specific parameters. We pursue this method in online Appendix C.8.

[33] To allow for respondent error, we consider it a reversal only if the allocations differ by two tokens or more. Results are very similar if we reduce the tolerance to just one token. In addition, Appendix Table C3 provides results where our preferred measure is replaced on the right-hand-side with the fraction of all five pairs of choices (including the one associated with the implemented interest rate) in which the respondent exhibited a present-biased static preference reversal. Coefficient estimates on fraction present-biased are, as expected, larger in magnitude than those of Table 5.

Table 5

Determinants of Revisions Towards Sooner

	Ordinary least-squares regressions					
Dependent variable: change in sooner allocation upon revisiting (MK)	(1)	(2)	(3)	(4)	(5)	(6)
Preferences under commitment						
Fraction present biased, non-implemented interest rates		196.385**	195.808**	196.736**	198.033**	214.286**
		(95.222)	(95.110)	(95.457)	(97.578)	(95.871)
Fraction of all tokens allocated to 'sooner'		−523.457***	−512.533**	−495.282**	−472.658**	−468.704**
		(144.727)	(183.875)	(185.305)	(208.907)	(228.647)
Indicator: days to first disbursement (targeted) ≤ 6		107.737**	111.270**	113.743**	113.971**	124.629**
		(50.162)	(50.508)	(50.763)	(50.890)	(51.040)
Indicator: more elastic in the far time frame		127.723*	131.860*	138.471*	137.956*	136.351*
		(76.722)	(77.815)	(77.303)	(77.797)	(78.877)
Fraction present biased × indicator: more elastic in the far time frame		−200.517	−204.079	−206.583	−207.747	−216.585
		(175.339)	(177.404)	(178.197)	(180.001)	(180.320)
Financial sophistication						
Fraction of decisions consistent with law of demand			0.813	21.584	20.991	6.693
			(177.867)	(178.276)	(178.423)	(178.688)
Words recalled			−1.384	−0.818	−0.737	0.437
			(19.274)	(19.169)	(19.287)	(19.275)
Raven's tests correct			−29.722	−30.415	−30.295	−22.245
			(28.502)	(28.697)	(28.608)	(29.117)
Financial literacy questions correct			16.379	14.878	14.671	25.197
			(28.644)	(28.452)	(28.672)	(28.771)
Shocks						
Death in the family (indic.)				62.129	63.379	55.954
				(203.497)	(203.721)	(202.864)
Shortfall in expected hh income (MK)				0.053	0.052	0.049
				(0.041)	(0.041)	(0.038)
Social pressure						
Spouse minus own allocation to sooner (MK)					0.013	0.044
					(0.064)	(0.081)
Number of relatives in the village					0.830	1.589
					(3.489)	(3.382)

Table 5
(*Continued*)

	Ordinary least-squares regressions					
Dependent variable: change in sooner allocation upon revisiting (MK)	(1)	(2)	(3)	(4)	(5)	(6)
Rate of return to waiting						
Implemented interest rate {0.1, 0.25, 0.5, 0.75, 1}	−143.004*	−230.279**	−227.048**	−217.948**	−216.638**	−222.626**
	(78.279)	(81.576)	(81.585)	(81.543)	(82.219)	(82.654)
Baseline characteristics						
Male	125.169**	102.197**	103.061**	105.891**	101.623*	57.796
	(49.851)	(49.639)	(51.948)	(51.810)	(53.367)	(61.461)
Age 35 or under	198.798**	183.041**	179.567**	177.550**	177.044**	283.013**
	(71.742)	(70.736)	(70.943)	(71.396)	(71.965)	(109.969)
36–56 years old	117.670**	107.302**	110.236**	113.512**	112.177**	178.055**
	(54.887)	(54.200)	(53.493)	(53.665)	(55.422)	(64.326)
Some primary school	−66.952	−81.393	−76.593	−74.929	−75.660	−32.505
	(70.212)	(68.760)	(72.070)	(70.397)	(70.858)	(72.125)
Primary school	−159.166*	−164.017*	−160.505*	−169.354*	−170.945*	−126.996
	(85.954)	(84.113)	(90.934)	(90.073)	(90.245)	(92.344)
More than primary school	−215.094**	−230.233**	−218.818*	−222.391*	−222.868*	−148.056
	(105.964)	(105.190)	(118.268)	(116.837)	(116.832)	(119.628)
Have adequate maize	35.744	21.233	25.514	22.494	20.949	3.560
	(56.325)	(55.263)	(56.060)	(56.891)	(57.838)	(57.848)
Total HH wealth	−0.122	−0.125	−0.125	−0.130	−0.128	−0.124
	(0.089)	(0.086)	(0.085)	(0.084)	(0.085)	(0.088)
Controls for:						
Spousal characteristics	–	–	–	–	–	Yes
R^2 (adj.)	0.02	0.05	0.04	0.05	0.04	0.05
N	661	661	661	661	661	661
p-val of F-test: Frac PB + $\mathbf{1}(\Delta\varepsilon f)$ + Frac PB × $\mathbf{1}(\Delta\varepsilon f) = 0$		0.29	0.30	0.38	0.29	0.26
p-value of F-test: financial sophistication variables jointly 0			0.84	0.84	0.84	0.86

Notes. Significance levels: *10%, **5%, ***1%. Robust standard errors in parentheses. Unit of observation is individual included in revisit sample. Spousal characteristics controls are: fraction present biased for all choices, indicators for age category, indicators for education category, word recall, raven's score, financial literacy score, and fraction of choices adhering to law of demand.

parametric measure of across-time-frame differences in the curvature of utility based on the average responsiveness to the interest rate of the share of consumption allocated to later for each time frame $f \in \{near, far\}$:

$$\bar{\varepsilon}_f = \frac{1}{4} \sum_{r=0.25}^{1.0} \varepsilon_{rf}.$$

Here, ε_{rf} is the change in the share of consumption allocated to later in time frame f associated with the incremental increase in the rate of return to r.[34] We use ε_{rf} instead of the elasticity of intertemporal substitution, $\mathrm{d}\ln[(c_{t+1})/c_t]/\mathrm{d}r$ ($1/\sigma$ or $1/\rho$ in example 5) because the latter is undefined for corner solutions and, in practice, the two measures are so well correlated that, among those with interior solutions, the two produce quantitatively very similar results. Then, we take the difference in the average responsiveness across time frames, $\triangle\bar{\varepsilon}_f \equiv \bar{\varepsilon}_{far} - \bar{\varepsilon}_{near}$. When $\triangle\bar{\varepsilon}_f$ is large it indicates that the respondent was more responsive to the rate of return, and thus exhibited less curvature in flow utility, in the far time frame.[35] If such respondents also exhibit present-biased preference reversals, those reversals would not be explained by changes in the marginal utility of consumption but instead point to time-inconsistent preferences.

The importance of hyperbolic discounting for revision could be understated if present-bias is positively correlated with an overall reluctance to delay consumption. If so, present-biased static preference reversals would be positively correlated with larger initial allocations to sooner that, by definition, leave less room for revisions towards sooner. We therefore also condition on a non-parametric measure of patience: fraction of tokens allocated to sooner, across nine baseline allocations (out of 10), excluding the implemented choice.

Column (2) of the Table shows initial results of Test 1. The results are consistent with the model outlined in Section 2 where respondents are heterogeneous in both β and in the time-dependence of flow utility. The coefficient on the main effect of fraction present biased is positive and statistically significantly different from zero at the 5% level. This effect, however, only exists for individuals that do not appear systematically more elastic in the 'far' time frame. Summing the coefficients on the main effect, the indicator for 'more elastic in the far time frame' $1(\triangle\bar{\varepsilon}_f > 0.1)$ and on the interaction of fraction present-biased with the indicator, we see that those who are more elastic in the far time frame are, on average, time-consistent (the sum of the coefficients is not statistically significant, p = 0.29).

[34] Thus, if ℓ_{rf} denotes the share of consumption allocated to later in time frame f when the rate of return is r, then:

$$\varepsilon_{r'f} = \frac{\ell_{r'f} - \ell_{rf}}{r' - r}.$$

The smallest incremental increase in the interest rate is 0.15, so ε_{rf} can range from ± 6.67.

[35] Among the respondents who were revisited, $\triangle\bar{\varepsilon}_f$ ranges from -2.10 to 2.33 with a median of 0.00 and a mean of 0.01. To reduce the confounding influence of implementation error in responses, we create an indicator variable equal to one if $\triangle\bar{\varepsilon}_f > 0.1$ and zero otherwise. This classifies 33% of the revisited sample as 'more elastic' in the later time frame. Using a continuous measure of the across time frame difference in the responsiveness to the interest rate yields very similar conclusions but with less precision.

Test 2 exploits the randomised revisit date. Column (2) also includes on the right-hand-side of the regression an indicator for the targeted lag to first disbursement being less than or equal to six days.[36] Here the prediction is robust to concerns about time-dependence of marginal utility. If individuals have hyperbolic preferences ($\beta < 1$), they will shift more towards the present if they are sufficiently close to the time of consumption. We chose an indicator of six days or less, which captures a third of the revisited sample, in order to balance concerns about power (which might argue for a linear target lag specification) against the prediction of a non-linear relationship between targeted lag and revision that comes from a model of quasi-hyperbolic time discounting.

The estimates in column (2) provide evidence consistent with quasi-hyperbolic time discounting among some respondents. The coefficient on the indicator for six or fewer days to first disbursement is positive and statistically significant at the 5% level. In addition, as expected, the non-parametric measure of general impatience is negatively correlated with revisions towards sooner. Inclusion of this control has little effect on other regression coefficients.[37]

3.4. *Other Motives for Revision*

In column 3 we add to the regression variables measuring financial sophistication and proxying for mistakes in initial allocations. We examine whether these indicators of error predict revisions and whether a correlation between these measures and preferences in stage one explain the latter's correlation with revisions. The coefficients on these variables are typically negative, suggesting that those with greater sophistication tend to revise towards later. But the standard errors on these estimates are large, and we cannot reject a null hypothesis of large effects (either positive or negative). A joint significance test yields a similar conclusion.

As discussed in subsection 3.2 there is a negative correlation between adherence to the law of demand and static preference reversals. However, including the measure of adherence to the law of demand has virtually no effect on the point estimates of the relationship between present-biased static preference reversals and revision behaviour. There is therefore no evidence that this link between stage one preference reversals and revisions is driven by a relationship between the preference reversals and mistakes.

In column (4) we add variables representing shocks experienced since the baseline survey. Coefficients on death in the family and on shortfall in expected income have the expected negative signs. Again, the standard errors are large and we cannot reject a

[36] Online Appendix C.3 shows that alternate (in particular, linear) specifications of the target lag yield similar results and that a highly flexible specification of the target lag suggests that the step-function we use at six days is a reasonable approximation.

[37] In results available upon request, we also estimate a specification that includes a triple interaction term allowing the effect of distance to first disbursement to differ by both fraction present-biased and the indicator for more being elastic in the far period. The statistical significance of the previously discussed coefficients does not change in this specification; the magnitude of the coefficient on the fraction present-biased increases somewhat. The coefficient on the triple interaction term is positive, consistent with a larger effect of distance to first disbursement among those who are more present-biased and more elastic in the far period but not statistically different from zero.

null hypothesis of large coefficients.[38] Inclusion of these shock variables has little impact on other regression coefficients.

In column (5), we add to the regression measures of social pressure. The first variable is one's spouse's allocation to sooner minus one's own, averaged across the nine baseline allocations (out of 10), excluding the implemented choice.[39] This variable should capture pressure to revise one's allocation towards sooner coming from one's spouse. Initial allocations were made without consultation between spouses but there was ample opportunity to express preferences regarding the implemented allocation (and, implicitly, alternatives) after the allocation was revealed and vouchers issued, and before the revisit. Moreover, even though the initial allocations were made privately, one choice from each spouse was selected for potential implementation and then a dice roll in the presence of both spouses determined which allocation was actually implemented.[40] The second variable is simply the number of relatives one reports having in the village, which should proxy for pressures to share with a wider social network. Both variables enter the regression positively, consistent with the pressure leading to less saving. Their magnitudes are precisely estimated to be economically small; we can reject a null hypothesis of large positive correlations with revisions towards sooner.

In column 6, we add to the set of regressors several characteristics of one's spouse choices and performance on tests in stage one (coefficients omitted for brevity).[41] There is no evidence that any of the results we have described so far are simply due to omitted spousal variables: their inclusion has little effect on other coefficients of interest.

In sum, the patterns in Table 5 provide some support for a model of quasi-hyperbolic discounting as an account of some respondents' behaviour. Test 1 shows that individuals whose stage one allocations exhibit more present-biased preference reversals – reversals that cannot easily be explained by changes in the marginal utility of consumption – revise more towards sooner. Test 2 shows that revisions towards sooner are also larger when individuals make their revision at a time sufficiently close to the funds disbursement date. We estimate quite precisely little effect of social pressure on the tendency to revise. Finally we find no evidence that variables representing financial sophistication or shocks have statistically significant or robust relationships with revision behaviour. Thus, the results provide no support for the idea that mistakes in initial allocations (which should be more prevalent for those with lower financial sophistication) are important determinants of revision over this horizon.

Examining the coefficients from column (6) of Table 5, we can assess their economic magnitude. A useful benchmark for this purpose is the impact of a 50-percentage point

[38] Deaths affect approximately 2% of households and shocks to income tend to be small. Households expected virtually no cash income over this period. Care should therefore be used in extrapolating these results to other settings subject to greater risk.

[39] As with the present-bias ratio, we exclude the implemented choice from this calculation to guard against a spurious positive relationship caused by random choice.

[40] Revisions towards the spousal allocation could happen unwillingly, as the result of pressure from the spouse (Ashraf, 2009; Schaner, 2015), or willingly, say on the basis of information provided by the spouse as to optimal actions.

[41] These variables are: fraction present biased across all choices, word recall, Raven's score, financial literacy score, and fraction of decisions consistent with law of demand.

reduction in the rate of return to waiting 30 days, which leads to a 111.31 MK increase in revisions towards sooner. In comparison, a one-standard-deviation (0.28) increase in the measure of present-bias is associated with 60.36 MK higher revisions towards sooner; making one's revision decision within six days of day $t = 61$ raises revisions towards sooner by 124.63 MK.[42]

4. Related Literature

There is a long tradition of evaluating time preferences from observational choices over time. 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), 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. More recently, observational data have been used to estimate potentially non-constant time-discount functions. This literature, which restricts itself to estimating quasi-hyperbolic discount functions, includes Laibson *et al.* (2007), Paserman (2008), and Fang and Silverman (2009). We depart from this literature by adopting experimental methods for eliciting intertemporal choices and working with non-parametric measures of patience and present-bias.

The experimental literature on time preference is large. Influential recent examples include Andersen *et al.* (2008), Benhabib *et al.* (2010), Andreoni and Sprenger (2012), Augenblick *et al.* (2015) and Halevy (2015). Frederick *et al.* (2002) provides a review. Our article is distinguished from the bulk of this literature by, among other things, our implementation of a lab-in-the-field experiment with a large and heterogeneous sample. We can thus examine the correspondence between subjects' experimental behaviour and their 'real world' characteristics and behaviours.

Our article thus joins the relatively recent trend to augment laboratory studies of time preference with experiments in the field, such as Harrison *et al.* (2005), Ashraf *et al.* (2006) and Tanaka *et al.* (2010). Two of these studies are closely related to ours.

[42] In the online Appendix, we provide the following additional analyses. First, we show in online Appendix C.3 that the indicator we use for the targeted lag to first disbursement is a reasonable approximation. Second, in online Appendix C.4 we show that no pattern similar to that shown by present-bias appears for an analogously-defined future-bias variable. In results available upon request, we find that the coefficients on the measures of present and future-bias are not statistically different from each other when included in the same regression, though the magnitude of the coefficient on the present-bias term remains almost 70% larger than that of the future-bias term. Third, in online Appendix C.5 we provide an analysis of attrition related to the randomised target lag, showing that while attrition is statistically significantly higher at lower target lags, the magnitude of this relationship is small enough that it would be highly implausible for our results related to the target lag to be driven purely by selection. Fourth, in online Appendix C.6 we estimate the specification of column (6), Table 5 separately for males and female respondents and find no strong evidence of gender differences in key coefficients. Fifth, in online Appendix C.7, we replicate Table 5 excluding individuals that are inconsistent in three or more pairs. One may think that these individuals do not understand the experiment thus contributing to measurement error. We find that most of the results hold and that the coefficients of interest are not larger in absolute value, suggesting that there is no attenuation bias. Finally, using a flexible '$\delta - \beta$' model we structurally estimate the individual discount factor β and include it as a regressor in the specification of Table 5. Online Appendix Table C9 contains the results. Online Appendix C.8 contains the details.

The first, Ashraf *et al.* (2006), fielded hypothetical time preference questions among Philippine respondents who were then later offered a commitment saving product. 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 article differs from this study by studying the link between incentivised intertemporal allocation decisions and revision of prior choices directly. We measure the extent of preference reversals, as well as the basic consistency of choice with rational economic models and thus provide a quantitative assessment of the mechanisms behind time inconsistency and the demand for commitment. The second related paper, Harrison *et al.* (2005), elicited time preferences among Danish respondents. A subset of respondents were later revisited and asked to perform the same time preference experiment again. Our experiment differs from Harrison *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 revisions of intertemporal plans and to shed light on the determinants of such revisions.

5. Conclusion

The consequences of sub-optimal intertemporal choices can be serious, especially among the poor in developing countries. We conducted an experiment among Malawian farmers to investigate why their intertemporal choices may appear not to serve their individual self-interest. More precisely, we provide the first field evidence on the causes and correlates of decisions to revise prior intertemporal choices made under commitment. The experiment allowed subjects to make an intertemporal allocation of substantial funds they would receive at two future times 30 days apart. This future 30-day period was timed to occur during a period of low income and low food stores, during which consumption smoothing of substantial amounts of future income is very difficult. Several weeks later, prior to the first disbursal of funds, we revisited study participants and allowed them to revise their previous allocations over the same 30-day period. We examine these revisions of allocations for evidence of self-control problems as well as other potential mechanisms behind intertemporal choice revision.

We provide a new evaluation of the importance of self-control problems in a developing context. We test, in particular, whether revisions of allocations towards the present are positively associated with measures of present-bias from an earlier baseline survey, or with the (randomly assigned) closeness in time to the first possible date of money disbursement.[43] These tests complement existing tests of self-control problems based on demand for commitment devices. In contrast to analyses of demand for commitment devices, our approach has the advantage of allowing even naïve individuals (who are not aware of their self-control problems) to contribute to estimates, since naïve as well as sophisticated respondents can display revision behaviour. In addition, analysis of revision behaviour avoids problems of low demand that may arise if commitment devices are poorly designed.

[43] This result is reminiscent of Kaur *et al.*'s (2015) finding that worker effort increases as a worker's randomly-assigned payday comes closer.

We find that only a minority of our sample exhibits present-biased static preference reversals that cannot be easily reconciled by predictable changes in the marginal utility consumption. But the correlation between these reversals and revision of prior choices towards sooner is relatively large. Consistent with a model of self-control problems, the correlation between the time to the first possible date of disbursement and revisions towards sooner is negative. We find no evidence that respondents' revisions tend to move in the direction of their spouses' preferences for such revision. Similarly, though with less precision, we find no evidence that mistakes or shocks predict revisions of prior intertemporal choices.

These results suggest cautious optimism about efforts to improve the lives of the very poor in developing countries via interventions that address their problems of self-control. Our results support the view that, if we privilege an individual's preferences at moments relatively far from the present, there may be important benefits of commitment for some people and the costs of such commitments, in terms of reduced flexibility, would be limited.

This view should be tempered, however, by two important caveats. First, our findings show that present-bias, as evidenced by static preference reversals, is far from ubiquitous in this population. Many of the participants in the experiment exhibited, at most, just a modest tendency to be present-biased. In addition, we provide evidence that some of the revisions towards the sooner allocation, consistent with present-bias, are more likely to reflect anticipated time-varying intertemporal elasticities of substitution, rather than time-varying discount rates. Policy design must take account of this heterogeneity; efforts to help some with their legitimate self-control problems must avoid saddling others with commitments they do not need. Second, and related, our experiment was conducted during the lean season when little income is generated. This makes ours an unlikely context for finding that income fluctuations influence revision. It is possible that income shortfalls may influence revision behaviour in other parts of the Malawian agricultural year, such as in the post-harvest months.

World Bank and BREAD
University of Maryland
Arizona State University and NBER
University of Michigan, NBER, and BREAD

Accepted: 31 January 2016

Additional Supporting Information may be found in the online version of this article:

Appendix A. The Experiment Subject Pool.
Appendix B. Variable Definitions.
Appendix C. Supplementary Analyses.
Appendix D. Simulations of Stochastic Choice.

References

Ainslie, G. (1992). *Picoeconomics: The Strategic Interaction of Successive Motivational States within the Person*, New York: Cambridge University Press.

Ameriks, J., Caplin, A. and Leahy, J. (2003). 'Wealth accumulation and the propensity to plan', *Quarterly Journal of Economics*, vol. 118(3), pp. 1007–47.
Anderson, S. and Baland, J.-M. (2002). 'The economics of roscas and intrahousehold resource allocation', *The Quarterly Journal of Economics*, vol. 117(3), pp. 963–95.
Andersen, S., Harrison, G., Lau, M. and Rustrom, E. (2008). 'Eliciting risk and time preferences', *Econometrica*, vol. 76(3), pp. 583–618.
Andreoni, J. and Sprenger, C. (2012). 'Estimating time preferences from convex budgets', *American Economic Review*, vol. 102(7), pp. 3333–56.
Ashraf, N. (2009). 'Spousal control and intra-household decision making: an experimental study in the Philippines', *American Economic Review*, vol. 99(4), pp. 1245–77.
Ashraf, N., Karlan, D. and Yin, W. (2006). 'Tying Odysseus to the mast: evidence form a commitment savings product in the Philippines', *Quarterly Journal of Economics*, vol. 121(2), pp. 635–72.
Augenblick, N., Niederle, M. and Sprenger, C. (2015). 'Working over time: dynamic inconsistency in real effort tasks', *Quarterly Journal of Economics*, vol. 130(3), pp. 1067–115.
Baland, J., Guirkinger, C. and Mali, C. (2011). 'Pretending to be poor: borrowing to escape forced solidarity in credit cooperatives in Cameroon', *Economic Development and Cultural Change*, vol. 60(1), pp. 1–16.
Bandura, A. (1997). *Self Efficacy: The Exercise of Control*, New York: W.H. Freeman.
Banerjee, A. V. and Duflo, E. (2011). *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*. New York: PublicAffairs.
Bauer, M., Chytilova, J. and Morduch, J. (2010). 'Behavioral foundations of microcredit', *American Economic Review*, vol. 102(2), pp. 1118–39.
Becker, G. (1962). 'Irrational behavior and economic theory', *Journal of Political Economy*, vol. 70(1), pp. 1–13.
Benhabib, J., Bisin, A. and Schotter, A. (2010). 'Present-bias, quasi-hyperbolic discounting, and fixed costs', *Games and Economic Behavior*, vol. 69(2), pp. 205–23.
Benjamin, D., Brown, S. and Shapiro, J. (2013). 'Who is 'behavioral'? cognitive ability and anomalous preferences', *Journal of the European Economic Association*, vol. 11(6), pp. 1231–55.
Beshears, J., Choi, J., Laibson, D., Madrian, B. and Sakong, J. (2015). 'Self control and commitment: can decreasing the liquidity of a savings account increase deposits?', mimeo, NBER.
Brune, L., Giné, X., Goldberg, J. and Yang, D. (2016). 'Commitments to save: a field experiment in rural Malawi', *Economic Development and Cultural Change*, vol. 64(2), pp. 187–220.
Bryan, G., Karlan, D. and Nelson, S. (2010). 'Commitment devices', mimeo, Yale University.
Chabris, C., Laibson, D. and Schuldt, J. (2008). 'Intertemporal choice', in (S. Durlauf and L. Blume, eds.), *The New Palgrave Dictionary of Economics*, 2nd edn, London: Palgrave Macmillan.
Duflo, E., Kremer, M. and Robinson, J. (2011). 'Nudging farmers to use fertilizer: theory and experimental evidence from Kenya', *American Economic Review*, vol. 101(6), pp. 2350–90.
Dupas, P. and Robinson, J. (2013). 'Why don't the poor save more? evidence from health savings experiments', *American Economic Review*, vol. 103(4), pp. 1138–71.
DellaVigna, S. (2009). 'Psychology and economics: evidence from the field', *Journal of Economic Literature*, vol. 47(2), pp. 315–72.
Ericson, K. and Noor, J. (2015). 'Delay functions as the foundations of time preference: testing for separable discounted utility', mimeo, Boston University.
Fang, H. and Silverman, D. (2009). 'Time-inconsistency and welfare program participation: evidence from the NLSY', *International Economic Review*, vol. 50(4), pp. 1043–77.
Finscope (2016). 'FinScope Malawi 2008', available at: http://www.finscope.co.za/documents/2009/Brochure_Malawi08.pdf (last accessed: 6 September 2016).
Frederick, S., Loewenstein, G. and O'Donoghue, T. (2002). 'Time discounting and time preference: a critical review', *Journal of Economic Literature*, vol. 40(2), pp. 351–401.
Fuchs, V. (1982). 'Time preference and health: an exploratory study', in (V. Fuchs, ed.), *Economic Aspects of Health*, pp. 93–120, Chicago, IL: University of Chicago Press.
Fudenberg, D. and Levine, D.K. (2012). 'Timing and self-control', *Econometrica*, vol. 80(1), pp. 1–42.
Gul, F. and Peendorfer, W. (2001). 'Temptation and self-control', *Econometrica*, vol. 69(6), pp. 1403–35.
Halevy, Y. (2015). 'Time consistency: stationarity and time invariance', *Econometrica*, vol. 83(1), pp. 335–52.
Harrison, G., Lau, M. and Rustrom, E. (2005). 'Dynamic consistency in Denmark: a longitudinal field experiment', mimeo, Georgia State University.
Hausman, J. (1979). 'Individual discount rates and the purchase and utilization of energy-using durables', *Bell Journal of Economics*, vol. 10(1), pp. 33–54.
Hoff, K. and Arijit, S. (2006). 'The kin system as a poverty trap?' in (S. Bowles, S. Durlauf and K. Hoff eds), *Poverty Traps*, pp. 95–115, Princeton University Press,
Jacobson, S. and Ragan, P. (2009). 'Learning from mistakes: what do inconsistent choices over risk tell us?', *Journal of Risk and Uncertainty*, vol. 38(2), pp. 143–58.
Jakiela, P. and Owen, O. (2016). 'Does Africa need a rotten kin theorem? experimental evidence from village economies', *Review of Economic Studies*, vol. 83(1), pp. 231–68.

Jamison, J., Karlan, D. and Schechter, L. (2008). 'To deceive or not to deceive: the effect of deception on behavior in future laboratory experiments', *Journal of Economic Behavior & Organization*, vol. 68(3-4), pp. 477–88.
Kaur, S., Kremer, M. and Mullainathan, S. (2015). 'Self-control at work', *Journal of Political Economy*, vol. 123 (6), pp. 1227–77.
Laibson, D. (1997). 'Golden eggs and hyperbolic discounting', *Quarterly Journal of Economics*, vol. 112(2), pp. 443–77.
Lawrance, E.C. (1991). 'Poverty and the rate of time preference: evidence from panel data', *Journal of Political Economy*, vol. 99(1), pp. 54–77.
Ligon, E., Thomas, J.P. and Worrall, T. (2002). 'Informal insurance arrangements with limited commitment: theory and evidence from village economies', *The Review of Economic Studies*, vol. 69(1), pp. 209–44.
Laibson, D., Repetto, A. and Tobacman, J. (2007). 'Estimating discount functions with consumption choices over the lifecycle', mimeo, University of Pennsylvania.
Loewenstein, G. and Elster, J. (1992). *Choice Over Time*, New York: Russell Sage.
Maranz, D. E. (2001), *African Friends and Money Matters: Observations from Africa (Publications in Ethnography, Vol. 37)*, SIL International.
Meier, S. and Sprenger, C. (2015). 'Temporal stability of time preferences', *Review of Economics and Statistics*, vol. 97(2), pp. 273–86.
O'Donoghue, T. and Rabin, M. (1999). 'Doing it now or later', *American Economic Review*, vol. 89(1), pp. 103–24.
Ozdenoren, E., Salant, S. and Silverman, D. (2012). 'Willpower and the optimal control of visceral urges', *Journal of the European Economic Association*, vol. 10(2), pp. 342–68.
Paserman, D. (2008). 'Job search and hyperbolic discounting: structural estimation and policy implications', Economic Journal, vol. 118(531), pp. 1418–52.
Pender, J. (1996). 'Discount rates and credit markets: theory and evidence from rural India', *Journal of Development Economics*, vol. 50(2), pp. 257–96.
Phelps, E.S. and Pollak, R.A. (1968). 'On second-best national saving and game-equilibrium growth', *Review of Economics Studies*, vol. 35, pp. 185–99.
Platteau, J.-P. (2000), *Institutions, Social Norms and Economic Development (Fundamentals of Development Economics)*, Routledge.
Schaner, S. (2015). 'Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings', *Applied Economics*, vol. 7(2), pp. 135–74.
Schultz, T. (1979). 'The economics of being poor. Prize lecture to the memory of Alfred Nobel', available at: http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/1979/schultz-lecture.html (last accessed: 6 September 2016).
Tanaka, T., Camerer, C. and Nguyen, Q. (2010). 'Risk and time preferences: experimental and household survey data from Vietnam', *American Economic Review*, vol. 100(1), pp. 557–71.
Thaler, R. (1991). 'Some empirical evidence on dynamic inconsistency', in (R. Thaler, ed.), *Quasi Rational Economics*, pp. 127–33, New York: Russell Sage.
Warner, J. and Pleeter, S. (2001). 'The personal discount rate: evidence from military downsizing', *American Economic Review*, vol. 91(1), pp. 33–53.