CrossMark

Be patient when measuring hyperbolic discounting: Stationarity, time consistency and time invariance in a field experiment☆

Wendy Janssens[a,b,c], Berber Kramer[b,d,*], Lisette Swart[a,c,e]

[a] Department of Economics, Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081HV Amsterdam, The Netherlands
[b] Amsterdam Institute for International Development, Pietersbergweg 17, 1105BM Amsterdam, The Netherlands
[c] Tinbergen Institute, Gustav Mahlerlaan 117, 1082MS Amsterdam, The Netherlands
[d] Markets, Trade and Institutions Division, International Food Policy Research Institute, 2033 K Street, Washington, DC, NW 20006, USA
[e] CPB Netherlands Bureau for Economic Policy Analysis, Bezuidenhoutseweg 30, 2594AV The Hague, The Netherlands

ARTICLE INFO

ABSTRACT

Hyperbolic discounting is one potential reason why savings remain low among the poor. Most 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 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 the same individuals violate both stationarity and time consistency. Our field experiment in Nigeria examines the extent to which this is the case. The experiment measured both stationarity and time consistency for the same participants. Violations of the two rarely coincide, especially among more liquidity-constrained participants. Thus, in a context of liquidity constraints, eliciting only one type of choice reversal is insufficient to identify hyperbolic discounting.

1. Introduction

For the poor, consumption smoothing is hindered by fluctuating cash flows and limited access to formal credit and insurance (Collins et al., 2009). This is compounded by a constrained ability to save (Dupas and Robinson, 2013). One potential reason for low savings is hyperbolic discounting, meaning that implicit discount rates are lower for tradeoffs in the more distant future than for tradeoffs in the near future (Frederick et al., 2002). A hyperbolic discounter violates *time consistency*, i.e. she prefers to invest towards increased future consumption when asked far in advance, but when asked right before investing the money, she opts for sooner but lower consumption. She also violates *stationarity*, meaning that she prefers for example $110 in 31 days over $100 in 30 days, but rather has $100 today instead of $110 tomorrow (Green et al., 1994; Kirby and Herrnstein, 1995). Empirical observations of either violation have been interpreted as

evidence of hyperbolic discounting.

However, hyperbolic discounting is an unambiguous explanation for such choice reversals only if the same person violates both stationarity and time consistency. A second yet often neglected explanation for choice reversals is a violation of *time invariance* (Halevy, 2015), which means that the marginal rate of substitution (MRS) changes over time.[1] 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 changes in the economic environment (Read et al., 2012), either through unanticipated shocks to household finances (Dean and Sautmann, 2016) or through anticipated changes in income (Epper, 2016). Crucially, when stationarity or time consistency are measured in isolation, one may wrongly interpret choice reversals caused by time-varying background wealth as evidence of hyperbolic discounting.

[1] Thus, violations of stationarity and time consistency do not necessarily imply any form or 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 et al., 2005). Further, note that violations of stationarity and time consistency do not necessarily imply any form of irrationality.

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, 2012a) 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 after the others 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. 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 by differences in the second and third allocation, elicited at different points in time regarding the same calendar dates. 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' (Halevy, 2015). If time invariance is satisfied, a hyperbolic discounter will violate both stationarity and time consistency. Observing either choice reversal is sufficient to infer hyperbolic discounting *only* in that case. If time invariance is not satisfied, observing a violation of either choice reversal is *not* sufficient to identify hyperbolic discounters. This paper sheds light on the severity of the potential misclassification.

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, and 62 percent of violations result from time-variant choices instead. Moreover, for nearly half of this subsample, the two choice reversals move in different directions with one present-biased and one future-biased violation. This is not just noise in decision-making; participants who violate time consistency but not stationarity have significantly less access to informal credit and lose more wealth over time than other participants. This suggests that violations of time invariance are in part due to liquidity constraints. As a result, when observed in isolation, choice reversals are not sufficient evidence of hyperbolic discounting. Instead, identifying hyperbolic discounters requires a longitudinal design eliciting both stationarity and time consistency.

This paper makes three unique contributions to the literature. First, the experiment links choice reversals to violations of time invariance. To our best knowledge, Halevy (2015) is the only study with similar analyses, but using a different subject pool (undergraduate students in economics). Giné et al. (2016) link measures of stationarity to time consistency for a subject pool that is more comparable to ours, but do not provide measures of time invariance. Meier and Sprenger (2015) measure stationarity and time invariance, but not time inconsistence. Other experimental studies either analyze violations of stationarity (e.g. Coller and Williams, 1999; Harrison et al., 2002; Carvalho et al., 2016) or of time consistency (e.g. Sayman and Öncüler, 2009; Read et al., 2012), without linking the two.

Second, we test whether liquidity constraints can explain why stationarity and time consistency often do not overlap, i.e. why participants violate time invariance.[2] Unlike Halevy (2015), we can identify liquidity-constrained participants using rich survey data on

participants' financial characteristics. We find that compared to other participants, those who are relatively most constrained are significantly more likely to violate time consistency without also violating stationarity. We identify liquidity-constrained participants as those with less access to credit and greater reductions in wealth, independent of whether the household explicitly reports a shock. This distinguishes our paper from Giné et al. (2016) and Dean and Sautmann (2016), whose empirical analyses focus mainly on *unanticipated* expenditure and income shocks, and from Carvalho et al. (2016), who link stationarity to the timing of *anticipated* income. Our measure captures both anticipated and unanticipated liquidity constraints.

Third, the finding that liquidity constraints result in violations of time invariance 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 et al., 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, supporting the theory that standard experimental measures of time preferences and stationarity capture financial constraints and changes in non-experimental wealth rather than innate discount rates (Dean and Sautmann, 2016; Epper, 2016).

These findings have potential implications for policies that promote savings in low-income settings. Savings play an important role in smoothing consumption, in particular for the poor with volatile cash flows and limited access to formal financial services (Collins et al., 2009). Hyperbolic discounting is one of the main theories used to explain low savings rates, but this theory is based on observed violations of *either* stationarity *or* time consistency. Our findings suggest that such violations are often driven by liquidity constraints. Hence, policies that aim at promoting savings among the poor should not only address hyperbolic discounting, but also consider the role of liquidity constraints when designing mechanisms to improve their ability to save.

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 the role of liquidity constraints. 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

---

Table 1
Three types of intertemporal choices.

Round 1: Near Future	$x_{1,1}$	❶ Ⓢ	Ⓛ		
Round 1: Distant Future	$x_{1,2}$	❶		Ⓢ	Ⓛ
Round 2: Near Future	$x_{2,2}$	❷ Ⓢ	Ⓛ		
Time	t	τ_1 p_{1S} p_{1L}		τ_2 p_{2S}	p_{2L}

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.

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 (see also 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. $x_{1,1} \neq x_{1,2}$.

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, $x_{1,2} \neq x_{2,2}$.

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, $x_{1,1} \neq x_{2,2}$.

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, 3 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).[3] We also assume that individuals lack access to financial markets, and cannot transfer background wealth ω from outside the experiment between the sooner and the later payment date. The three voucher allocations optimize the following three target functions:[4]

Table 2
Defining three types of violations.

Type of violation				Decision round	Payment round
Violation of stationarity	$x_{1,1}$	$\neq$	$x_{1,2}$	R1	R1 vs. R2
Violation of time consistency	$x_{1,2}$	$\neq$	$x_{2,2}$	R1 vs. R2	R2
Violation of time invariance	$x_{1,1}$	$\neq$	$x_{2,2}$	R1 vs. R2	R1 vs. R2

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

---

[3] 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.

[4] Our experiment implements only one of these three allocations, so that the first-round allocation regarding first-round payment dates does not influence background wealth in allocations regarding second-round payment dates.

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

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

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

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, and v_S and v_L the value of vouchers allocated to the sooner and later payment date, respectively. 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.

An individual's background wealth on the sooner versus the later payment date in round t is indicated as ω_{tS} and ω_{tL}, respectively. For simplicity, we hold background wealth on the later payment date constant, $\omega_{1L} = \omega_{2L}$, and focus on ceteris paribus effects of changes in background wealth on the sooner payment date. We do so because background wealth at the later payment date, ω_{tL}, is not known at decision moment τ_t, in contrast to ω_{tS}. As a result, analyzing changes in expected future background wealth is analytically less tractable. In addition, we measured background wealth empirically around the time of the sooner but not the later payment date. In the context of our experiment, the effects of a ceteris paribus change in background wealth on the sooner payment date are hence more relevant from an empirical perspective. Allowing for changes in future background wealth would either dampen or amplify choice reversals, depending on whether the difference in background wealth on the sooner versus later payment date is reduced or widened.[5]

We will now discuss how present bias, $\beta < 1$, and changes in background wealth, ω_{tS}, lead to choice reversals. 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 $x_{2,2} < x_{1,2}$, because for them, $\beta < 1$. In this case, measuring either stationarity or time consistency is sufficient to infer hyperbolic discounting.

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

In this case, 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, this choice reversal is future-biased. 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 choice reversal. 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. Dean and Sautmann (2016) highlight this point in relation to the effect of unexpected income shocks on standard experimental measures of time preferences; and Epper (2016) uses a similar argument when attributing stationarity violations to anticipated increases in income. We therefore formulate the following hypotheses for individuals who lack access to credit:

1) If changes in background wealth are *anticipated*,

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

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

2) If changes in background wealth are *unanticipated*,

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

Hypothesis 2b. ... *increasing* background wealth on sooner payment dates ($\mathbb{E}_{\tau_1}\omega_{1S} < \mathbb{E}_{\tau_2}\omega_{2S}$) 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. 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.

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 (2012a)'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.[6] 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' 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

---

[5] Epper (2016) and Noor (2009) develop models for intertemporal choice in which currently liquidity-constrained individuals violate stationarity without violating time consistency because they expect an increase in income in a not so distant future, i.e. a change in ω_{1L} compared to ω_{1S}. The mechanisms behind the results in our conceptual framework – changing ω_{2S} compared to ω_{1S} - are comparable to the mechanisms in these two papers.

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

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 et al., 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, 2012a).

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.[7] Fig. 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 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).[8] 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'.[9]

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 second-round payment dates was 0.9.[10] 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
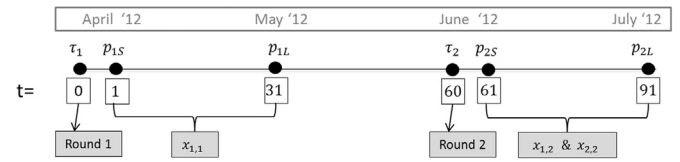

Fig. 1. Timeline of the Study. Circles represent two different decision moments, $t = \tau_1$ for a first round at t=0 and $t = \tau = 2$ for a second round sixty days later. During these decision moments, people allocate vouchers to a sooner and a later payment date. Squares represent the payment dates 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.

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 1000 and 2000 NGN, and they earned 1862 NGN on average. These stakes are fairly high, as the maximum possible payment of 2000 NGN 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.

Allowing participants to revise their choices unexpectedly has advantages and disadvantages (Halevy, 2015). One advantage is that all choices in the first round are truly incentivized. If instead participants knew they would be allowed to revise their choice regarding second-round payment dates, we would have incentivized only choices regarding the near-future payment dates, $x_{1,1}$ and $x_{2,2}$. We wanted to rule out this potential explanation for violations of stationarity and time consistency. Further, in this way, the option to allocate money between the two payment dates was unanticipated in both rounds, limiting the effect of the experimental payments on financial behavior prior to the decision.

At the same time, paying the second-round decision without prior disclosure raises a number of concerns. First, participants may think they are expected to revise their choices during the second round. If this were an important consideration, participants would violate time consistency more often than stationarity, deviate from the initial allocation by just a few vouchers, and violate at similar rates in present- and future-biased directions. We do not observe such choice patterns. Second, the revisit might decrease participants' trust in future payments. However, when asked to motivate their second-round choices, participants only mentioned how they would use the experimental payments and never pointed at a lack of trust (see Online Appendix Fig. B2). Hence, we have no indication that second-round decisions were driven by the option to revise allocations unexpectedly.[11]

3.3. Description of the participant sample

The experiment targeted 303 individuals who participated in the baseline survey in March 2012. Of those, 286 persons (94.4 percent)

---

[7] 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).

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

[9] 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 (2012a, 2012b) and Giné et al. (2016) do not find any evidence of order effects.

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

[11] In the period between the first and the second round, participants were visited on a weekly basis by enumerators for the larger study in which we embedded the experiment. This most likely increased their trust in receiving future payments. Also, in most communities, some of the participants had received payments in the first round, further signaling experimenter trustworthiness to others.

Table 3
Description of participant characteristics.

	All participants		Revisited sample		Difference
	N	Mean	N	Mean	in means
	(1)	(2)	(3)	(4)	(5)
Age	286	40.58	240	40.17	−0.406
Male	286	0.395	240	0.379	−0.016
No formal education	285	0.596	240	0.596	−0.001
Number of participants within the household	286	2.825	240	2.850	0.025
Main income from farming	286	0.378	240	0.367	−0.011
Main income from business	286	0.402	240	0.417	0.015
Main income from other	286	0.098	240	0.096	−0.002
No main source of income	286	0.122	240	0.121	−0.002
Planning to harvest before July	286	0.066	240	0.071	−0.005
Financial wealth R1 (in 1000 NGN)	286	63.25	240	71.15	7.902[**]
Financial wealth R2 (in 1000 NGN)	277	50.84	240	50.02	−0.82
More access to informal credit	284	0.426	238	0.445	0.019
Satisfies monotonicity	286	0.087	240	0.087	0.000

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 she 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 except four variables (age, the number of children and two financial balances) are binary indicators. Significance levels in Column (5) for variable y are based on a t-test for $\hat{\beta} = 0$, where $\hat{\beta}$ is the difference between revisited and non-revisited participants, estimated using a linear regression with standard errors clustered at the household level. [†] $p < 0.10$, [*] $p < 0.05$.
    [**] $p < 0.01$.

participated in the first round of our experiment. For 256 participants (89.5 percent of 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 286 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. On average, 2.8 members per household participated in the experiment. We therefore cluster standard errors in all analyses at the household level. The majority of participants never entered the formal school system. The two predominant sources of income among participants are farming (37.8 percent) and business (40.2 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 2000 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 need 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 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, 2012a). 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) 200 NGN vs. 150 NGN, (2) 150 NGN vs. 120 NGN, and (3) 120 NGN vs. 100 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 1440 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 (2012a), and better understanding than more comparable participants in Giné et al. (2016).

Following Chakraborty et al. (2015), we also test 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 (see Online Appendix Table B1). Such violations are hence 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. The total budget remained constant at 10 vouchers of 200 NGN, which also implies that we cannot test whether participants choose to save more as the experimental stakes increase; a magnitude effect that Epper (2016) attributes to liquidity constraints.

The within-subject analyses only include participants for whom all three choices depicted in Table 1 were elicited. For the 16 dropouts and the randomly selected participants who received their allocation for first-round payment dates, we cannot observe violations of time invariance or time consistency. Columns (3) to (5) compare the 240 participants revisited during the second round with the full sample. Attrition is not related to 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 finds that violations of time invariance that account for this discrepancy can
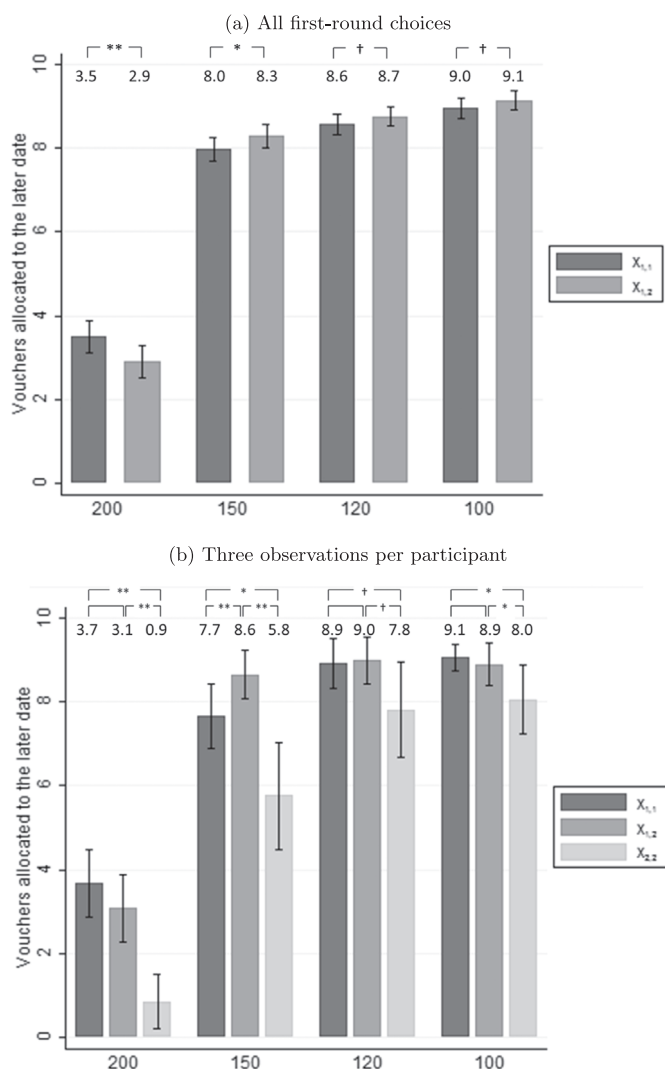
Fig. 2. Number of vouchers allocated to the later date by choice type and voucher value. Significance levels reported for choice type comparisons ($x_{1,1}$ vs. $x_{1,2}$, $x_{1,2}$ vs. $x_{2,2}$ and $x_{1,1}$ vs. $x_{2,2}$) are estimated from a regression of the number of vouchers allocated to the later payment date on choice type, with standard errors clustered at the household level. † $p < 0.10$, * $p < 0.05$, ** $p < 0.01$. (a) All first-round choices, (b) Three observations per participant.

partly be explained by changes in participants' wealth.

4.1. Description of choices

Fig. 2 summarizes the number of vouchers that participants allocate to the later payment date. Panel (a) includes all eight first-round choices for each participant, showing allocations regarding first-round payment dates (the light-gray bar, $x_{1,1}$) versus second-round payment dates (the dark-gray bar, $x_{1,2}$). The figure separates choices by the value of vouchers allocated to the *sooner* payment date ('sooner vouchers'). These 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, when sooner vouchers are worth 200 NGN, the return on waiting is the lowest. In that case, participants indeed allocate most vouchers to the sooner payment date, leaving on average 3.5 and 2.9 vouchers for the later date in choices regarding first- and second-round payment dates, respectively. When sooner vouchers are worth 150 NGN, participants have a higher return on waiting and allocate about five additional vouchers to the later date ($p < 0.01$). Compared to these choices, when vouchers are worth 120 NGN, participants allocate an

additional half voucher to the later date ($p < 0.01$). Reducing the value of sooner vouchers even further to 100 NGN has a very similar effect. Hence, as the return on waiting increases, participants allocate more vouchers to the later payment date; consistent with monotonicity.

Panel (a) also compares first-round choices regarding payment dates in the near future, 'tomorrow versus in one month' ($x_{1,1}$), with choices regarding payment dates in the distant future, 'in two months versus in three months' ($x_{1,2}$). A significant difference in allocations between these two choices implies a rejection of stationarity. When there is no return on waiting, that is, when sooner vouchers are worth 200 NGN, participants allocate *more* vouchers to the later payment date in choices regarding the near future, $x_{1,1}$, than in choices regarding the more distant future, $x_{1,2}$ ($p < 0.001$), violating stationarity in a future-biased direction. When there is a positive return on waiting, that is, sooner vouchers are worth 150, 100 or 120 NGN, participants allocate *fewer* vouchers to the later payment date in choices regarding payments in the near future. These present-biased violations of stationarity are consistent with hyperbolic discounting, but they are small in absolute terms.[12]

In order to investigate time consistency and time invariance, Panel (b) adds second-round choices. For comparison, this panel focuses on choices for which we observe all three choice types, omitting 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. Second-round choices for payment dates in the near future, $x_{2,2}$, are different from both first-round allocations, $x_{1,1}$ and $x_{1,2}$, independent of the value of vouchers allocated to the sooner payment date. We hence 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. In addition, these violations are of much larger magnitude than violations of stationarity ($x_{1,2}$ versus $x_{1,1}$).

We conclude that the aggregate data violate time consistency and time invariance, and that these violations are much more pronounced than violations of stationarity. This suggests that time-inconsistent behavior is linked more closely to violations of time invariance than stationarity, and the overlap between violations of time consistency and stationarity appears limited. The remainder of this section will further investigate the overlap between stationarity, time consistency and time invariance, including only 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.

4.2. Classification of participants

Fig. 3 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). Participants in Groups 3, 4 and 5 violate time consistency. Nearly half of them – 19.6 out of 43.8 percent in our sample – coincide with a violation of time invariance, without stationarity being violated (Group 3). Likewise, among the 43.3 percent of participants who violate stationarity (Groups 2, 4 and 5), 19.2 percent violates time invariance without violating time consistency (Group 2). Only 24.2 percent of participants violates both stationarity and time consistency (Groups 4 and 5). Finally, 62 percent of all violations (Groups 2 to 5) are either stationary but time-inconsistent or non-stationary but time-consistent,

---

[12] Allocations for sooner vouchers worth 200 NGN (with no return on waiting) thus seem to follow a different pattern than allocations with sooner vouchers 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 paper.
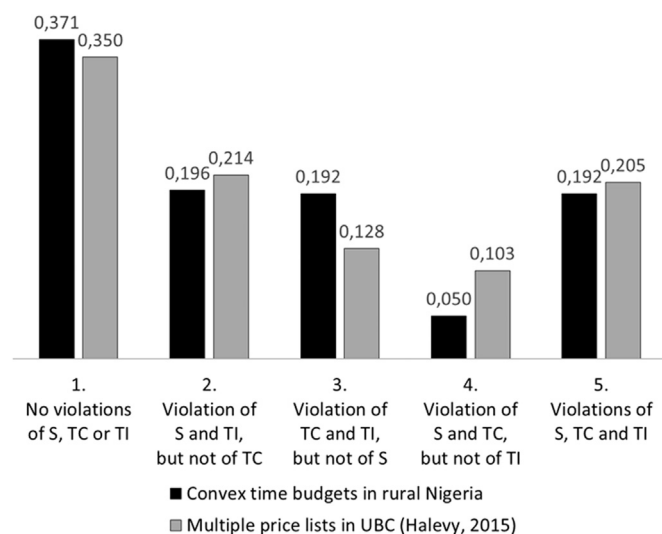
Fig. 3. Comparing distribution of participants to distribution in Halevy (2015). S: Stationarity ($x_{1,1} = x_{1,2}$); TC: Time Consistency ($x_{1,2} = x_{2,2}$); TI: Time Invariance ($x_{1,1} = x_{2,2}$). The percentages listed here from Halevy (2015) are based on Column 1 from Table II on page 345 of Halevy Y., Time consistency: stationarity and time invariance, *Econometrica* **83** (1), 2015, 335–352.

coinciding with time-variant choices instead (Groups 2 and 3). Thus, the share of participants that violates stationarity is very similar to the share that violates time consistency, but overlap between these two groups is limited.

An important question is to what extent these findings differ from those in Halevy (2015), 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, 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 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). 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 4 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 0.541 and 0.548, respectively. The correlation between violations of stationarity and time consistency in Halevy (2015) is 0.325, which is somewhat higher than the correlation in our sample, but still lower than his correlations between stationarity and time invariance (0.384), and between time consistency and time invariance (0.537).

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 choices are not time invariant. To analyze overlap in the *direction* of these violations, 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, which is a mere 22.7 percent of those with at least one present-biased violation. 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).[13] In Panel C, only a small 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. 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.[14] In this way, the results represent an upper bound to the number of violations. These robustness checks show qualitatively similar patterns as Column (1). We also rule out the possibility that the low correlation between violations of time consistency and stationarity is only driven by random noise (see A.1), or by features of the experimental design (see A.2).

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. We discussed earlier that in a context with limited access to credit, discrepancies between stationarity and time consistency may arise from changing background wealth. The remainder of this section will test to what extent liquidity constraints can explain our results.

4.3. Behavioral mechanisms: liquidity constraints

Section 2 formulated the hypotheses that in a context with limited access to credit, 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 relatively 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)

---

[13] 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 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.

[14] In both Columns (3) and (4), violations of stationarity, time consistency and time invariance are defined in the same way as Column (1).

Table 4
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 S, but not of TC	0.192	0.138	0.140	0.058
Violation of TC but not of S	0.196	0.258	0.270	0.113
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 S, but not of TC	0.142	0.058	0.170	0.075
PB violation of TC, but not of S	0.213	0.229	0.320	0.196
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 S, but not of TC	0.154	0.100	0.220	0.171
FB violation of TC, but not of S	0.088	0.050	0.200	0.154
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 PB violations of S and TI compares PB stationarity violations to time-invariance violations where participants become *more* patient, because when choices satisfy time consistency but violate stationarity in a PB direction, choices will have become more patient. The correlation of PB violations of TC and TI on the other hand compares PB time-consistency violations to time-invariance violations where participants become *less* patient, following a similar line of reasoning. For FB correlations, the patterns are exactly reversed: the correlations between FB violations of S and TI use time-invariance violations where participants become *less* patient, while the correlations between FB violations of TC and TI use time-invariance violations where participants become *more* patient.

and a relatively 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). We label participants satisfying both criteria as *more* liquidity-constrained.

Fig. 4 compares their behavior (light-grey bar) with that of the remainder of participants who are *less* liquidity-constrained (dark-grey bars); that is, participants with either more access to informal credit, or without a large reduction in wealth. The figure classifies participants into four categories: those who violate neither stationarity nor time consistency (Group 1), those who violate stationarity but not time consistency (Group 2), those who violate time consistency but not stationarity (Group 3), and those who violate both stationarity and time consistency (Groups 4 and 5). Panels (a) and (b) focus on present-biased and future-biased violations, respectively.

Hypothesis 1a states that anticipated decreases in wealth are associated with future-biased violations of stationarity but not time consistency. This correlation should be stronger for *more* liquidity-constrained participants, who by definition experienced a relatively large loss but have less access to informal credit to smooth consumption over time. The proportion of participants in Group 2 in Fig. 4b is indeed larger for the more liquidity-constrained compared to the less liquidity-constrained sample, although the difference is not statistically significant.

To further delve into the differences between more and less liquidity-constrained participants, Table 5 splits the sample of *less* constrained participants (whose behavior is summarized in the dark-grey bars in Fig. 4) into three groups: participants with more access to credit without a large reduction in wealth in Column (1); participants with more access to credit and with a large reduction in wealth in Column (2); and participants with less access to credit but without large wealth reduction in Column (4). Column (5) presents violation patterns for *more* liquidity-constrained participants (represented by the light-grey bars in Fig. 4, with both less

access to credit and a large reduction in wealth). Panel A only includes violations in a present-biased direction and Panel B in a future-biased direction. Columns (3), (6), (7) and (8) test whether differences between the various subsamples are significant. As an additional test of H1a, Panel B of Table 5 compares the proportion of participants who violate stationarity but not time consistency in a future-biased direction between the four subsamples and confirms that all differences are in the expected direction but lack statistical significance.

Hypothesis 2a states that unanticipated decreases in background wealth are associated with present-biased violations of time consistency, but not stationarity. Consistent with this hypothesis, 47 percent of more liquidity-constrained participants violate time consistency but not stationarity in a present-biased direction (Fig. 4a), which is a significantly larger proportion than the 18 percent in the less constrained sample. Panel A of Table 5 shows that this pattern is most pronounced for participants who have *both* less access to credit *and* who have experienced a large loss, as differences between participants with and without a large loss or between participants with more and less access to credit are not significant (Columns (3), (7), and (8)).[15]

Hypotheses 1b and 2b predict violation patterns for participants whose background wealth increases. As a corollary, more liquidity-

---

[15] As a robustness check, we also estimated a logit model for the relation between liquidity constraints and violating only time consistency (but not stationarity) in a present-biased direction, using various definitions of 'large' reductions in wealth. Our main findings are insensitive to the percentage at which a reduction in wealth is defined as 'large' (within reasonable bounds), and to controlling for potential confounds of the relation between liquidity constraints and choice reversals (see Online Appendix Table B4). 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.
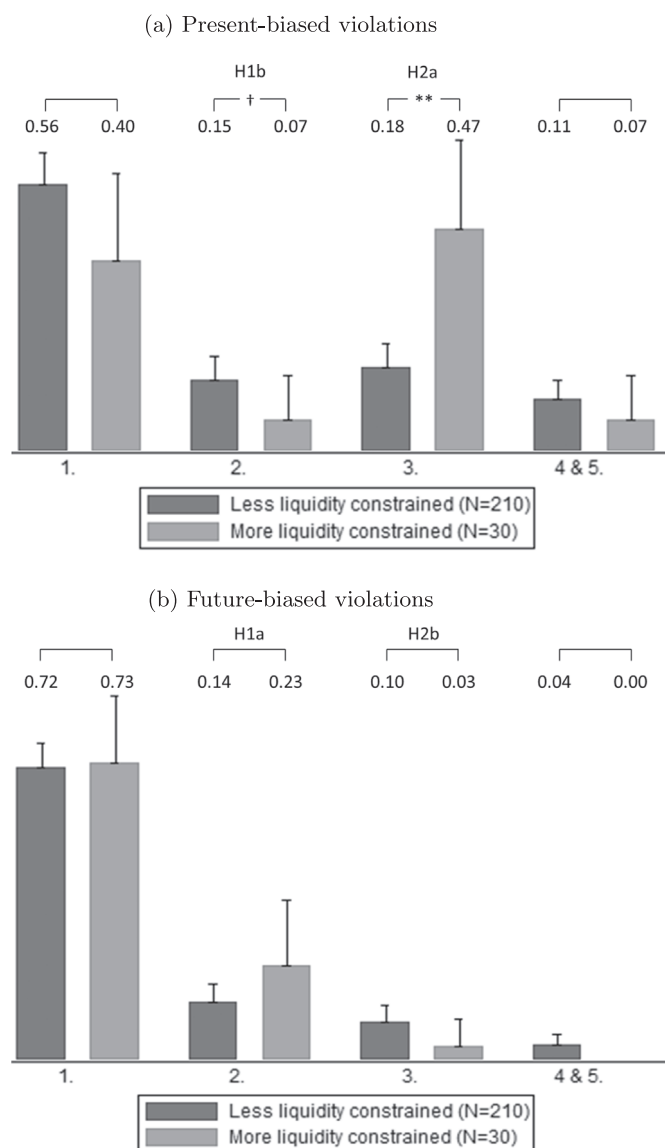
(a) Present-biased violations

(b) Future-biased violations

Fig. 4. Violation patterns for more versus less liquidity constrained participants. Significance levels for comparisons of more versus less constrained participants are estimated from a regression of falling into Group *X* on participant type, with standard errors clustered at the household level. † $p < 0.10$, * $p < 0.05$, ** $p < 0.01$. (a) Present-biased violations, (b) Future-biased violations.

constrained participants should display these violation patterns *less* often than the less constrained sample. Hypothesis 1b states that anticipated increases in wealth are associated with present-biased violations of stationarity, but not time consistency. Hence, in Fig. 4a more liquidity-constrained participants should be classified into Group 2 less often compared to less liquidity-constrained participants. Indeed, 15 percent of less constrained participants violate stationarity but not time consistency in a present-biased direction, while this is the case for only 7 percent of more constrained participants. This difference is driven by the subsample with more access to credit, as shown in Panel A of Table 5. Present-biased violations of stationarity but not time consistency are especially prevalent among participants with more access to credit and without a large loss, i.e. the least constrained subsample.[16]

Hypothesis 2b states that unanticipated increases in wealth are associated with future-biased violations of time consistency, but not stationarity. This implies that in Fig. 4b, more constrained participants should fall less often into Group 3 than other participants. The proportion violating time consistency but not stationarity in a future-biased direction is small among both samples, and is indeed the lowest among more liquidity-constrained participants. Although this difference is not statistically significant, Panel B of Table 5 shows that among participants with less access to credit, those with a large loss are significantly less likely to fall into Group 3 when focusing on violations in a future-biased direction ($p < 0.10$).[17]

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. Thus, by measuring either stationarity or time consistency, one may well measure the extent to which a participant is liquidity-constrained, instead of measuring whether he or she is a hyperbolic discounter.

5. Conclusion

Hyperbolic discounting is one potential reason for low savings, limiting consumption smoothing among low-income households. Most evidence of hyperbolic discounting is inferred from (cross-sectional) static choice experiments 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 choice reversals 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 intertemporal choice is time-invariant.

A field experiment in rural Nigeria analyzed to what extent violations of stationarity and time consistency result from violations of time invariance, and whether these violations are related to liquidity constraints. Using convex time budgets, we asked participants during a first round to allocate a future gift 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, we asked the same participants to allocate their gift between 'tomorrow' and 'in one month from now', referring to the same calendar dates as in the second choice. A participant with different allocations in the first and the second choice violates stationarity; different allocations in the

---

(footnote continued)

individuals with less access to credit, those with large gains are more likely to satisfy both time consistency and stationarity, and less likely to violate time consistency but not stationarity in a present-biased direction. Gains improve their ability to smooth consumption over time, reducing the need to violate time consistency in a present-biased direction. Further, conditional on having less access to credit, individuals with large gains are less likely to violate stationarity, but not time consistency, in a future-biased direction. This is consistent with the theory that these individuals anticipate a less binding liquidity constraint in the distant future than those without large gains.

[17] Online Appendix Table B5 investigates heterogeneity in violation patterns by participants' main source of income. The table compares farming and non-farming participants. We find no heterogeneity in present-biased violation patterns. By contrast, in terms of future-biased violations, a large reduction in wealth is associated with a higher probability of violating especially stationarity among farming participants, and a lower probability of time consistency violations among non-farmers. We however note that for most of our participants, household income depends heavily on the agricultural season, even when the participant himself or herself does not report farming as main source of income.

---

[16] Online Appendix Table B3 presents the results for people who experienced a large *increase* in net wealth instead of a large loss. Large gains are associated with significantly different violation patterns only among individuals with less access to credit, suggesting that again the interaction of credit access and changes in wealth influences decisions. For

Table 5
Distribution of violations of stationarity and time consistency.

	More access to credit			Less access to credit			More vs Less access	
	No large loss (1)	Large loss (2)	Diff. (2)-(1) (3)	No large loss (4)	Large loss (5)	Diff. (5)-(4) (6)	No large loss (4)-(1) (7)	Large loss (5)-(2) (8)
Panel A. Counting only violations in a present-biased direction								
No PB violations of either S or TC	0.500	0.600	0.100	0.608	0.400	−0.208[†]	0.108	−0.200
PB Violation of S, but not of TC (cf. H1b)	0.211	0.033	−0.177[**]	0.137	0.067	−0.071	−0.073	0.033
PB Violation of TC, but not of S (cf. H2a)	0.171	0.300	0.129	0.137	0.467	0.329[**]	−0.034	0.167
PB Violations of both S and TC	0.118	0.067	−0.051	0.118	0.067	−0.051	−0.000	0.000
Panel B. Counting only violations in a future-biased direction								
No FB violations of either S or TC	0.763	0.733	−0.030	0.686	0.733	0.047	−0.077	0.000
FB Violation of S, but not of TC (cf. H1a)	0.132	0.200	0.068	0.137	0.233	0.096	0.006	0.033
FB Violation of TC, but not of S (cf. H2b)	0.079	0.033	−0.046	0.127	0.033	−0.094[†]	0.049	0.000
FB Violations of both S and TC	0.026	0.033	0.007	0.049	0.000	−0.049[*]	0.023	−0.033
Number of observations	76	30		102	30			

S: Stationarity. TI: Time invariance. TC: Time consistency. 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; and as having a large loss when wealth reduces by more than 80 percent between the first and the second round (the quarter with the largest relative reduction). Significance levels are based on a *t*-test for $\widehat{\beta} = 0$, where $\widehat{\beta}$ is the difference across choice types in vouchers allocated to the later payment date, estimated with standard errors clustered at the household level.

[†] $p < 0.10$.
[*] $p < 0.05$.
[**] $p < 0.01$.

second and the third choice imply a violation of time consistency; and a difference between the first and the third choice is a violation of time invariance.

We find that 43.4 percent of participants violate time consistency and a similar 43.8 percent violates stationarity. However, 62 percent of these participants do not violate both stationarity and time consistency; and among those who do violate both, choice reversals often move in opposite directions, frequently violating time consistency in a present-biased direction while violating stationarity in a future-biased direction. Only 10.4 percent of participants violate both stationarity and time consistency in a present-biased direction, which is potentially explained by hyperbolic discounting. Other participants' choice reversals coincide with violations of time invariance, making hyperbolic discounting a less likely explanation. 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).

Our results indicate that especially individuals who are relatively more liquidity constrained (i.e. those with greater reductions in wealth and less access to credit) are more likely to violate time consistency but not stationarity in a present-biased direction, and that they are more likely to violate stationarity but not time consistency in a future-biased direction. These findings are in line with our theoretical predictions for participants facing unanticipated and anticipated liquidity constraints, respectively.[18] We conclude that liquidity constraints can drive a wedge between stationarity and time consistency. When disposable income fluctuates over time and individuals lack sound financial instruments to smooth consumption, one cannot infer hyperbolic discounting from observing either a stationarity or a time consistency violation alone. Hence, experiments that aim to measure hyperbolic discounting cannot rely on measuring either time consistency or stationarity, unless they can rule out violations of time invariance.

The literature highlights two potential explanations for violations of time invariance in the presence of liquidity constraints. First, the poor are generally exposed to large uninsured risks, resulting in frequent expenditure shocks and volatile incomes (Dercon, 2002). Those with less access to credit are unable to absorb unanticipated shocks to their financial situation, and cannot fully smooth consumption in the presence of anticipated changes either. Hence, both anticipated and unanticipated changes in one's background wealth can shift the marginal rate of substitution (Dean and Sautmann, 2016; Epper, 2016). This, in turn, results in violations of time invariance.

Second, recent theories on the psychology of poverty suggest that scarcity can reduce decision-making quality by making households focus on the present (Mullainathan and Shafir, 2013). A financial shock creating scarcity may hence temporarily increase present bias, amplifying violations of time invariance. In low-income settings, the overlap between violations of time consistency and stationarity may therefore not be as strong as it is in a context with less volatile incomes and expenditures, especially for the most liquidity-constrained households. Thus, our findings are not necessarily contradictory of hyperbolic discounting and present bias; however, the role of such preferences in decision-making may become less dominant as households start facing more binding liquidity constraints. More research is needed to test these two explanations for violations of time invariance.

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 hyperbolic discounting and present bias. Commitment devices help individuals who are aware of this bias to commit to their earlier plans (Ashraf et al., 2006; Bryan et al., 2010; Thaler and Benartzi, 2004). An important question, however, is how 'hard' commitments should be for a savings device to generate the highest value; when individuals cannot tailor their savings plans to fluctuations in their financial means, or when they cannot withdraw their savings in times of financial pressure, they have less flexibility to smooth consumption in the presence of changing circumstances. In theory, sophisticated hyperbolic discounters will hence trade off their desired level of commitment with their perceived need for flexibility.

Based on our findings, we conjecture that flexibility becomes relatively more important when individuals have limited access to credit and insurance, especially when incomes or expenditures are highly volatile. For such individuals, the hard constraints imposed by commitment savings devices may be too strong. In support of these theoretical conjectures, recent empirical work finds greater impacts of soft commitment savings devices with flexible rules compared to

---

[18] As such, our results provide a potential explanation for future-biased violations of stationarity, which are a commonly observed puzzle in choice experiments.

impacts of more binding alternatives (Dupas and Robinson, 2013; Karlan and Linden, 2016). More research is needed to understand which types of commitment technologies are most effective under varying circumstances in generating demand, stimulating usage, and reaching intended goals. Nonetheless, it is crucial that such commitment devices - although context-dependent - help the poor achieve their savings and investment goals without compromising their ability to also smooth consumption.

Appendix A. Alternative explanations for time invariance violations

Section 4.3 attributes the wedge between violations of time consistency and stationarity to liquidity constraints. This appendix analyzes the potential role of random noise in decision-making and experimental factors.

A.1. Behavioral mechanisms II: random noise in decision-making

The remarkable similarity of our findings with those of Halevy (2015), despite stark differences in the experimental designs, raises the question whether the observed patterns result from random noise in decision-making. A trembling hand results in violations of time invariance and can hence potentially express itself in similar ways in different experiments.

To analyze whether random noise can generate choice patterns comparable to those in our experiment, we simulate choices under the assumption that participants violate time invariance only because of random noise, not because of changes in background wealth. To that end, we first estimate a random utility model using maximum likelihood.[19] Estimated parameters include discount factors δ and β, a CRRA parameter ρ and the level of noise μ. Using these estimates, we then simulate participants' allocations, abstracting from changes in background wealth, so that only random noise can explain why a participant violates time invariance (and hence violates either time consistency or stationarity).

Formally, 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 Eqs. (1)–(3). We assume that background wealth at both the sooner and the later payment date is zero in each of the three equations. 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:

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

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). If background wealth is constant in Eqs. (1)–(3) (Section 2), the optimization problem is equivalent for first- and second-round allocations regarding near-future payment dates, $x_{1,1}$ and $x_{2,2}$. Thus, these allocations would be the same in the absence of random noise, and simulated choices can violate time invariance only due to noise.

Following Harrison et al. (2013), we estimate the model using maximum likelihood.[20] We assume that instantaneous utility is of 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:

$$CDF(a; v_S, s, \sigma) = \sum_{x=0}^{a} P(x_{s,\sigma} = x)$$
(A.2)

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.[21]

Appendix Table A1 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.

---

[19] Apesteguia and Ballester (2015) note that when using this method, the estimated choice probabilities will not satisfy monotonicity with respect to preference parameters. They show that this problem is of economic relevance especially when estimating risk aversion parameters; when estimating discount rates, this problem occurs only with substantial delays between the sooner and the later payment dates (e.g. two years), well beyond the delays observed in experiments that are comparable to ours.

[20] 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 (2012a) propose using Tobit regressions to estimate the model, but the Tobit model makes a number of theoretical assumptions that are inconsistent with the set-up of convex time budget tasks (Harrison et al., 2013).

[21] 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.

Table A1
Outcomes realized in the experiment and in Monte Carlo simulations.

Discounting	(1) Actual	Linear utility ($\rho = 0$)				CRRA utility ($\rho \neq 0$)			
		Exponential $\rho = 0$, $\beta = 1$		Hyperbolic $\rho = 0$, $\widehat{\beta} = 0.82$		Exponential $\widehat{\rho} = -0.77$, $\beta = 1$		Hyperbolic $\widehat{\rho} = -0.71$, $\widehat{\beta} = 0.72$	
		(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
μ		$\widehat{0.09}$	0.24	$\widehat{0.09}$	0.24	0.09	$\widehat{0.24}$	0.09	$\widehat{0.24}$
Proportion violating time consistency									
Violation: $x_{1,2} \neq x_{2,2}$	0.44	0.84	0.90	0.86	0.90	0.26	0.73	0.31	0.77
Present-biased: $x_{1,2} > x_{2,2}$	0.32	0.42	0.45	0.57	0.52	0.13	0.36	0.18	0.47
Proportion violating stationarity									
Violation: $x_{1,2} \neq x_{1,1}$	0.43	0.84	0.90	0.86	0.90	0.26	0.73	0.30	0.77
Present-biased: $x_{1,2} > x_{1,1}$	0.25	0.42	0.45	0.57	0.52	0.13	0.36	0.18	0.48
Proportion violating time invariance									
Violation: $x_{1,1} \neq x_{2,2}$	0.58	0.84	0.90	0.81	0.89	0.26	0.72	0.27	0.72
Less patient: $x_{1,1} > x_{2,2}$	0.37	0.42	0.45	0.40	0.44	0.13	0.36	0.13	0.36
Correlations between violations of stationarity and time consistency									
Violation	0.21	0.04	0.00	0.09	0.01	0.37	0.16	0.42	0.20
Present-biased	0.13	0.44	0.37	0.49	0.39	0.86	0.57	0.88	0.62

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 ∧. Columns (2)–(5) assume linear utility ($\rho = 0$), while Columns (6)–(9) use the estimated risk aversion parameter ($\widehat{\rho}$). Columns (2)–(3) and (6)–(7) assume exponential discounting ($\beta = 1$), while Columns (4)–(5) and (8)–(9) use the estimated present bias parameter ($\widehat{\beta}$). All models are estimated for two levels of noise: the estimated noise level assuming linear utility ($\widehat{\mu} = 0.09$) and the estimated noise level assuming CRRA ($\widehat{\mu} = 0.24$). When linear utility is assumed (Columns (2)–(5)), the estimated degree of noise is 0.09, so there is a hat in Column (2) and (4). There are no hats in Column (3) and (5), because there the assumed degree of noise is the degree that is estimated when assuming CRRA utility. Similarly, when CRRA utility is assumed in Columns (6)–(9), the estimated degree of noise is 0.24. As a result, there are hats in Columns (7) and (9), but not in Columns (6) and (8) which assume the degree of noise based on the estimation assuming linear utility.

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)–(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 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.

A.2. 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. 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 et al., 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 et al., 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, 2012a). 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 Fig. 2. 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 et al. (2016) 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 Fig. 2, 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 as part of a larger project, 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. Self-reported financial planning improved from baseline to midline 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 (see Online Appendix Table B6). 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 using their experimental payments. Hence we do not need to censor the marginal rate of substitution elicited from participants' choices by the market interest rate (Cubitt and Read, 2007, Chabris et al., 2008; Andersen et al., 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.

Appendix B. Supplementary data

Supplementary data associated with this article can be found in the online version at http://dx.doi.org/10.1016/j.jdeveco.2016.12.011.

References

Andersen, S., Harrison, G., Lau, M., Rutström, E., 2014. Discounting behavior: a reconsideration. Eur. Econ. Rev. 71, 15–33.

Andreoni, J., Sprenger, C., 2012a. Estimating time preferences from convex budgets. Am. Econ. Rev. 102 (7), 3333–3356.

Andreoni, J., Sprenger, C., 2012b. Risk preferences are not time preferences. Am. Econ. Rev. 102 (7), 3357–3376.

Andreoni, J., Kuhn, M., Sprenger, C., 2015. On measuring time preferences. J. Econ. Behav. Organ. 116, 451–464.

Apesteguia, J., Ballester, M., 2015. Monotone stochastic choice models: the case of risk and time preferences. Barcelona GSE Working Paper 859.

Ashraf, N., Karlan, D., Yin, W., 2006. Tying odysseus to the mast: evidence from a commitment savings product in the Philippines. Q. J. Econ. 121 (2), 635–672.

Augenblick, N., Niederle, M., Sprenger, C., 2015. Working over time: dynamic inconsistency in real effort tasks. Q. J. Econ. 130 (3), 1067–1115.

Bryan, G., Karlan, D., Nelson, S., 2010. Commitment devices. Annu. Rev. Econ. 2 (1), 671–698.

Carvalho, L., Meier, S., Wang, S., 2016. Poverty and economic decision-making: evidence from changes in financial resources at payday. Am. Econ. Rev. 106 (2), 260–284.

Chabris, C., Laibson, D., Schuldt, J., 2008. Intertemporal choice. Palgrave Dict. Econ..

Chakraborty, A., Calford, E., Fenig, G., Halevy, Y., 2015. External and Internal Consistency of Choices Made in Convex Time Budgets. Unpublished manuscript.

Coller, M., Williams, M., 1999. Eliciting individual discount rates. Exp. Econ. 2 (2), 107–127.

Collins, D., Morduch, J., Rutherford, S., Ruthven, O., 2009. Portfolios of the Poor: How the World's Poor Live on $2 a Day. Princeton University Press, New Jersey.

Cubitt, R., Read, D., 2007. Can intertemporal choice experiments elicit time preferences for consumption? Exp. Econ. 10 (4), 369–389.

Dean, M., Sautmann, A., 2016. Credit Constraints and the Measurement of Time Preferences. Unpublished manuscript.

Dercon, S., 2002. Income risk, coping strategies, and safety nets. World Bank Res. Obs. 17 (2), 141–166.

Dupas, P., Robinson, J., 2013. Why don't the poor save more? Evidence from health savings experiments. Am. Econ. Rev. 103 (4), 1138–1171.

Epper, T., 2016. Income expectations, limited liquidity, and anomalies in intertemporal choice. University of St. Gallen, School of Economics and Political Science Discussion Paper No. 2015-19.

Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference: a critical review. J. Econ. Lit. 40 (2), 351–401.

Giné, X., Goldberg, J., Silverman, D., Yang, D., 2016. Revising commitments: field evidence on adjustment of prior choices. Econ. J.. http://dx.doi.org/10.1111/ecoj.12378.

Green, L., Fristoe, N., Myerson, J., 1994. Temporal discounting and preference reversals in choice between delayed outcomes. Psychon. Bull. Rev. 1 (3), 383–389.

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

Halevy, Y., 2014. Some Comments on the Use of Monetary and Primary Rewards in the Measurement of Time Preferences.

Harrison, G., Lau, M., Williams, M., 2002. Estimating individual discount rates in Denmark: a field experiment. Am. Econ. Rev. 92 (5), 1606–1617.

Harrison, G., Lau, M., Rutstrom, E., 2005. Dynamic consistency in Denmark: A longitudinal field experiment. Unpublished Manuscript.

Harrison, G., 2008. Maximum likelihood estimation of utility functions using stata. UCF Working Paper 06-12.

Harrison, G., Lau, M., Rutström, E., 2013. Identifying time preferences with experiments: Comment. CEAR Working Paper.

Hoddinott, J., Hoel, J., Schwab, B., 2016. Self-control exertion and the expression of time preference: experimental results from ethiopia. J. Econ. Psychol. 52, 136–146.

Janssens, W., Kramer, B., Van der List, M., Pap, D., 2013. Financial and Health Diaries: A Yearlong Weekly Panel of Farming Households in Nigeria and Kenya. Unpublished raw data.

Karlan, D., Linden, L., 2016. Loose knots: Strong versus weak commitments to save for education in Uganda. Yale University Economic Growth Center Discussion Paper No. 1037.

Kirby, K., Herrnstein, R., 1995. Preference reversals due to myopic discounting of delayed reward. Psychol. Sci. 6 (2), 83–89.

Krupka, E., Stephens, M., 2013. The stability of measured time preferences. J. Econ. Behav. Organ. 85, 11–19.

Laibson, D., 1997. Golden eggs and hyperbolic discounting. Q. J. Econ. 112 (2), 443–477.

Meier, S., Sprenger, C., 2015. Temporal stability of time preferences. Rev. Econ. Stat. 97 (2), 273–286.

Mullainathan, S., Shafir, E., 2013. Scarcity: why Having too Little Means so Much. Macmillan.

Noor, J., 2009. Hyperbolic discounting and the standard model: eliciting discount functions. J. Econ. Theory 144 (5), 2077–2083.

Read, D., Frederick, S., Airoldi, M., 2012. Four days later in Cincinnati: longitudinal tests of hyperbolic discounting. Acta Psychol. 140 (2), 177–185.

Sayman, S., Öncüler, A., 2009. An investigation of time inconsistency. Manag. Sci. 55 (3), 470–482.

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

Sprenger, C., 2015. Judging experimental evidence for dynamic inconsistencies. Am. Econ. Rev. 105 (5), 280–285.

Thaler, R., 1981. Some empirical evidence on dynamic inconsistency. Econ. Lett. 8 (3), 201–207.

Thaler, R., Benartzi, S., 2004. Save more tomorrow using behavioral economics to increase employee saving. J. Political Econ. 112 (1), 164–187.