Are the poor so present-biased?

IFS Working Paper W18/24

Rachel Cassidy

Are the poor so present-biased?*

Rachel Cassidy†

October 15, 2018

Abstract

Estimates of "present-bias" among the poor may be exaggerated if poor individuals are credit-constrained and expect to have greater liquidity in the future. I conduct an experiment in rural Pakistan which provides causal evidence of this effect. I use windfalls to generate fully exogenous variation in subjects' liquidity constraints. I show that fluctuating liquidity has a significant and sizeable effect on measures of time-inconsistency, which does not operate via cognitive functioning. Importantly, I establish that the causation runs from tighter liquidity constraints to appearing "present-biased" — rather than truly present-biased individuals making choices which lead to tighter liquidity constraints.

*I thank Nava Ashraf, Abigail Barr, Stefan Dercon, Yoram Halevy, David Kunst, Glenn Harrison, Sean Higgins, Simon Quinn, Kate Vyborny and Chris Woodruff for their comments. I also thank Tazeem Khan, Hashim Zaman and all at NRSP, as well as Mahniya Zafar, Ahmed Ayub, Mehreen Irshad and Tawfiq Hamid for fieldwork assistance. All errors are my own. This project was funded by the International Growth Centre (IGC) Pakistan. The IGC aims to promote sustainable growth in developing countries by providing demand-led policy advice based on frontier research. The views expressed here do not necessarily reflect those of IGC Pakistan nor the IGC as a whole. My work as a PhD student was supported by the Economic and Social Research Council [grant number ES/J500112/1]. Ethical approval for this project was given by the University of Oxford's CUREC [ECONCIA13-032]. The pre-analysis plan for this project was registered at http://www.socialscienceregistry.org/trials/1343. The Online Appendix can be found at https://rachelmcassidy.com/research/.

†Institute for Fiscal Studies. Email: rachel.cassidy@ifs.org.uk

1 Introduction

While poverty may impair decision-making (Banerjee and Mullainathan, 2010; Mani et al., 2013), some of the apparently irrational behaviour observed among the poor may have a rational explanation. For example, while lack of investment in fertilizer may partly be driven by present-bias (Duflo et al., 2011), it may also be driven by heterogeneous returns (Suri, 2011). Researchers attempting to measure the poor's time preferences directly have found that many individuals are unwilling give up a smaller, immediate cash payment in order to wait and receive a larger payment a few weeks later, but are relatively more willing to wait when one or both payments are dated further into the future (Ashraf et al., 2006; Tanaka et al., 2010). This has been interpreted as evidence of widespread present-bias.[1] Yet in some circumstances, such choices may represent a rational, time-consistent response to liquidity constraints. Specifically, if an individual expects to have higher liquidity in the long run, and cannot smooth this into the present because she is credit constrained, then she may act as if she has a higher marginal rate of intertemporal substitution in the short run even if she has perfectly time-consistent preferences (Noor, 2009; Dean and Sautmann, 2014; Epper, 2015; Gerber and Rohde, 2015). Researchers may therefore overestimate the extent of present-bias among populations who are credit-constrained and face time-varying liquidity constraints.

This paper provides causal evidence that poor individuals who expect their liquidity constraints to ease in the future are indeed more likely to appear "present-

[1] In wealthier settings, critics have argued that subjects may arbitrage experimental payments, although the evidence on arbitrage is mixed (Coller and Williams, 1999; Harrison et al., 2002; Cubitt and Read, 2007; Augenblick et al., 2015; Andreoni et al., 2017). Tasks involving time-dated effort or consumption have therefore been used; although subjects may also arbitrage goods via trading or stockpiles (Ubfal, 2016). The argument in this paper is that while concerns about arbitrage imply that *not* finding "present-bias" over money is insufficient to conclude that a subject is *not* truly present-biased; time-varying liquidity constraints imply that *finding* "present-bias" over money is insufficient to conclude that a subject *is* truly present-biased.

biased", and that this effect is sizeable. I conduct an experiment in rural Pakistan, in which I randomise the timing of experimental windfalls such that some subjects are anticipating to receive liquidity in a few weeks' time. I show that these individuals are significantly more likely to appear "present-biased" over money (i.e. more impatient about payments in the immediate future than payments in a few weeks' time), and significantly less likely to appear "future-biased" (i.e. vice versa). This indicates that subjects take their liquidity constraints in account when making such choices — violating the so-called "narrow framing" assumption, which is required for such choices to identify time preferences (unless subjects' liquidity does not vary over time). Moreover, this evidence demonstrates a broader confound which is also present in attempts to identify present-bias from observational data on financial or consumption behaviour; or indeed from effort choices, if subjects also expect their time constraints or cognitive bandwidth constraints to vary over time.

The experimental design is summarised as follows. Subjects participate in two sessions of activities, held two weeks apart. I randomise at the individual level whether a subject receives a windfall — a fixed sum, framed as the experimental participation fee — during the first session, or is given a receipt during the first session to claim the windfall during the second session. Subjects who receive the participation fee during the first session thereby experience an immediate and unanticipated positive liquidity shock, easing whatever liquidity constraints they currently face. Subjects who are told during the first session that they will receive the participation fee during the second session thereby experience an anticipated positive liquidity injection, dated two weeks into the future. Given that subjects are highly credit-constrained and hold low savings, this exacerbates any liquidity constraints they currently face, as they are unable to smooth this future liquidity into the present.

The results show that even this relatively modest, experimentally-induced change in liquidity constraints can indeed affect measures of "time-inconsistency". Specifically, being told that they will have to wait until the second session to receive the windfall makes subjects 31 percent more likely to appear "present-biased", and 21 percent less likely to appear "future-biased". This temporary experimental treatment by definition cannot alter subjects' "deep" underlying time preferences. Nor does the effect of the treatment on choices appear to operate via effects on subjects' trust, cognitive functioning, risk preferences, optimism, which I also measure and control for in robustness checks. The results therefore provide support for the hypothesis that subjects take their background liquidity constraints (including windfalls) into account when making such choices, and thus that such choices do not identify their time preferences.

To test whether the results also extend to quasi-experimental variation in subjects' natural liquidity constraints, I also randomly schedule each village's first session either before or after the main onset of the wheat harvest. I find that subjects interviewed prior to the harvest are 32 percent less likely to look "time-consistent" than subjects interviewed after the onset of the harvest. However, the effect is mainly driven by subjects being much more likely to appear "future-biased" prior to the harvest. A possible explanation is that subjects interviewed prior to the harvest face a high degree of uncertainty about harvest yields, and thus defer receipt of payments in behaviour akin to precautionary savings. Alternatively, especially given that subjects are women, it may be that they prefer to defer payments until after the harvest when they will be easier to conceal from their spouse or relatives. Either way, again such choices do not identify their time preferences.

This paper is most closely related to the work of Carvalho et al. (2016), who show that poor individuals in the US appear "present-biased" over money just before but not after pay-day. This may be evidence of the effect hypothesised above: that time-consistent, credit constrained subjects try to smooth liquidity into the period before pay-day, and thus spuriously appear "present-biased" before pay-day. On the other hand, the fact that there is a jump in liquidity at the end of the month may in part be driven by the fact that some individuals are truly present-biased, and so fail to smooth their pay-check to the end of the month. These individuals may thus correctly show up as present-biased at the end of the month, but spuriously appear "time-consistent" after pay-day because they are able to arbitrage experimental payments.

The major contribution of this paper is to offer that the causation runs from tighter liquidity constraints to appearing "present-biased" — rather than truly present-biased individuals making choices which lead them to have tighter liquidity constraints. I am able to do so because the windfall experiment generates purely exogenous variation in both the timing *and size* of jumps in subjects' liquidity. As a further contribution, the results from the wheat harvest show that similar effects to those observed by Carvalho et al. (2016) arise from a very different source of quasi-experimental variation in subjects' liquidity, in a very different setting that is commonly faced by the poor in developing countries.

Other recent studies have also highlighted how attempts to measure poor individuals' "time preferences" may be confounded by anticipated and unanticipated changes in their liquidity over time. Dean and Sautmann (2014) show that reduced savings and unfavourable expenditure shocks are correlated with individuals appearing "impatient" over monetary payments in a panel dataset in Mali. Yet again, causation may run from time preferences to shocks, if subjects who are truly impatient are the ones who fail to build up savings, or to invest in risk-mitigating technologies or insurance. Balakrishnan et al. (2015) find that subjects in the lab in Kenya appear to partially integrate their background expenditure when answering time preference tasks; although controlling for this does not dramatically change their estimates of present-bias. In contrast, Ambrus et al. (2015) find that correcting for subjects' anticipated income growth removes all estimated "present-bias" from their sample of lab subjects in Iceland. Yet again this does not prove that causation runs from income expectations to appearing "present-biased": it could instead be that truly present-biased subjects (who are by definition optimistic that their future self will be more patient) are also the ones who have the most optimistic income expectations.

The evidence in this paper may also help to explain two stylized facts in the literature on measuring time-inconsistency. The first is that many individuals often appear "future-biased", i.e. more "patient" about the immediate future. Such behaviour is hard to interpret as a widespread preference for delayed gratification. In contrast, a rational explanation is that savings-constrained subjects may make apparently "future-biased" choices if they expect to be *more* liquidity-constrained in the future. Second, several studies have attempted to measure time-inconsistency by examining whether subjects revise their choices about future-dated payments when asked again, nearer to the payment dates (Halevy, 2015; Giné et al., 2016; Janssens et al., 2017). Many individuals are found to exhibit "present-bias" on the standard measure described above, but not on this revision measure, or vice versa.[2] If such measures capture liquidity constraints, then the explanation for this is straightforward : the revision measure may simply capture subjects who have experienced a negative liquidity shock after they made their first choice.

[2] Giné et al. find that standard "present-biased" choices are correlated with revisions in the direction of present-bias, supporting the idea that these measures capture some truly present-biased individuals. The authors find no such correlation between standard "future-biased" choices and revisions in the direction of "future-bias", and instead find that "future-biased" revisions are correlated with households experiencing a (non-experimental) positive income shock.

2 Experimental Design

Sample and setting: I conducted the experiment in a rural district in northern Punjab, Pakistan. The subjects consisted of 530 women across 53 villages (ten women per village), who were randomly sampled from the microfinance clients of the NGO National Rural Support Program (NRSP).[3] Each subject participated in two sessions lasting approximately an hour: a baseline on day one, and a follow-up two weeks later on day fifteen. The day fifteen sessions contain exactly the same activities as the day one sessions, as described below. Their primary purpose is to enable half of respondents to be paid their participation fee after a two-week delay, which is the main experimental treatment. The fact that time preferences are measured again on day fifteen also allows me to test a number of secondary hypotheses. Five subjects were absent at follow-up, and so are dropped from all analyses in case they anticipated their attrition at baseline, leaving a final sample of 525 women.

Session structure: Both sessions and all experimental payments took place at the subject's home, to ensure that subjects made decisions in a familiar setting, and to minimise subjects' transaction costs. The sessions began with surveys: a baseline survey at the start of the day one session only; and a short income and expenditure survey at the start of both sessions. In the first session, subjects were then told about when exactly they would receive their participation fee (see below). In both sessions, the enumerator then led the subject through an identical set of time preference activities (see also below) and control activities to measure

[3] See Online Appendix for full details of the sampling and randomisation procedure. The rationale for sampling microfinance clients is that they have had extensive interaction with NRSP and are thus likely to trust NRSP to deliver future-dated payments. To the extent that group-liability microfinance loans can themselves be considered commitment savings products (Bauer et al., 2012; Afzal et al., 2017), these women might be expected to have higher rates of sophisticated present-bias than a random sample of the population.

subjects' risk preferences, cognitive functioning, and optimism (see Section 5).[4] To incentivize truthful and considered responses, subjects were randomly paid for one of their choices in the time or the risk preference tasks, determined by the subject drawing balls from a bag at the end of each session.

Windfall experiment: Subjects received a participation fee of 1000 PKR (approximately 10 USD) for taking part in the two sessions. This is equivalent to a day's income for the average household in the sample, and thus represents a non-trivial windfall. That said, it is well within average range of uncertainty that subjects report facing over household income in the next four weeks (3920 PKR), and thus is not so large a windfall as to generate unnatural behaviour. The main experiment randomised at the individual level whether a subject received her participation fee before the activities on day one, or before the activities on day fifteen.[5] If a subject was paid before the day one activities, she was given the 1000 PKR and a receipt bearing NRSP's logo. If she was going to be paid on day fifteen, she was given an almost identical IOU also bearing NRSP's logo, although she was assured that she would be paid on day fifteen even if she lost the IOU. This was to maximise trust, and to ensure that all subjects received the same amount of paperwork at the same point before the day one activities. The purpose of revealing the timing of the windfall payment after the baseline survey but before the activities was to ensure that it could not influence responses about covariates, but could influence responses in the activities, mimicking the effects

[4] Figure A.4 in the Online Appendix illustrates the full sequence of activities. The activity payment protocol can also be found in the Online Appendix. Subjects are pre-randomised by computer to receive either the block of time preference activities first or the block of control activities first; and to receive either the near or the far frame first within the time preference activities. Results are robust to controlling for this; see Online Appendix.

[5] When gaining consent to participate, subjects were informed that they would receive a 1000 PKR combined participation fee for both sessions, but not when this fee would be paid. When the timing of the payment was revealed, it was framed neutrally as being due to administrative reasons. Subjects were told that, for the same reasons, other participants in their village might be paid either on day one or on day fifteen, as randomly determined by computer.

that expectations of tighter or looser liquidity constraints might have on subjects' responses to such activities.

Survey timing experiment: Cross-cutting the main windfall experiment, I conducted a quasi-experiment which exploited the timing of the harvest to generate natural variation in subjects' liquidity constraints. The order in which villages were visited was randomised, and the baseline survey was timed such that the first half of the villages received their day one sessions before the approximate harvest start date, while the second half received their day one sessions after this date.[6] Subjects in villages assigned to the first half of the study thereby received their baseline interview in a period of tighter liquidity constraints, with the expectation that their liquidity constraints would ease in the near future once harvesting had begun (around the time of their follow-up session). Meanwhile, subjects in villages randomly assigned to the second half received their baseline (and follow-up) session after the main onset of the harvest. They were therefore interviewed at a time when liquidity constraints had begun to relax, and also when uncertainty about harvest yields had been largely resolved.

Time preference activities: The time preference activities were multiple price lists,[7] which have been widely used in both the developing-country and developed-country literature.[8] The enumerator conducted the activity with respondents

[6] In advance of the experiment, local NRSP staff predicted that the modal date for farmers to start harvesting wheat would be 25th April. The final sessions were concluded a week before the beginning of Ramadan, which began at the start of June.

[7] As an additional time preference activity, I also attempted to implement a present-equivalent task (Abdellaoui et al., 2013). However, from piloting and enumerator observation this task appeared very poorly understood by subjects, and the data show a very high degree of inconsistent switching behaviour.

[8] Recent literature has seen a shift towards use of "convex time budget" tasks (Andreoni and Sprenger, 2012a) to allow estimation of the discount rate and the curvature of the utility function from one task. However, such tasks appeared too complex for subjects in piloting. Balakrishnan et al. (2015) find that responses to convext time budget tasks are highly correlated with responses in multiple price lists.

using calendars and artificial money as a visual aid (see Online Appendix), and then recorded her answers on a tablet. In the near frame, the enumerator asked the subject to choose between 400 PKR (approximately 4 USD, and one day of the legal minimum wage in Pakistan) today and increasing amounts in two weeks' time. In the far frame, the enumerator asked the subject to choose between 400 PKR in two weeks' time and increasing amounts in four weeks' time. For each choice, the enumerator placed the corresponding cash amounts on calendars to demonstrate. The data suggest very high levels of understanding, with fewer than 2% of subjects exhibiting inconsistent choices (at some point preferring the later payment, but then switching back to preferring the sooner payment when a larger later payment was offered) in either the baseline or the follow-up session.

Subjects' responses to the time preference activities are used to construct a number of measures, in line with other studies in the literature. A subject's *near-frame switch-point* is defined as the lowest amount of Rupees at which she switches from preferring 400 PKR now to preferring that amount of Rupees two weeks from now. Her *far-frame switch-point* is defined as the lowest amount of Rupees at which she switches from preferring 400 PKR in two weeks to preferring that amount of Rupees in four weeks. She is labelled *"present-biased", "future-biased", or "time-consistent"* if her near-frame switch-point on a given day is respectively greater than, less than, or the same as her far-frame switch-point on that day. Intuitively, she is labelled "present-biased" in this way if she requires a larger return to wait for money in the near future than in the more distant future (and vice versa for "future-bias"). Since subjects perform the activity on both day one and on day fifteen, all of these measures are constructed twice.[9]

[9] Looking at revisions of prior choices on day fifteen, she is labelled *dynamically "present-biased", "future-biased", or "time-consistent"* if her near-frame switch-point on day fifteen is respectively greater than, less than, or the same as her far-frame switch-point on day one. The experimental treatment described in this paper is not predicted to cause changes in this dynamic measure; however, it is described in the Online Appendix.

The time discounting activities deliberately did not carry a front-end delay — that is, the earlier payment in the near frame was dated "today" rather than a future date. This was to maximise the chance of observing true present-bias: if subjects' non-constant discounting applies only to the very immediate future, then front-end delays may lead to under-estimation of true present-bias. On the other hand, front-end delays have been used to mitigate concerns that subjects may not trust experimenters to deliver future-dated payments. In particular, if subjects view future-dated payments as having a positive probability of default, and combine this with some kind of convex weighting probability weighting or taste for certainty, then this may lead them disproportionately to prefer payments dated today and thus spuriously to appear "present-biased" (Halevy, 2008; Andreoni and Sprenger, 2012b). Trust issues could also potentially confound the treatment effect, if receiving the participation fee on day one rather than day fifteen makes subjects have more trust that future-dated payments from the time preference activities will be made.To mitigate these concerns, all subjects were also given IOU vouchers bearing NRSP's logo for the activity payments, and the script introducing the task emphasised that the enumerator would return to the household to deliver any future-dated payments selected for payment.[10] Section 5 also presents various analyses to check that trust does not confound the treatment effects.

2.1 Experimental predictions

The following are the key experimental predictions and the intuition behind them. A full derivation can be found in the Online Appendix.

Participation fee timing experiment: If a subject is told she will receive the windfall on day fifteen (as opposed to day one), then this constitutes information

[10]Andersen et al. (2014) find that after establishing trust in a similar way, introduction of a front-end delay has a negligible effect on estimated discount rates.

that she will receive a positive liquidity shock in two weeks' time (as opposed to a positive liquidity shock now). All else equal, she is thus more likely to be anticipating increased liquidity in the near future.

1. If she exhibits narrow framing, not considering her liquidity constraints, then this should have no effect on the likelihood that she appears "present-biased" or "future-biased" on day one.

2. If she does not, or only partially, exhibits narrow framing, she should be more likely to appear "present-biased" on day one and less likely to appear "future-biased" on day one.

Survey timing experiment: If a subject receives her day one session before the main onset of the harvest (as opposed to after), then she is more likely to be anticipating an increase in liquidity in the near future.

3. If she exhibits narrow framing, this should have no effect on the likelihood that she appears "present-biased" or "future-biased" on day one.

4. If she does not exhibit narrow framing, she should be more likely to appear "present-biased" on day one and less likely to appear "future-biased" on day one.

Interaction of experiments: If a subject receives her day one session before the main onset of the harvest (as opposed to after), then she is also less likely to be able to smooth the liquidity shock generated by the windfall via her own liquidity or informal borrowing.

5. If she exhibits narrow framing, the interaction of the two treatments should have no effect on the likelihood that she appears "present-biased" or "future-biased" on day one.

6. If she does not exhibit narrow framing, the effect described in experimental prediction 2 should be stronger if she is interviewed prior to the harvest.

3 Data

3.1 Liquidity constraints

Table 1 presents pre-specified descriptive statistics from the baseline survey. Subjects appear to have limited scope for saving and borrowing, especially at short notice, and thus are unlikely to be able fully to smooth consumption or experimental payments over time. Just 23% of subjects report that their household has a bank account. Moreover, whilst 97% report having access to formal borrowing — which makes sense given that the sample is drawn from individuals who have been NRSP microfinance clients — just 3% report that they have access to informal borrowing from family, friends or neighbours. Whilst average savings are sizeable, 64% of subjects report having no cash savings at home (not shown). This comprises 37% of subjects who have no savings at all, and a further 27% who only have illiquid savings in ROSCAs or livestock. As proxies of trust that future payments will take place, subjects report high levels of trust that a female NRSP representative would keep a future appointment,[11] and indeed that they themselves would keep a future appointment with NRSP, both averaging over 4 on a Likert scale from 1-5.[12]

[11] All enumerators were female, and piloting showed that some subjects were unwilling to answer this question unless the representative was specified to be a female.

[12] The measure of trust in oneself may capture a subject's belief about her own level of personal organisation and reliability: it is highly correlated with self-reported measures from the baseline survey that the subject is good at keeping track of her money, and that she uses one or more types of commitment device to save. It may also be a euphemistic way for her to report whether she trusts NRSP staff to come back: it also has a highly significant (although small) correlation with the number of years that the subject has been an NRSP client. Either way, it should proxy the subject's belief about the probability that a future-dated payment would be successfully delivered to her by an NRSP representative.

A key descriptive statistic for the quasi-experiment is whether liquidity is indeed increasing over the survey period. Figure A.1 in Appendix A.2 plots fractional polynomial fits of subjects' total income for the last two weeks, and expected income for the next two weeks, by the date on which they received the baseline survey. The first panel shows that average income received in the last two weeks is smoothly increasing over the entire survey period. There is no sharp discontinuity around late April when the bulk of harvesting begins; but this likely reflects that some farmers start harvesting earlier and others later, and that other forms of economic activity already begin to take place in the run-up to the harvest. The second panel shows that expected income is also increasing, except for the final weeks of the survey period when harvesting has ended. Overall, it therefore appears that subjects are on a path of increasing income, at least until the very end of the survey period.

3.2 Balance

The column "mean diff." in Table 1 is generated by regressing the variable in question on a dummy for receiving the participation fee on day one, with standard errors robust to individual heteroskedasticity. Since the baseline survey was conducted prior to subjects learning about the timing of their participation fee, "mean difference" therefore constitutes a test of balance for each variable. Only "harvests wheat" is unbalanced at the 10% level, and this variable is already included in the vector of pre-specified controls for regression analyses. A test that the mean differences are not jointly different from zero also cannot be rejected.

Table A.1 in Appendix A.1 presents the same baseline survey variables split by whether subjects are in a village which randomly received its baseline interview prior to or after the main onset of the harvest. In this case the column "mean diff." is not a test of pre-treatment balance, since subjects clearly knew when

Table 1: Balance – windfall timing

	Mean	Pay day 1 Mean	Pay day 15 Mean	Mean Diff.	Total N	Pay day 1 N	Pay day 15 N
Liquidity							
HH income (100,000 PKR)	2.70	2.71	2.69	0.02	523	260	263
Savings (100,000 PKR)	0.44	0.41	0.47	-0.06	525	260	265
Bank account	0.23	0.24	0.22	0.02	525	260	265
Could borrow	0.98	0.98	0.99	-0.01	525	260	265
Could borrow formal	0.97	0.97	0.97	-0.01	525	260	265
Could borrow informal	0.03	0.03	0.03	0.00	525	260	265
Harvest							
Harvests wheat	0.37	0.33	0.40	-0.07*	525	260	265
Demographics							
Muslim	0.88	0.88	0.87	0.02	525	260	265
Education (years)	2.24	2.24	2.23	0.01	525	260	265
Housewife	0.75	0.74	0.75	-0.02	525	260	265
Age	37.85	38.23	37.48	0.74	525	260	265
Married	0.88	0.88	0.88	-0.00	525	260	265
HH size	6.30	6.27	6.33	-0.06	524	260	264
HH head	0.06	0.05	0.07	-0.02	525	260	265
HH decisions (index 0-1)	0.11	0.11	0.10	0.01	525	260	265
Trust							
Trust NRSP (1-5)	4.22	4.24	4.20	0.04	525	260	265
Trust self (1-5)	4.42	4.45	4.40	0.05	525	260	265

Notes: All variables are taken from the baseline survey, conducted at the start of the day one session prior to revelation of treatment status. Treatment status — receiving the participation fee on day one or on day fifteen — is computer-randomised prior to session. Mean diff. represents difference in means across the two treatment arms. *, ** and *** indicate significance of this difference at the 10%, 5% and 1% levels respectively, as estimated from a regression of the variable of interest on the treatment indicator, with standard errors robust to individual heteroskedasticity. 100,000 PKR $\approx$ 1,000 USD.

responding to the baseline survey whether their interview was taking place before or after the harvest. Thus most of the observed significant differences likely reflect real differences before and after the main onset of the harvest, rather than a failure of randomisation.For example, "trust" is significantly lower before the harvest, but this may reflect the fact that women interviewed just before the harvest think of "future appointments" as taking place during the busy harvest time. To control for such effects, all variables showing a significant difference are added to the vector of pre-specified controls in regressions involving the survey timing treatment (as per the pre-analysis plan).

3.3 Measures of "time-inconsistency"

Figures A.2 and A.3 in Appendix A.2 show the full distributions of near-frame switch-points on day one and far-frame switch-points on day one, by participation fee timing and by whether the day one interview is before or after the main onset of the harvest. Across all frames it is clear that, whilst many subjects are prepared to wait for payments between 400 and 500 Rs after two weeks rather than 400 Rs at the earlier date, a substantial fraction of subjects are unwilling to wait until offered much larger amounts. In each frame, 15-20% of subjects (labelled as ">1000") are unwilling to wait two weeks even when offered the maximum delayed payment of 1000 Rs. This may imply that their true marginal rate of intertemporal substitution is exceedingly high. However, it may also signal either lack of understanding, or unwillingness to accept any deferred payments because of lack of trust. The results below are therefore presented with and without these subjects.

Turning to apparent "time-inconsistency", the analysis in this paper focuses on the measures of inconsistency on day one, in line with the experimental predictions above. Table 2 displays the relationship between subjects' apparent "in-

Table 2: Inconsistency, day one & day fifteen

Day 1 static choices	Day 15 static choices "Present-biased"	"Time-consistent"	"Future-biased"	Total
"Present-biased"	41	55	40	136
	(7.9%)	(10.5%)	(7.7%)	(26.1%)
"Time-consistent"	51	119	32	202
	(9.8%)	(22.8%)	(6.1%)	(38.7%)
"Future-biased"	51	84	49	184
	(9.8%)	(16.1%)	(9.4%)	(35.2%)
Total	143	258	121	522
	(27.4%)	(49.4%)	(23.2%)	(100.0%)

Notes: "Present-biased", "time-consistent" and "future-biased" are dummy variables indicating that a subject's near-frame switch-point in the multiple price list activity on that day was respectively greater than, less than or the same as her far-frame switch-point on that day. Subjects are classified as either "present-biased", "time-consistent" or "future-biased". Subjects received the same activity order and frame order on day fifteen as they had received on day one.

consistency" across the near and the far frame on day one, and across the near and the far frame on day fifteen. Several patterns emerge. First, the level of apparently "time-inconsistent" decision-making is high, with less than half of the sample exhibiting the same choices across the near and the far frames on either day. The size of the discrepancies across time frames (not shown) suggests that these inconsistencies across frames are unlikely to be due to indifference or noise: for those subjects who are classified as "time-inconsistent" on day one, subjects on average differ by four choice bands across the near and the far frame. Second, in line with the studies mentioned earlier, the proportions of subjects appearing "future-biased" is always similar in magnitude to the proportion appearing "present-biased". Third, there is no evidence to suggest that "present-bias" on day one is driven by lack of trust that enumerators will return with payments: on day fifteen, subjects have already seen enumerators return once and thus should trust them more, but subjects exhibit the same amount of "present-bias".[13] Fi-

[13] The proportion of subjects making "future-biased" decisions decreases over sessions. This may be indicative of a learning effect, if subjects made "future-biased" decisions through mis-

nally, just 7.9% of subjects appear "present-biased" on this measure on both dates, again suggesting that these measures do not identify time preferences (at least not stable ones).

4 Results

4.1 Participation fee timing

Table 3 presents the results of the windfall experiment. The estimating equation is the following logit model:[14]

$$Pr\left[Y_{idvt}=1|treat_{iv},\boldsymbol{X_{iv}}\right]=\Lambda\left[\alpha_1 treat_{iv}+\boldsymbol{\gamma}\boldsymbol{X_{iv}}+\eta_{vdt}\right], \tag{1}$$

where Y_{ivtd} is a binary outcome variable for subject i in village v on session day $d \in \{1, 15\}$ at calendar date t.[15] The dummy variable $treat_{iv}$ is an indicator equal to one if subject i was assigned to being paid the participation fee on day fifteen rather than day one. $\boldsymbol{X_{iv}}$ is a vector of pre-specified time-invariant individual controls from the baseline survey: household income, savings, and possession of a bank account; ability to borrow in the next two months (formally or informally) if needed; whether the household harvests wheat; household religion; the subject's education; and her occupation (housewife or other). η_{vdt} is a village fixed effect, included because randomisation was stratified by village (Bruhn and McKenzie, 2009). Standard errors are robust to individual-level heteroskedasticity, since randomisation was at the individual level. $\hat{\alpha}_1$ represents the average treatment effect,

understanding on day one but not on day fifteen. Alternatively, given that far fewer day fifteen sessions take place before the harvest, the increase in "time-consistent" decision-making on day fifteen could also reflect the sample being on average less liquidity-constrained.

[14] In line with the pre-analysis plan, ordinary least squares specifications were also estimated. All results are virtually identical; available on request.

[15] For completeness, the Online Appendix also reports results where the dependent variables are the continuous measures of a respondent's "switch-point" in each time frame.

since assignment to treatment and actual treatment status coincided by construction.[16] Since village fixed effects are included, $\hat{\alpha}_1$ is identified by comparing within each village the five subjects who were randomly paid their participation fee on day one to the five subjects who were randomly paid on day fifteen.

Experimental prediction 2 is that $\hat{\alpha}_1$ should be positive when the dependent variable is appearing "present-biased" on day one, and negative when the dependent variable is appearing "future-biased" on day one. Columns (1) and (3) of Table 3 show that the results are consistent with these predictions: being told on day one that the participation fee will be paid on day fifteen, as opposed to receiving the participation fee on day one, increases the probability of a subject appearing "present-biased" by 7.1 percentage points (p-value 0.081) and decreases the probability of her appearing "future-biased" by 8.0 percentage points (p-value 0.080). There is no effect on the probability of appearing "time-consistent", which is consistent with there being no clear theoretical prediction: the treatment should shift some subjects from appearing "present-biased" into appearing "time-consistent", but shift others out of appearing "time-consistent" and into appearing "future-biased".

Columns (4), (5) and (6) of Table 3 restrict the sample to subjects who at some point switch to choosing a later payment in both the near frame and the far frame, on both day one and day fifteen. For "switchers", deferring the participation fee to day fifteen increases the probability of appearing "present-biased" by 11.1 percentage points (p-value 0.078) and decreases the probability of appearing "future-biased" by 15.8 percentage points (p-value 0.012). Thus the point estimates are more significant for this subsample and the effect sizes appear larger, although not significantly so. The effects observed in the full sample in columns

[16]Treatment-arm-specific instructions automatically displayed on the enumerator's survey tablet, including payment and voucher procedures which had to be followed and verified by the enumerator before she could proceed.

Table 3: Treatment effects on day one inconsistency – windfall timing

	(1)	(2)	(3)	(4)	(5)	(6)
	"Present-biased"	**"Time-consistent"**	**"Future-biased"**	**"Present-biased"**	**"Time-consistent"**	**"Future-biased"**
	Day one	**Day one**	**Day one**	**Day one**	**Day one**	**Day one**
	All	**All**	**All**	**Switchers**	**Switchers**	**Switchers**
	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)
Participation fee timing						
Pay on day 15	0.071*	0.009	-0.080*	0.111*	0.050	-0.158**
	(0.041)	(0.048)	(0.046)	(0.063)	(0.070)	(0.067)
Controls	✓	✓	✓	✓	✓	✓
Village f.e.'s	✓	✓	✓	✓	✓	✓
Observations	493	503	504	260	272	294
Control mean	0.231	0.381	0.388	0.221	0.350	0.429

Notes: Standard errors in parentheses. *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively. All dependent variables are measured after revelation of participation fee treatment status. "Present-biased" ["time-consistent"]("future-biased") is a dummy indicating a near-frame switch-point greater than [equal to] (less than) the far-frame switch-point in the multiple price list activity. Subjects are classified as either "present-biased", "time-consistent" or "future-biased". "Switchers" are subjects who at some point switch to accepting a deferred payment in both the near frame and the far frame, on both day one and day fifteen. Pre-specified controls are: household income, savings, and possession of a bank account; ability to borrow in the next two months (formally or informally) if needed; whether the household harvests wheat; household religion; respondent's education; and her occupation (housewife or other). Reported effects represent the marginal effects at the mean. N=525 for the full sample, and N=294 for switchers. Sample sizes differ where village fixed effects perfectly predict the outcome variable.

(1) and (3) are therefore clearly not driven by "non-switchers", i.e. subjects who may have had severe trust concerns, or who may have been trying to signal their poverty to NRSP.

4.2 Survey timing

Table 4 shows the results of the survey timing experiment. The estimating equation is the following logit model:

$$Pr\left[Y_{idvt} = 1 | \textit{pre-harvest}_v, \boldsymbol{X_{iv}}\right] = \Lambda\left[\alpha_1 \textit{pre-harvest}_v + \gamma \boldsymbol{X_{iv}}\right], \qquad (2)$$

where $\textit{pre-harvest}_v$ is a dummy indicating that the subject's village received its baseline interviews prior to the main onset of the wheat harvest.[17] Standard errors are clustered at the village level, since survey dates were randomised at the village level.[18]

In terms of prediction 4, columns (1)-(3) of Table 4 show that subjects randomly interviewed just prior to the harvest are not significantly more likely to appear "present-biased". However, there is a dramatic reduction of 14.6 percentage points in the proportion who appear "time-consistent". This is driven by the fact that, in contrast to prediction 4, subjects interviewed prior to the main harvest are 12.9 percentage points *more* likely to appear "future-biased". This cannot be straightforwardly explained by subjects anticipating the easing of

[17] This "pre-post" specification was registered in the pre-analysis plan to allow comparability with the effects of the participation fee timing treatment. Given the smoothness of the increase in income over the survey period shown in Figure A.1 in Appendix A.2, I also re-estimate Equation 2 with the day of the baseline survey period (1-48) on the right-hand side. Subjects are on average 0.5 pp more likely to appear "time-consistent" and 0.5 pp less likely to appear "future-biased" for each day later into the baseline survey period; see Online Appendix.

[18] As a conservative measure, Equation 2 is estimated on the first 48 villages only, as the remaining five villages were brought into the sample after the initial randomisation. Point estimates become slightly larger and even more significant if the final five villages are included; results available on request.

Table 4: Treatment effects on day one inconsistency – survey timing

	(1)	(2)	(3)	(4)	(5)	(6)
	“Present-biased”	**“Time-consistent”**	**“Future-biased”**	**“Present-biased”**	**“Time-consistent”**	**“Future-biased”**
	Day one	**Day one**	**Day one**	**Day one**	**Day one**	**Day one**
	All	**All**	**All**	**All**	**All**	**All**
	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)
Survey timing						
Pre-harvest	0.016	-0.146***	0.129***	0.045	-0.147**	0.100*
	(0.037)	(0.055)	(0.049)	(0.046)	(0.064)	(0.055)
Pre-harvest*Harvests wheat				-0.083	0.004	0.082
				(0.087)	(0.097)	(0.094)
Harvests wheat	-0.004	-0.004	0.009	0.039	-0.006	-0.036
	(0.047)	(0.053)	(0.049)	(0.068)	(0.078)	(0.076)
Controls	✓	✓	✓	✓	✓	✓
Observations	476	476	476	476	476	476
Control mean	0.244	0.456	0.300	0.244	0.456	0.300

Notes: Standard errors in parentheses. *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively. All dependent variables are measured after revelation of participation fee treatment status. “Present-biased” [“time-consistent”](“future-biased”) is a dummy indicating a near-frame switch-point greater than [equal to] (less than) the far-frame switch-point in the multiple price list activity. Subjects are classified as either “present-biased”, “time-consistent” or “future-biased”. “Pre-harvest” indicates a day one session prior to 25th April 2016. “Harvests wheat” is a dummy variable equal to one if the subject’s household will harvest wheat at this harvest. The sample excludes the five villages which were included at the end of the survey period to boost sample size. Pre-specified controls are: household income, savings, and possession of a bank account; ability to borrow in the next two months (formally or informally) if needed; whether the household harvests wheat; household religion; respondent’s education; and her occupation (housewife or other). Additional controls due to imbalance on survey timing are: trust in NRSP to keep a future appointment; trust in oneself to keep future appointments; and decision-making power within the household. Reported effects represent the marginal effects at the mean. N=476.

liquidity constraints after the harvest, as this should make them appear "present-biased".

One possible explanation, not predicted in the pre-analysis plan, is that subjects interviewed prior to the harvest face more uncertainty: about the timing of the harvest, their income from the harvest and from related economic activity, and indeed whether they will be able to repay any supplier credit and microfinance loans that they have used for inputs. They may therefore defer receipt of experimental payments in the near future, in behaviour akin to precautionary savings. To check for evidence consistent with this, I re-estimate Equation 2 with the difference between a subject's maximum and minimum expected income over the next four weeks as the dependent variable (see Online Appendix). This measure of income uncertainty is omdeed strongly positively correlated with *pre-harvest*: whilst subjects interviewed after the main onset of the harvest report an average expectation range of 3441 PKR, those interviewed prior to the harvest report a range which is 1358 PKR larger.

However, columns (4)-(6) of Table 4 show that the effects of the increase in "future-bias" is not concentrated among subjects whose households will actually harvest wheat: the point estimates on the interaction term *pre-harvest*harvests wheat* are small and never significant. This may suggest that uncertainty before the harvest is equally important for individuals whose households do not harvest wheat, for example because the harvest catalyses other forms of economic activity such as hiring of casual labour.On the other hand, it may be suggest that apparent "future-bias" before the harvest is driven by something else. Another possible explanation is that subjects prefer to defer payments until after the harvest because they will be easier to conceal from other family members when there is more cash in the household. I test this by interacting the pre-harvest treatment with various proxies of intra-household demands. The increase in "future-bias"' prior to the harvest appears to be driven by women in households with a larger number of individuals (see Online Appendix); although again this analysis was not pre-specified.[19]

4.3 Interaction of treatments

Next, I re-estimate Equation 2 adding interaction terms between the survey timing treatment and the windfall timing treatment. Table 5 presents the results. Column (1) shows no evidence in support of experimental prediction 6, that the effect of deferring the windfall on "present-bias" should be concentrated among subjects interviewed prior to the harvest: the interaction term is insignificant, and the main treatment effect is not significantly different from the estimate in column (1) of Table 3 but loses significance.

Turning to columns (4)-(6), there is still no significant evidence that being interviewed pre-harvest increases measured "present-bias" (experimental prediction 4), although again the point estimate carries the predicted sign. The coefficient on the interaction term also carries the predicted sign but is marginally insignificant (p-value 0.101). The large effect of *pre-harvest* on reducing "time-inconsistency" is unchanged compared to Table 4, and the interaction term is insignificant. However, column (6) shows that the spike in measured "future-bias" prior to the harvest comes entirely from subjects who receive the windfall on day one. It therefore appears that subjects use their choices to "save" the participation fee, for one of the motives described above.

[19] A further alternative explanation is that subjects are truly present-biased, but that receiving a windfall prior to the harvest affects their propensity to become sophisticated and use the experimental time-preference payments as a commitment device. Evidence from the baseline survey is mixed: "future-bias" is not correlated with a subject's reports that uses one or more commitment devices to save; but is negatively correlated with subjects reporting that they attempt to use commitment devices to save but fail.

Table 5: Treatment effects – interaction of windfall & survey timing

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	“Present-biased”	**“Time-consistent”**	**“Future-biased”**	**“Present-biased”**	**“Time-consistent”**	**“Future-biased”**	**“Present-biased”**	**“Time-consistent”**	**“Future-biased”**
	Day one	**Day one**	**Day one**	**Day one**	**Day one**	**Day one**	**Day one**	**Day one**	**Day one**
	All	**All**	**All**	**All**	**All**	**All**	**All**	**All**	**All**
	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)
Participation fee timing									
Pay on day 15	0.030	0.071	-0.105*				0.020	-0.015	-0.007
	(0.043)	(0.060)	(0.061)				(0.051)	(0.072)	(0.070)
Survey timing									
Pre-harvest				-0.030	-0.176***	0.198***	-0.019	-0.184**	0.194***
				(0.047)	(0.063)	(0.059)	(0.055)	(0.075)	(0.068)
Interaction of timings									
Pay on day 15*Pre-harvest	0.049	-0.104	0.055	0.088	0.060	-0.140**	0.067	0.075	-0.133
	(0.050)	(0.073)	(0.069)	(0.054)	(0.069)	(0.064)	(0.075)	(0.101)	(0.096)
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	476	476	476	476	476	476	476	476	476
Control mean	0.225	0.472	0.303	0.225	0.472	0.303	0.225	0.472	0.303

Notes: Standard errors in parentheses. *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively. All dependent variables are measured after revelation of participation fee treatment status. “Present-biased” (“future-biased”) is a dummy indicating a near-frame switch-point greater than (less than) the far-frame switch-point in the multiple price list activity. Subjects are classified as either “present-biased”, “future-biased” or “time-consistent” (not shown). “Pre-harvest” (“post-harvest”) indicates a day one session prior to (later than) 25th April 2016. “Harvests wheat” is a dummy variable equal to one if the subject’s household will harvest wheat at this harvest. The sample excludes the five villages which were included at the end of the survey period to boost sample size. Pre-specified controls are: household income, savings, and possession of a bank account; ability to borrow in the next two months (formally or informally) if needed; whether the household harvests wheat; household religion; respondent’s education; and her occupation (housewife or other). Additional controls due to imbalance on survey timing are: trust in NRSP to keep a future appointment; trust in oneself to keep future appointments; and decision-making power within the household. Reported effects represent the marginal effects at the mean. N=476.

4.4 Heterogeneity by existing liquidity constraints

The same reasoning behind experimental prediction 6 implies that the effect of windfall should be stronger for subjects who are in general more liquidity-constrained, as they are less able to smooth the timing of the participation fee themselves. To test for this, Table A.2 in Appendix A.1 shows the interaction of the windfall timing treatment with household savings, whether the household would be able to borrow in the next two months, whether the household owns a bank account, and household income. None of the interaction terms except the interaction with "could not borrow" are significant, although lack of power means large effects cannot be ruled out.Nonetheless, the direct effect of these variables is in line with the hypothesised explanation of "time-inconsistency" based on liquidity constraints. Specifically, the marginal effect of the control *could not borrow* on appearing "present-biased" is positive, large and significant (34.3 percentage points, p-value 0.033). In contrast, the marginal effect of the control *has a bank account* is negative (9.8 percentage points, p-value 0.083). Moreover, subjects with higher household income are significantly more likely to appear "future-biased" (4.3 percentage points for every $1,000 USD increase in household income, p-value 0.015).

5 Robustness

Trust: To test directly whether a subject's trust that future meetings and payments would take place affects her measured "present-bias", I re-estimate Equation 1 controlling for the two different proxies of trust from the baseline survey. As Table A.3 in Appendix A.1 shows, neither *trust NRSP* nor *trust self* is a significant predictor of appearing "present-biased" or indeed "future-biased".However, this does not indicate that they are simply a poor proxy of subjects' beliefs that future payments will take place. Columns (3) and (4) of Table A.3 show that an increase in *trust self* significantly reduces a subject's switch-point in both the near and the far frame, thus making her appear significantly more "patient" in the near and the far frame. The results are therefore in line with concerns that subjects who perceive a lower probability of future payments taking place may more "impatient".However, the reason that this has no net effect on measured "present-bias" or "future-bias" is that the effect of trust on increasing "patience" is virtually identical across the near and the far frame.[20]

Arbitrage: Another possible explanation for why giving subjects who already receive the windfall on day one or are interviewed after the harvest exhibit lower levels of "present-bias" is if under these treatment conditions subjects already have enough liquidity to fund their desired level of current consumption. If so, experimental payments dated "now" do not translate into consumption "now", and so the monetary tasks become unable to measure true present-bias (O'Donoghue and Rabin, 2015). Instead, individuals should chose time-dated payments to maximise experimental earnings, which in this context means that they should choose any later payment greater than 400 Rs (i.e. any positive return) and should do so across both frames.[21] However, the data do not support this: even the median subject interviewed after the harvest prefers an earlier payment of 400 Rs rather than a later payment of up to 600 Rs in the near frame, and up to 650 Rs in the far frame.

[20]This suggests that subjects may weight the subjective probability of experimenter default similarly across the near and far frame, rather than exhibiting a "certainty effect" or broader convex probability weighting.

[21]No subject reports receiving a financial return on any form of savings other than livestock. It is therefore safe to assume that saving a marginal 400 Rupees outside of the experiment always produces a return inferior to the positive returns offered within the time preference activities.

Cognitive functioning: Another possible concern is that deferring the participation fee, or being interviewed prior to the harvest, might lead subjects to make more inconsistent choices by increasing financial stress and thereby reducing cognitive bandwidth. To test for this, I re-estimate Equations 1 and 2 controlling for a subject's score on the digit span test, maths test, and time taken to complete the numerical Stroop test. Like the time preference activities, these cognitive tests were elicited after the revelation of the participation fee timing; thus including these controls captures any simultaneous effects that the treatments may have had on subjects' cognitive functioning. Columns (1) and (2) of Table A.4 in Appendix A.1 report the results for the participation fee treatment (results for the survey timing treatment can be found in the Online Appendix). The coefficients on the two treatments appear unchanged and are never significantly different from the main specifications; although the coefficient on *pay on day fifteen* at times becomes marginally insignificant since the inclusion of additional controls again weaken power.[22]

Risk preferences: Receiving the participation fee on day one could also move a subject to a less curved part of her within-period utility function. If she projects this onto her future utility function (via so-called "projection bias") this might increase her propensity to take the later payment in the time-preference activities, if the later payment is seen as risky. To check for this, I re-estimate Equations 1 and 2 adding a proxy for subjects' risk preferences: a subject's aggregate certainty premium in risk preference questions elicited after the revelation of the windfall timing. The correlation between the risk premium and "time-inconsistency" is zero or very small, and including risk controls does not significantly change the

[22]Mathematical performance itself is strongly negatively correlated with "present-bias" and strongly positively correlated with "future-bias". This may suggest that subjects with higher cognitive ability are more patient, or that higher-ability subjects understand the opportunity to save at a good return through the time preference activity.

estimated treatment effects of the windfalls or the survey timing.

Optimism: Similarly, it could be that receiving the participation fee on day one makes subjects feel more optimistic that uncertain future events will go in their favour, which may change her responses to time-preference questions aside from the pure effect of easing liquidity constraints.[23] To test for this, I also re-estimate Equations 1 and 2 adding an extra control proxying subjects' optimism, also elicited after the windfall timing was revealed. Specifically, optimism was measured as the difference between a subject's stated belief of her probability of winning a draw and the objective, given probability of winning that draw. As columns (5) and (6) or Table A.4 shows, the coefficient on *optimism* is a tightly-estimated zero, and including this measure does not significantly change the treatment effect of the windfall or the survey timing.

Alternative specifications I also re-estimate the main treatment effects re-constructing the experimental measures of "time-inconsistency" in various pre-specified ways. Results show that results are robust to treatment of the experimental sample and data; see Online Appendix for details. Going beyond the pre-analysis plan, I also estimate post-double LASSO models including all of the covariates in table 1 as potential controls. Results are reported in the Online Appendix, and are insignificantly different from the results presented here.

6 Conclusion

This paper provides what appears to be the first evidence that anticipated changes in liquidity constraints can cause poor individuals' to appear "present-biased" —

[23]Subjects' mood could also impact their discount rates directly, for example Haushofer and Fehr (2014) show that experimentally-induced negative affect states increase discount rates.

or indeed "future-biased" — over money. It is therefore possible that researchers using such tasks have over-estimated the extent of present-bias exhibited by the poor in everyday financial decision-making. Indeed, this may offer a rational explanation for some of the mixed evidence on demand for and use of commitment savings products (John, 2017), in contrast with the idea that the poor are too naïve to adopt and use commitment.

Recent efforts to measure subjects' time preferences have also used effort tasks, or data on real-life effort choices such as when to file one's tax returns (Martinez et al., 2017). However, such tasks may be vulnerable to a similar critique: an individual who is anticipating (correctly or optimistically) to be less time-constrained or bandwidth-constrained in the future might spuriously appear "present-biased" in her effort choices, unless the task studied perfectly equalises the time expense and the cognitive demand for subjects in both time frames. It thus remains an empirical question — and may be context-dependent — as to whether any bias in effort tasks is greater than the bias in monetary tasks. In light of this, developing alternative approaches to measuring time-inconsistency (Toussaert, 2015) thus remains an important topic for future research.

References

Abdellaoui, M., H. Bleichrodt, O. l'Haridon, and C. Paraschiv (2013). Is there one unifying concept of utility? An experimental comparison of utility under risk and utility over time. *Management Science 59*(9), 2153–2169.

Afzal, U., G. d'Adda, M. Fafchamps, S. Quinn, and F. Said (2017). Two sides of the same rupee? Comparing demand for microcredit and microsaving in a framed field experiment in rural Pakistan. *The Economic Journal*.

Ambrus, A., T. Asgeirsdottir, J. Noor, and L. Sándor (2015). Compensated Discount Functions – An Experiment on Integrating Rewards with Expected Income. *Working paper*.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2006). Elicitation using multiple price list formats. *Experimental Economics 9*(4), 383–405.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2013). Discounting behaviour and the magnitude effect: evidence from a field experiment in Denmark. *Economica 80*(320), 670–697.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2014). Discounting behavior: A reconsideration. *European Economic Review 71*, 15–33.

Anderson, M. L. (2012). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American Statistical Association*.

Andreoni, J., C. Gravert, M. Kuhn, S. Saccardo, and Y. Yang (2017). It's Just Money: Do Experimental Subjects Arbitrage Intertemporal Monetary Rewards? *Working paper*.

Andreoni, J. and C. Sprenger (2012a). Estimating time preferences from convex budgets. *The American Economic Review 102*(7), 3333–3356.

Andreoni, J. and C. Sprenger (2012b). Risk preferences are not time preferences. *The American Economic Review 102*(7), 3357–3376.

Ashraf, N., D. Karlan, and W. Yin (2006). Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines. *The Quarterly Journal of Economics 121*(2), 635–672.

Augenblick, N., M. Niederle, and C. Sprenger (2015). Working over time: dynamic inconsistency in real effort tasks. *The Quarterly Journal of Economics 1067*, 1115.

Augenblick, N. and M. Rabin (2015). An Experiment on Time Preference and Misprediction in Unpleasant Tasks. *Working paper*.

Balakrishnan, U., J. Haushofer, and P. Jakiela (2015). How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments. *IZA Discussion Paper*.

Banerjee, A. and S. Mullainathan (2010). The shape of temptation: Implications for the economic lives of the poor. *NBER Working Papers*.

Banerjee, A. V. and S. Mullainathan (2008). Limited attention and income distribution. *The American Economic Review*, 489–493.

Bauer, M., J. Chytilova, and J. Morduch (2012, April). Behavioral Foundations of Microcredit: Experimental and Survey Evidence from Rural India. *American Economic Review 102*(2), 1118–39.

Bruhn, M. and D. McKenzie (2009). In pursuit of balance: Randomization in practice in development field experiments. *American Economic Journal: Applied Economics 1*(4), 200–232.

Carvalho, L. S., S. Meier, and S. W. Wang (2016). Poverty and economic decision-making: Evidence from changes in financial resources at payday. *The American Economic Review 106*(2), 260–284.

Coller, M. and M. B. Williams (1999). Eliciting individual discount rates. *Experimental Economics 2*(2), 107–127.

Cubitt, R. P. and D. Read (2007). Can intertemporal choice experiments elicit time preferences for consumption? *Experimental Economics 10*(4), 369–389.

Dean, M. and A. Sautmann (2014). Credit constraints and the measurement of time preferences. *Working paper*.

Deaton, A. (1991). Saving and liquidity constraints. *Econometrica 59*(5), 1221–1248.

Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2010). Are risk aversion and impatience related to cognitive ability? *The American Economic Review 100*(3), 1238–1260.

Duflo, E., M. Kremer, and J. Robinson (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from kenya. *The American Economic Review 101*(6), 2350–2390.

Epper, T. (2015). Income Expectations, Limited Liquidity, and Anomalies in Intertemporal Choice. *Working paper*.

Gerber, A. and K. I. Rohde (2015). Eliciting discount functions when baseline consumption changes over time. *Journal of Economic Behavior & Organization 116*, 56–64.

Giné, X., J. Goldberg, D. Silverman, and D. Yang (2016). Revising commitments: Field evidence on the adjustment of prior choices. *The Economic Journal*.

Halevy, Y. (2008). Strotz meets Allais: Diminishing impatience and the certainty effect. *The American Economic Review*, 1145–1162.

Halevy, Y. (2015). Time consistency: Stationarity and time invariance. *Econometrica 83*(1), 335–352.

Harris, C. and D. Laibson (2001). Dynamic choices of hyperbolic consumers. *Econometrica 69*(4), 935–957.

Harrison, G. W., M. I. Lau, and M. B. Williams (2002). Estimating individual discount rates in Denmark: A field experiment. *American Economic Review*, 1606–1617.

Haushofer, J. (2014). The Cost of Keeping Track. *Working paper*.

Haushofer, J. and E. Fehr (2014). On the psychology of poverty. *Science 344*(6186), 862–867.

Hirshleifer, S., D. McKenzie, R. Almeida, and C. Ridao-Cano (2015). The impact of vocational training for the unemployed: experimental evidence from Turkey. *The Economic Journal*.

Janssens, W., B. Kramer, L. Swart, et al. (2017). Be patient when measuring Hyperbolic Discounting: Stationarity, Time Consistency and Time Invariance in a Field Experiment. *Journal of Development Economics*.

John, A. (2017). When Commitment Fails – Evidence from a Field Experiment. *Working paper*.

Mani, A., S. Mullainathan, E. Shafir, and J. Zhao (2013). Poverty impedes cognitive function. *Science 341*(6149), 976–980.

Martinez, S.-K., S. Meier, and C. Sprenger (2017). Procrastination in the field: Evidence from tax filing. *UC San Diego Working Paper*.

McKenzie, D. J. (2017). Identifying and spurring high-growth entrepreneurship: experimental evidence from a business plan competition. *American Economic Review (forthcoming)*.

Noor, J. (2009). Hyperbolic discounting and the standard model: Eliciting discount functions. *Journal of Economic Theory 144*(5), 2077–2083.

O'Donoghue, T. and M. Rabin (2015). Present bias: Lessons learned and to be learned. *The American Economic Review 105*(5), 273–279.

Suri, T. (2011). Selection and comparative advantage in technology adoption. *Econometrica 79*(1), 159–209.

Tanaka, T., C. F. Camerer, and Q. Nguyen (2010). Risk and time preferences: linking experimental and household survey data from Vietnam. *The American Economic Review 100*(1), 557–571.

Toussaert, S. (2015). Eliciting temptation and self-control through menu choices: a lab experiment. *Working Paper*.

Ubfal, D. (2016). How general are time preferences? Eliciting good-specific discount rates. *Journal of Development Economics 118*, 150–170.

A Appendices

A.1 Tables

Table A.1: Descriptive statistics by survey timing

	Mean	Pre-harvest Mean	Post-harvest Mean	Mean Diff.	Total N	Pre-harvest N	Post-harvest N
Liquidity							
HH income (100,000 PKR)	2.70	2.64	2.77	-0.12	523	285	238
Savings (100,000 PKR)	0.44	0.42	0.47	-0.05	525	286	239
Bank account	0.23	0.23	0.23	-0.00	525	286	239
Could borrow	0.98	0.98	0.99	-0.01	525	286	239
Could borrow formal	0.97	0.95	0.99	-0.03	525	286	239
Could borrow informal	0.03	0.06	0.00	0.05***	525	286	239
Harvest							
Harvests wheat	0.37	0.40	0.33	0.07	525	286	239
Demographics							
Muslim	0.88	0.85	0.90	-0.05	525	286	239
Education (years)	2.24	2.55	1.87	0.68**	525	286	239
Housewife	0.75	0.72	0.78	-0.06*	525	286	239
Age	37.85	37.48	38.30	-0.83	525	286	239
Married	0.88	0.87	0.90	-0.03	525	286	239
HH size	6.30	6.29	6.31	-0.02	524	286	238
HH head	0.06	0.07	0.05	0.02	525	286	239
HH decisions (index 0-1)	0.11	0.13	0.08	0.04*	525	286	239
Trust							
Trust NRSP (1-5)	4.22	4.16	4.29	-0.12***	525	286	239
Trust self (1-5)	4.42	4.39	4.46	-0.08***	525	286	239

Notes: All variables are taken from the baseline survey, conducted at the start of the day one session. Treatment status — receiving the baseline interview before or after the onset of the harvest — is computer-randomised prior to the start of the entire survey. Mean diff. represents difference in means across the two treatment arms. *, ** and *** indicate significance of this difference at the 10%, 5% and 1% levels respectively, as estimated from a regression of the variable of interest on the treatment indicator, with standard errors clustered at the village level. For time-varying controls, this is not a test of balance, but rather shows how these baseline covariates differ when the survey is conducted before or after the harvest. HH decisions, Trust NRSP and Trust self are added to the controls (the others imbalanced covariates are already included in the vector of controls).

Table A.2: Treatment effects of windfall timing – heterogeneity by existing liquidity constraints

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	"Present-biased"	**"Future-biased"**	**"Present-biased"**	**"Future-biased"**	**"Present-biased"**	**"Future-biased"**	**"Present-biased"**	**"Future-biased"**
	Day one	**Day one**	**Day one**	**Day one**	**Day one**	**Day one**	**Day one**	**Day one**
	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)	Mfx / (s.e.)
Participation fee timing								
Pay on day 15	0.071	-0.050	0.074*	-0.089**	0.064	-0.045	0.146	-0.114
	(0.046)	(0.051)	(0.041)	(0.044)	(0.045)	(0.053)	(0.089)	(0.102)
Liquidity								
Savings (100,000 PKR)	0.015	-0.017	0.015	-0.048	0.015	-0.054	0.016	-0.053
	(0.045)	(0.044)	(0.030)	(0.032)	(0.030)	(0.034)	(0.030)	(0.034)
Could not borrow	0.343**	-0.282	0.389**	-3.190***	0.341**	-0.276	0.344**	-0.285
	(0.160)	(0.311)	(0.170)	(0.237)	(0.160)	(0.304)	(0.164)	(0.313)
Bank account	-0.098*	-0.009	-0.098*	-0.007	-0.117	0.063	-0.099*	-0.004
	(0.056)	(0.061)	(0.057)	(0.058)	(0.084)	(0.078)	(0.057)	(0.061)
HH income (100,000 PKR)	-0.016	0.044**	-0.017	0.043**	-0.016	0.043**	-0.001	0.037
	(0.016)	(0.018)	(0.017)	(0.017)	(0.016)	(0.018)	(0.022)	(0.024)
Heterogeneity								
Pay on day fifteen*Savings	-0.000	-0.072						
	(0.054)	(0.062)						
Pay on day fifteen*Could not borrow			-0.147	3.392***				
			(0.332)	(0.458)				
Pay on day fifteen*Bank account					0.035	-0.155		
					(0.109)	(0.118)		
Pay on day fifteen*HH income							-0.028	0.012
							(0.030)	(0.034)
Controls	✓	✓	✓	✓	✓	✓	✓	✓
Village f.e.'s	✓	✓	✓	✓	✓	✓	✓	✓
Observations	493	504	493	504	493	504	493	504
Control mean	0.231	0.388	0.231	0.388	0.231	0.388	0.231	0.388

Notes: Standard errors in parentheses. *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively. All dependent variables are measured after treatment status is revealed. "Present-biased" ("future-biased") is a dummy indicating a near-frame switch-point greater than (less than) the far-frame switch-point in the multiple price list activity. Subjects are classified as either "present-biased", "future-biased" or "time-consistent" (not shown). Measures of liquidity constraints are taken from the baseline survey. "HH income" and "savings" are denominated in Pakistani Rupees (100 PKR = 1 USD) winsorized at the 95% level. "Could not borrow" is a dummy equal to one if the subject could not borrow from a formal or informal source in the next two months if they needed to. "Bank account" is a dummy variable equal to one if the household has a bank account. Reported effects represent the marginal effects at the mean. N=525. Sample sizes differ where liquidity variables are missing or village fixed effects perfectly predict the outcome variable.

Table A.3: Treatment effects of windfall timing – trust controls

	(1) **"Present-biased"** **Day one** **All**	(2) **"Future-biased"** **Day one** **All**	(3) **Near-frame switch-point** **Day one** **All**	(4) **Far-frame switch-point** **Day one** **All**
	Mfx / (s.e.)	Mfx / (s.e.)	OLS	OLS
Participation fee timing				
Pay on day 15	0.070*	-0.078*	12.810	3.444
	(0.041)	(0.046)	(20.069)	(19.483)
Trust				
Trust NRSP (1-5)	0.011	0.022	-5.811	-8.068
	(0.037)	(0.039)	(16.506)	(17.946)
Trust self (1-5)	-0.022	0.030	-79.506***	-72.607***
	(0.048)	(0.053)	(23.872)	(22.849)
Controls	✓	✓	✓	✓
Village f.e.'s	✓	✓	✓	✓
Observations	493	504	523	523
Control mean	0.231	0.388	666.654	679.154

Notes: Standard errors in parentheses. *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively. All dependent variables are measured after revelation of participation fee timing treatment status. "Present-biased" ("future-biased") is a dummy indicating a near-frame switch-point greater than (less than) the far-frame switch-point in the multiple price list activity. Subjects are classified as either "present-biased", "future-biased" or "time-consistent" (not shown). Switch-points are in Pakistani Rupees (100 PKR $\approx$ 1 USD). "Trust NRSP" and "Trust self" are 1-5 Likert-scale responses ("how strongly do you agree or disagree with the following statement?") to the following statements: "if a female representative of NRSP made an appointment to see me about a different study, they would be unlikely to cancel or change that appointment"; "if I made an appointment to see someone, for example a female representative of NRSP involved in a different study, I would be unlikely to cancel or change that appointment". Pre-specified controls are: household income, savings, and possession of a bank account; ability to borrow in the next two months (formally or informally) if needed; whether the household harvests wheat; household religion; respondent's education; and her occupation (housewife or other). Reported effects in columns (1) and (2) represent the marginal effects at the mean. N=525. Sample sizes in columns (1) and (2) differ where village fixed effects perfectly predict the outcome variable. The coefficient on *trust self* in column (3) has an FDR-adjusted q-value of 0.002 taken across the two proxies of trust (Anderson, 2012). The coefficient on *trust self* in column (4) has an FDR-adjusted q-value 0.002 taken across the two proxies of trust.

Table A.4: Treatment effects of windfall timing – additional controls

	(1) **"Present-biased"** **Day one** Mfx / (s.e.)	(2) **"Future-biased"** **Day one** Mfx / (s.e.)	(3) **"Present-biased"** **Day one** Mfx / (s.e.)	(4) **"Future-biased"** **Day one** Mfx / (s.e.)	(5) **"Present-biased"** **Day one** Mfx / (s.e.)	(6) **"Future-biased"** **Day one** Mfx / (s.e.)
Participation fee timing						
Pay on day 15	0.058	-0.076	0.071*	-0.081*	0.071*	-0.081*
	(0.041)	(0.047)	(0.041)	(0.046)	(0.041)	(0.046)
Cognitive functioning						
Digit span 1 (score 1-7)	0.021	0.003				
	(0.025)	(0.026)				
Maths 1 (score 1-8)	-0.023***	0.030***				
	(0.008)	(0.009)				
Stroop time 1	-0.001	0.001				
	(0.001)	(0.001)				
Risk preferences						
Certainty premium 1 (100 PKR)			-0.005	0.002		
			(0.004)	(0.004)		
Optimism						
Probablity optimism 1					-0.000	0.000
					(0.000)	(0.000)
Controls	✓	✓	✓	✓	✓	✓
Village f.e.'s	✓	✓	✓	✓	✓	✓
Observations	488	499	493	504	493	504
Control mean	0.231	0.388	0.231	0.388	0.231	0.388

Notes: Standard errors in parentheses. *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively. All dependent variables are measured after revelation of participation fee timing treatment status, as are cognitive functioning, risk preferences and optimism. "Present-biased" ("future-biased") is a dummy indicating a near-frame switch-point greater than (less than) the far-frame switch-point in the multiple price list activity. Subjects are classified as either "present-biased", "future-biased" or "time-consistent" (not shown). "Certainty premium" is a PKR value aggregated across five sets of certainty-equivalent questions which involved different probabilities. "Optimism" is the difference between a subject's subjective belief of her own probability of winning a draw and the objective, given probability of winning that draw, aggregated across five draws with different probabilities. Pre-specified controls are: household income, savings, and possession of a bank account; ability to borrow in the next two months (formally or informally) if needed; whether the household harvests wheat; household religion; respondent's education; and her occupation (housewife or other). Reported effects represent the marginal effects at the mean. N=525. Sample sizes differ where village fixed effects perfectly predict the outcome variable.

A.2 Figures

Figure A.1: Income in the last two & next two weeks, by baseline survey date

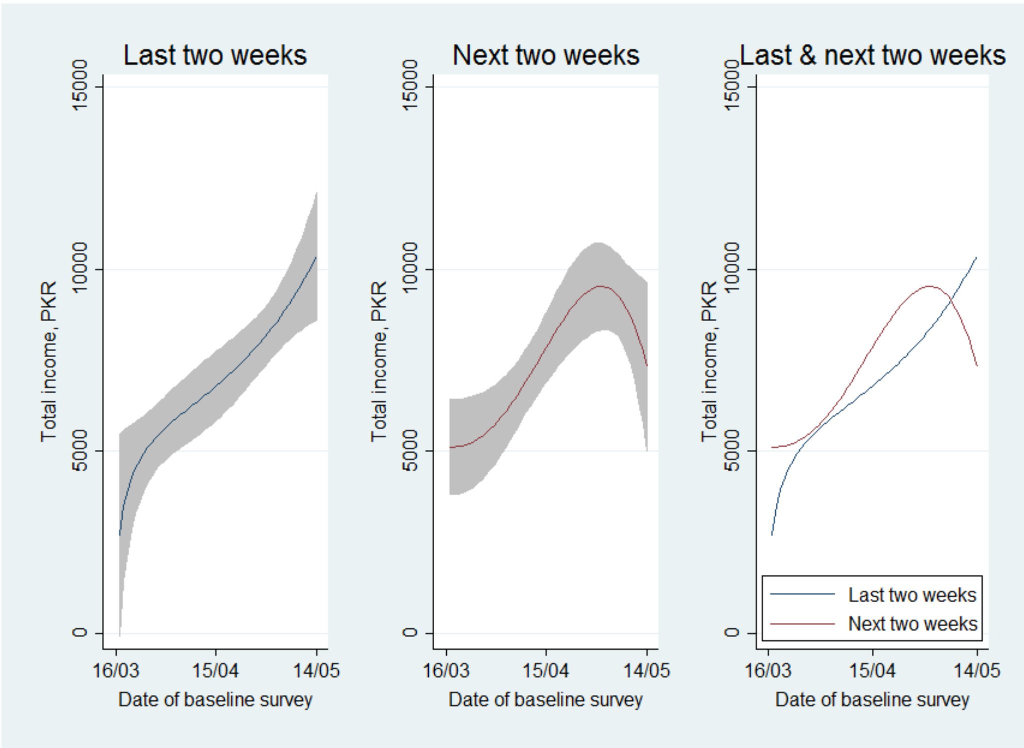

Notes: Plots display the predicted value and 95% confidence intervals for a fractional polynomial fit of the reported income values, after winsorizing at the 95th percentile. Confidence intervals are suppressed in the third panel to allow visibility of both lines simultaneously. Values are as measured on day one during the baseline survey. Respondents are first asked to list all household income sources, then to report income and expected income values separately by source. The survey tablet calculates the totals, which the enumerator verifies with the respondent.

Figure A.2: Day one switch-points by windfall timing

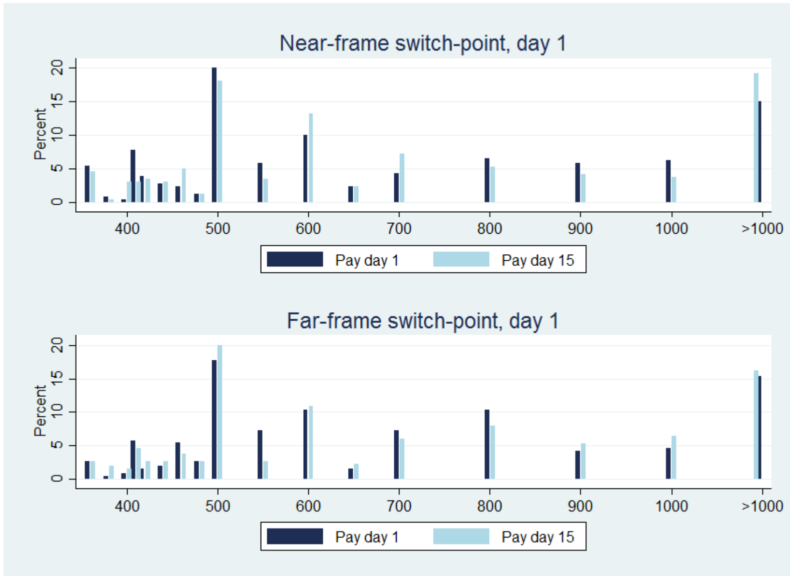

Figure A.3: Day one switch-points by survey timing

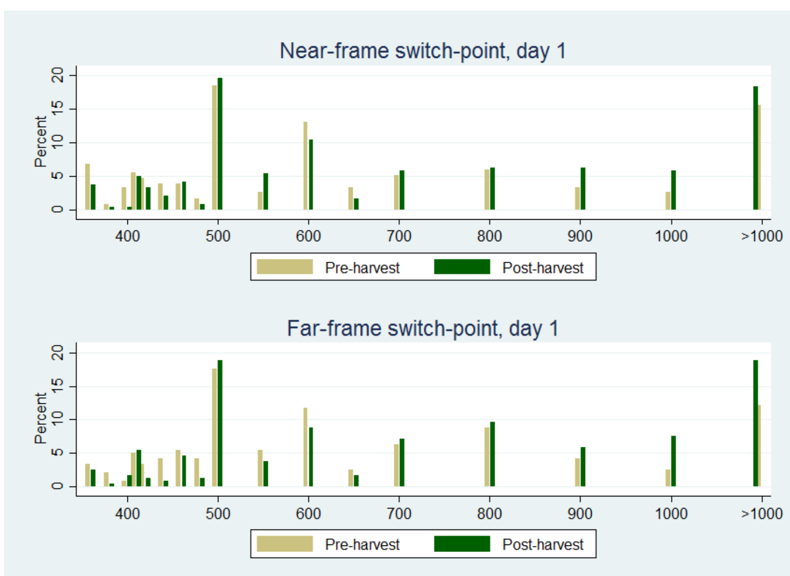

Notes: All variables are measured on day one. Switch-points are in Pakistani Rupees (100 PKR $\approx$ 1 USD). In each frame, subjects were asked to choose between 400 PKR on the earlier date or the amount shown on the later date. The switch-point in each frame is equal to the first value at which the subject chose to receive the payment on the later date. In the near frame, the earlier date was today and the later date was two weeks from today. In the far frame, the earlier date was two weeks from today and the later date was four weeks from today.