TI 2015-097/V
Tinbergen Institute Discussion Paper

Be patient when measuring Hyperbolic Discounting: Stationarity, Time Consistency and Time Invariance in a Field Experiment

Wendy Janssens
Berber Kramer
Lisette Swart

Faculty of Economics and Business Administration, VU University Amsterdam, and Tinbergen Institute, the Netherlands.

Be Patient When Measuring Hyperbolic Discounting: Stationarity, Time Consistency and Time Invariance in a Field Experiment

Wendy Janssens* Berber Kramer† Lisette Swart*

April 14, 2016

Abstract

Most experimental evidence of hyperbolic discounting is based on violations of either stationarity or time consistency. Stationarity is violated when intertemporal choices differ for trade-offs in the near versus the more distant future. Time consistency on the other hand is violated if the optimal allocation for specific dates changes over time. Both types of choice reversals may however also result from time-varying discount rates. Hyperbolic discounting is an unambiguous explanation for choice reversals only if individuals simultaneously violate both stationarity and time consistency. Our field experiment examines the extent to which this is the case. At different points in time, the same participants allocated a future gift over sooner-smaller and later-larger rewards with varying front-end delays. We find that most violations of time consistency do not coincide with violations of stationarity. Random noise in decision-making alone does not explain this finding. Instead, we find a significant association between individual violations and changes in household wealth, in particular for participants with less access to credit. We conclude that in a context of liquidity constraints, eliciting violations of either stationarity or time consistency alone is insufficient to identify hyperbolic discounting.

Keywords: Time preferences, hyperbolic discounting, temporal stability, liquidity constraints

JEL Codes: C93; D03; D14; D90; G02.

---

*Tinbergen Institute and Department of Economics, VU University Amsterdam, w.janssens@vu.nl and l.swart@vu.nl

†Markets, Trade and Institutions Division, International Food Policy Research Institute (IFPRI), b.kramer@cgiar.org. An earlier version circulated under the title 'Time Inconsistent Behavior Under Incomplete Markets: Results from a Fields Experiment in Nigeria.' We gratefully acknowledge funding from the PharmAccess Foundation and VU University Amsterdam (VU-IRPP Grant). We thank the Nigerian team for excellent data collection. Chris Elbers, Jan Willem Gunning, Glenn Harrison, Shachar Kariv, Charles Sprenger and two anonymous referees provided valuable comments, as well as participants in a number of seminars and conferences.

1 Introduction

People often revise their initial choices regarding decisions to borrow, invest or perform a tedious task. For example, someone may prefer to invest towards increased future consumption when asked far in advance, but change her mind right before investing the money, and opt for sooner but lower consumption. Such dynamic choice reversals are called violations of time consistency. Another typical choice reversal concerns a violation of stationarity, for example when someone prefers $110 in 31 days over $100 in 30 days, but rather has $100 today instead of $110 tomorrow (Green, Fristoe and Myerson, 1994; Kirby and Herrnstein, 1995). Both violations are consistent with hyperbolic time preferences, meaning that implicit discount rates are lower for tradeoffs in the more distant future than for tradeoffs in the near future (Frederick, Loewenstein and O'Donoghue, 2002). Stationarity violations are typically observed in (cross-sectional) static choice experiments. As such, they are the easiest and hence most common way to test for hyperbolic discounting.

However, a second yet often neglected explanation for violations of either time consistency or stationarity is a violation of time invariance, which means that subjective discount rates change over time (Halevy, 2015). For example, one month ago someone preferred $110 a day later over $100 the same day, but when asked again today she prefers $100 immediately over $110 tomorrow. This difference may – among others – be due to noise and uncertainty, to changes in underlying preferences, or to changes in incentives and the economic environment (Read, Frederick and Airoldi, 2012). [1]

Crucially, hyperbolic discounting is an unambiguous explanation for such choice reversals only if the same person violates both stationarity and time consistency. When stationarity or time consistency are measured in isolation, one may wrongly ascribe choice reversals that are in fact caused by

[1] Thus, violations of stationarity and time consistency do not necessarily imply any form of irrationality. Violations of stationarity may also result from distrust in experimenters sending future payments and from uncertainty around future preferences and states. Experiments use small front-end delays to minimize the influence of these confounding factors (Harrison, Lau and Rutström, 2005)

time-varying preferences to hyperbolic time preferences.

This paper therefore analyzes to what extent *stationarity* and *time consistency* overlap by means of a field experiment in rural Nigeria. The experiment elicited three convex time budget allocations (Andreoni and Sprenger, 2012*a*) using a longitudinal design adapted from Giné et al. (2016). Participants distributed a future gift over a sooner-smaller and a later-larger reward. Sooner and later rewards arrived 'tomorrow' and 'in one month' for the first allocation, 'in two months' and 'in three months' for the second allocation made on the same day, and 'tomorrow' and 'in one month' for the third allocation. The third allocation was made two months later and hence concerned the same calendar dates as the second allocation, but the time until the two payment dates was the same as in the first allocation.

The experiment elicited each of these three allocations for 240 participants. Building on Halevy (2015), rejecting stationarity requires different choices in the first and second allocation, elicited on the same day with varying front-end delays. Time consistency is rejected when the second and third allocation, elicited at different points in time regarding the same calendar dates, differ. Finally, time invariance is violated when a participant chooses differently in the first and third allocation, elicited on different days but both framed as an allocation over 'tomorrow' and 'in one month'. In theory, if time invariance is satisfied, a hyperbolic discounter will violate both stationarity and time consistency and observing either non-stationarity or time inconsistency is sufficient to infer hyperbolic discounting *only* in that case.

We find that violations of time consistency and stationarity often do not overlap. While 43.4 percent of participants violates time consistency, only 24.2 percent violates *both* time consistency and stationarity. Moreover, nearly half of this subsample violates time consistency and stationarity in different directions with one present- and one future-biased violation. We show that this is not just noise in decision making; violating time consistency but not stationarity is correlated with reductions in wealth from the first to the second decision moment, in particular for individuals with less access to informal credit. This suggests that violations of time invariance are in part due to liquidity constraints. We conclude that identifying hyperbolic discounters cannot be done on the basis of cross-sectional static choice experiments alone. Instead, this requires a longitudinal design eliciting both stationarity and time consistency.

This paper makes three unique contributions to the literature. First, the experimental design links violations of time consistency and stationarity to violations of time invariance for individuals in a context with limited access to financial markets. In a context with sound financial instruments, allocations involving monetary rewards are potentially influenced by the interest rate at which participants can save and borrow outside the experiment (Chabris, Laibson and Schuldt, 2008). When access to financial markets is restricted, as is the case for our subject pool, changes in consumption are likely to follow small changes in income very closely (Halevy, 2014), so that intertemporal allocations of monetary rewards are more closely related to the true underlying time preferences.

To our best knowledge, Halevy (2015) is the only existing choice experiment that also links violations of time consistency and stationarity to violations of time invariance, but using a different subject pool (undergraduate students in economics), who may have better options to save and borrow outside the experiment. Giné et al. (2016) carried out an experiment similar to ours in a context with missing financial markets, and Augenblick, Niederle and Sprenger (2015) elicit time preferences using effort rather than monetary rewards, but both do not explicitly test to what extent violations are related to violations of time invariance. Other experimental studies either analyze violations of stationarity (e.g. Coller and Williams, 1999; Harrison, Lau and Williams, 2002) or of time consistency (e.g. Sayman and Öncüler, 2009; Read, Frederick and Airoldi, 2012), without linking the two.

Second, we analyze why stationarity and time consistency do not always overlap, focusing on the role of potential liquidity constraints. Unlike Halevy (2015), we can identify liquidity-constrained participants as those with less access to credit and larger reductions in wealth. Compared to other participants, they are significantly more likely to violate time consistency without also violating stationarity. This suggests that limited overlap between time consistency and stationarity is not driven solely by random noise, but also by liquidity constraints. Our approach differs from Giné et al. (2016) in that we distinguish between participants with more versus less access to credit, and that we include wealth changes independent of whether the household explicitly reports a wealth shock. We thereby capture both anticipated and unanticipated gains and losses.

Third, discrepancies between stationarity and time consistency imply that the patience level of participants changes over time. As such, our longitudinal design relates to the literature on the temporal stability of time preferences. Identifying temporal stability (or time invariance) requires a longitudinal design in which the experimental methodology and the subject pool are fixed (Frederick, Loewenstein and O'Donoghue, 2002). The main incentivized field experiment with such a design, Meier and Sprenger (2015), finds that any observed temporal instability can be explained by random noise. By contrast, Krupka and Stephens (2013) use a panel with hypothetical choices collected during a period of high inflation and find that elicited discount rates are correlated to economic factors such as the inflation rate and household income, suggesting that temporal instability of expressed time preferences is not purely random. This is more consistent with our findings.

This paper is structured as follows. The next section outlines a conceptual framework to interpret the relation between stationarity, time consistency, and time invariance. Section 3 describes the experiment. Section 4 presents our results and discusses potential behavioral mechanisms. Section 5 concludes.

2 Conceptual Framework

To show why violations of stationarity as measured in most (cross-sectional) time preference experiments do not necessarily overlap with time inconsistent behavior, this section first outlines the types of intertemporal allocations considered in the experiment. We then describe how one can infer violations of stationarity, time consistency and time invariance from these allocations, and discuss conditions under which one can identify hyperbolic discounting. Finally, we formulate hypotheses on how liquidity constraints resulting from changes in background wealth may lead to non-overlapping violations of stationarity and time consistency.

Consider a consumer allocating a gift of g vouchers over two future payment dates. She allocates x vouchers to a later date, denoted p_L, and the remaining $g-x$ vouchers to a sooner date, p_S. Each voucher allocated to the later date is worth v_L. Vouchers allocated to the sooner date are worth v_S and are never worth more than vouchers allocated to the later date, $v_S \leq v_L$.

Allocations are made at the start of two distinct rounds, at decision moments τ_1 and τ_2. The consumer allocates her vouchers between a sooner and later payment date in the first round, $\{p_{1S}, p_{1L}\}$, and between a sooner and later payment date in the second round, $\{p_{2S}, p_{2L}\}$. In both rounds, the sooner payment date immediately follows the decision moment associated with that round, τ_1 and τ_2, respectively. This yields the three intertemporal allocations $x_{1,1}$, $x_{1,2}$ and $x_{2,2}$ depicted in Table 1.

The first allocation, $x_{1,1}$, is made at the start of the *first* round (at $t = \tau_1$) regarding the payment dates during the *first* round, $\{p_{1S}, p_{1L}\}$, which are both in the near future. The second allocation, $x_{1,2}$, also concerns a choice made at the *first* decision moment, but concerns the payment dates during the *second* round, $\{p_{2S}, p_{2L}\}$, which are in the distant future. The third allocation, $x_{2,2}$, is made at the start of the *second* round and concerns the payment dates during the *second* round, $\{p_{2S}, p_{2L}\}$. This allocation hence concerns the same payment dates as the second choice, but these payment dates are again in the near future, as in the first choice.

Table 2 illustrates how these three allocations combined elicit violations of time consistency, stationarity, and time invariance, comparable to Halevy (2015). Stationarity is violated when otherwise similar intertemporal choices (with respect to the delay between p_S and p_L) depend on the front-end delay, i.e. the amount of time between the decision moment and the sooner payment date. In our experiment the delay between payment dates is the same across allocations. We thus observe a violation of stationarity when the two first-round decisions $x_{1,1}$ and $x_{1,2}$ are not identical, i.e. $x_{1,1} \neq x_{1,2}$.

Table 1: Three Types of Intertemporal Choices

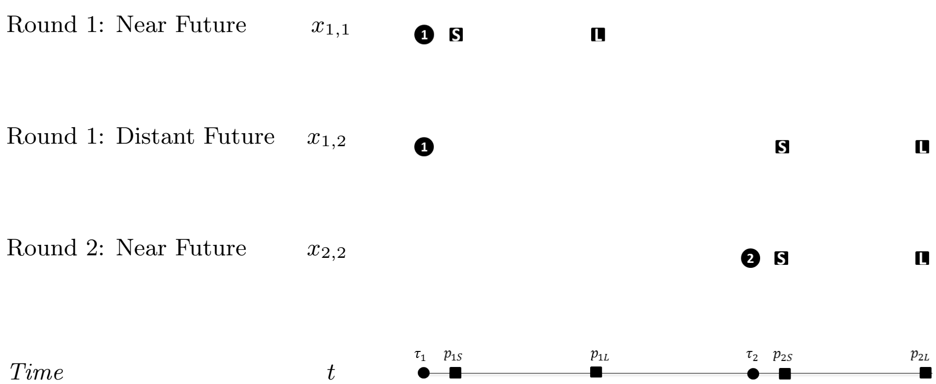

Circles represent two different decision moments, $t = \tau_1$ for a first round and $t = \tau = 2$ for a second round. During these decision moments, people allocate g vouchers to a sooner and a later payment date. Squares represent the payment dates at which consumers can choose to receive the future gift. The sooner date is labeled 'S', and the later date is labeled 'L'. These payment dates are either in the period following the first round (for the first choice, $x_{1,1}$) or in the period following the second round (for the second and third choice, $x_{1,2}$ and $x_{2,2}$). The first and third choice concern payout dates in the near future. For the second choice, made in the first round regarding the payment dates in the second round, payout dates are in the distant future.

Table 2: Defining Three Types of Violations

Violation of stationarity	$x_{1,1}$	$\neq$	$x_{1,2}$
Violation of time invariance	$x_{1,1}$	$\neq$	$x_{2,2}$
Violation of time consistency	$x_{1,2}$	$\neq$	$x_{2,2}$

$x_{i,j}$ represents the number of vouchers (out of a maximum of ten) that a participants allocates to the later payment date (rather than the sooner) at decision moment i for payment payment dates j.

A violation of time invariance implies that the timing of the *decision moment* influences the intertemporal choice when the front-end delay remains the same. In other words, in an otherwise similar choice, a person becomes more or less patient depending on when she takes the decision. This can result from random noise in decision-making, changes in wealth or changes in the underlying structural time preferences. The experiment therefore tests whether the first-round allocation over first-round payment dates, $x_{1,1}$, differs from the second-round allocation over second-round payment dates, $x_{1,1} \neq x_{2,2}$.

Time consistency is violated when a person's allocation between two payment dates at fixed points in time is affected by the time span between the decision moment and the two payment dates. In our experiment, we observe a violation of time consistency when first-round allocations regarding the second-round payment dates, $x_{1,2}$, are not the same as second-round allocations regarding the same payment dates, $x_{1,2} \neq x_{2,2}$.

These three violations are closely linked. Halevy (2015) proves that if one of them occurs, we must observe at least one other violation. An individual's allocations $\{x_{1,1}, x_{1,2}, x_{2,2}\}$ can hence be categorized into one of five collectively exhaustive groups:

1. $x_{1,1} = x_{1,2} = x_{2,2}$. In this group, choices are identical regardless of front-end delay and decision moment, thereby satisfying time consistency, stationarity and time invariance.

2. $x_{1,1} \neq x_{1,2} = x_{2,2}$. In this group, allocations for second-round payment dates do not depend on the decision moment, thereby satisfying time consistency, $x_{1,2} = x_{2,2}$. However, these two allocations differ from the first-round allocation regarding first-round payment dates, $x_{1,1}$, violating stationarity and time invariance.

3. $x_{1,1} = x_{1,2} \neq x_{2,2}$. In the first round, this group makes identical decisions independent of the timing of payment dates, thereby satisfying stationarity, $x_{1,1} = x_{1,2}$. However, in the second round, this group chooses a different allocation, violating time consistency and time invariance.

4. $x_{1,1} = x_{2,2} \neq x_{1,2}$. In this group, allocations regarding near-future payment dates do not depend on when the decision is made, thereby satisfying time invariance, $x_{1,1} = x_{2,2}$. This group however chooses a different allocation regarding distant-future payment dates, violating stationarity and time consistency.

5. $x_{1,1} \neq x_{1,2} \neq x_{2,2}$. In this group, individuals choose different allocations in each type of choice, thereby violating time consistency, stationarity, and time invariance.

Thus, as long as time invariance is satisfied (Groups 1 and 4), a violation of stationarity coincides with - and can be interpreted as - a violation of time consistency. However, when time invariance is violated, the two do not necessarily coincide, and a violation of stationarity cannot be interpreted as a violation of time consistency (Groups 2, 4 and 5).

To illustrate how these concepts relate to hyperbolic discounting, assume a two-period discounted utility framework with time-separable utility and - for tractability - quasi-hyperbolic discounting (also referred to as $\beta\delta$-discounting, Laibson, 1997).[2] We also assume that individuals lack access to financial markets, and cannot transfer background wealth from outside the experiment between the sooner and the later payment date. Later, we will discuss how relaxing this assumption alters the predictions. The three voucher allocations optimize the following three target functions:[3]

$$\max_{x_{1,1}} \quad \mathbb{E}_{\tau_1}\left[u\left(v_S(g - x_{1,1}); \omega_{1S}\right) + \beta\delta u\left(v_L x_{1,1}; \omega_L\right)\right] \tag{1}$$

$$\max_{x_{1,2}} \quad \mathbb{E}_{\tau_1}\left[\beta u\left(v_S(g - x_{1,2}); \omega_{2S}\right) + \beta\delta u\left(v_L x_{1,2}; \omega_L\right)\right] \tag{2}$$

$$\max_{x_{2,2}} \quad \mathbb{E}_{\tau_2}\left[u\left(v_S(g - x_{2,2}); \omega_{2S}\right) + \beta\delta u\left(v_L x_{2,2}; \omega_L\right)\right] \tag{3}$$

[2]Strictly speaking, quasi-hyperbolic discounting distinguishes the present (today) from the future (tomorrow and any later day). Given that our soonest payment takes place the next day, we need to assume that tomorrow will still be considered as the (extended) present by the participants, so that $\beta = 1$ for payments tomorrow. This will be the case when adopting a more general hyperbolic discount function.

[3]Only one of these three allocations is implemented, so that the first-round allocation regarding first-round payment dates does not influence background wealth in allocations regarding second-round payment dates.

where $\mathbb{E}_{\tau_t}$ represents expectations at the time of round t, $u(\cdot)$ instantaneous consumption utility, g the experimental gift, $x_{t,j}$ the number of vouchers allocated to the later date of payment dates j in round t, v_S and v_L the value of vouchers allocated to the sooner and later payment date, respectively, ω_{jS} an individual's background wealth on the sooner payment date in round t, and ω_L background wealth on the later payment date. For simplicity, we hold background wealth on the later payment date constant and focus on ceteris paribus effects of changes in background wealth on the sooner payment date.[4] Further, $0 < \delta < 1$ represents an exponential discount factor for the later relative to the sooner payment date and $0 < \beta \leq 1$ a present-bias parameter by which all instantaneous utilities for future payments are discounted.

First assume that background wealth is stable over time, so that $\mathbb{E}_{\tau_1}\omega_{1S} = \mathbb{E}_{\tau_2}\omega_{2S}$. In that case, the first-round allocation regarding first-round payment dates equals the second-round allocation regarding second-round payment dates, $x_{1,1} = x_{2,2}$, and time invariance is not violated. If $\beta \neq 1$, allocations will violate both stationarity, $x_{1,1} \neq x_{1,2}$, and time consistency, $x_{1,2} \neq x_{2,2}$. Thus, under the assumption of time invariance, a violation of stationarity implies a violation of time consistency and vice versa, and (quasi-) hyperbolic discounters will violate both stationarity and time consistency in a present-biased direction, $x_{1,1} < x_{1,2}$ and $x_{1,2} < x_{2,2}$, because for them, $\beta < 1$.

Now assume that background wealth at the sooner payment date changes over time, $\mathbb{E}_{\tau_1}\omega_{1S} \neq \mathbb{E}_{\tau_2}\omega_{2S}$. If participants cannot save or borrow, they cannot smooth background wealth over time. As a result, allocations regarding same-round payment dates, $x_{1,1}$ and $x_{2,2}$, will differ, and the participant will violate time invariance. Specifically, if $\mathbb{E}_{\tau_1}\omega_{1S} < \mathbb{E}_{\tau_2}\omega_{2S}$, a participant will allocate fewer vouchers to the later payment date in the first round than in the second round in allocations regarding same-round payment dates. In other words, a participant becomes more patient over time.

[4] ω_{tL} is not known at decision moment τ_t, in contrast to ω_{tS}. As a result, allowing expected future background wealth to vary over decision rounds leads to less clear predictions.

If $\mathbb{E}_{\tau_1}\omega_{1S} > \mathbb{E}_{\tau_2}\omega_{2S}$, participants become less patient over time.

In addition, even in the absence of hyperbolic discounting, $\beta = 1$, a participant will violate either stationarity or time consistency, depending on whether the change in wealth is already anticipated in the first round. If anticipated, such that $\mathbb{E}_{\tau_1}\omega_{1S} \neq \mathbb{E}_{\tau_1}\omega_{2S} = \mathbb{E}_{\tau_2}\omega_{2S}$, stationarity is violated, $x_{1,1} \neq x_{1,2} = x_{2,2}$. In case of reduced wealth (and patience) from the first to the second round, stationarity is violated in a future-biased direction. If unanticipated, $\mathbb{E}_{\tau_1}\omega_{1S} = \mathbb{E}_{\tau_1}\omega_{2S} \neq \mathbb{E}_{\tau_2}\omega_{2S}$, time consistency is violated, $x_{1,1} = x_{1,2} \neq x_{2,2}$, and a reduction in wealth results in a present-biased violation of time consistency. Thus, when time invariance is violated, observing non-stationary or time-inconsistent choices is not necessarily indicative of hyperbolic discounting, but may instead capture consumption smoothing.

An important assumption in the discussion above is that individuals lack access to credit and savings. In perfectly functioning financial markets, participants would not need to violate time invariance in order to smooth consumption. We therefore formulate the following hypotheses:

If individuals lack access to credit and changes in background wealth are anticipated,

Hypothesis 1a *...* decreasing *background wealth ($\mathbb{E}_{\tau_1}\omega_{1S} > \mathbb{E}_{\tau_2}\omega_{2S}$) is associated with future-biased violations of stationarity, but not of time consistency,*

Hypothesis 1b *...* increasing *background wealth ($\mathbb{E}_{\tau_1}\omega_{1S} < \mathbb{E}_{\tau_2}\omega_{2S}$) is associated with present-biased violations of stationarity, but not of time consistency,*

If changes in background wealth are unanticipated, ...

Hypothesis 2a *...* decreasing *background wealth ($\mathbb{E}_{\tau_1}\omega_{1S} > \mathbb{E}_{\tau_2}\omega_{2S}$) is associated with present-biased violations of time consistency, but not of stationarity.*

Hypothesis 2b ... increasing *background wealth* ($\mathbb{E}_{\tau_1}\omega_{1S} < \mathbb{E}_{\tau_2}\omega_{2S}$) *is associated with future-biased violations of time consistency, but not of stationarity.*

In conclusion, we argue that when time invariance is violated, one can only infer hyperbolic discounting from observing both time consistency and stationarity. Halevy (2015) shows that this requires longitudinal designs with allocations at different decision moments for payment dates at different points in time. The majority of existing time preference experiments however elicit only violations of stationarity, using cross-sectional designs with one decision moment regarding different payment dates. Systematic violations of time invariance due to predictable or unpredictable changes in the economic environment may confound the conclusions from these experiments, and this paper sheds light on the severity of the potential misclassification.

3 Experimental methods and procedures

3.1 Design

To test whether violations of time consistency empirically overlap with violations of stationarity, we conducted an artefactual field experiment in rural Nigeria. The experiment elicited participants' intertemporal allocations using Andreoni and Sprenger (2012*a*)'s convex time budget method. Participants received ten vouchers to divide between two future payment dates, with the later date exactly one month after the sooner date. Vouchers allocated to the later payment date were always worth 200 NGN.[5] Vouchers allocated to the sooner payment date were worth either 200, 150, 120 or 100 NGN.

Participants allocated their budgets between the two payment dates in three different incentivized scenarios: (i) a first-round allocation dividing the ten vouchers between payment dates soon after the first round, 'tomorrow'

[5]At the time of the experiment, 100 NGN (Nigerian Naira) was worth approximately 0.62 USD.

and 'one month from now' (yielding choice $x_{1,1}$); (ii) a first-round allocation dividing the vouchers between payment dates in a more distant future, '2 months from now' and '3 months from now' (yielding $x_{1,2}$); and (iii) a second-round allocation conducted two months later for the same payment dates, and hence framed again as 'tomorrow' and 'one month from now' (yielding $x_{2,2}$). Thus, within subjects, we varied (a) the delay between the decision moment and the payment dates; and (b) the timing of the decision moment itself. As such, the experiment elicits measures of stationarity, time invariance and time consistency, as shown in Table 1.

Note that in choices regarding the near future, the earliest payment date was tomorrow. Due to this front-end delay, we are unable to identify pure quasi-hyperbolic discounting, which assumes structurally different discounting of the present versus the future. We opted for a small delay before the first payment for two reasons. First, paying participants the same day was logistically difficult. Second, delaying the payment by one day helped avoid possible confounds such as differential transaction costs between payment dates or trust issues (Chabris, Laibson and Schuldt, 2008). Sozou (1998) showed that the perceived risk of default of the experimenter differs between immediate payments and any future payments, but that the perceived difference in risk between different payment moments in the future is negligible. An increasing number of studies therefore avoids immediate payments and we followed this approach (for additional references and a detailed discussion, see Andreoni and Sprenger, 2012*a*).

3.2 Procedures

Participants were recruited from a sample of farming households in Kwara State, Nigeria, who were interviewed weekly about their health and finances from March 2012 to May 2013.[6] Figure 1 illustrates a timeline of this experiment. In March 2012, a baseline survey collected individual characteristics for all household members. In April 2012, we conducted the first round

[6]This is the Health and Financial Diaries study implemented by the Amsterdam Institute for International Development in collaboration with the PharmAccess Foundation and the University of Ilorin Teaching Hospital (Janssens et al., 2013).

of the experiment. Enumerators visited the households and interviewed all adult household members in private following a script with the experimental instructions (see the Online Appendix).[7] They first elicited choices regarding the second-round payment dates, framed as payments 'in three months' versus 'in two months' from now, followed by a break with survey questions. After this intentional break, which served to reduce potential efforts to appear consistent across choices, enumerators elicited choices regarding the first-round payment dates, framed as payments 'in one month' from now versus 'tomorrow'.[8]

Figure 1: Timeline of the Study

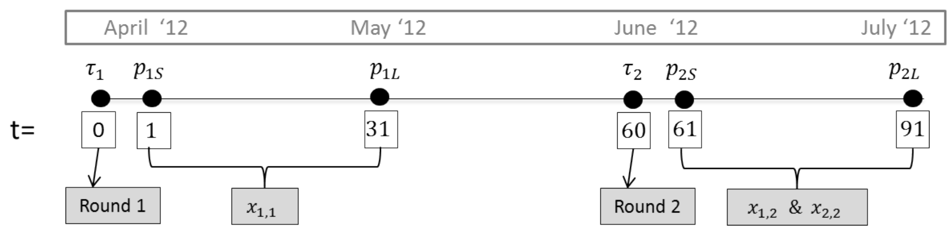

Circles represent two different decision moments, $t = \tau_1$ for a first round at t=0 and $t = \tau = 2$ for a second round sixty days later. During these decision moments, people allocate vouchers to a sooner and a later payment date. Squares represent the payment dates at which consumers can choose to receive the future gift. The sooner date is labeled 'S', and the later date is labeled 'L'. These payment dates are either in the period following the first round (for the first choice, $x_{1,1}$) or in the period following the second round (for the second and third choice, $x_{1,2}$ and $x_{2,2}$). The first and third choice concern payout dates in the near future. For the second choice, made in the first round regarding the payment dates in the second round, payout dates are in the distant future.

Allocations regarding first- and second-round payment dates were both made for the four different values of vouchers allocated to the sooner date. To ensure incentive compatibility, we randomly selected one of these allocations for each participant for actual payout. To retain a large enough sample for the second round, the probability of selecting a choice regarding

[7]We targeted the household head, their spouses, and other adult household members not enrolled in school.

[8]To enhance understanding of the time preference games, enumerators used a wooden board with two bowls representing the sooner and the later payment date, and small vouchers that people had to divide over the two bowls. The order of the questions was not randomized. Order effects are expected to be limited, since Andreoni and Sprenger (2012*a*,*b*) and Giné et al. (2016) do not find any evidence of order effects.

second-round payment dates was 0.9.[9] Participants did not know the exact probabilities. They were told that the computer would randomly select one question and that this would be one of the eight questions they were about to answer.

The ten percent of participants for whom a first-round choice was selected for payment received their payments according to their initial allocation. By contrast, those who were to be paid during the second round were revisited unexpectedly two months later, in June, just before their 'sooner' payment date. They received the opportunity to revise their earlier choice that was selected for payment. The enumerator clearly showed them their initial choice given the selected voucher values for second-round payment dates, $x_{1,2}$, and asked them to indicate their preferred allocation once more. They were paid according to this new allocation rather than the initial choice. Participants were reassured that they could leave their allocation as it was or change it to whatever allocation they preferred.

On payment dates, enumerators returned to every participant with a payout on that day and exchanged vouchers valid on that particular day for cash. The experimental design allowed participants to earn between 1,000 and 2,000 NGN, and they earned 1,862 NGN on average. These stakes are fairly high, as the maximum possible payment of 2,000 Naira is equivalent to approximately three days of work among the employed participant sample. Further, concerns about a lack of participant trust in receiving the experimental pay-outs are limited, as participants were part of a larger ongoing study for which they were being interviewed by the same research team on a regular basis.

3.3 Description of the participant sample

The experiment targeted 303 individuals who participate in the baseline survey in March 2012. Of those, 293 persons (96.7 percent) participated in the first round of our experiment. For 256 participants (87.4 percent of

[9] This probability was less than 1 to ensure incentive compability of choices regarding first-round payment dates.

first-round participants), the experiment selected a choice regarding payment dates following the second round, and among them, 240 (93.8 percent) participated in the second round. For the remaining sixteen first-round participants, we did not observe second-round allocations because a few participants moved away from the study area, and one participant had passed away. His family members were hence mourning and did not participate in the second round either.

Table 3 presents summary statistics for all participants in the experiment. Columns (1) and (2) show the number of observations and the mean for all 293 participants who completed the first round of the experiment. The average age of the participants is just over 40 years of age and around forty percent of participants are male. The majority of participants never entered the formal school system. The two predominant sources of income among participants are farming (36.9 percent) and business (39.6 percent).

Since businesses are often related to farming, participants' financial situation depends heavily on the agricultural season. The experiment was conducted in the period between planting and harvest. At baseline, only seven percent of the farmers expected to harvest before July, when the later payment date of the second round was due. Since farmers incur expenditures to harvest their produce and generally prefer to wait until market prices increase instead of selling their harvest right away, the harvest time is a liquidity-constrained period. Participants may well take this into account in allocations regarding second-round payment dates. Consistent with this idea, average wealth (calculated as the balance of all financial assets and liabilities within a household) reduces from the start of the first to the second round by a sizable 10,000 NGN, which is 17 percent of wealth at baseline, and five times the maximum experimental payout of 2,000 NGN.

Furthermore, only 42.6 percent of participants have relatively easy access to informal credit. The other participants either cannot borrow 20,000 NGN in case of an emergency, or needs to borrow from three or more different people to raise this amount. Given that the vast majority of our sample is unbanked, such a limited ability to borrow from one's informal network suggests that this person has limited access to credit in general. This group

Table 3: Description of the participant sample

	(A) All participants		(B) Revisited participants		
	N	Mean	N	Mean	Diff. in means
	(1)	(2)	(3)	(4)	(5)
Age	293	40.31	240	40.17	-0.143
Male	293	0.396	240	0.379	-0.017
No formal education	292	0.589	240	0.596	0.007
Main income from farming	293	0.369	240	0.367	-0.002
Main income from business	293	0.396	240	0.417	0.021
Main income from other	293	0.106	240	0.096	-0.01
No main source of income	293	0.130	240	0.121	-0.009
Planning to harvest before July	293	0.066	240	0.071	-0.005
Financial wealth in round 1	293	61,262	240	71,153	9,891***
Financial wealth in round 2	277	50,841	240	50,020	-821
More access to informal credit	291	0.426	238	0.445	-0.019
Satisfies monotonicity	293	0.904	240	0.912	0.008

Financial wealth is calculated as the balance of all financial assets and liabilities within a household (the sum of current bank account balances, formal and informal savings, loans and credits receivable, subtracted by outstanding credits and loans). An individual is classified as having less access to credit if they cannot borrow 20,000 NGN in case of an emergency, or needs to borrow from three or more different people to raise this amount. We do not present standard deviations because all variables, apart from four (age, the number of children and two financial balances) are binary indicators. Reported p-values are based on t-tests with standard errors (shown in parentheses) clustered by household. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

will be liquidity constrained if experiencing a significant reduction in wealth.

A large body of literature discusses the possible effects of limited understanding on conclusions drawn from time preference experiments (see for example Andreoni and Sprenger, 2012*a*). If a participant does not fully understand the task or its implications, her decisions will not accurately represent her underlying time preferences. Enumerators devoted a significant amount of time to explain the convex time budget task. To test whether poor understanding can nevertheless have introduced noise in the allocations, leading to violations of stationarity, time consistency or time invariance, we test a simple monotonicity condition. When the return on waiting increases, participants should never allocate *fewer* vouchers to the later payment date.

To test whether participants satisfied this monotonicity concept, we compare allocations when sooner vouchers are worth (1) 100 NGN vs. 120 NGN, (2) 120 NGN vs. 150 NGN, and (3) 150 NGN vs. 200 NGN; for both near-future ($x_{1,1}$) and distant-future ($x_{1,2}$) allocations. Using these six comparison pairs, 219 of the 240 participants in the final sample (91.3 percent) never violate monotonicity. Further, of the 1,440 pairs (6 pairs times 240 participants), 1410 pairs (97.9 percent) satisfy monotonicity, suggesting similar levels of understanding as university students participating in Andreoni and Sprenger (2012*a*), and better understanding than more comparable participants in Giné et al. (2016).

Following Chakraborty et al. (2015), Appendix Table B1 tests for demand monotonicity by comparing the number of interior versus corner allocations in the entire data set. The percentage of choice sets violating demand monotonicity never exceeds 11.5 percent and does not increase in the number of interior choices in a choice set. Thus, demand monotonicity violations are not a major concern in our data. Chakraborty et al. (2015) perform three additional tests to analyze the internal and external consistency of data from convex time budgets: they test for the weak axiom of revealed preferences, wealth monotonicity, and impatience monotonicity. We do not have experimental variation to perform these three tests.

This paper restricts the sample in the main analysis to participants for whom all three choices depicted in Table 1 were elicited. For the 16 dropouts and the randomly selected participants who were not revisited for the time preference game in the second round, we cannot observe violations of time invariance or time consistency. Columns (3) to (5) compare the 240 participants who were revisited during the second round with the full sample and confirm that attrition from the first to the second round was not driven by observable characteristics. The only variables that differ significantly between the full sample and the revisited sample are household size and financial wealth at baseline. Wealth in the full sample did not reduce as much as in the revisited subsample, in part because most non-revisited participants received their experimental pay-out in the first round.

4 Results

This section describes the experimental results, starting with a description of how participants allocate their future gift over time. Next, we exploit our within-subject design to identify how frequently violations of time consistency overlap with violations of stationarity. Finally, this section analyzes whether violations of time invariance that account for this discrepancy can be explained by random noise or by other, more systematic, factors including changes in participants' wealth and features of the experimental design. We only include the randomly selected value of vouchers allocated to the sooner payment date, for which all three choices ($x_{1,1}$, $x_{1,2}$ and $x_{2,2}$) are observed.

4.1 Description of choices

Figure 2 plots the cumulative distribution of the number of vouchers allocated to the later payment date and Table 4 presents corresponding summary statistics. In this and subsequent tables, reported p-values are based on a t-test for differences in means with standard errors clustered at the household level. Panel (a) illustrates how the return on waiting, varied between subjects, affects allocations in the experiment. Panel (b) presents the distribution by type of choice, which was varied within subjects.

Panel (a) of Figure 2 illustrates the cumulative distributions when vouchers allocated to the *sooner* payment date ('sooner vouchers') are worth 200 NGN, 150 NGN, 120 NGN or 100 NGN. Since vouchers allocated to the *later* payment date are worth a fixed 200 NGN, the return on waiting decreases in the value of sooner vouchers. Thus, for participants whose sooner vouchers are worth 200 NGN, the return on waiting is the lowest. They allocate most vouchers to the sooner payment date, leaving on average 2.54 vouchers for the later date (see Panel (a) in Table 4). Participants whose sooner vouchers are worth 150 NGN have a higher return on waiting and allocate on average nearly five additional vouchers to the later date ($p < 0.01$). Compared to this subsample, participants with vouchers worth 120 NGN allocate an additional 1.22 vouchers to the later date ($p < 0.01$). Reducing the value of sooner vouchers to 100 NGN does not significantly increase the number of vouchers allocated to the later date any further.[10]

Panel (b) shows allocations by type of choice. The solid line describes first-round choices when the payment dates are in the near future, 'tomorrow versus in one month' ($x_{1,1}$). The dashed line describes first-round choices with payment dates in the distant future, 'in two months versus in three months' ($x_{1,2}$). The dotted line represents the choice elicited during the second round nearly two months later, framed again as an allocation regarding payment dates in the near future ($x_{2,2}$). The distributions of the two first-round choices $x_{1,1}$ and $x_{1,2}$ are very similar; in first-round choices regarding the near future, $x_{1,1}$, participants allocate on average 7.27 (out of ten) vouchers to the later payment date, while they allocate on average 7.30 vouchers to the later payment date in choices regarding the more distant future, $x_{1,2}$ (see Table 4). We fail to reject the hypothesis that $x_{1,1} = x_{1,2}$ ($p = 0.853$). Hence, the aggregate data satisfy stationarity.

[10]Allocations for sooner vouchers worth 200 NGN (with no return on waiting) thus seem to a different pattern than allocations where sooner vouchers are worth less than 200 NGN (yielding a positive return on waiting). We replicated all tables presented in this paper for the sample excluding sooner vouchers worth 200 NGN (available upon request). Omitting participants with sooner vouchers worth 200 NGN results in qualitatively similar patterns to the ones presented in this section.

Figure 2: Distribution of vouchers allocated to later date

(a) By value of vouchers allocated to the sooner date

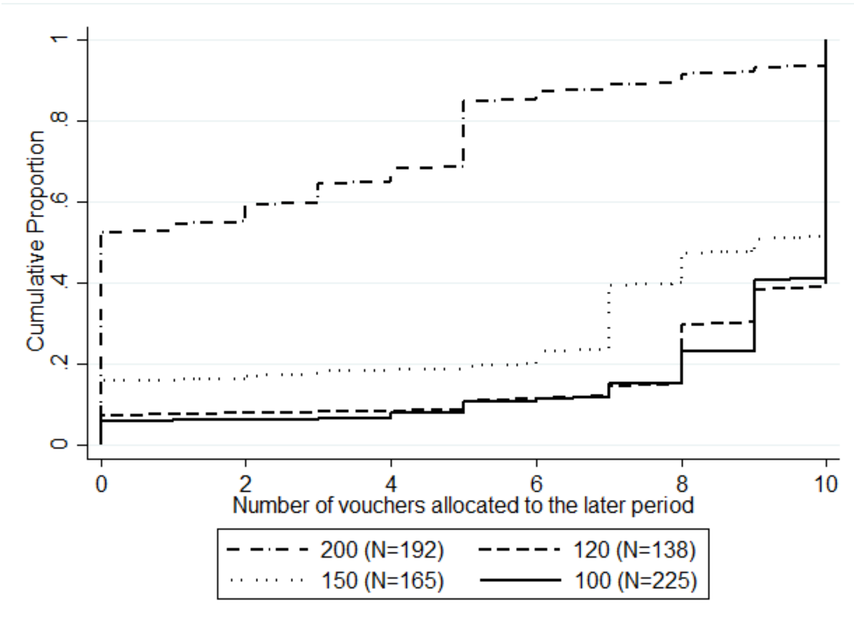

(b) By choice type (incl. vouchers of 200 NGN)

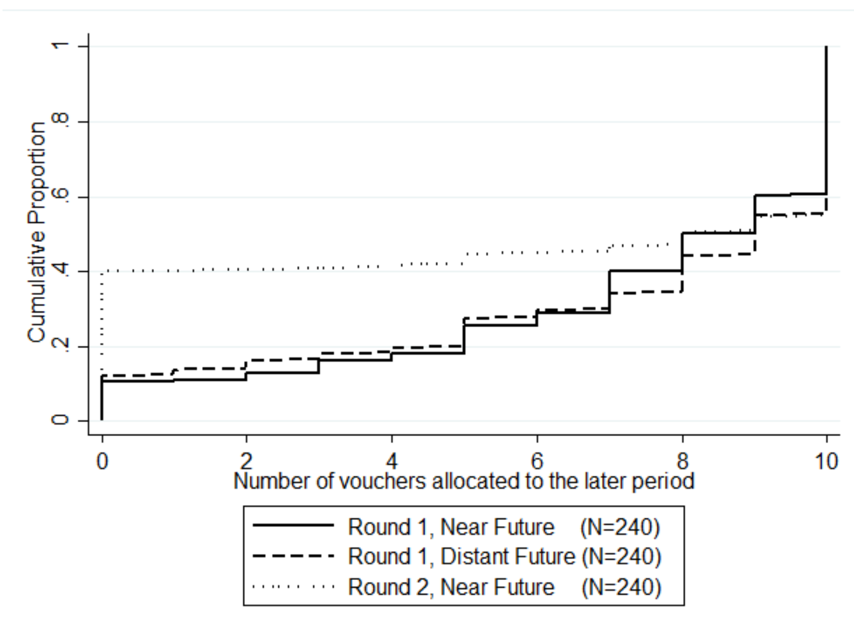

Table 4: Distribution of vouchers allocated to later payment date

	Summary statistics				p-value equal means		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. By sooner voucher value (varied between-subjects)							
	200	150	120	100	200 = 150	150 = 120	120 = 100
Mean	2.54	7.35	8.57	8.66			
Std Dev	3.23	3.59	2.74	2.59	0.000	0.001	0.746
Observations	192	165	138	225			
Panel B. By type of choice (varied within-subjects)							
	All	$x_{1,1}$	$x_{1,2}$	$x_{2,2}$	$x_{1,1} = x_{1,2}$	$x_{1,1} = x_{2,1}$	$x_{1,1} = x_{2,2}$
Mean	6.71	7.27	7.30	5.56			
Std Dev	3.98	3.29	3.53	4.72	0.853	0.000	0.000
Observations	720	240	240	240			
Panel C. Excl. sooner vouchers worth 200 NGN							
	All	$x_{1,1}$	$x_{1,2}$	$x_{2,2}$	$x_{1,1} = x_{1,2}$	$x_{1,1} = x_{2,1}$	$x_{1,1} = x_{2,2}$
Mean	8.23	8.58	8.83	7.27			
Std Dev	3.03	2.18	2.11	4.12	0.101	0.000	0.001
Observations	528	176	176	176			

The value of vouchers allocated to the later payment date is fixed at 200 NGN, while the value of vouchers allocated to the sooner payment date varies from 200 NGN to 100 NGN. $x_{i,j}$ represents the number of vouchers (out of a maximum of ten) that a participants allocates to the later payment date (rather than the sooner) at decision moment i for payment payment dates j. As such, stationarity tests whether the first-round allocation for payment dates in the near future ($x_{1,1}$) is identical to the first-round allocation for payment dates in the more distant future ($x_{1,2}$). Time invariance tests whether allocations for near-future payment dates are identical regardless of whether they were made in the first round ($x_{1,1}$) or in the second round ($x_{2,2}$). Finally, time consistency tests whether the first-round allocation for second-round payment dates ($x_{1,2}$) is identical to the second-round allocation for second-round payment dates ($x_{2,2}$). Reported p-values are based on t-tests with standard errors (shown in parentheses) clustered by household. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

By contrast, second-round choices for payment dates in the near future, $x_{2,2}$, are different from the two first-round allocations, $x_{1,1}$ and $x_{1,2}$. The average number of vouchers allocated to the later payment date in the second-round choice is on average 5.57, which is significantly lower than both choices made in the first round ($p < 0.01$). Because participants tend to revise their first-round allocations for the distant-future payment dates in the second round, we reject the hypothesis that $x_{1,2} = x_{2,2}$ ($p < 0.01$), implying a violation of time consistency in the aggregate. Moreover, given that allocations regarding near payment dates are not constant across the two rounds, we reject the hypothesis that $x_{1,1} = x_{2,2}$ ($p < 0.01$), implying an aggregate violation of time invariance.

Our analyses omit two types of first-round choices that do not have a second-round equivalent: choices for voucher values that were not selected for payment, and choices made by participants who were not revisited during the second round (because they were selected to be paid during the first round, or because the participant dropped out of the study). Also for these non-selected first-round choices and non-revisited participants, we cannot reject stationarity (see Appendix Table B2). This reinforces the conclusion drawn from Table 3 that attrition does not bias our results.

We conclude that the aggregate data violate time consistency and time invariance, but not stationarity. This implies that individual choice patterns will also violate time consistency and time invariance, while stationarity may or may not be satisfied. Thus, time-inconsistent behavior appears to be linked more closely to violations of time invariance than stationarity, and the overlap between violations of time consistency and stationarity appears limited.

4.2 Classification of participants

Stationarity in the aggregate data is necessary but not sufficient for individual-level stationarity. We may observe aggregate stationarity simply because some participants choose present-biased allocations whereas others choose future-biased allocations that cancel each other out on average. Figure 3 therefore indicates the proportion of participants with a violation of time consistency (represented by the grey shaded bar), a violation of stationarity (represented by the bar with horizontal lines), and violations of time invariance (represented by the bar with vertical lines).

Figure 3: Proportions of participants violating the three concepts

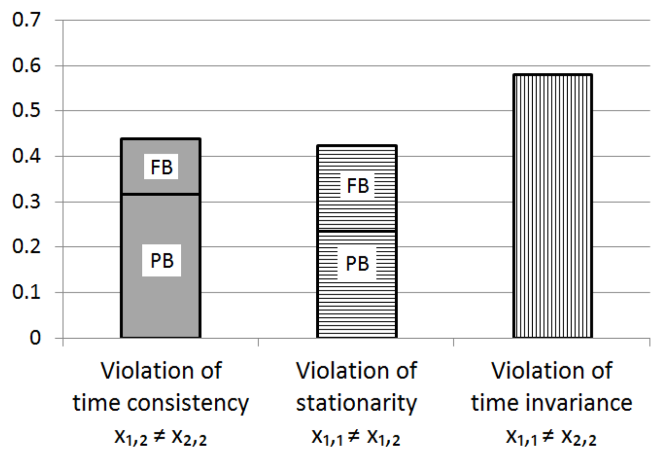

Violations of time consistency (stationarity) are divided into present-biased violations, $x_{1,2} > x_{2,2}$ ($x_{1,2} > x_{1,1}$), noted by PB, and future-biased violations, $x_{1,2} < x_{2,2}$ ($x_{1,2} < x_{1,1}$), noted by FB. Violations of time invariance cannot be classified as either present-biased or future-biased and are hence not divided into these two categories.

The left bar shows that 43.8 percent of the participants violates time consistency ($x_{1,2} \neq x_{2,2}$). Although we do not observe violations of stationarity in the aggregate data, the middle bar indicates that stationarity is violated ($x_{1,1} \neq x_{1,2}$) by an almost similar percentage of participants as time consistency. As long as choices satisfy time invariance, the left and middle bar would represent the same group of participants. However, the right bar indicates that choices regarding near-future payment dates change from the first to the second round for more than half (58.0 percent) of all participants. Their violations of time invariance imply that stationarity and time consistency do not necessarily overlap.

The figure also illustrates why, despite very similar percentages of participants that violate stationarity and time consistency, only time consistency is violated at an aggregate level. The figure divides violations of stationarity and time consistency into present-biased violations (the lower areas marked 'PB') and future-biased violations (the upper areas marked 'FB'). Present-biased violations of stationarity and time consistency occur when a participant allocates more vouchers to the later date in choices regarding *distant-future* payment dates, $x_{1,2} > x_{1,1}$ and $x_{1,2} > x_{2,2}$. Future-biased violations of stationarity and time consistency, on the other hand, occur when a higher number of vouchers is allocated to the later date in choices regarding *near-future* payment dates, $x_{1,2} < x_{1,1}$ and $x_{1,2} < x_{2,2}$. While over 70 percent of time inconsistencies are present-biased, present-biased violations account for only 55 percent of non-stationarities.

Figure 4 divides our participants into one of the five collectively exhaustive groups discussed in Section 2 (the dark bars) and compares our sample with the Halevy (2015) sample (the gray bars). In our sample, nearly half of all violations of time consistency - 19.6 out of 43.8 percent - coincide with a violation of time invariance, without stationarity being violated. Likewise, among the 43.3 percent of participants who violate stationarity, 19.2 percent violates time invariance without violating time consistency. Only 24.2 percent of participants violates both stationarity and time consistency.

An important question is to what extent these findings differ from Halevy (2015)'s findings, the only existing experiment that relates time-inconsistent behavior to violations of stationarity and time invariance. There are large differences in methodology and subject pool between Halevy's study and our own: Halevy (2015) used a multiple price list, while we use a convex time budget method with visual aids, and Halevy's sample consisted of undergraduate students at the University of British Columbia in Canada while the participants in our experiment are adults living in poor households in rural Nigeria with limited access to financial markets.

Interestingly, the percentages of participants belonging to the different groups is remarkably similar between the two studies. In both studies, the majority of participants does not violate stationarity, time consistency and time invariance (Group 1) and very few participants violate both time consistency and stationarity without also violating time invariance (Group 4).

Figure 4: Comparing distribution of participants to distribution in Halevy (2015)

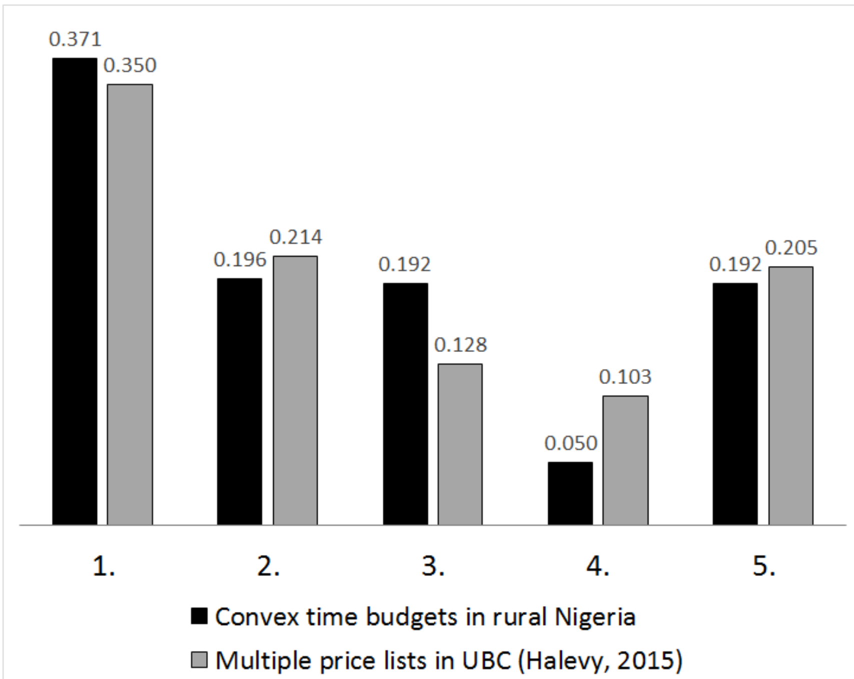

Group 1 consists of individuals satisfying time consistency, stationarity and time invariance ($x_{1,1} = x_{1,2} = x_{2,2}$); Group 2 consists of individuals satisfying time consistency, but violating stationarity and time invariance ($x_{1,1} \neq x_{1,2} = x_{2,2}$); Group 3 consists of individuals satisfying stationarity, but violating time consistency and time invariance ($x_{1,1} = x_{1,2} \neq x_{2,2}$); Group 4 consists of individuals satisfying time invariance, but violating time consistency and stationarity ($x_{1,1} = x_{2,2} \neq x_{1,2}$); and finally Group 5 consists of individuals violating time consistency, stationarity and time invariance ($x_{1,1} \neq x_{1,2} \neq x_{2,2}$).
The percentages listed here from Halevy (2015) are based on Column 1 from Table II on page 345 of Halevy, Yoram, 2015. "Time Consistency: Stationarity and Time Invariance." *Econometrica,* 83(1): 335 - 352.

The share of participants in the other three groups is also very similar. Thus, despite differences in design and subject pool between the two experiments, violations of time consistency, stationarity and time invariance arise in very similar ways.

Table 5 describes in more detail how well stationarity and time consistency overlap in our experiment. The first column of Panel A defines violations the same way as Table 2 does: any difference between two allocations results in a violation of stationarity, time consistency or time invariance. The correlation between violations of stationarity and time consistency is 0.212, which is substantially lower than the correlations between violations of stationarity and time invariance, or time consistency and time invariance, which are 0.548 and 0.541, respectively.

The model presented in Section 2 predicts that hyperbolic discounters violate stationarity and time consistency in a present-biased direction. These violations can however move in opposite directions when time invariance is violated. To analyze how often present-biased violations of stationarity and time consistency overlap, Panel B presents statistics for present-biased violations only, treating future-biased violations of stationarity (or time consistency) as an observation satisfying stationarity (or time consistency). Similarly, Panel C specifically analyzes future-biased violations, treating present-biased violations of stationarity (or time consistency) as an observation satisfying stationarity (or time consistency).

The first column of Panel B shows that only 10.4 percent of all participants violates both stationarity and time consistency in a present-biased direction. The correlation between present-biased violations of stationarity and time consistency is 0.131, which is again substantially lower than both the correlation between present-biased violations of stationarity and time invariance (0.517) and the correlation between present-biased violations of time consistency and time invariance (0.738).[11] In Panel C, only a small

[11] The second correlation compares present-biased violations of stationarity to violations of time invariance where participants become *more* patient. A participant who allocates 6 vouchers to 'in one month' (and the remaining 4 to 'tomorrow') and 8 vouchers to 'in 3 months' (and the remaining 2 to 'in 2 months') violates stationarity in a present-biased direction. If this person is time consistent (i.e. allocates 8 vouchers to the later payment

Table 5: Distribution of violations of stationarity and time consistency

	Violation if allocation differs by more than ... vouchers:		Excl participants with ≥ 2 identical corners	Counting 2 identical corners as a violation
	> 0	> 1		
	(1)	(2)	(3)	(4)
Panel A. Counting violations in both PB and FB directions				
No violations of S or TC	0.371	0.492	0.120	0.050
Violation of TC but not of S	0.196	0.258	0.270	0.113
Violation of S, but not of TC	0.192	0.138	0.140	0.058
Violations of both S and TC	0.242	0.113	0.470	0.779
Correlation violations S, TC	0.212	0.095	0.087	0.283
Correlation violations S, TI	0.541	0.424	0.423	0.493
Correlation violations TC, TI	0.548	0.621	0.584	0.627
Panel B. Counting only violations in a PB direction				
No PB violations of S or TC	0.542	0.646	0.330	0.154
PB violation of TC, but not of S	0.213	0.229	0.320	0.196
PB violation of S, but not of TC	0.142	0.058	0.170	0.075
PB violations of both S and TC	0.104	0.067	0.180	0.575
Correlation PB violations S, TC	0.131	0.197	0.021	0.369
Correlation PB violations S, TI	0.517	0.371	0.261	0.518
Correlation PB violations TC, TI	0.738	0.730	0.864	0.689
Panel C. Counting only violations in a FB direction				
No FB violations of S or TC	0.725	0.825	0.540	0.254
FB violation of TC, but not of S	0.088	0.050	0.200	0.154
FB violation of S, but not of TC	0.154	0.100	0.220	0.171
FB violations of both S and TC	0.033	0.025	0.040	0.421
Correlation FB violations S, TC	0.084	0.179	-0.120	0.332
Correlation FB violations S, TI	0.449	0.357	0.353	0.621
Correlation FB violations TC, TI	0.597	0.633	0.819	0.744
Number of observations	240	240	100	240

PB: Present-biased. FB: Future-biased. S: Stationarity. TI: Time invariance. TC: Time consistency. Time invariance cannot be classified as either present- or future-biased. The correlation of present-biased violations of S and TI compares present-biased stationarity violations to time-invariance violations where participants become *more* patient, because when choices satisfy time consistency but violate stationarity in a present-biased direction, choices will have become more patient. The correlation of present-biased violations of TC and TI on the other hand compares present-biased time-consistency violations to time-invariance violations where participants become *less* patient, following a similar line of reasoning. For future-biased correlations, the patterns are exactly reversed: the correlations between future-biased violations of S and TI use time-invariance violations where participants become *less* patient, while the correlations between future-biased violations of TC and TI use time-invariance violations where participants become *more* patient.

share of the participants violates both time consistency and stationarity in a future-biased direction, but more than twenty percent violates either stationarity or time consistency in a future-biased direction.

Columns (2) - (4) explore whether noise from a trembling hand or the presence of corner allocations can account for the limited overlap between time consistency and stationarity. Column (2) relaxes the definition of stationarity, time consistency and time invariance violations to allocations differing by at least two vouchers to investigate the effects of a trembling hand. Column (3) excludes all participants who selected two or more identical corner allocations in which participants allocate all vouchers to one of the two payment dates, since these reveal participants' preferences only weakly: their preferred allocations may violate time consistency, stationarity, or time invariance, but this is not observed.[12] As a final robustness check, Column (4) assumes that choices involving two identical corner allocations (i.e. all choices from which we cannot infer whether a concept is violated) do in fact represent a violation. In this way, the results represent an upper bound to the number of violations. These robustness checks show qualitatively similar patterns as Column (1).

In sum, our experiment provides evidence that violations of time consistency and stationarity often do not coincide. Violations of time invariance correlate much better with violations of stationarity and time consistency. To the extent that time consistency and stationarity do not overlap, observed behavior may well be driven by other mechanisms than hyperbolic discounting. These mechanisms may or may not be systematic. The remarkable similarity with Halevy's findings, combined with the stark differences in the experimental designs, raises the question whether the observed patterns simply result from random noise in decision-making that expresses itself in similar ways in very different experiments. The next subsection sheds light

date in the second-round choice), she becomes *more* patient. Following a similar line of reasoning, the correlation of present-biased violations of time consistency and time invariance on the other hand compares present-biased violations of time consistency to violations of time invariance where participants become *less* patient.

[12] Violations of stationarity, time consistency and time invariance are defined in the same way as Column (1).

on this question.

4.3 Behavioral mechanisms I: Random noise in decisions

One possible reason for observed violations of time invariance may simply be noise. A trembling hand may explain why participants change their allocation regarding near-future payment dates from the first to the second round. This will also result in violations of stationarity and time consistency, so that a nonzero correlation between any two violations can result at least partially from random noise.

We therefore estimate discount factors δ and β, a CRRA parameter ρ and the level of noise μ using maximum likelihood, and simulate allocations for the estimated parameters, assuming that violations of time invariance are driven only by random noise in decision-making. Appendix A describes the estimation and simulation procedures in more detail. Table 6 compares our experimental findings in Column (1) with simulated outcomes in Column (2)-(9). Due to the large number of corner allocations, utility is estimated to be convex, $\rho < 0$. Since this assumption is disputable, Columns (2) - (5) present simulation results imposing linear utility ($\rho = 0$), while utility is convex in Columns (6) - (9). Columns (2) - (3) and (6) - (7) impose exponential discounting, $\beta = 1$, so that violations of stationarity and time consistency are driven solely by noise. Columns (4) - (5) and (8) - (9) assume quasi-hyperbolic discounting, $\beta < 1$. As a result, time invariance violations are still due to noise, while violations of stationarity and time consistency may also arise from a present bias. We present simulated outcomes assuming two different levels of noise, μ, which are estimated using maximum likelihood with linear and CRRA utility respectively.

The bottom section shows the correlation between violations of stationarity and time consistency. The *actual* correlation is 0.21 in Column (1). This correlation is fairly high compared to the different simulated correlations in Columns (2) to (9), although it does not always fall outside the simulated confidence intervals. The last row highlights the fact that overlapping violations of stationarity and time consistency often occur in a future-biased or

Table 6: Outcomes realized in the experiment and in Monte Carlo simulations

		Linear utility ($\rho = 0$)				CRRA utility ($\rho \neq 0$)			
Discounting		Exponential $\rho = 0$, $\beta = 1$		Hyperbolic $\rho = 0$, $\hat{\beta} = 0.82$		Exponential $\hat{\rho} = -0.77$, $\beta = 1$		Hyperbolic $\hat{\rho} = -0.71$, $\hat{\beta} = 0.72$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
μ	Actual	$0.\hat{0}9$	0.24	$0.\hat{0}9$	0.24	0.09	$0.\hat{2}4$	0.09	$0.\hat{2}4$
Proportion violating time consistency									
Violation: $x_{1,2} \neq x_{2,2}$	0.44	0.85	0.90	0.85	0.90	0.42	0.78	0.35	0.75
Present-biased: $x_{1,2} > x_{2,2}$	0.32	0.43	0.45	0.56	0.52	0.31	0.48	0.17	0.38
Proportion violating stationarity									
Violation: $x_{1,2} \neq x_{1,1}$	0.43	0.85	0.90	0.85	0.90	0.43	0.78	0.35	0.75
Present-biased: $x_{1,2} > x_{1,1}$	0.25	0.43	0.45	0.56	0.52	0.31	0.49	0.17	0.38
Proportion violating time invariance									
Violation: $x_{1,1} \neq x_{2,2}$	0.58	0.85	0.90	0.85	0.90	0.35	0.77	0.35	0.75
Less patient: $x_{1,1} > x_{2,2}$	0.37	0.43	0.45	0.43	0.45	0.17	0.39	0.17	0.38
Correlations between violations of stationarity and time consistency									
Violation	0.21	0.06	0.01	0.06	0.00	0.47	0.14	0.38	0.15
Present-biased	0.13	0.31	0.32	0.31	0.32	0.54	0.33	0.41	0.32

For a detailed description of the simulated model, see Appendix A. Choices $x_{1,1}$, $x_{1,2}$ and $x_{2,2}$ represents the number of vouchers allocated to the later payment date in the near-future first-round allocation, distant-future first-round allocation, and near-future second-round allocation, respectively. Symbol ρ is the CRRA coefficient, β a present bias by which instantaneous utility for future payments is discounted, and μ indicates the degree of noise. Parameters estimated by means of maximum likelihood are marked $\hat{.}$. Columns (2) - (5) assume linear utility ($\rho = 0$), while Columns (6) - (9) use the maximum likelihood estimate of the risk aversion parameter ($\hat{\rho}$). Furthermore, Columns (2) - (3) and (6) - (7) assume exponential discounting ($\beta = 1$), while Columns (4) - (5) and (8) - (9) use the maximum likelihood estimate of the quasi-hyperbolic discounting parameter ($\hat{\beta}$). Finally, all models are estimated for two levels of noise: the maximum likelihood estimated noise level assuming linear utility ($\hat{\mu} = 0.09$) and the maximum likelihood estimated noise level assuming CRRA ($\hat{\mu} = 0.24$).

opposing direction, as the actual correlation is significantly lower than the simulated correlations when zooming in on present-biased violations. Thus, given the modeled noise structure, random noise alone cannot account for the low correlation between (present-biased) violations of stationarity and time consistency.

4.4 Behavioral mechanisms II: Liquidity constraints

Liquidity constraints resulting from reductions in wealth among those with limited access to credit provide an alternative explanation for why time consistency and stationarity do not overlap. Section 2 formulated hypotheses on how changes in wealth affect stationarity and time consistency in a context with limited access to credit. In this context, anticipated wealth changes result in violations of stationarity but not time consistency, while unanticipated wealth changes result in violations of time consistency but not stationarity.

To test these hypotheses, we analyze whether violations of stationarity and time consistency emerge differently for participants with less access to informal credit (who either cannot raise NGN 20,000 at all, or need to borrow from three or more persons to raise this amount) and a large reduction in wealth (measured as a wealth reduction of more than 80 percent between the first and the second round, i.e. the quarter with the largest relative reduction). Table 7 reports the distribution of participants over the different types of violations, distinguishing between participants with more versus less access to credit, and between participants with versus without a large reduction in household wealth.[13]

The first column reports results for the entire sample (identical to the first column of Table 5), while Column (2) restricts the sample to the people who experienced a large wealth loss. Columns (3) and (4) present results for the same two samples, but zoom in on participants with more access to credit. By contrast, Columns (5) and (6) restrict the sample to individuals

[13] Appendix Table B3 presents the results for people who experienced a large increase of their net wealth. The results illustrate that large gains do not affect choices differently for participants with more versus less access to credit.

Table 7: Distribution of violations of stationarity and time consistency

	All		More access to credit		Less access to credit	
	All	Had a Large Loss	All	Had a Large Loss	All	Had a Large Loss
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Counting violations in both PB and FB directions						
No violations of S or TC	0.371	0.383	0.377	0.433	0.371	0.333
Violation of TC but not of S	0.196	0.267	0.179	0.233	0.205	0.300
Violation of S, but not of TC	0.192	0.117	0.217	0.133	0.167	0.100
Violations of both S and TC	0.242	0.233	0.226	0.200	0.258	0.267
Correlation violations S, TC	0.212	0.245	0.191	0.238	0.250	0.247
Correlation violations S, TI	0.541	0.478	0.507	0.424	0.578	0.538
Correlation violations TC, TI	0.548	0.710	0.443	0.605	0.635	0.809
Panel B. Counting only violations in a PB direction						
No PB violations of S or TC	0.542	0.500	0.528	0.600	0.561	0.400
PB Violation of TC but not of S	0.213	0.383	0.208	0.300	0.212	0.467
PB Violation of S, but not of TC	0.142	0.050	0.160	0.033	0.121	0.067
PB Violations of both S and TC	0.104	0.067	0.104	0.067	0.106	0.067
Correlation PB violations S, TC	0.131	0.089	0.106	0.208	0.173	-0.026
Correlation PB violations S, TI	0.517	0.527	0.615	0.557	0.426	0.523
Correlation PB violations TC, TI	0.738	0.808	0.662	0.675	0.794	0.935
Panel C. Counting only violations in a FB direction						
No FB violations of S or TC	0.725	0.733	0.755	0.733	0.697	0.733
FB Violation of TC but not of S	0.088	0.033	0.066	0.033	0.106	0.033
FB Violation of S, but not of TC	0.154	0.217	0.151	0.200	0.159	0.233
FB Violations of both S and TC	0.033	0.017	0.028	0.033	0.038	0.000
Correlation FB violations S, TC	0.084	0.054	0.102	0.169	0.068	-0.102
Correlation FB violations S, TI	0.449	0.455	0.483	0.432	0.429	0.482
Correlation FB violations TC, TI	0.597	0.552	0.487	0.695	0.685	0.557
Number of observations	240	60	106	30	132	30

S: Stationarity. TI: Time invariance. TC: Time consistency. PB: Present-biased. FB: Future-biased. Time invariance cannot be classified as either present- or future-biased. The correlation of present-biased violations of S and TI compares present-biased stationarity violations to time-invariance violations where participants become *more* patient, because when choices satisfy time consistency but violate stationarity in a present-biased direction, choices will have become more patient. The correlation of present-biased violations of TC and TI on the other hand compares present-biased time-consistency violations to time-invariance violations where participants become *less* patient, following a similar line of reasoning. For future-biased correlations, the patterns are exactly reversed: the correlations between future-biased violations of S and TI use time-invariance violations where participants become *less* patient, while the correlations between future-biased violations of TC and TI use time-invariance violations where participants become *more* patient. An individual is classified as having less access to credit if they cannot borrow 20,000 NGN in case of an emergency, or needs to borrow from three or more different people to raise this amount.

who have less access to credit. Column (6) is therefore of explicit interest as it focuses on the most liquidity constrained participants: the individuals with less access to credit who faced substantial declines in wealth.

Hypothesis 1a states that anticipated decreases in wealth are associated with future-biased violations of stationarity but not time consistency. The proportion of the participants making such choices in Panel C is indeed larger among people who faced a large loss in Column (2) than it is in the full sample in Column (1). As predicted, this effect is somewhat more pronounced for participants with less access to credit in Column (6).

Hypothesis 2a states that unanticipated decreases in background wealth are associated with present-biased violations of time consistency, but not stationarity. Approximately 38 percent of the participants who faced a large loss violated time consistency but not stationarity in a present-biased direction, a much larger proportion than the 21.3 percent in the full sample. This effect is particularly strong for participants with less access to credit in Column (6) at 46.7 percent, compared to 30.0 percent of participants in Column (4) who could use informal credit to cope with the wealth loss. In the former group, very few present-biased violations are related to a violation of stationarity. Instead, present-biased violations of time consistency are highly correlated with violations of time invariance (0.935). Among less constrained participants, this correlation is lower (0.675).

Appendix Table B4 presents the marginal effects of relative reductions in wealth, but now modeling reductions in wealth as a continuous variable. Results are qualitatively similar. Thus, the finding that violations of time consistency but not stationarity are more common among more liquidity-constrained participants is robust to varying thresholds in the definition of a 'large' loss.

Hypotheses 1b and 2b predict violation patterns for participants whose background wealth increases. To be in line with these predictions, participants who faced a relatively large decrease in wealth, should display these violation patterns *less* often than the full sample. Hypothesis 1b states that anticipated increases in wealth are associated with present-biased violations of stationarity, but not time consistency. The results in Panel B show that

14.2 percent of participants in the full sample violate stationarity but not time consistency in a present-biased direction, while this is the case for only 5.0 percent of the participants who faced a large loss. Hypothesis 2b states that unanticipated increases in wealth are associated with future-biased violations of time consistency, but not stationarity. As shown in Panel C, the proportion violating time consistency but not stationarity in a future-biased direction among participants who faced a large decline in wealth is again lower than the proportion in the full sample.

In sum, these findings suggest that the observed non-overlapping violations of time consistency and stationarity are not purely random. Violations of either time consistency or stationarity are correlated with wealth changes and access to informal credit, consistent with theoretical predictions.

4.5 Behavioral Mechanisms III: Experimental Factors

Finally, we discuss whether key features of the experimental design may have caused discrepancies between violations of stationarity and time consistency. A first potential limitation is that intertemporal allocations are elicited using convex time budgets, which have been questioned to yield valid choice patterns (Chakraborty et al., 2015). Our design does not include multiple price lists and does not vary experimental wealth or the delay between the sooner and the later payment date, so that we cannot test the weak axiom of revealed preferences, or test for wealth and impatience monotonicity, robustness checks proposed by (Chakraborty et al., 2015). Nevertheless, we showed in Section 3.3 that only few participants violate demand monotonicity, suggesting they have understood the convex time budget task. Further, convex time budgets have been shown to have equal or better predictive validity compared to double multiple price lists (Andreoni, Kuhn and Sprenger, 2015), supporting the validity of our method to elicit violations of time consistency and stationarity.

A second factor potentially biasing time preference experiments is a lack of trust in the experimenters among participants (Thaler, 1981; Chabris, Laibson and Schuldt, 2008; Sprenger, 2015). Participants who do not trust experimenters to return with their money on future dates will make different decisions for allocations where 'today' is one of the payment dates than for allocations with only future payment dates, regardless from whether this person is a hyperbolic discounter. To avoid this we only included payment dates in the future, so that the soonest payment date is no longer immediate. Several recent studies adopt this approach and fail to reject stationarity at the aggregate level, suggesting that this indeed does eliminate changes in trust as a potential confound (e.g. Sprenger, 2015; Giné et al., 2016; Andreoni and Sprenger, 2012*a*). Further indications that limited levels of trust have not influenced participants' allocations follow from the observation that even when there is no return on waiting (i.e. sooner vouchers are worth 200 NGN so that they do not 'lose' any money by allocating vouchers to the sooner date) participants allocate some vouchers to the later date in the first-round choice regarding near-future payment dates, as shown in Table 4. Participants thus appear to have trusted the experimenter to hold on to their money, suggesting that discrepancies between stationarity and time consistency violations are not driven by a lack of trust.

Decision fatigue or limited attention spans offer a third potential reason for the lack of overlap between violations of stationarity and time consistency in our experiment. In the second round, participants only had to allocate their vouchers once, while they were presented with eight different choices in the first round. As a result, participants may have paid more attention to their choice in the second round than they did in their first-round choices. Hoddinott, Hoel and Schwab (2014) find that when participants become fatigued, they behave more impatiently. Thus, if fatigue were a problem in our sample, we would expect participants to choose less patient allocations in their last decisions in the first round. However, the last first-round choices were near-future allocations. As shown in Table 4, participants actually do not behave more impatiently in these choices compared to their first set of choices, which concerned the distant-future payment dates. Fatigue does not explain why time consistency and stationarity do not overlap.

A fourth possible confound is that participants were interviewed about their cash inflows and outflows on a weekly basis for the larger project the experiment was embedded in, potentially increasing awareness of a present bias among hyperbolic discounters. In other words, hyperbolic discounting may have presented itself in a more sophisticated way during the second round than during the first round, affecting second-round choices, $x_{2,2}$. To test whether frequent interviewing affected the level of participants' financial sophistication, the project randomly selected a number of control households to be interviewed only at baseline and during an endline survey one year later. Table 8 compares respondents interviewed with high frequency with the control group in terms of a number of self-reported financial planning variables measured at baseline and endline. Financial planning reportedly improved among all respondents, also those in the control group, and we do not observe stronger improvements in financial planning for participants interviewed on a weekly basis. It is hence unlikely that frequent interviewing explains the discrepancy between violations of stationarity and time consistency.

A fifth factor potentially weakening the correlation between stationarity and time consistency is the elicitation method of the second-round choice. In the second round of the experiment, participants are shown their first-round allocation for the distant future and asked whether they would like to revise this choice. This procedure, adopted from Giné et al. (2016), could potentially increase the probability of observing time-consistent choices that nevertheless violate stationarity and time invariance. However, comparing our results to Halevy (2015) (who did not present the second round as a revision and did not show participants their first-round allocations during the second round), we observe a similar proportion of participants who satisfy time consistency but violate stationarity and time invariance. We hence conjecture that presenting the second-round choice as a revision does not confound our results.

Finally, the limited presence of formal financial institutions in the region where our research was conducted may also have influenced our findings. Most participants in our experiment do not have access to formal financial instruments to save or borrow against future payments at a fixed, salient interest rate. Therefore, participants cannot engage in financial arbitrage

Table 8: Effect of diary participation on the ability to plan financially

	(1) Ability to plan finances	(2) Type of planner	(3) Sticks to plans	(4) Only saves leftovers	(5) # times runs out of money
Participates	-0.012	0.124**	-0.115**	0.009	-0.031
	(0.094)	(0.063)	(0.054)	(0.017)	(0.167)
Endline	0.710***	0.182***	-0.168***	0.014	0.535***
	(0.115)	(0.052)	(0.046)	(0.017)	(0.146)
Participates *X* Endline	0.107	-0.110	0.097	-0.014	0.073
	(0.156)	(0.087)	(0.076)	(0.028)	(0.222)
Male respondent	-0.048	0.023	-0.033	-0.010	-0.152
	(0.053)	(0.041)	(0.031)	(0.016)	(0.097)
Age	0.007***	-0.001	0.001	0.000	-0.000
	(0.002)	(0.001)	(0.001)	(0.001)	(0.004)
Household size	-0.004	0.000	-0.001	-0.001	0.004
	(0.011)	(0.008)	(0.007)	(0.002)	(0.020)
Has 1 to 6 yrs of education	0.072	-0.055	0.057	0.002	0.154
	(0.084)	(0.049)	(0.041)	(0.016)	(0.129)
Has 7 or more yrs of education	0.318***	-0.036	0.054	0.018	0.217
	(0.078)	(0.058)	(0.048)	(0.019)	(0.151)
Community effects	Yes	Yes	Yes	Yes	Yes
N	1110	1125	1125	1125	1115
R-squared	0.285	0.031	0.039	0.006	0.072

In Column (1), the dependent variable is the first factor of a principal component analysis for financial planning ability; in Column (2), the dependent variable is a categorical scale of saving behavior, with value 1 for someone who plans well and sticks to these plans, 3 for someone who only saves when money is left over, and 2 for someone in the middle; in Column (3), it is a dummy variable that indicates whether someone plans well and sticks to these plans; in Column (4), the dependent variable is a dummy variable that indicates whether someone only saves when money is left over; and in Column (5), the dependent variable is categorical variable that indicates the number of times a person runs out of money in a month. Reported p-values are based on t-tests with standard errors (shown in parentheses) clustered by household. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

using their experimental payments. Hence we do not need to censor subjective discount rates by the market interest rate (Cubitt and Read, 2007; Chabris, Laibson and Schuldt, 2008; Andersen, Harrison and Lau, 2014).

In sum, limited trust in the experimenters, decision-making fatigue, improved financial awareness due to high-frequency data collection among a subsample of participants, presenting the second-round choice as a revision, and arbitrage do not seem to drive our results. We conclude that violations of stationarity and time consistency often do not overlap, and this can be explained partly by changes in participants' financial situation.

5 Conclusion

An increasing number of scholars is interested in eliciting experimental measures of hyperbolic discounting. Most studies do so by means of a (cross-sectional) static choice experiment in which participants choose whether to receive a sooner-smaller or later-larger payment, with payment dates in either the near future or in a more distant future. Such experiments elicit violations of *stationarity*. Alternatively, one can elicit violations of *time consistency* by means of a longitudinal design in which participants choose at different points in time whether to receive a sooner-smaller versus later-larger payment, keeping the payment dates fixed. Both violations of time consistency and stationarity are commonly interpreted as evidence for hyperbolic discounting. These violations may however also be driven by violations of *time invariance*, meaning that participants express different preferences regarding near-future payment dates depending on when they make their decisions. Hyperbolic discounting can be inferred from stationarity or time consistency violations only when participants choose identical allocations for near-future payment dates independent of the timing of these decisions.

A field experiment in rural Nigeria analyzed to what extent violations of stationarity and time consistency overlap, or whether they are instead related to violations of time invariance. Using convex time budgets, participants were asked during a first round to allocate vouchers between 'tomorrow' and 'one month from now', as well as between 'two months from now' and 'three months from now', and during a second round two months later to allocate vouchers between 'tomorrow' and 'in one month from now', the same calendar dates as in the second choice. A difference between the first and the second choice is labeled a violation of stationarity; a difference between the second and the third choice is labeled a violation of time consistency; and a difference between the first and the third choice is labeled a violation of time invariance.

Although 43.4 percent of participants violate time consistency and a similar 43.8 percent violate stationarity, these violations are mainly attributed to violations of time invariance. The correlation between stationarity and time consistency is weak. Only 10.4 percent of participants violates both stationarity and time consistency in a present-biased direction and could be qualified as hyperbolic discounters.

These findings are both qualitatively and quantitatively similar to results from a laboratory experiment with undergraduate students at the University of British Colombia in Vancouver (Halevy, 2015). We subsequently investigate potential reasons for why time consistency and stationarity do not overlap. We find that random noise alone cannot explain their low correlation. This conclusion is further reinforced by the observation that changes in wealth predict violations of time consistency but not stationarity (or vice versa), especially for liquidity constrained individuals. Thus, when income or expenditures fluctuate over time, systematic changes in wealth can drive a wedge between stationarity and time consistency. In that case, one cannot infer hyperbolic discounting from observing either a stationarity or a time consistency violation.

At the aggregate level, participants violate time consistency and time invariance but not stationarity. The finding that stationary is not rejected in the aggregate is in line with many recent studies with monetary rewards that carefully take into account trust issues (e.g. Sprenger, 2015; Giné et al., 2016; Andreoni and Sprenger, 2012*a*; Augenblick, Niederle and Sprenger, 2015; Carvalho, Meier and Wang, 2014). Increasing participants' trust in the experimenter is one explanation for why aggregate-level stationarity is satisfied. Our study provides a second explanation for the finding of aggregate stationarity. Participants who anticipate losses in wealth are more likely to violate stationarity in a future-biased direction, which offsets present-biased violations among other participants.

These results have direct consequences for the design of studies that aim to elicit empirical measures of hyperbolic discounting. In order to identify hyperbolic discounters, experiments need to measure both violations of time consistency and of stationarity, unless violations of time invariance can be ruled out. We suggest one channel through which time preferences may change over time, i.e. wealth changes, but more research is needed on the (in)stability of time preferences.

Our findings have important implications for the design of financial tools, such as commitment savings devices. Choice reversals in one-time experiments have often led to the conclusion that time inconsistent behavior is driven by present bias, and that people need commitment devices and other nudges to be able to commit to their earlier plans (Ashraf, Karlan and Yin, 2006; Bryan, Karlan and Nelson, 2010; Thaler and Benartzi, 2004). Such commitment devices tie individuals to the mast, but for those individuals who did not anticipate a change in their financial circumstances, this may actually harm their welfare. In regions with large fluctuations in income and expenditures, commitment is not necessarily welfare-improving.

References

Andersen, Steffen, Glenn Harrison, and Elisabet Lau, Morten nd Rutström. 2014. "Discounting behavior: A reconsideration." *European Economic Review*, 71: 15–33.

Andreoni, James, and Charles Sprenger. 2012*a*. "Estimating time preferences from convex budgets." *American Economic Review*, 102(7): 3333–3356.

Andreoni, James, and Charles Sprenger. 2012*b*. "Risk preferences are not time preferences." *American Economic Review*, 102(7): 3357–3376.

Andreoni, James, Michael Kuhn, and Charles Sprenger. 2015. "On measuring time preferences." *Journal of Economic Behavior and Organization*, 116: 451–464.

Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. "Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines." *Quarterly Journal of Economics*, 121(2): 635–672.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2015. "Working over time: Dynamic inconsistency in real effort tasks." *Quarterly Journal of Economics*, 130(3): 1067–1115.

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

Carvalho, Leandro, Stephan Meier, and Stephanie Wang. 2014. "Poverty and economic decision making: Evidence from changes in financial resources at payday." *Unpublished manuscript.*

Chabris, Christopher, David Laibson, and Jonathon Schuldt. 2008. "Intertemporal choice." *Palgrave Dictionary of Economics.*

Chakraborty, Anujit, Evan Calford, Guidon Fenig, and Yoram Halevy. 2015. "External and Internal Consistency of Choices made in Convex Time Budgets." *Unpublished manuscript.*

Coller, Maribeth, and Melonie Williams. 1999. "Eliciting individual discount rates." *Experimental Economics*, 2(2): 107–127.

Cubitt, Robin, and Daniel Read. 2007. "Can intertemporal choice experiments elicit time preferences for consumption?" *Experimental Economics*, 10(4): 369–389.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time discounting and time preference: A critical review." *Journal of Economic Literature*, 40(2): 351–401.

Giné, Xavier, Jessica Goldberg, Dan Silverman, and Dean Yang. 2016. "Revising commitments: Field evidence on the adjustment of prior choices." *Unpublished manuscript.*

Green, Leonard, Nathanael Fristoe, and Joel Myerson. 1994. "Temporal discounting and preference reversals in choice between delayed outcomes." *Psychonomic Bulletin and Review*, 1(3): 383–389.

Halevy, Yoram. 2014. "Some comments on the use of monetary and primary rewards in the measurement of time preferences."

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

Harrison, Glenn. 2008. "Maximum likelihood estimation of utility functions using STATA." *UCF Working Paper 06-12.*

Harrison, Glenn, Morten Lau, and Elisabet Rutström. 2005. "Dynamic consistency in Denmark: A longitudinal field experiment." *Unpublished Manuscript.*

Harrison, Glenn, Morten Lau, and Elisabet Rutström. 2013. "Identifying time preferences with experiments: Comment." *CEAR Working Paper.*

Harrison, Glenn, Morten Lau, and Melonie Williams. 2002. "Estimating individual discount rates in Denmark: A field experiment." *American Economic Review*, 92(5): 1606–1617.

Hoddinott, John, Jessica Hoel, and Benjamin Schwab. 2014. "Self-control fatigue, cognitive function, and the expression of time preference: Experimental results from Ethiopia." *Unpublished Manuscript.*

Janssens, Wendy, Berber Kramer, Marijn Van der List, and David Pap. 2013. "Financial and Health Diaries: A yearlong weekly panel of farming households in Nigeria and Kenya." *Unpublished raw data.*

Kirby, Kris, and Richard Herrnstein. 1995. "Preference reversals due to myopic discounting of delayed reward." *Psychological Science*, 6(2): 83–89.

Krupka, Erin, and Melvin Stephens. 2013. "The stability of measured time preferences." *Journal of Economic Behavior and Organization*, 85: 11–19.

Laibson, David. 1997. "Golden eggs and hyperbolic discounting." *Quarterly Journal of Economics*, 112(2): 443–477.

Meier, Stephan, and Charles Sprenger. 2015. "Temporal stability of time preferences." *Review of Economics and Statistics*, 97(2): 273–286.

Read, Daniel, Shane Frederick, and Mara Airoldi. 2012. "Four days later in Cincinnati: Longitudinal tests of hyperbolic discounting." *Acta Psychologica*, 140(2): 177–185.

Sayman, Serdar, and Ayse Öncüler. 2009. "An investigation of time inconsistency." *Management Science*, 55(3): 470–482.

Sozou, Peter. 1998. "On hyperbolic discounting and uncertain hazard rates." *Proceedings of the Royal Society of London*, 265(1): 2015–2020.

Sprenger, Charles. 2015. "Judging experimental evidence for dynamic inconsistencies." *American Economic Review*, 105(5): 280–285.

Thaler, Richard. 1981. "Some empirical evidence on dynamic inconsistency." *Economics Letters*, 8(3): 201–207.

Thaler, Richard, and Shlomo Benartzi. 2004. "Save more tomorrow: Using behavioral economics to increase employee saving." *Journal of Political Economy*, 112(1): 164–187.

Appendices

A Random noise in decisions: details of the model

In Section 4.3, we compare our experimental findings with simulated predictions from a model in which violations of time invariance are driven only by random noise in decision-making. In this section, the details of the model are discussed in more detail.

If $s \in \{1, 2\}$ indicates the decision round, $\sigma \in \{1, 2\}$ the payment round, and $v_S \in \{100, 120, 150, 200\}$ the value of vouchers allocated to the sooner payment date, let the intertemporal utility from allocating x vouchers to the later payment date be $U_{s,\sigma}(x; v_S)$, which is defined in Equations (1) - (3). The probability that a participant allocates x vouchers to the later payment date can be written as the ratio of utility from this allocation to the utility summed over all ten possible allocations, so that choices with higher utility have a higher probability of being selected:

$$P(x_{s,\sigma} = x; v_S) = \frac{U(x_{s,\sigma}; v_S)^{\frac{1}{\mu}}}{\sum_{z=0}^{10} U(z_{s,\sigma}; v_S)^{\frac{1}{\mu}}} \tag{4}$$

where $\mu > 0$ is a parameter specifying the degree of noise (if μ is infinitesimal, then there is no noise, and as μ is going to infinity, decision making becomes an entirely random process). Since in Equations (1) - (3) in Section 2, the optimization problem is equivalent for first- and second-round allocations regarding near-future payment dates, $x_{1,1}$ and $x_{2,2}$, these allocations would be the same in the absence of random noise. In our simulations, time invariance can hence be violated only due to noise.

Following Harrison, Lau and Rutström (2013), we estimate the model using maximum likelihood.[14] We assume that instantaneous utility is of

[14] Doing so, we build on STATA routines carefully explained in Harrison (2008). This estimation procedure has the ability to properly analyze corner allocations as well. This is particularly important in our data set, since many participants do not choose an interior allocation, but allocate all vouchers to one of the payment dates. Such corner allocations represent censored decisions, which potentially biases estimates in linear regressions. Andreoni and Sprenger (2012*a*) propose using Tobit regressions to estimate the model, but

the CRRA type, $u(c) = c^{1-\rho}/(1-\rho)$, where ρ is the coefficient of relative risk aversion and consumption c is equal to $c_S = v_S(10 - x)$ on the sooner payment date and $c_L = 200x$ on the later payment date. The estimated (quasi-)hyperbolic discounting parameter β, the noise parameter μ, and the risk aversion parameter ρ, together with a participant's voucher value v_S, the decision round s, and the round in which payments occur σ, yield estimates of the cumulative probability that a participant allocates $x_{s,\sigma}$ vouchers to the later payment date:

$$CDF(a; v_S, s, \sigma) = \sum_{x=0}^{a} P(x_{s,\sigma} = x) \tag{5}$$

These cumulative distribution functions are in turn used to simulate the allocation every participant chooses in each of the three choices, which allows us to calculate correlations between violations of stationarity, time consistency, and time invariance.[15]

the Tobit model makes a number of theoretical assumptions that are inconsistent with the set-up of convex time budget tasks (Harrison, Lau and Rutström, 2013).

[15]The routine follows a procedure described in Meier and Sprenger (2015). First, every participant-choice observation is assigned a random number $z_{ihs\sigma}$ from a uniform distribution $U \sim [0, 1]$. Second, the random number is compared with the cumulative probabilities. If the random number satisfies $CDF((a–1); V_{ih}, s\sigma) \leq z_{ihs\sigma} < CDF(a; V_{ih}, s\sigma)$, with $CDF(-1; V_{ih}, s\sigma) = 0$ for a decision moment s, payment dates σ and voucher value V, then the simulated number of vouchers allocated to the later payment date is $\sim X_{ihs\sigma} = a$. Third, the routine calculates a number of summary statistics: the percentage of participants for whom we observe violations of time invariance, of time consistency, and of stationarity; the direction of these violations; and the correlation between these different violations. This routine is repeated 999 times for each observation to derive 95% confidence intervals for each statistic and p-values for the realized statistics in the experiment.

B Appendix Tables

Table B1: Violations of demand monotonicity by number of interior choices

Number of interior choices in a choice set	Number of choice sets	Violations of demand monotonicity: Number of choice sets	Violations of demand monotonicity: Proportion of choice sets
0	187	0	0.0000
1	49	3	0.0612
2	26	3	0.1154
3	75	5	0.0667
4	143	15	0.1049
Total	480	26	0.0542

Table includes all eight allocations that participants made in the first round: four allocations regarding near-future payment dates and four allocations regarding distant-future payment dates. Demand monotonicity implies that the amount allocated to the later payment date is weakly increasing in the return on waiting, and is tested by comparing pairs of allocations within a set of four allocations ('choice set') where only the return on waiting increases. A choice set violates demand monotonicity if any of these pairs violates demand monotonicity.

Table B2: Comparing implemented and non-implemented first-round choices

Choice		Implemented choices			Non-implemented choices			Difference
Timing of payment dates	Value of sooner voucher	N	Mean	Std. dev.	N	Mean	Std. dev.	in means
Near	200	65	3.62	3.18	221	3.46	3.39	0.15
Near	150	55	7.65	2.86	231	8.04	2.16	-0.38
Near	120	46	8.91	2.00	240	8.48	2.10	0.43
Near	100	75	9.05	1.37	211	8.91	2.25	0.13
Far	200	65	3.05	3.22	221	2.86	3.35	0.19
Far	150	55	8.64	2.15	231	8.19	2.32	0.45
Far	120	46	8.98	1.88	240	8.70	2.07	0.28
Far	100	75	8.88	2.22	211	9.20	1.86	-0.32

The value of vouchers allocated to the later payment date is fixed at 200 NGN, while the value of vouchers allocated to the sooner payment date varies from 200 NGN to 100 NGN. Mean represents the average number of vouchers (out of ten) that a participants allocates to the later payment date (rather than the sooner) at decision moment i for payment payment dates j. Reported p-values are based on t-tests with standard errors (shown in parentheses) clustered by household. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B3: Distribution of violations of stationarity and time consistency for people who experienced a large gain

	All		More access to credit		Less access to credit	
	All	Had a Large Gain	All	Had a Large Gain	All	Had a Large Gain
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Counting violations in both PB and FB directions						
No violations of S or TC	0.371	0.517	0.377	0.400	0.371	0.600
Violation of TC but not of S	0.196	0.133	0.179	0.120	0.205	0.143
Violation of S, but not of TC	0.192	0.200	0.217	0.280	0.167	0.143
Violations of both S and TC	0.242	0.150	0.226	0.200	0.258	0.114
Correlation violations S, TC	0.212	0.237	0.191	0.199	0.250	0.252
Correlation violations S, TI	0.541	0.628	0.507	0.519	0.578	0.721
Correlation violations TC, TI	0.548	0.495	0.443	0.199	0.635	0.721
Panel B. Counting only violations in a PB direction						
No PB violations of S or TC	0.542	0.633	0.528	0.520	0.561	0.714
PB Violation of TC but not of S	0.213	0.117	0.208	0.120	0.212	0.114
PB Violation of S, but not of TC	0.142	0.150	0.160	0.240	0.121	0.086
PB Violations of both S and TC	0.104	0.100	0.104	0.120	0.106	0.086
Correlation PB violations S, TC	0.131	0.257	0.106	0.164	0.173	0.341
Correlation PB violations S, TI	0.517	0.671	0.615	0.749	0.426	0.551
Correlation PB violations TC, TI	0.738	0.689	0.662	0.561	0.794	0.791
Panel C. Counting only violations in a FB direction						
No FB violations of S or TC	0.725	0.850	0.755	0.840	0.697	0.857
FB Violation of TC but not of S	0.088	0.050	0.066	0.040	0.106	0.057
FB Violation of S, but not of TC	0.154	0.083	0.151	0.080	0.159	0.086
FB Violations of both S and TC	0.033	0.017	0.028	0.040	0.038	0.000
Correlation FB violations S, TC	0.084	0.134	0.102	0.345	0.068	-0.075
Correlation FB violations S, TI	0.449	0.427	0.483	0.369	0.429	0.484
Correlation FB violations TC, TI	0.597	0.418	0.487	0.180	0.685	0.685
Number of observations	240	60	132	35	106	25

S: Stationarity. TI: Time invariance. TC: Time consistency. PB: Present-biased. FB: Future-biased. Time invariance cannot be classified as either present- or future-biased. The correlation of present-biased violations of S and TI compares present-biased stationarity violations to time-invariance violations where participants become *more* patient, because when choices satisfy time consistency but violate stationarity in a present-biased direction, choices will have become more patient. The correlation of present-biased violations of TC and TI on the other hand compares present-biased time-consistency violations to time-invariance violations where participants become *less* patient, following a similar line of reasoning. For future-biased correlations, the patterns are exactly reversed: the correlations between future-biased violations of S and TI use time-invariance violations where participants become *less* patient, while the correlations between future-biased violations of TC and TI use time-invariance violations where participants become *more* patient. An individual is classified as having less access to credit if they cannot borrow 20,000 NGN in case of an emergency, or needs to borrow from three or more different people to raise this amount.

Table B4: Logit modeling present-biased violations of time consistency but not stationarity, $x_{1,1} \geq x_{1,2}$ & $x_{2,2} \neq x_{1,2}$

	(1)	(2)	(3)	(4)
Panel A. Marginal Effects				
Credit Constrained	0.011	-0.055	-0.000	-0.023
	(0.052)	(0.063)	(0.053)	(0.071)
Large Loss (dummy)	0.209***	0.104		
	(0.049)	(0.068)		
Large Loss x Credit Constrained		0.191*		
		(0.098)		
Loss (continuous)			0.148***	0.126**
			(0.043)	(0.050)
Loss x Credit Constrained				0.044
				(0.091)
Controls	Yes	Yes	Yes	Yes
Panel B. Full Results of the Logit				
Credit Constrained	0.084	-0.427	-0.001	-0.175
	(0.395)	(0.484)	(0.395)	(0.527)
Large Loss (dummy)	1.575***	0.800		
	(0.417)	(0.542)		
Large Loss x Less access to credit		1.468**		
		(0.747)		
Loss (continuous)			1.099***	0.937**
			(0.336)	(0.393)
Loss x Less access to credit				0.326
				(0.676)
Age	0.015	0.019	0.017	0.017
	(0.017)	(0.018)	(0.017)	(0.017)
Number of children in the household	-0.096	-0.100	- 0.099	-0.101
	(0.062)	(0.062)	(0.065)	(0.066)
Male	1.952***	1.915***	1.780***	1.765***
	(0.609)	(0.641)	(0.588)	(0.596)
Level of education: some primary school	-0.527	-0.416	-0.480	-0.453
	(0.679)	(0.680)	(0.646)	(0.652)
Level of education: more than primary school	-0.481	-0.461	-0.330	-0.343
	(0.631)	(0.615)	(0.613)	(0.604)
Source of income: farming	-1.205**	-1.1599*	-1.059*	-1.037*
	(0.569)	(0.623)	(0.566)	(0.573)
Source of income: other or nothing	-1.628***	-1.670***	-1.682***	-1.665***
	(0.547)	(0.552)	(0.570)	(0.558)
Poor household	-0.128	-0.146	-0.061	-0.058
	(0.441)	(0.442)	(0.489)	(0.490)
Violates monotonicity	0.709	0.513	0.496	0.444
	(0.635)	(0.677)	(0.675)	(0.684)
Regional Effects	Yes	Yes	Yes	Yes
Number of observations	236	236	236	236
Mean Dep. Var	0.212	0.212	0.212	0.212
Pseudo R-squared	0.188	0.201	0.178	0.179

An individual is classified as having suffered a large loss if the household wealth declined by more than 80 percent. An individual is classified as having less access to credit if they cannot borrow 20,000 NGN in case of an emergency, or needs to borrow from three or more different people to raise this amount. A household is classified as poor if the level of net wealth at baseline is below the median, i.e. below 50,000 NGN. Reported p-values are based on t-tests with standard errors (shown in parentheses) clustered by household. * $p < 0.10$, ** $p < 0.05$, ***