

Balakrishnan, Uttara; Haushofer, Johannes; Jakiela, Pamela

Working Paper

How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments

IZA Discussion Papers, No. 9653

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Balakrishnan, Uttara; Haushofer, Johannes; Jakiela, Pamela (2016) : How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments, IZA Discussion Papers, No. 9653, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<http://hdl.handle.net/10419/130326>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 9653

**How Soon Is Now?
Evidence of Present Bias from
Convex Time Budget Experiments**

Uttara Balakrishnan
Johannes Haushofer
Pamela Jakiela

January 2016

How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments

Uttara Balakrishnan

University of Maryland

Johannes Haushofer

Princeton University

and Busara Center for Behavioral Economics

Pamela Jakiela

University of Maryland

and IZA

Discussion Paper No. 9653

January 2016

IZA

P.O. Box 7240

53072 Bonn

Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments*

Empirically observed intertemporal choices about money have long been thought to exhibit present bias, i.e. higher short-term compared to long-term discount rates. Recently, this view has been called into question on both empirical and theoretical grounds, and a spate of recent findings suggest that present bias for money is minimal or non-existent when one allows for curvature in the utility function and transaction costs are tightly controlled. However, an alternative interpretation of many of these findings is that, in the interest of equalizing transaction costs across earlier and later payments, small delays were introduced between the time of the experiment and the soonest payment. We conduct a laboratory experiment in Kenya in which we elicit time and risk preference parameters from 291 participants, using convex time budgets and tightly controlling for transaction costs. We make the soonest payments truly immediate, using the Kenyan mobile money system M-Pesa to make real-time transfers to subjects' phones. We find strong evidence of present bias, with estimates of the present bias parameter ranging from 0.901 to 0.937. This result suggests that present bias for money does in fact exist, but only for truly immediate payments.

JEL Classification: C91, D90, O12

Keywords: discount rate, present bias, experiment, mobile money

Corresponding author:

Pamela Jakiela
Department of Agricultural and Resource Economics
University of Maryland
2200 Symons Hall,
College Park, MD 20742
USA
E-mail: pjakiela@umd.edu

* We are grateful to Channing Jang, James Vancel, and the staff of the Busara Center for Behavioral Economics for excellent research assistance, and to Pascaline Dupas, Ray Fisman, Jess Goldberg, Shachar Kariv, Maggie McConnell, Owen Ozier, Charlie Sprenger, and to numerous conference and seminar participants for helpful comments. This research was supported by Cogito Foundation Grant R-116/10 and NIH Grant R01AG039297 to Johannes Haushofer.

1 Introduction

How people trade off immediate and delayed consumption is a question of fundamental importance in economics (von Böhm-Bawerk 1890, Fisher 1930). The canonical economic model of time preferences is the discounted utility model, first proposed by Samuelson (1937); in it, all future payments are discounted by a constant factor each period, leading to exponential discounting.¹ In the second half of the 20th century, the discounted utility model was called into question by the finding that empirically observed discounting behavior, both in animals and humans, did not correspond to the predictions of exponential discounting; in particular, short-term discount rates were found to be higher than long-term discount rates (Ainslie 1975, Thaler 1991). These findings led to the development of alternative models of intertemporal tradeoffs in which agents are present-biased in the sense that they overweight immediate payoffs relative to those that occur in the future.² In economics, the most widely used example is the quasi-hyperbolic model, first proposed by Phelps and Pollak (1968) and adapted to the case of time preferences by Laibson (1997) and O’Donoghue and Rabin (1999).³ The quasi-hyperbolic model has since been used to explain empirical phenomena ranging from retirement saving (Laibson, Repetto, Tobacman, Hall, Gale, and Akerlof 1998) to gym attendance (Acland and Levy 2015, DellaVigna and Malmendier 2006). Present bias is important in a range of policy settings because it predicts preference reversals: agents who exhibit present bias will make consumption and savings plans that they fail to carry out; more generally, present-biased agents tend to invest less than they intend to in goods that yield long-run benefits (e.g. education and exercise), and to consume more than they intend to when goods are associated with future costs (e.g. unhealthy foods).

In recent years, just as present bias has begun receiving widespread attention from policy-makers (cf. World Bank 2015), many scholars have come to question the experimental evidence documenting violations of the discounted utility model. On the one hand, as many have pointed out, it is not clear that we should observe present bias in decisions about money — even if humans are present-biased. Utility is defined over consumption, so if subjects are able to borrow and save, intertemporal tradeoffs over dated money payoffs should depend on market interest rates, not individual preferences (Coller and Williams 1999, Augenblick, Niederle, and Sprenger 2015, Dean and Sautmann 2014). Experimental economists, in contrast, have long argued that experimental subjects “narrowly bracket” their decisions in the lab, viewing dated monetary payoffs as though they were a consumption plan, and numerous experimental studies have supported this view (Rabin and Weizsacker 2009). However, recent evidence has called narrow bracketing assumption into

¹In other words, consumption that occurs t periods in the future is discounted by a factor δ^t , where $\delta \leq 1$ and does not vary over time. See Frederick, Loewenstein, and O’Donoghue (2002) for an overview of the development of discounted utility model and its use in economics.

²In the discounted utility model, agents care more about immediate payoffs than about payoffs that occur k days in the future, but only as much as they care more about payoffs at time t than payoffs at time $t + k$.

³Within psychology, the most widely used model of present bias is the modified hyperbola (Kirby 1997). In that model, utility takes the form: $U(c_t) = \frac{1}{1+kt}u(c_t)$.

question. For example, Augenblick, Niederle, and Sprenger (2015) find evidence of present bias in effort tasks, but no sign of present bias in decisions about money. Dean and Sautmann (2014) find that intertemporal tradeoffs in their lab-in-the-field experiment are associated with both expenditure shocks and savings, suggesting that narrow bracketing fails in their data. These results have sparked a lively debate, with some scholars arguing that choices in time preference experiments are driven primarily by liquidity constraints and interest rates outside the lab (Dean and Sautmann 2014, Epper 2015, Carvalho, Meier, and Wang forthcoming), while others maintain that there is little evidence that agents integrate moderately-sized monetary payoffs into their optimal lifetime consumption plan through smoothing and arbitrage (Halevy 2014, Halevy 2015).

Paralleling this rising chorus of theoretical objections, there is mounting concern that standard experimental designs used to measure time preferences may be confounded. For example, Frederick, Loewenstein, and O’Donoghue (2002) point out that most experimental studies documenting present bias among humans ask subjects to choose between smaller, immediate cash payoffs — which are typically given out at the end of the experimental session — and larger, delayed payoffs. If subjects are not sure that they will actually receive the later payments, or if collecting delayed payments involves larger transaction costs (because, for example, subjects would need to return to the lab to pick up a check), they may appear present-biased when in fact they are not. Another concern is that many experiments assume that utility is linear in money; such an assumption will lead to over-estimates of the degree of present bias if subjects are risk averse — because the utility difference between larger future payments and smaller immediate payments is not as large as the dollar difference between the payoff amounts (Andersen, Harrison, Lau, and Rutström 2008).

In an attempt to address many of these methodological issues, Andreoni and Sprenger (2012) introduce a novel experimental design — the convex time budget (CTB) experiment. In a CTB experiment, a subject divides an endowment between two time periods subject to a budget constraint and an interest rate that makes the delayed payout date relatively attractive. Because subjects are not restricted to the endpoints of the budget line, this method allows for separate estimation of the time preference parameters and the curvature of the utility function (given sufficient data). Under the right conditions, CTB experiments also allow for explicit tests of the hypothesis that subjects are arbitraging between lab and non-lab savings technologies (as we would expect if they were integrating experimental payments into an optimal forward-looking consumption plan). Andreoni and Sprenger (2012) conduct CTB experiments in a university lab setting that allows them to take a number of steps to equalize transaction costs and uncertainty across time periods. Importantly, they make same-day payments using the same technology as delayed payments (checks in campus mailboxes). After introducing such protocols, they find no evidence of present bias among university undergraduates, casting further doubt on the existence of present bias over money payments.

One concern with many studies focused on equalizing transaction costs across immediate delayed payments is that the steps taken to do so also introduce a small front-end delay. For example, as

discussed above, Andreoni and Sprenger (2012) make “immediate” payments by placing a personal check in each subject’s mailbox before the close of the business day.⁴ Thus, “immediate” payments may not always be accessible immediately. If subjects do, in fact, have preferences consistent with the quasi-hyperbolic model, it is possible that they may view such almost-immediate payments as “later” rather than “now” — in which case, some of the recent failures to reject the discounted utility model may be attributable to the use of under-powered experimental tests.

We report the results of a series of convex time budget experiments that we conducted at the Busara Center for Behavioral Economics in Nairobi, Kenya. We make two important methodological departures from previous CTB studies such as Andreoni and Sprenger (2012) and Giné, Goldberg, Silverman, and Yang (2012). First, and most importantly, we make the soonest available payment truly immediate by capitalizing on Kenya’s mobile money payment system, M-Pesa, which makes it possible to send payments to participants in real-time. Specifically, same-day payouts were delivered to participants through their phones before they left the experimental session. M-Pesa payments are widely accepted throughout Kenya, and could be converted to cash by walking into a shop across the street from the experimental lab. Thus, we are able to equalize transaction costs and uncertainty by delivering both immediate and delayed payments through M-Pesa — while making “now” more immediately accessible than in previous CTB experiments.⁵

Our second methodological innovation is the development of a user-friendly touchscreen computer interface that allows us to collect a large data set of 48 CTB and 24 multiple price list decisions from every subject — while working in a population that is substantially less affluent, educated, and elite than standard subject pools of students at top universities in the U.S. and Europe. This allows us to estimate preference parameters at the individual level (including the curvature of the utility function, eliminating one of the confounds discussed above), and to explore the association between liquidity constraints and estimated preference parameters — in a population of economically-independent adults characterized by substantial heterogeneity in terms of socioeconomic status and involvement in the credit market. Moreover, the stakes in our experiment were large in terms of purchasing power: the median total payout was more than four times the median level of daily expenditure. Thus, subjects had every incentive to think carefully about their decisions, and our design provides a powerful test of the extent of arbitrage between lab and non-lab savings vehicles.

We report four main results. First, and most importantly, our results suggest a substantial degree of present bias over money. Our preferred empirical specifications suggest that delayed payoffs are discounted by between 6.3 and 9.9 percent relative to immediate payoffs (i.e. estimated β parameters range from 0.901 to 0.937 in specifications that control for background consumption).

⁴Other studies in a similar vein involve even greater delays. For example, in Giné, Goldberg, Silverman, and Yang (2012), the soonest payments occur one day after decisions are made.

⁵As discussed below, all subjects in our experiment had received payments from the Busara Center in the past, and the overwhelming majority reported an extremely high degree of confidence that all payments would arrive on schedule.

Our design allows us to control for risk aversion and individual-level variation in background consumption (proxied by average daily expenditures), both of which are empirically important — ignoring either factor could bias one’s estimate of the level of present bias. When we estimate parameters at the individual level, we find that just over 60 percent of our subjects have estimated β_i parameters below 1. This is consistent with our pooled estimates in that it suggests that present bias is common, but highlights the importance of individual heterogeneity in preferences. We find that the overwhelming majority of subjects who display a statistically significant degree of (present or future) bias at the individual level are present-biased; only a handful display a significant level of future bias.

Our second finding is that individual time preference parameters are not significantly related to measures of liquidity constraints, suggesting that such constraints are not a plausible alternative account of our findings. Neither having a job nor having a substantial amount saved (in a bank savings account) predicts the degree of patience or present-bias. Moreover, subjects do not display any tendency to shift experimental payouts toward days when they anticipate having limited cash-on-hand. Thus, it is highly unlikely that we are falsely ascribing to present bias patterns of behavior that are actually driven by liquidity constraints.

Third, we find that subjects who are not liquidity-constrained do not engage in the sorts of arbitrage we would expect if they were integrating their experimental payments into an optimal forward-looking consumption and savings plan. The overwhelming majority of subjects who hold liquid savings choose interior allocations in CTB decision problems which offer gross interest rates over 100 percent — well above those available through the credit market. Thus, they do not fully exploit the investment opportunities offered by the experiment, even though doing so would increase the net present value of their income (and therefore consumption) stream.

Fourth, we show that time preference parameters estimated using data from our CTB task are strongly and significantly correlated with analogous parameters estimated using data from a multiple price list task (without controlling for risk aversion). This finding contrasts with the results reported in Andreoni and Sprenger (2012), though our implementation of our CTB and MPL tasks was much more similar in terms of the visual display and computer interface.⁶ In our view, this finding helps to confirm the validity of the MPL design, though the CTB task generates a much richer data set and provides substantially more precise parameter estimates.

Taken together, our results demonstrate that present bias over money is not simply an artifact of experimental design flaws in previous studies; we find strong evidence of present bias in an experiment that controls for risk aversion, using protocols that equalize transaction costs and payment modalities across all possible payout dates. Moreover, our study provides clear evidence that subjects — specifically, a diverse sample of adults in a lower middle-income country — are not arbitraging between lab and outside savings vehicles; we also find no evidence that choices in

⁶Andreoni and Sprenger (2012) use a purpose-built computer interface to conduct their CTB experiments, but collect MPL data using a standard pencil-and-paper questionnaire. As discussed below, our CTB and MPL computer interfaces are extremely similar.

our experiment are driven by liquidity constraints. Taken together, the results support the view that individual choices in time preference experiments are driven by time preferences, not market interest rates.

The remainder of the paper is structured as follows. Section 2 describes the design and implementation of the study. Section 3 presents our theoretical framework and derives testable predictions. Section 4 presents our experimental main results. Section 5 discusses the relationship and other recent time preference experiments, and presents additional analysis designed to test for both arbitrage (between lab and non-lab savings technologies) and narrow bracketing of experimental payments. Section 6 concludes.

2 Experimental Design and Procedures

2.1 Experimental Design

We employ the convex time budget (CTB) design first utilized by Andreoni and Sprenger (2012). In a CTB experiment, each subject divides a budget across two payoff dates subject to the present-valued budget constraint:

$$c_t + \frac{c_{t+k}}{(1+r)} = m. \quad (1)$$

In this framework, t denotes the front-end delay, the number of days between the experiment and the earlier payoff date; and k denotes the delay between the earlier and later payoff dates. The CTB design has a number of advantages over more traditional discrete choice approaches to eliciting time preferences. First, choices from continuous budget sets contain more information than discrete (typically binary) choices — in each decision, a utility-maximizing subject reveals her most preferred allocation relative to a continuum of alternatives, not just a single less-preferred option. This additional information allows us to estimate both risk and time preferences parameters using data from a single experiment.⁷ Finally, CTB experiments can provide evidence that intertemporal tradeoffs in the lab are not driven by market interest rates and individual liquidity constraints (as opposed to time preferences) — if subjects who are not liquidity-constrained choose interior allocations when the interest rates offered through the experiment exceed those that would be available outside the experiment.⁸

Subjects in our experiment faced a total of 48 CTB decision problems, each of which was presented using a user-friendly touchscreen computer interface.⁹ This design generates an extremely

⁷As Andersen, Harrison, Lau, and Rutström (2008) point out, risk averse subjects will appear more impatient than they actually are if one ignores the issue of diminishing marginal utility when estimating discount rates.

⁸In other words, if subjects with access to credit markets were treating opportunities to save in the lab as part of a broader financial portfolio, they should fully exploit the above-market interest rates offered through the experiment (unless they are liquidity-constrained). See Coller and Williams (1999) for an early discussion of the issue. Meier and Sprenger (2010) find little evidence that experimentally-measured discount rates are predicted by liquidity constraints outside of the lab.

⁹The experimental interface was programmed using z-tree (Fischbacher 2007). Complete instructions and screenshots of the computer interface are included in the Online Appendix.

rich data set and allows us to estimate preference parameters at the individual level. The CTB decisions included in our experiment were organized into eight sets of six decision problems. The earlier and later payoff dates were fixed within each set. The front end delay was either 0, 14, or 28 days after the experimental session, and the delay between payoffs was either 14 or 28 days.¹⁰ Within each decision set, the maximum earlier payoff was fixed at either 400 or 600 Kenyan shillings.¹¹ Thus, each decision set corresponded to a triple: the front-end delay, the delay between payoffs, and the maximum earlier payoff.¹² The eight decision sets were presented in a random order.

Within each decision, the maximum later payoff depended on the gross interest rate, $1 + r$: reducing the earlier payoff (c_t) by one shilling meant increasing the later payoff (c_{t+k}) by $1 + r$ shillings. Within each set of decisions, subjects faced six gross interest rates: 1.1, 1.25, 1.75, 2, 3, and 4. Gross interest rates always appeared in increasing order within a decision set to minimize the potential for confusion. The Online Appendix lists the front-end delay, delay between payoffs, budget size, and gross interest rate for each of the 48 CTB decisions included in our experiment.

At the end of the CTB portion of the experiment, subjects completed a standard Multiple Price List (MPL) task that included 24 decision problems. In each MPL decision problem, a subject chooses between a smaller, earlier payoff and a larger, later payoff — so the MPL choice can be viewed as the restriction of a CTB decision problem to the endpoints of the budget line. MPL decision problems were organized into four sets of six decisions. Within each set, the earlier and later payoff dates and the earlier payoff amount were fixed; the later payoff amount increased over the course of the decisions within a set, with the later payoff amounts corresponding to the six gross interest rates included in the CTB decision problems (1.1, 1.25, 1.75, 2, 3, and 4). In MPL experiments such as this, we expect all but the most impatient subjects to eventually switch to preferring the delayed payment as the implied interest rate increases.¹³ In the four sets of MPL decisions included in our experiment, the front-end delay was either 0 or 14 days and the delay between payoffs was either 14 or 28 days. Thus, the MPL tasks covered a subset of the payoff dates, budget sizes, and interest rates included in the CTB experiment.

At the end of the experiment, one of the 72 decision problems (48 CTB decisions and 24 MPL decisions) was randomly chosen to determine final payoffs. This randomization was done separately for each subject, guaranteeing that all information on the timing and size of experimental payoffs

¹⁰We follow Andreoni and Sprenger (2012) in ensuring that all payoff dates occur on the same day of the week to eliminate any end-of-week confounds. As discussed below, we also made sure that payoff dates did not fall on holidays or the last day of the month.

¹¹These budgets are equivalent to approximately 4.08 and 6.12 USD, respectively, using the average exchange rate over the period during which the experiments took place. These endowments are large in purchasing power terms: the median level of daily expenditures in our sample is 150 Kenyan shillings (1.53 USD)

¹²We presented all possible combinations of the three front-end delays and the two delays between payoffs for the budget size (i.e. maximum earlier payoff) of 400 Kenyan shillings (4.08 USD). In addition, we included two decision sets in which the budget size was increased to 600 Kenyan shillings (6.12 USD). In these decisions, the front-end delay was either 0 or 14 days and the delay between payments was fixed at 14 days.

¹³In fact, relatively patient subjects may prefer the delayed payment in all six MPL decisions within a set.

remained private. In addition to their payoffs from the experiment, subjects received a fixed show-up fee which was evenly divided between the earlier and later payoff dates — so every subject, including those who chose corner solutions in the CTB decisions and those whose payoffs were determined by an MPL decision, received two dated payoffs.¹⁴ We describe the procedures used to deliver payments to subjects in detail below.

2.2 Experimental Procedures

The experiment was conducted between April and July of 2015 at the Busara Center for Behavioral Economics in Nairobi, Kenya. Subjects were drawn from two of Nairobi’s informal settlements, Kawangware and Kibera. Our sample includes data from 291 adult subjects.¹⁵ Table 1 reports summary statistics on the subjects in our sample.

Experimental sessions were conducted in a dedicated computer lab at the Busara Center.¹⁶ Instructions were presented orally in Swahili, one of Kenya’s official languages and a local *lingua franca*.¹⁷ Our user-friendly touchscreen interface was programmed using z-tree (Fischbacher 2007), and was intended to be easily comprehensible by subjects with limited levels of formal education. As discussed above, the dates of the earlier and later payoffs were fixed within each decision set in both the CTB and MPL portions of the experiment; these were announced aloud before subjects began making decisions within a given set. The dates also appeared in large font on the computer screen for each decision — the earlier payoff date on the left and the later payoff date on the right. In the CTB decision problems, the maximum possible payoffs at each date appeared directly below the dates on the screen. Subjects shifted money from the earlier to the later payoff date (or vice versa) by sliding their finger along a brightly colored touchscreen bar. The amounts allocated to the earlier and later payoff dates initially displayed as zeros; these amounts updated every time a subject touched the colored bar. After arriving at a desired allocation, the subject touched an “OK” button to confirm her choice and proceeded to the next decision. During the MPL task, the screen displayed the amounts associated with each payoff date directly below the date, and subjects touched a colored button to select their preferred payment and date. Full

¹⁴This approach is also taken by Andreoni and Sprenger (2012). Haushofer (2014) presents a theoretical model suggesting that a mental cost of keeping track of time-dated payments may act as an additional (cognitive) transaction cost, pushing subjects toward corner solutions and immediate payoffs when the show-up fee is paid (in its entirety) on the day of the experiment.

¹⁵MPL decision data is missing for six subjects because of a computer malfunction.

¹⁶No experimental sessions were held on Fridays or weekends to avoid any potential end-of-week effects. When considering a potential date for a session, we verified that no payout dates associated with that (potential) session fell on holidays or any other day that would lead to a foreseeable change in the desire for cash on hand (for example, the day when school fees are due).

¹⁷Rigorous translation procedures were used to ensure fidelity to the intended meaning of the instructions. The investigators first worked with experienced members of the Busara Center staff to refine the English instructions to produce a version that would make sense when translated into colloquial Kenyan Swahili. This multilingual team then translated the English and reviewed the translation as a group. The translated instructions were then sent to a separate team (not involved in administering the experiments); this team produced a back-translation to English that was then checked for equivalence to the original English text. English instructions are included in the Online Appendix; Swahili versions are available from the authors upon request.

experimental instructions and screenshots of the computer interfaces used for the CTB and MPL decision problems are included in the Online Appendix.

At the end of the experiment, the computer randomly selected one decision for payment. Payments from the chosen decision were added to a show-up fee, which was divided evenly between the earlier and later payoff dates (from the decision that was chosen to determine the final payout). Thus, all subjects received two dated payments. Subjects completed a short sociodemographic survey and then learned the dates and amounts of their final payoffs before departing from the lab. If the decision chosen to determine final payoffs involved an immediate payment, the payment was sent to the subject's phone before she left the session.

All payments in our experiment — including payments made on the day of the experiment — were made using the M-Pesa mobile money technology. M-Pesa is a money transfer service operated by Kenya's largest mobile phone company, Safaricom. Users can send and receive transfers and make direct payments to firms using their phones, and they can also withdraw cash from their M-Pesa accounts at over 80,000 M-Pesa agents throughout the country. All subjects in our experiment have active M-Pesa accounts, and all had received transfers from the Busara Center via M-Pesa prior to the experiment. Payments were timed to occur no more than two hours after the start of the experimental session — for example, if the experimental session started at ten o'clock in the morning, the instructions clearly stated that all payments would be sent no later than noon on the relevant payment date. This meant that subjects receiving payments on the day of the experiment would receive them before they departed from the Busara Center (while they were completing the sociodemographic survey). The implication is that transaction costs for immediate and delayed payments were equalized — all payments were sent from a trusted source (the Busara Center) via the familiar M-Pesa technology — but same day payments were truly immediate in the sense that they could be spent as soon as the experiment was over.

This payoff method addresses an important concern with many experimental studies of intertemporal preferences: behavior that appears present-biased may in fact be driven by differential transaction costs or uncertainty — for example, if subjects realize that they will need to cash a check or return to the lab to collect cash if they choose a delayed (rather than immediate) payment (Frederick, Loewenstein, and O'Donoghue 2002). We follow several recent studies (cf. Andersen, Harrison, Lau, and Rutström 2008, Andreoni and Sprenger 2012, Giné, Goldberg, Silverman, and Yang 2012) in equalizing the financial and cognitive transaction costs associated with immediate and delayed payments. For example, we adopt Andreoni and Sprenger's (2012) approach of dividing the show-up fee into two equal, dated payments; and we use a payment technology (M-Pesa) with which subjects are extremely familiar. Importantly, however, our study differs from Andersen, Harrison, Lau, and Rutström (2008), Andreoni and Sprenger (2012), and Giné, Goldberg, Silverman, and Yang (2012) because the steps taken to equalize transaction costs do not necessitate any delay in payments — subjects receive any payments made on the day of the experiment before

they depart from the lab, and are able to begin spending those payments immediately.¹⁸

3 Theoretical Framework

Each subject divides an endowment of $m > 0$ between two accounts: one associated with an earlier payoff date ($t \geq 0$ days in the future) and one associated with a later payoff date ($t + k > 0$ days in the future). We assume subject i maximizes her (additively separable) utility

$$u(c_t, c_{t+k}) = \begin{cases} u(c_t + \omega_t) + \beta \delta^k u(c_{t+k} + \omega_{t+k}) & \text{if } t = 0 \\ u(c_t + \omega_t) + \delta^k u(c_{t+k} + \omega_{t+k}) & \text{if } t \neq 0 \end{cases} \quad (2)$$

subject to the budget constraint

$$c_t + \frac{c_{t+k}}{(1+r)} = m. \quad (3)$$

ω_t denotes background consumption (consumption from outside the experiment) in period t ; this will be equal to zero if subjects narrowly bracket their decisions in the experiment (Rabin and Weizsacker 2009). The tangency condition is given by

$$\frac{u'(c_t + \omega_t)}{u'(c_{t+k} + \omega_{t+k})} = \beta \delta^k (1+r). \quad (4)$$

When individual preferences are dynamically consistent, the optimal c_t^* depends on the size of the budget (m), the gross interest rate (r), and the delay between payoffs (k), but not on the front end delay (t). Substituting the budget constraint into the tangency condition demonstrates that this result holds whenever ω_t is constant across periods and $\beta = 1$: the optimal c_t^* given m , r , and k does not depend on t . However, the exponential discounting model is the only model that generates dynamically consistent choices. In the equation above, when $\beta < 1$, the optimal c_t^* given m , r depends on t : the optimal allocation to the earlier period is higher when $t = 0$ than for all $t > 0$. Such changes in the optimal c_t^* as t changes are termed static preference reversals.¹⁹

If we assume that utility takes the constant relative risk aversion (CRRA) form such that

$$u(c_t) = \frac{c_t^{1-\rho}}{1-\rho}, \quad (5)$$

¹⁸Because subjects had experience receiving mobile money payments from the Busara Center via M-Pesa, there is little reason to be concerned that they doubted that their payments would arrive on time. When asked (at the end of the experiment) whether they thought that both of their experimental payments would arrive on time, 98 percent of subjects answered in the affirmative.

¹⁹Our experiment is not designed to distinguish between the quasi-hyperbolic model of present bias and alternative models, for example, the hyperbolic model described in Kirby (1997), that allow the discount factor to vary as a continuous function of the front-end delay in the neighborhood of 0. We consider only three levels of front-end delay: 0, 14, and 28 days. Our experiment is therefore not powered to test whether small front-end delays elicit intermediate levels of present bias. However, the results reported in Andreoni and Sprenger (2012) are not consistent with this.

then Equation 4 can be rewritten as

$$\frac{c_t}{c_{t+k}} = \left[\beta \delta^k (1+r) \right]^{-1/\rho}. \quad (6)$$

We can also solve for the demand function for c_t :

$$c_t = \frac{[\beta \delta^k (1+r)]^{-1/\rho} (1+r)m - \omega_t + [\beta \delta^k (1+r)]^{-1/\rho} \omega_{t+k}}{1 + [\beta \delta^k (1+r)]^{-1/\rho} (1+r)} \quad (7)$$

which reduces to

$$c_t = \frac{[\beta \delta^k (1+r)]^{-1/\rho} (1+r)m}{1 + [\beta \delta^k (1+r)]^{-1/\rho} (1+r)} \quad (8)$$

when $\omega_t = \omega_{t+k} = 0$. The parameters β , δ , and ρ can then be estimated by non-linear least squares (NLS). The limitation of the NLS approach is that it does not adjust for censoring of c_t at 0 and m . Andreoni and Sprenger (2012) suggest a complementary approach: taking logs of the Equation 6 allows one to estimate the model parameters in a two-limit Tobit framework that corrects for censoring.²⁰

4 Analysis

4.1 Comprehension and Consistency

An important question in all preference elicitation experiments is whether subjects make coherent choices that are consistent with utility maximization. Such concerns are particularly salient in our context, because our subjects have less education and experience with computers than typical experimental subject pools composed of university students. We take two approaches to assessing the consistency of subjects' choices. First, we test for choices for consistency with the Generalized Axiom of Revealed Preference (GARP); GARP provides a direct test of whether individual choices can be rationalized by a utility function that is continuous, increasing, concave, and piecewise linear (Afriat 1967, Varian 1982, Varian 1983). The majority of subjects in our experiment do not violate GARP. However, because of the small number of intersecting budget lines in our experiment, the power of our GARP test is somewhat limited. We therefore adopt a second approach: testing the extent of adherence to the law of demand. Taken together, both approaches suggest that subjects understood the experiment and made consistent decisions that can be viewed through the lens of utility maximization.

²⁰Similar procedures can be used to derive the demand function assuming a constant absolute risk aversion (CARA) utility function, which takes the form $u(c_t) = -e^{-\alpha c_t}$. Given this utility function, demand for c_t is given by:

$$c_t = \frac{1}{(2+r)} \left[\frac{\ln(\beta)}{-\alpha} + \frac{k \ln(\delta)}{-\alpha} + \frac{\ln(1+r)}{-\alpha} + m(1+r) + w_t - w_{t+k} \right]$$

As a robustness check, we also report estimates of δ and β that assume consumption utility takes the CARA form.

4.1.1 Rationality

One of the most important questions one can ask about individual decision data is whether choices are consistent with the utility maximization. When budgets are linear, revealed preference theory offers a direct test of rationality: choices can be rationalized by a utility function that is well-behaved (in the sense of being continuous, increasing, concave, and piecewise linear) if and only if they satisfy GARP (Afriat 1967).

Our experimental design contrasts with many time preference experiments because — by varying both the (early-valued) budget size and the gross interest rate for fixed pairs of earlier and later payoff dates (i.e. fixed values of t and $t + k$) — we confront subjects with sets of intersecting budget lines that create the possibility of violating GARP. Specifically, our experiment includes two sets of 12 intersecting budget lines in which the earlier and later payoff dates are fixed but the budget size and interest rate vary across decision problems.²¹ Our tests of consistency with GARP are based on choices in these 24 decision problems.

Unfortunately, the power of revealed preference tests based on relatively small numbers of decisions (i.e. intersecting budget lines) is often quite low (Choi, Fisman, Gale, and Kariv 2007b, Andreoni, Gillen, and Harbaugh 2013). Our design creates, in essence, two separate tests of GARP, each of which involves 12 decision problems; though we are able to add the number of violations across the two sets of intersecting budget lines, two GARP tests involving only 12 decisions may not generate sufficient power to be confident that one will detect violations of rationality. To assess the power of our GARP test, we follow the standard approach, which builds on Becker (1962) and Bronars (1987), generating a population of 1,000 simulated subjects who choose points at random from each budget line (according to a uniform distribution). 86.9 percent of these simulated subjects violate GARP at least once, suggesting that our revealed preference test does, in fact, have a reasonable level of power. The median number of violations in the sample of simulated subjects is 8. In contrast, 63.2 percent of subjects in our experiment never violate GARP, and only 17.2 percent have 8 or more violations. Figure 1 presents histograms of the distributions of GARP violations in our actual and simulated samples. Though the power of our GARP test is lower than in some recent experiments (cf. Fisman, Jakiela, Kariv, and Markovits 2015), the evidence suggests that our subjects are substantially closer to consistency with utility maximization than could occur at random.

4.1.2 Adherence to the Law of Demand

To further gauge the extent to which subjects in our experiment made meaningful and consistent choices, we follow Giné, Goldberg, Silverman, and Yang (2012) in examining “basic consistency”

²¹Decision Sets 1 and 4 in the CTB portion of our experiment both involved a front-end delay of 14 days and a delay between payments of 14 days, but the maximum earlier payout in Decision Set 4 was 600 (rather than 400) Kenyan shillings (i.e. 6.12 USD as opposed to 4.08 USD). Similarly, Decision Sets 3 and 7 both involved a front-end delay of 0 days and a delay between payments of 14 days, but the maximum earlier payout in Decision Set 7 was 600 (as opposed to 400) Kenyan shillings.

— a measure of adherence to the law of demand. The idea underlying basic consistency is that, for a fixed t and k (i.e. fixed earlier and later payoff dates), an increase in the gross interest rate is equivalent to a decrease in the price of consumption in the later period. So, if we consider two interest rates, r' and r'' such that $r'' > r'$, the amount allocated to the later period should be at least as large under r'' as under r' .

Subjects in our experiment made 8 sets of 6 CTB decisions. Within each set, the budget size (i.e. the maximum earlier period payoff), t , and k were fixed. Each set of decisions included 6 gross interest rates: 1.1, 1.25, 1.75, 2, 3, and 4. There are therefore 15 possible pairs of interest rates in each set of decisions. For each pair, we generate an indicator for a basic consistency violation that is equal to one if the allocation to the later account is strictly higher under the lower of the two interest rates. We then calculate individual-level frequency of such violations. One minus the frequency of basic consistency violations provides an index of the level of adherence to the law of demand. The median rate of basic consistency is 0.93, suggesting that subjects understood the experiment and were able to implement purposeful choices using the computer interface.

For comparison, we again follow the approach suggested by Bronars (1987), generating a population of simulated subjects who choose points from each budget line randomly (according to a uniform distribution). The median rate of basic consistency among simulated subjects is only 0.72, and only 6.1 percent have basic consistency indices above 0.8 (versus 71.6 percent of actual subjects). Figure 2 compares the distribution of the basic consistency index in our (actual) sample to the simulated distribution. It is clear that a large majority of subjects make choices that are much more consistent than what would occur by chance.

4.2 Aggregate Analysis

With consistency established, we now examine the intertemporal tradeoffs made by subjects in our experiment. Figure 3 presents the average percentage of the budget allocated to the earlier payoff date; the sample is restricted to decisions where the early-valued budget was 400 Kenyan shillings (4.08 USD) to ensure balance across delay lengths. Subjects in our experiment allocate more to the earlier payoff date when the front-end delay is 0 relative to cases where the earlier payoff is not immediate: subjects allocate 47.7 percent of their endowment to the earlier payoff date when earlier payments are immediate, versus 45.5 percent of the budget when the front-end delay is greater than 0 (p-value 0.005). Figure 4 plots the percent of the (present-valued) budget allocated to the earlier payoff date as a function of front-end delay, delay (between payments), and interest rate. Again, the graph suggests some degree of present bias: the amount allocated to the early payment is almost always at least weakly higher when the front-end delay is zero as opposed to 14 or 28 days (it is strictly higher at 8 of 12 data points). Interestingly, the tendency toward present bias appears more pronounced at the higher interest rates.²²

²²This suggests one potential explanation for differences between our results and those reported in both Andreoni and Sprenger (2012) and Augenblick, Niederle, and Sprenger (2015): the highest gross interest rates are substantially

Next, we more formally test for the presence of present bias by estimating Equations 6 and 7, pooling the data from all subjects.²³ Results are reported in Table 2. In Columns 1 through 4, we estimate parameters via non-linear least squares for a range of different values of background consumption. In Column 1 of Table 2, we restrict background consumption to 0 in all periods. We estimate a daily discount rate of 0.992; the parameter is quite precisely estimated, and is significantly different from 1 at the 99 percent confidence level (p-value < 0.001). The estimated β parameter is 0.895; it is significantly different from 1 at the 99 percent confidence level (p-value < 0.001). In Column 2, we allow background consumption to differ from zero but assume that it is constant across periods; we then estimate the level of background consumption as one of the model parameters. Again, we find that the estimated β and δ parameters are significantly different from 1 at the 99 percent confidence level (p-values < 0.001), though the estimated β parameter jumps to 0.920. The estimated background consumption parameter, ω_t , is 276 Kenyan shillings (2.82 USD), and is significantly different from 0 (p-value < 0.001). For comparison, the mean self-reported daily expenditure across all subjects was 204 Kenyan shilling (2.08 USD), and the median was 150 Kenyan shillings (1.53 USD). In Column 3, we estimate separate background consumption parameters for the earlier and later periods (though these, of course, vary across decisions). Again, the estimated β and δ parameters are significantly different from 1 (p-values < 0.001); the estimated β is 0.937. Interestingly, estimated background consumption is significantly higher in the later period than in the earlier period (p-value < 0.001), in spite of the fact that the calendar dates of the earlier and later payouts varied across the 28 experimental sessions. Point estimates suggest that background consumption is approximately 245 Kenyan shillings (2.50 USD) at time t versus 315 Kenyan shillings (3.21 USD) at time $t+k$. Finally, in Column 4, we replace background consumption with self-reported daily consumption expenditure; thus, background consumption varies across individuals but not across periods. The estimated β is 0.910, and it remains significantly different from 1 (p-value < 0.001).

In Columns 5 and 6, we estimate parameters via two-limit Tobit (by taking logs of the tangency condition described by Equation 6). The advantage of the Tobit approach is that it adjusts for censoring of the amount allocated to the earlier period at 0 and m ; the disadvantage is that background consumption cannot be estimated as one of the parameters.²⁴ In Column 5, we set background consumption at 0.01. In Column 6, we use self-reported daily expenditure. In both cases, the estimated β is significantly less than 1 (p-values < 0.001). Point estimates are comparable to those generated via NLS. Thus, across all specifications, we find substantial evidence of present bias: subjects discount delayed monetary payments by roughly 10 percent relative to truly immediate monetary payments. Interestingly, the estimated δ parameter is not significantly different from 1 in the specifications that adjust for censoring. The coefficient is, in fact, estimated to be almost exactly 1, suggesting that subjects tendency to discount future payoffs is driven primarily

lower. For example, the highest interest rate considered in Augenblick, Niederle, and Sprenger (2015) is 1.43.

²³Standard errors are clustered at the subject level in all specifications.

²⁴See Andreoni and Sprenger (2012) for discussion.

by present bias.

4.2.1 Robustness Checks

CARA Utility. To address concerns that our results might be driven by our assumption that consumption utility takes the CRRA, we also report NLS estimates of β and δ derived under the assumption that utility takes the CARA form (Table 3). In Column 1, we restrict background consumption to 0; in Column 2, we estimate the difference in background consumption across periods, $\omega_{t+k} - \omega_t$, as one of the model parameters.²⁵ β estimates are slightly higher under the assumption of CARA utility (relative to CRRA utility), but remain significantly different from 1 (p-values ≤ 0.001). Thus, choices in our CTB experiment with truly immediate payments provide strong evidence of present bias over money, and our results are not dependent on the assumption that consumption utility takes the CRRA form.

Falsification Test. A more general mis-specification concern, suggested by Halevy (2015), is that consumption utility may not be stable over time. One interpretation of this concern is that, in the quasi-hyperbolic model, the β parameter is the only possible way of deviating from exponential discounting; deviations from $\beta = 1$ might then be observed (though not always in the direction of present bias) because they offer the mis-specified quasi-hyperbolic model an additional degree of freedom when explaining individual choice patterns. We explore this possibility through a simple falsification test: we re-estimate our main specifications, but replace β with a parameter λ that enters the discount factor (only) when the front-end delay is 28 days, the longest front-end delay we consider. In other words, we replicate the estimation reported in Table 2 assuming utility takes the form:

$$u(c_t, c_{t+k}) = \begin{cases} u(c_t + \omega_t) + \delta^k u(c_{t+k} + \omega_{t+k}) & \text{if } t \neq 28 \\ u(c_t + \omega_t) + \lambda \delta^k u(c_{t+k} + \omega_{t+k}) & \text{if } t = 28 \end{cases} \quad (9)$$

If our main results are driven by underlying instability in the utility function across payout dates, we might expect that the λ parameter might also be consistently different than one. In results reported in the Online Appendix, we show that this is not the case. λ is sometimes greater than 1 and sometimes less than 1, but we can never reject the hypothesis that $\lambda = 1$. Thus, our results do not appear to be driven by a more general pattern of instability in discounting or utility across all payout dates.

Sample Restrictions. Finally, in the Online Appendix, we report three additional robustness checks: we restrict the sample to (i) those subjects who never violate GARP, (ii) those subjects with basic consistency indices above 0.85 (i.e. those who display a degree of adherence to the law of demand that is never observed among simulated subjects who choose random points on each

²⁵ ω_{t+k} and ω_t are not separately identified under the assumption of CARA utility.

budget line), and (iii) subjects who choose interior allocations more than half the time (suggesting that their choices are driven by preferences rather than interest rates or liquidity constraints). Our main result — that β is significantly less than 1 — is extremely robust across these different subsamples of the data.

4.3 Individual-Level Analysis

We next examine decisions in our experiment at the individual level. As discussed above, an advantage of our experimental design is that it allows us to estimate β and δ at the subject level without needing to assume that the distribution of individual-level parameters takes a specific functional form.²⁶ We estimate subject-level β_i , δ_i , and ρ_i parameters via non-linear least squares while controlling for self-reported background consumption (average daily expenditure).²⁷ Histograms of the estimated individual-level β_i and δ_i parameters are presented in Figure 5. The median individual-level β_i is 0.933, and 169 subjects (60.6 percent) have estimated β_i parameters below 1. Only 74 subjects (26.5 percent) have individual β_i parameters that are significantly different from 1 (at the 95 percent confidence level), but the overwhelming majority of these individuals (61 subjects) are present-biased (i.e. have estimated β_i parameters significantly below 1). Thus, present-biased subjects outnumber future-biased subjects, both in terms of the distribution of point estimates and the frequency of statistically significant deviations from exponential discounting.

Though we observe considerable variation in patience, our subjects tend to discount the future sharply: the median δ_i is 0.992, which implies an annual discount rate above 95 percent. Our results also suggest that present bias and impatience are (weakly) positively correlated (Spearman’s $\rho = 0.176$).

We also observe substantial heterogeneity in risk aversion. The median estimated individual-level CRRA coefficient is 0.936; 21.7 percent of subjects have estimated ρ_i parameters below 0.5, while 47.0 percent of subjects have estimated ρ_i parameters above 1.²⁸ The distribution of estimated ρ_i parameters is quite skewed: 13 subjects have estimated CRRA coefficients above 20.

One alternative account of present bias parameter estimates less than one is that participants face liquidity constraints (in the present), and that these cause subjects to take immediate payments whenever possible. To assess the degree to which such constraints might explain our results, in

²⁶In other words, our functional form assumptions impose restrictions on individual preferences (most notably, we assume that discounting is either quasi-hyperbolic or exponential), but we impose no restrictions on the relationship between individual preference parameters (across subjects). See Choi, Fisman, Gale, and Kariv (2007a), Fisman, Kariv, and Markovits (2007), Choi, Kariv, Müller, and Silverman (2014), and Fisman, Jakiela, and Kariv (2014) for examples of similar estimation approaches that estimate preference parameters at the individual level. This approach contrasts with the estimation strategies used in Andersen, Harrison, Lau, and Rutström (2008), Von Gaudecker, van Soest, and Wengström (2011), and Jakiela and Ozier (forthcoming), who estimate mixed logit models of individual choices.

²⁷Thus, we estimate the specification reported in Column 4 of Table 2 separately for each subject. We are unable to estimate individual parameters for 12 of our 291 subjects. Five subjects always allocated their entire endowment to the earlier payoff date, and five always allocated their entire endowment to the later payoff date. In addition, estimation does not converge for two of the remaining 281 subjects.

²⁸A histogram of the estimated ρ_i parameters is included in the Online Appendix.

Table 5, we report the results of median regressions of the estimated individual-level parameters (β_i , δ_i , and ρ_i) on individual characteristics, including factors likely to be associated with liquidity constraints — regular employment (an indicator for having regular wage work or a self-employment activity that one does every week) and liquid savings (an indicator for having more than 1,000 Kenyan shillings, i.e. 10.20 USD, in the bank). Neither employment nor having more than 1,000 Kenyan shillings in the bank is associated with the estimated model parameters, though we do find that women are significantly more patient and more risk averse (but not more present-biased).²⁹ Together, these results suggest that liquidity constraints are unlikely to account for the pattern of results described above, and in particular the evidence for present bias.

A further question of interest is whether the time preference parameters obtained with the CTB approach are similar to those obtained using the traditional multiple price list (MPL). In Figure 6, we compare the individual-level estimates of the model parameters derived from our CTB experiment to estimates constructed using the data from our MPL task. To construct our MPL estimates of model parameters, we follow the procedures outlined in Andreoni and Sprenger (2012); since our experiment did not include an MPL task designed to elicit risk preferences, we construct estimates of β and δ assuming linear utility (in the MPL task).³⁰ For the 65 subjects who did not make internally consistent choices in the MPL tasks, we follow the standard practice of using the count of the number of times they chose the earlier payoff to assign them to β and δ categories.³¹ The figure shows that CTB and MPL parameter estimates are highly correlated, though the CTB task generates a richer data set and allows for a wider range of parameter values. In Table 5, we report the results of median regressions of the estimated CTB parameters on the analogous MPL parameter estimates. In Columns 1 and 2, we include all subjects for whom individual-level CTB parameter estimates are available; in Columns 3 and 4, restrict the sample to those subjects who made consistent choices in the MPL task. The MPL estimates of both β and δ are significantly related to the CTB estimates of the same parameters in all specifications (all p-values < 0.001).

5 Discussion

Over the last few years, several theoretically-sophisticated, technically-rigorous experiments have sparked a lively debate about the use of lab experimental methods to measure intertemporal trade-offs (cf. Andreoni and Sprenger 2012, Giné, Goldberg, Silverman, and Yang 2012, Augenblick,

²⁹A number of previous studies have also found that women are more risk averse than men; see Croson and Gneezy (2009) for discussion.

³⁰Specifically, we construct estimated of δ using choices in MPL tasks where the earlier payoff was not immediate, and then use the δ estimates to calculate β (which is identified by the change in the frequency of choosing the earlier payoff when the earlier payoff is immediate, relative to MPL tasks where the earlier payoff is not immediate).

³¹As discussed above, subjects who implement decisions in the MPL task without error should choose the earlier payout between zero and six times, and then choose the later payment for the remainder of the six decisions in the set (which have the highest interest rates).

Niederle, and Sprenger 2015, Dean and Sautmann 2014, Halevy 2015, Epper 2015). This body of work raises two critical questions. First, do lab experiments with money payoffs measure time preferences? Utility is defined over consumption, not money, and subjects have access to a range of credit products. So, sophisticated subjects may treat dated experimental payoffs as a (nother) form of credit, integrating their (present-discounted) experimental income into an optimal forward-looking consumption plan — in which case, choices in the lab will reflect market interest rates, not individual preferences. Second, even if one assumes that tradeoffs in experiments with money payoffs are driven by time preferences, should static preference reversals be interpreted as evidence of present bias? As discussed at length in Giné, Goldberg, Silverman, and Yang (2012) and Halevy (2015), this interpretation assumes that subjects have stable utility functions and discount rates — that one’s willingness to make tradeoffs between dated payoffs that arrive 1, 2, or 3 (or 100, 200, or 300) days from “now” does not depend on the calendar date upon which one is asked. Of course, the assumption that preferences are stable is standard in economics; however, even subjects with stable preferences may seek to shift money toward points in time when they expect the marginal utility of money to be high (because, for example, they anticipate needing to make a large payment or being short of income), particularly if credit markets are imperfect.³²

Providing definitive answers to both of these questions is beyond the scope of this paper. However, our experiment does allow for explicit tests of a number of the hypotheses under scrutiny. In what follows, we present these additional pieces of analysis and compare our results to those of other recent studies. We first test whether subjects behave in the manner predicted by standard models of intertemporal optimization, exploiting the high interest rates offered through the experiment to increase their present-discounted income stream. We then explore the issue of narrow bracketing by testing whether subjects shift experimental payments toward dates when they expect to have limited liquidity. We conclude by summarizing our results and suggesting directions for further research.

5.1 Do Subjects Engage in Arbitrage?

Economists typically assume that people have access to perfect credit markets, and that their intertemporal tradeoffs are part of an optimal forward-looking consumption plan characterized by a set of Euler equations. However, most experimental economists — particularly those using lab

³²Stigler and Becker (1977) argue that economists should not accept heterogeneity in preferences (“tastes”) as an explanation for differences in choice behavior across individuals or periods; instead, they suggest that human beings are homogeneous, and that economists should seek price and income based explanations for differences in behavior. However, the hypothesis that individual preferences are heterogeneous (across individuals) has now been confirmed in a large number of controlled lab experiments (cf. Andreoni and Miller 2002, Choi, Fisman, Gale, and Kariv 2007a, Fisman, Kariv, and Markovits 2007, Choi, Kariv, Müller, and Silverman 2014, Fisman, Jakiela, and Kariv 2014). Whether future work will also rule out the hypothesis that preferences are stable over time — at least after one properly controls for variation in income, prices, and other arguments that enter into the utility function — remains to be seen; however, we take this assumption as a natural point of departure for economic research, and focus on explanations for variation in revealed preferences (within individual over time) that are related to predictable variation in income and anticipated expenditures.

experimental methods to study discounting — assume that subjects engage in narrow bracketing, viewing their dated lab experimental payoffs in isolation, as though they were consumption plans (Rabin and Weizsacker 2009). Coller and Williams (1999) were some of the first to question this assumption, highlighting the fact that experimental subjects with access to credit should engage in arbitrage, choosing immediate payments when the gross interest rate within the experiment is lower than their return on savings, and choosing delayed payments when the experimental return is lower than their cost of borrowing. More recently, Dean and Sautmann (2014) provide evidence that the intertemporal tradeoffs made by subjects in their (lab-in-the-field) experiment are related to liquidity constraints and accumulated savings outside of the lab — suggesting that the narrow bracketing assumption fails in their data. Building on Harris and Laibson (2001), Dean and Sautmann (2014) propose a model in which sophisticated, potentially credit-constrained consumers make optimal forward-looking plans that integrate their experimental payments into their consumption streams. If the assumptions of their model hold, then lab experiments cannot be used to measure individual time preferences; instead, they measure the marginal rate of substitution across periods, which depends on the market interest rate and one’s anticipated income and expenditures in all future periods.

Of course, it is entirely possible that both of these extreme modeling assumptions — narrow bracketing with zero background consumption and full dynamic optimization — fail to provide a psychologically realistic description of the ways human beings actually make financial decisions (both inside and outside the lab). If savings are illiquid or there are substantial transaction costs to borrowing or lending, people are unlikely to engage in the types of arbitrage proposed by Coller and Williams (1999), and the tradeoffs made in the lab (over moderate amounts of money) are likely to have little to do with the market interest rate.³³ Moreover, even outside the lab, there is ample evidence that people do not smooth moderately-sized one-off payments over their lifetimes (cf. Halevy 2014).

CTB experiments allow for explicit tests of the extent to which subjects engage in these types of arbitrage between lab and non-lab savings vehicles: because the implicit interest rates offered in experiments are typically well above those available through the credit market, subjects who integrate lab payments into fully-optimal forward-looking consumption plans should only choose immediate consumption if they are liquidity constrained; those who hold liquid savings should allocate their entire endowment to the later payoff date when the (implicit) lab interest rate exceeds the return on their savings (because this maximizes the net present value of their income stream).³⁴

We test this prediction in our data by dividing the sample into subjects known to have liquid

³³The same may be true if there are large cognitive costs in forming an optimal forward-looking consumption and savings plan (or re-optimizing in response to a moderately-sized shock).

³⁴In addition, those who can borrow at an interest rate below that offered by the experiment should also allocate their entire endowment to the later payoff date. However, it is difficult to collect reliable information on subjects’ access to credit (and the interest rates that they would face) outside of the lab — particularly if borrowing would take place outside of the formal credit market. We therefore focus on the subset of subjects who hold moderate amounts of money in savings accounts since it is apparent that these individuals are not liquidity constrained.

savings (the 28 percent of subjects who have more than 1,000 Kenyan shillings, or 10.20 USD, in a savings account) and those who might be liquidity constrained (those who do not have 1,000 Kenyan shillings in a savings account). Figure 7 plots the proportion of subjects within each category who allocate their entire endowment to the later payoff date. The figure makes it clear that, though corner solutions are common, those subjects who are not liquidity constrained do not appear to be engaging in arbitrage by consistently cashing in on the extremely high interest rates offered in the experiment. Across all interest rates, subjects with substantial savings in the bank allocate their entire endowment to the later payoff date 25.5 percent of the time (while those without 1,000 shillings in the bank do so 21.2 percent of the time). Focusing in on extremely high gross interest rates above 2 (i.e. guaranteed returns above 100 percent over a 2 to 4 week time horizon), which are extremely unlikely to be available outside the lab, we still find that those who are not liquidity constrained only allocate the entire endowment to the later payoff 32.9 percent of the time. Thus, our data make it very clear that subjects are not engaging in the types of arbitrage that we would expect if lab payments were integrated into an optimal forward-looking consumption and savings plan.³⁵

5.2 Do Subjects Engage in Narrow Bracketing?

Next, we turn our attention to the issue of narrow bracketing. We began exploring this issue in Section 4, when we tested the hypothesis that background consumption was equal to zero in our pooled analysis. In Column 2 of Table 2, we estimated background consumption as one of the parameters of our CRRA model; in Column 3 of Table 2, we estimated separate background consumption parameters for the earlier and later payoff dates. We found that the estimated background consumption parameter was significantly different from 0, and that we were able to reject the hypothesis that background consumption was the same in the earlier and later periods. Interestingly, we are unable to reject the hypothesis that the single background consumption parameter reported in Column 2 of Table 2 is significantly different than the average level of self-reported daily expenditure (p-value 0.184). When we estimate a specification that allows background consumption to differ across periods (in Column 3), we are unable to reject the hypothesis that consumption in the earlier period is equal to mean self-reported expenditures (p-value 0.462), but we can just reject the hypothesis that background consumption in the later period is equal to mean self-reported expenditure (p-value 0.075). Importantly, neither accounting for background consumption nor allowing it to differ across periods impacts our estimated β coefficient dramatically; it is significantly below 1 in all specifications. However, these results clearly reject the most extreme form of narrow bracketing: background consumption is not equal to zero; in fact, it is quite close to the amount spent in a day by the typical subject. Choices in our experiment also suggest that subjects expect

³⁵Interestingly, our results differ from those of Andreoni and Sprenger (2012) in this regard: 70 percent of the CTB decisions observed in their experiment are corner solutions, suggesting that arbitrage may be a more reasonable explanation of behavior in that context.

background consumption to be systematically higher in the future.

Since background consumption is not equal to zero and differs across periods, it is important to ask whether choice patterns that appear present-biased might, in fact, result from subjects' attempts to shift payoffs toward dates when they anticipate that the marginal utility of consumption will be higher (for example, because they will have relatively little income or cash on hand), as suggested by Dean and Sautmann (2014) and Halevy (2015). More generally, if subjects are shifting payoffs toward days when they expect the marginal utility of income to be high, static preferences reversals may or may not constitute evidence of dynamic inconsistency (Halevy 2015). To explore this issue, we added a series of questions about anticipated liquidity to the post-experiment survey. For each of the five possible payout dates in the experiment, we asked subjects whether they expected to have more cash on hand than normal, less cash on hand than normal, or approximately the typical amount of cash on hand. These questions were only administered during the last eight experimental sessions, so we have data from 137 of our 291 subjects. We observe substantial within-subject variation in anticipated liquidity: only 29.9 percent of subjects expected to have more cash on hand than is typical at all payoff dates, and no subjects expected to have less cash than is typical on all payoff dates. As a result, we can use our self-reported measure of expected liquidity to test whether subjects shift income from the experiment toward dates when they expect to have relatively little cash on hand.

In Table 6, we regress the fraction of the endowment allocated to the earlier payoff date on controls for the delay between payments, the interest rate, and the endowment size plus an indicator for having a front-end delay of zero days. We report OLS specifications controlling for individual fixed effects in Columns 1 and 2 and Tobit specifications that adjust for censoring of the dependent variable at 0 and 1 in Columns 3 and 4. In all specifications, we restrict the sample to those subjects who answered the date-specific liquidity questions (results are nearly identical in the full sample). As expected (given our main results), the percent of the endowment allocated to the earlier payoff date is significantly higher when the earlier payoff is immediate. In Columns 2 and 4, we add an indicator for having lower expected liquidity (cash on hand) at the later payoff date, relative to the earlier payoff date. Including the control for anticipated liquidity has little impact on the parameter measuring present bias, the indicator for choices in which the front-end delay is zero. Moreover, the estimated coefficient on the relative liquidity measure does not suggest that subjects shift funds to dates when they expect to be liquidity-constrained. In the fixed effects specifications, the coefficient is close to zero and not statistically significant; in the Tobit specifications, it is positive and weakly significant — suggesting that, if anything, subjects shift experimental payments away from dates when they expect to have relatively less liquidity. We therefore find no evidence that subjects reduce their allocation to the earlier payoff date when they expect marginal utility to be higher at the later payoff date; and we can rule out the possibility that such predictable changes in marginal utility explain the observed level of present bias in our data.

5.3 Summary and Directions for Further Research

Our results — and, in our view, the results of many of the related studies listed above — point to a few broad conclusions. First, neither the standard model of intertemporal tradeoffs (in which lab payments are integrated into an optimal forward-looking consumption and savings plan) nor the simplest form of narrow bracketing (with background consumption equal to 0) provides an adequate description of the ways actual subjects make intertemporal tradeoffs in the lab. We should therefore devote greater attention to the development of experiments explicitly designed to test potential refinements of existing models, in order to build a more accurate positive description of the decision-making process. Experiments intended to better understand the way humans use mental accounting and heuristics to simplify the process of organizing their financial lives seem likely to be particularly fruitful.

Second, our results suggest that present bias over money is important when the earliest payouts are truly immediate; our findings are consistent with models such as the quasi-hyperbolic model (Laibson 1997, O’Donoghue and Rabin 1999) that predict a sharp distinction between immediate and slightly delayed payoffs. However, our experiment is not intended to trace out the relationship between individual discount factors and the level of the front-end delay over the minutes and hours after decisions are made. As such, we are also unable to fully distinguish between the quasi-hyperbolic model and alternatives that do not suggest a clear discontinuity between immediate and nearly immediate payoffs (cf. Kirby 1997). We leave that to future work.

Finally, and perhaps most speculatively, we believe that there might be important lessons to be learned from an attempt to test the extent to which the widely diverse results reported in the literature can be explained by differences across studies in either experimental parameters (e.g. the interest rate) or contextual variables (e.g. credit market conditions or financial literacy). As we observed earlier, our results suggest that present bias is most pronounced at high interest rates; and our experiment offers a substantially broader range of interest rates than several recent studies which do not find evidence of present bias. Moreover, while considerable energy has gone in to testing the relationship between, for example, liquidity constraints and intertemporal tradeoffs in the lab *within* a particular study, work to date has not explored the potential links between the variation in observed results *across* studies the widely divergent financial situations of different subject pools. Yet, the potential for such analysis exists, since rigorous time preference experiments have now been conducted in an extremely wide range of settings, with subjects ranging from Malawian subsistence farmers to undergraduates at some of the most prestigious universities in North America. In our view, these conclusions suggest a focus on experiments in broad samples of the population that focus on explicitly testing the links between credit constraints, portfolio holdings, and income expectations — both inside and outside the lab — and intertemporal tradeoffs.

6 Conclusion

We conduct convex time budget experiments with truly immediate repayment, exploiting Kenya’s mobile money technology M-Pesa to send same-day payments before subjects leave the lab while still keeping transaction costs precisely equal across all payout dates. We find substantial evidence of present bias. Estimates of the present bias parameter, β , range from 0.901 to 0.937. We also observe substantial heterogeneity across subjects. Variables associated with likely liquidity constraints (or the lack thereof) do not predict individual-level time preference parameters.

As discussed above, our results differ from a number of other recent findings. First, we document substantial present bias. Our focus on truly immediate payment may explain the observed differences between our findings and those of, for example, Andreoni and Sprenger (2012) and Giné, Goldberg, Silverman, and Yang (2012). However, Dean and Sautmann (2014) also fail to find evidence of present bias in a setting with immediate repayment. Second, we can reject the hypothesis that subjects narrowly bracket their decisions in the experiment and assume background consumption is equal to zero. Background consumption appears to be very close to the typical level of daily expenditures made by our subjects. However, we find no evidence that subjects shift lab payouts toward dates when they expect to be liquidity constrained. Finally, we are able to rule out the possibility that subjects are arbitraging between lab and non-lab savings vehicles. Subjects who hold amounts of liquid savings larger than the immediate lab payouts considered in our study do not exploit the high rates of interest offered by the experiment, and therefore fail to maximize the net present value of their income stream. This result differs from the findings of Dean and Sautmann (2014); Andreoni and Sprenger (2012) also report that 70 percent of the CTB decisions in their experiment were corner solutions, suggesting that their data is more consistent with arbitrage between the lab and the credit market.

We have argued that one of the key differences between our experiment and other studies is our focus on truly immediate payment. There are, however, several other differences that should be noted. Most notably, relative to university students at elite schools in the North America, our subjects — like those of Giné, Goldberg, Silverman, and Yang (2012) and Dean and Sautmann (2014) — are living closer to their liquidity and credit constraints. While university students may have low incomes, particularly if one ignores family and government transfers, they have high and stable consumption streams; they are able to save at almost no cost, and to borrow at reasonable rates if it were absolutely necessary. Moreover, students may find the task of re-optimizing their financial plan more feasible, both because they are highly numerate and because their portfolio is likely limited to a modest range of formal credit products with explicit terms. It is therefore reasonable to expect that university students may be more likely to arbitrage between the lab and the credit market than low-income adults in developing countries.

The development of the quasi-hyperbolic model of time preferences was motivated by the desire to provide a better positive explanation of the ways humans actually made intertemporal tradeoffs.

However, since the advent of this tractable model of present bias, much of the energy has shifted toward exploring its predictions in field settings and testing between the quasi-hyperbolic model and the standard model. The new wave of theoretically-motivated time preference experiments, and the controversies generated by the conflicting results of these studies, highlights the need for more experimental work directed towards testing and refining the behavioral economic model of intertemporal decisionmaking, with a specific focus on not only the extent of present bias, but the ways subjects do or do not integrate controlled choices in the lab — or isolated decisions outside of the lab — into a larger, long-term optimization problem. Such an integrated framework has the potential to allow for a more complete reconciliation of existing results, and will also help policymakers seeking to nudge citizens and consumers toward achieving their long-term savings and consumption goals.

References

- ACLAND, D., AND M. R. LEVY (2015): “Naivete, Projection Bias, and Habit Formation in Gym Attendance,” *Management Science*, 61(1), 146–160.
- AFRIAT, S. N. (1967): “The Construction of Utility Functions from Expenditure Data,” *International Economic Review*, 8(1), 67–77.
- AINSLIE, G. (1975): “Specious reward: A behavioral theory of impulsiveness and impulse control,” *Psychological Bulletin*, 82(4), 463–496.
- ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2008): “Eliciting Risk and Time Preferences,” *Econometrica*, 76(3), 583–618.
- ANDREONI, J., B. GILLEN, AND W. HARBAUGH (2013): “The Power of Revealed Preference Tests: Ex-Post Evaluation of Experimental Design,” Working Paper.
- ANDREONI, J., AND J. MILLER (2002): “Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism,” *Econometrica*, 70(2), 737–753.
- ANDREONI, J., AND C. SPRENGER (2012): “Estimating Time Preferences from Convex Budgets,” *American Economic Review*, 102(7), 3333–56.
- AUGENBLICK, N., M. NIEDERLE, AND C. SPRENGER (2015): “Working Over Time: Dynamic Inconsistency in Real Effort Tasks,” *Quarterly Journal of Economics*, 130(3), 1067–1115.
- BECKER, G. S. (1962): “Irrational Behavior and Economic Theory,” *The Journal of Political Economy*, 70(1), 1–13.
- BRONARS, S. G. (1987): “The Power of Nonparametric Tests of Preference Maximization,” *Econometrica: Journal of the Econometric Society*, pp. 693–698.
- CARVALHO, L. S., S. MEIER, AND S. W. WANG (forthcoming): “Poverty and economic decision making: evidence from changes in financial resources at payday,” *American Economic Review*.
- CHOI, S., R. FISMAN, D. GALE, AND S. KARIV (2007a): “Consistency and Heterogeneity of Individual Behavior under Uncertainty,” *American Economic Review*, 97(5), 1921–1938.
- (2007b): “Revealing Preferences Graphically: An Old Method Gets a New Toolkit,” *American Economic Review*, 97(2), 153–158.
- CHOI, S., S. KARIV, W. MÜLLER, AND D. SILVERMAN (2014): “Who Is (More) Rational?,” *American Economic Review*, 104(6), 1518–1550.
- COLLER, M., AND M. WILLIAMS (1999): “Eliciting Individual Discount Rates,” *Experimental Economics*, 2(2), 107–127.
- CROSON, R., AND U. GNEEZY (2009): “Gender Differences in Preferences,” *Journal of Economic Literature*, 47(2), 1365–1384.
- DEAN, M., AND A. SAUTMANN (2014): “Credit Constraints and the Measurement of Time Preferences,” Working Paper 2014-1, Brown University, Department of Economics.
- DELLAVIGNA, S., AND U. MALMENDIER (2006): “Paying Not to Go to the Gym,” *American Economic Review*, pp. 694–719.
- EPPER, T. (2015): “Income Expectations, Limited Liquidity, and Anomalies in Intertemporal Choice,” Discussion paper, Universitat St. Gallen Discussion Paper 2015-19.
- FISCHBACHER, U. (2007): “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, 10(2), 171–178.
- FISHER, I. (1930): “The Theory of interest rates,” *New York*.

- FISMAN, R., P. JAKIELA, AND S. KARIV (2014): “The Distributional Preferences of Americans,” NBER working paper no. 20145.
- FISMAN, R., P. JAKIELA, S. KARIV, AND D. MARKOVITS (2015): “The Distributional Preferences of an Elite,” *Science*, 349(5254), aab0096.
- FISMAN, R., S. KARIV, AND D. MARKOVITS (2007): “Individual Preferences for Giving,” *American Economic Review*, 97(5), 1858–1876.
- FREDERICK, S., G. LOEWENSTEIN, AND T. O’DONOGHUE (2002): “Time discounting and time preference: A critical review,” *Journal of economic literature*, pp. 351–401.
- GINÉ, X., J. GOLDBERG, D. SILVERMAN, AND D. YANG (2012): “Revising Commitments: Field Evidence on the Adjustment of Prior Choices,” Working Paper 18065, National Bureau of Economic Research.
- HALEVY, Y. (2014): “Some Comments on the Use of Monetary and Primary Rewards in the Measurement of Time Preferences,” Working Paper.
- (2015): “Time Consistency: Stationarity and Time Invariance,” *Econometrica*, 83(1), 335–352.
- HARRIS, C., AND D. LAIBSON (2001): “Dynamic Choices of Hyperbolic Consumers,” *Econometrica*, 69(4), 935–957.
- HAUSHOFER, J. (2014): “The Cost of Keeping Track,” Working Paper.
- JAKIELA, P., AND O. OZIER (forthcoming): “Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies,” *Review of Economic Studies*.
- KIRBY, K. N. (1997): “Bidding on the future: evidence against normative discounting of delayed rewards.,” *Journal of Experimental Psychology: General*, 126(1), 54.
- LAIBSON, D. (1997): “Golden Eggs and Hyperbolic Discounting,” *The Quarterly Journal of Economics*, 112(2), 443–478.
- LAIBSON, D., A. REPETTO, J. TOBACMAN, R. E. HALL, W. G. GALE, AND G. A. AKERLOF (1998): “Self-Control and Saving for Retirement,” *The Quarterly Journal of Economics*, 1998(1), 91–196.
- MEIER, S., AND C. SPRENGER (2010): “Present-Biased Preferences and Credit Card Borrowing,” *American Economic Journal: Applied Economics*, 2(1), 193–210.
- O’DONOGHUE, T., AND M. RABIN (1999): “Doing It Now or Later,” *American Economic Review*, 89(1), 103–124.
- PHELPS, E. S., AND R. A. POLLAK (1968): “On Second-Best National Saving and Game-Equilibrium Growth,” *The Review of Economic Studies*, 35(2), 185–199.
- RABIN, M., AND G. WEIZSACKER (2009): “Narrow Bracketing and Dominated Choices,” *The American Economic Review*, 99(4), 1508–1543.
- SAMUELSON, P. A. (1937): “A Note on Measurement of Utility,” *The Review of Economic Studies*, 4(2), 155–161.
- STIGLER, G. J., AND G. S. BECKER (1977): “De Gustibus Non Est Disputandum,” *American Economic Review*, 67(2), 76–90.
- THALER, R. H. (1991): “Some Empirical Evidence on Dynamic Inconsistency,” *Quasi rational economics*, 1, 127–136.
- VARIAN, H. R. (1982): “The Nonparametric Approach to Demand Analysis,” *Econometrica*, 50(4), 945–972.
- (1983): “Non-Parametric Tests of Consumer Behaviour,” *Review of Economic Studies*, 50(1), 99–110.

VON BÖHM-BAWERK, E. (1890): *Capital and interest: A critical history of economical theory*. Macmillan and Co.

VON GAUDECKER, H.-M., A. VAN SOEST, AND E. WENGSTRÖM (2011): “Heterogeneity in Risky Choice Behavior in a Broad Population,” *American Economic Review*, 101(2), 664–694.

WORLD BANK (2015): “World Development Report 2015: Mind, Society, and Behavior,” Report, International Bank for Reconstruction and Development.

Table 1: Summary Statistics

Variable:	Mean	S.D.	Median	Min.	Max.	Obs.
Female	0.58	0.49	1	0	1	291
Age	32.55	9.57	30	18	62	279
Completed primary school	0.96	0.21	1	0	1	291
Completed secondary school	0.56	0.50	1	0	1	291
Married or cohabitating	0.59	0.49	1	0	1	291
Employed or self-employed	0.29	0.45	0	0	1	291
Has bank account with at least 1,000 shillings	0.28	0.45	0	0	1	291
Average daily expenditure (in shillings)	204.20	225.94	150	0.14	2857.14	291
Expenditure on previous Monday (in shillings)	325.06	553.49	200	2	7000	291
Subject describes self as very-patient	0.65	0.48	1	0	1	291
Trusts that lab payments will be sent on time	0.98	0.15	1	0	1	291

Table 2: NLS and Tobit Estimates of Model Parameters Assuming CRRA Utility

Parameter	NLS (1)	NLS (2)	NLS (3)	NLS (4)	Tobit (5)	Tobit (6)
β	0.895*** (0.018)	0.920*** (0.016)	0.937*** (0.016)	0.910*** (0.019)	0.856*** (0.034)	0.901*** (0.035)
δ	0.992*** (0.001)	0.991*** (0.001)	0.996*** (0.001)	0.991*** (0.001)	1.000*** (0.003)	1.001*** (0.002)
ρ	0.578*** (0.023)	0.980*** (0.076)	0.975*** (0.082)	0.996*** (0.079)	0.727*** (0.003)	1.417*** (0.002)
$\omega_t = \omega_{t+k}$	0 —	276.445*** (54.288)		$\bar{\omega}_i$ —	0.01 —	$\bar{\omega}_i$ —
ω_t			244.916*** (55.275)			
ω_{t+k}			314.921*** (62.044)			
$H_0: \beta = 1$	0.000	0.000	0.000	0.000	0.000	0.004
$H_0: \delta = 1$	0.000	0.000	0.000	0.000	0.864	0.679
$H_0: \omega_t = \omega_{t+k}$			0.000			
Observations	13968	13968	13968	13968	13968	13968
Subjects	291	291	291	291	291	291

Robust standard errors clustered at the subject level. $\bar{\omega}_i$ indicates self-reported average daily expenditure, which varies across subjects.

Table 3: NLS Estimates of Model Parameters Assuming CARA Utility

Parameter	NLS (1)	NLS (2)
β	0.938*** (0.013)	0.958*** (0.013)
δ	0.990*** (0.001)	0.995*** (0.001)
α	0.001*** (0.001)	0.001*** (0.001)
$\omega_{t+k} - \omega_t$		86.353*** (19.241)
$H_0: \beta = 1$	0.000	0.001
$H_0: \delta = 1$	0.000	0.000
Observations	13968	13968
Subjects	291	291

Robust standard errors clustered at the subject level. Columns 3 and 4 are regressions of the tangency condition and solution function respectively.

Table 4: Median Regressions of Individual-Level Parameter Estimates

<i>Estimated parameter:</i>	β_i (1)	δ_i (2)	ρ_i (3)
Female	0.009 (0.044)	0.01*** (0.003)	0.376*** (0.101)
Age	0.002 (0.002)	0.0001 (0.0001)	0.012** (0.005)
Married or cohabitating	0.019 (0.044)	0.002 (0.003)	0.121 (0.1)
Completed secondary school	-0.081* (0.043)	0.003 (0.003)	-0.074 (0.099)
Employed or self-employed	-0.013 (0.047)	0.001 (0.003)	0.101 (0.108)
Has 1000 KSH in a bank account	0.013 (0.047)	-0.001 (0.003)	0.071 (0.107)
Constant	0.927*** (0.086)	0.977*** (0.005)	0.271 (0.196)
Observations	279	279	279

Robust standard errors in parentheses.

Table 5: Median Regressions of Associations between CTB and MPL Parameter Estimates

<i>Sample restriction:</i> <i>Estimated parameter:</i>	ALL SUBJECTS		CONSISTENT SUBJECTS	
	β_i (1)	δ_i (2)	β_i (3)	δ_i (4)
MPL estimate of parameter	0.177*** (0.064)	0.269*** (0.06)	0.275*** (0.074)	0.412*** (0.065)
Constant	0.766*** (0.068)	0.727*** (0.059)	0.682*** (0.078)	0.585*** (0.064)
Observations	273	273	210	210


Robust standard errors in parentheses. In Columns 3 and 4, the sample is restricted to subjects who make consistent decisions in the MPL tasks — choosing the early payoff at low interest rates, and switching (once) to preferring the later payoff at higher interest rates.

Table 6: Regressions of Percent of Budget Allocated to Earlier Payoff Date

<i>Specification:</i>	FES	FES	TOBIT	TOBIT
	(1)	(2)	(3)	(4)
Front end delay: 0 days	0.035*** (0.011)	0.034*** (0.011)	0.05*** (0.018)	0.055*** (0.019)
Delay between payments: 28 days	0.024* (0.013)	0.024* (0.013)	0.053** (0.022)	0.051** (0.022)
Interest rate: 125 percent	-0.089*** (0.017)	-0.089*** (0.017)	-0.156*** (0.032)	-0.156*** (0.032)
Interest rate: 175 percent	-0.188*** (0.025)	-0.188*** (0.025)	-0.331*** (0.055)	-0.331*** (0.055)
Interest rate: 200 percent	-0.229*** (0.029)	-0.229*** (0.029)	-0.397*** (0.063)	-0.397*** (0.063)
Interest rate: 300 percent	-0.249*** (0.032)	-0.249*** (0.032)	-0.442*** (0.07)	-0.442*** (0.07)
Interest rate: 400 percent	-0.264*** (0.033)	-0.264*** (0.033)	-0.466*** (0.073)	-0.466*** (0.073)
Endowment size: 600 shillings	0.013 (0.014)	0.013 (0.014)	0.021 (0.024)	0.021 (0.024)
Lower expected liquidity at later payoff date	.	-0.006 (0.006)	.	0.036* (0.02)
Constant	0.587*** (0.021)	0.587*** (0.021)	0.656*** (0.044)	0.654*** (0.044)
Observations	6576	6576	6576	6576

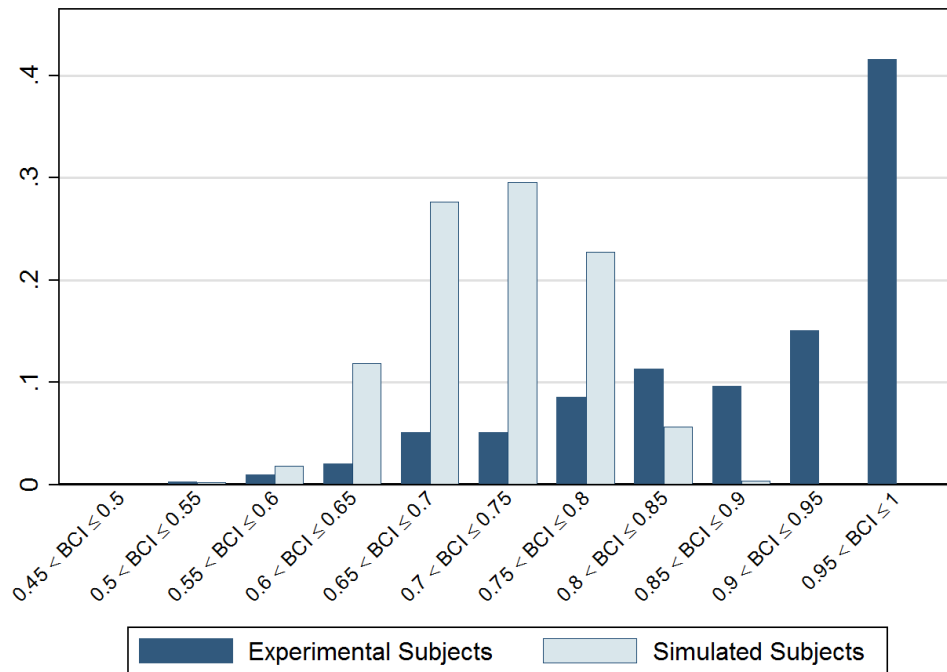
Robust standard errors clustered at the subject level in all specifications.

Figure 1: Frequency of GARP Violations for Actual vs. Simulated Subjects




The figure reports the total number of violations of the Generalize Axiom of Revealed Preference for experimental subjects and four 1,000 thousand simulated subjects who randomize uniformly on each budget line. Violations are calculated for the two pairs of (earlier and later) dates that were presented at both the 400 shilling (early-valued) budget size and the 600 shilling (early-valued) budget size (i.e. decision sets 1, 3, 4, and 7 from the CTB experiment), creating intersecting budget lines that allow for revealed preference tests.

Figure 2: The Basic Consistency Index for Actual vs. Simulated Subjects



To calculate the Basic Consistency Index (BCI), we consider all pairs of interest rates, r' and r'' , such that $r'' > r'$; a subject's decisions satisfy basic consistency if the amount allocated to the later period is at least as large under r'' as under r' . The BCI is fraction of all possible pairs of choices that satisfy this property.


Figure 3: Percent of Budget Allocated to Earlier Account



The sample is restricted to CTB decisions in which the early-valued budget was 400 Kenyan shillings (4.08 USD).

Figure 4: Percent of Budget Allocated to Earlier Account, by Interest Rate

Panel A: Delay Between Payments = 28 Days



Panel B: Delay Between Payments = 14 Days





Figure 5: Histograms of Individual-Level Parameter Estimates

Panel A: Estimated $\hat{\beta}_i$ Parameters



Panel B: Estimated $\hat{\delta}_i$ Parameters




Figure 6: Scatter Plots of Individual-Level Parameters: CTB vs. MPL Estimates






Figure 7: Do Subjects Engage in Arbitrage? The Frequency of (Later Payoff Date) Corner Solutions



Potentially liquidity-constrained subjects are those who do not have at least KES 1,000 (USD 9.80) in a bank savings account. The figure plots the fraction of decisions (conditional on the gross interest rate) that are corner solutions where the subject allocates her entire endowment to the later payoff date.

Online Appendix: not for print publication

Experimental Instructions

This is a study of the different ways people make decisions about money. Please listen carefully to the instructions that are being read to you. Tell us if you have any questions or if there is anything that you do not understand.

Each of you will receive two payments for participating in this study, an EARLIER payment and a LATER payment. You will receive the EARLIER payment on an EARLIER date, and you will receive the LATER payment on a LATER date. These two payments will be sent via M-Pesa by XXXX o'clock on each date. So, if you receive a payment today, it will be sent to you before you leave the Busara Center. On the day that you are scheduled to receive one of your payments, we will send you a text message reminder that the payment is coming. After you receive the text message, your payment will be sent via M-Pesa. If you do not receive one of your payments, you should immediately contact the staff at Busara by flashing +254(0) 704851141

Each of you will receive 300 shillings just for participating in the study. You will receive this money in two equal amounts of 150 shillings. These two payments of 150 shillings will be sent to you via M-Pesa on the two different dates (the EARLIER date and the LATER date). You'll receive your payments before XXXX o'clock on each date.

In addition to the 300 shillings, you will be paid money based on one of your decisions in this study. You will make 72 decisions in this study. The study has two parts. In the first part of the study you will make 48 decisions; in the second part you will make 24 decisions. After you've made all 72 decisions, the computer will choose one of the decisions to be the payment decision. We will use the payment decision to determine how much money you are paid in this study. All 72 decisions have the same chance of being chosen as the payment decision. So, you should make each decision as if it were the payment decision. At the end of the study, your computer will choose the payment decision. Because of the payment decision, you may receive more than 150 shillings on the EARLIER date, the LATER date, or both. We will add any money from the payment decision to your two payments of 150 shillings. So, you will receive at least 150 shillings by XXXX o'clock on each of the two dates (EARLIER and LATER). Both of these payments will be sent via M-Pesa.

In this part of the study, you will make 48 decisions about how to divide money between two times, one is on an EARLIER date and one is on a LATER date. So, you will decide how much money you want to be sent to you at the EARLIER date and how much you want at the LATER date. The easiest way to explain how you will make decisions in this study is to show you an example.

For each decision that you will make, you will see a screen like the one that you have in front of you now. These two boxes show you when you will receive the two payments. The box on the left shows the date of the EARLIER payment, and the box on the right shows the date of the LATER payment. In this example, you will receive the EARLIER payment by XXXX o'clock today (before you leave the Busara Center), and you will receive the LATER payment in two weeks later - on XXXX.

The boxes also show you the amount you will receive on each date. The box on the left shows the amount of money in the EARLIER payment, and the box on the right shows the amount of the LATER payment. The amounts are in red. You will be paid this money in addition to the 150 shillings that you will receive on each date (EARLIER and LATER) for participating in this study. For now, you can see that the two amounts are 0.

Next, notice a thick blue and green line on the center of the screen. Touch anywhere on that thick line. You should see a black pointer above the thick line and two buttons at the bottom part of the screen. Now you also see that the two amounts are not zero anymore.

You can change the position of the pointer moving it from right to left, or left to right by touching anywhere on the thick line. You will see (in the two boxes) that you are moving money between the EARLIER payment and the LATER payment by touching different parts on the thick line. Thus, touching more towards the right side moves more money to the LATER payment from the EARLIER payment.

Please try this yourself. Touch anywhere on this thick line. As you touch more towards the left side,

the amount of the EARLIER payment increases and the amount of the LATER payment decreases. As you touch more towards the right side, the amount of the EARLIER payment decreases and the amount of the LATER payment increases. Touching more towards the right side moves more money to the LATER payment from the EARLIER payment.

In this part of the study, you will make 48 decisions. Each decision will be different. In the boxes on the screen, you will see the highest amounts of money that you can receive at the EARLIER and LATER dates in that decision. The highest amounts of money will not be the same in all decisions. Now practice touching different parts of the line and changing the size of the payments. The highest amounts of money will not be the same in all decisions. The highest amount that you can receive on the LATER date is always the same or more than the highest amount that you can receive on the EARLIER date.

Today, you will make decisions in 8 different rounds. Each round will have 6 decisions. The dates of the EARLIER and the LATER payments might be different in each of the rounds. At the beginning of each round, we'll announce when the EARLIER and LATER payments will take place for all the decisions in that round. This information is also shown in the boxes on your screen. Within a round, the date of the EARLIER payment is the same for all decisions and the date of the LATER payment is the same for all decisions.

In this part of the study, you will make 48 decisions. You will indicate your decision by touching anywhere you want on the line. To confirm your decision, you will be touching the "OK" button. After each round, you will see a screen with the words "Please wait for the study to continue".

As we said earlier, the computer will choose one of your 72 decisions to be the payment decision. All 48 decisions in the first part of the study have the same chance of being chosen, so you should think carefully about each decision.

We will start with a practice round. This practice round will be the same as the 8 later rounds. In this practice round, you will make 6 decisions about how you want to divide money between and EARLIER date and a LATER date. This round is just for practice; these decisions will not be the payment decision. Remember: this round is just for practice; these decisions will not be the payment decision.

Now we will start the second part of the study. You will make 24 decisions in this part of the study. In each decision, you will choose between a smaller EARLIER payment and a larger LATER payment. All 24 decisions in this part of the study have the same chance of being chosen as the payment decision, so you should think carefully about each decision. The computer will choose one of the decisions to be the payment decision at the end of the study. So, please listen carefully to the instructions being read to you. Tell us if you have any questions or if there is anything that you do not understand.

The easiest way to explain how you will make decisions in this study is by showing you an example. For each decision that you will make, you will see a screen like the one in front of you now. These two boxes show you when you will receive the two payments. In this example, you will receive the EARLIER payment on XXXX, and you will receive the LATER payment two weeks LATER on XXXX. The box on the left shows the amount of money you will receive in the EARLIER payment, and the box on the right shows the amount of money you will receive in the LATER payment. The amounts are in red colour. In this example, the EARLIER payment is XXXX, and the LATER payment is XXXX.

Each decision will be different. In each of the decisions in this study, you will see a screen like this one. You will decide if you want to receive the EARLIER payment or the LATER payment. There is no right or wrong answer - you just choose the payment that you prefer more. To choose the EARLIER payment, touch the green "CHOOSE THE EARLIER PAYMENT" button on the left side of the screen. To choose the LATER payment, touch the yellow "CHOOSE THE LATER PAYMENT" button on the right side of the screen. So, to make a decision, you touch the button which is below the payment that you prefer. When you touch the button, its colour will turn red. To confirm your decision, you touch the "OK" button at the bottom of you screen. You can change your decision before you touch the "OK" button.

Today, you will make decisions in 4 different rounds. Each round will have 6 decisions. The dates of the EARLIER and the LATER payments will be different in every round. At the beginning of each round, well announce when the EARLIER and LATER payments will take place for all the decisions in that round. This message is also shown in the boxes on your screen.

Within a round, the date of the EARLIER payment is the same for all decisions and the date of the

LATER payment is the same for all decisions. After each round, you will see a screen with the words "Please wait for the study to continue." At the end of the study, the computer will choose one of your decisions to be the payment decision. All 24 decisions in this part of the study also have the same chance of being chosen, so you should think carefully about each decision.

You will receive the payment that you chose - the EARLIER payment or the LATER payment - from the payment decision. If you chose the EARLIER payment and the EARLIER date is today, you will receive your payment before leaving Busara Center.

The computer will now choose one of your 72 decisions for payment.

Additional Tables and Figures

Table A1: Falsification Test: Allowing Variation in Discounting when Front-End Delay is 28 Days

Parameter	NLS (1)	NLS (2)	NLS (3)	NLS (4)	Tobit (5)	Tobit (6)
λ	1.013*** (0.022)	0.979*** (0.017)	0.988*** (0.017)	0.965*** (0.022)	1.047*** (0.047)	1.021*** (0.040)
δ	0.991*** (0.001)	0.990*** (0.001)	0.997*** (0.001)	0.990*** (0.001)	0.997*** (0.003)	0.999*** (0.002)
ρ	0.588*** (0.024)	1.071*** (0.088)	1.067*** (0.094)	1.009*** (0.081)	0.743*** (0.003)	1.437*** (0.002)
$\omega_t = \omega_{t+k}$	0 —	334.355*** (63.363)		$\bar{\omega}_i$ —	0.01 —	$\bar{\omega}_i$ —
ω_t			300.571*** (65.254)			
ω_{t+k}			385.868*** (71.570)			
$H_0: \lambda = 1$	0.572	0.200	0.480	0.113	0.310	0.596
$H_0: \delta = 1$	0.000	0.000	0.000	0.000	0.257	0.764
$H_0: \omega_t = \omega_{t+k}$			0.000			
Observations	13968	13968	13968	13968	13968	13968
Subjects	291	291	291	291	291	291

Robust standard errors clustered at the subject level. $\bar{\omega}_i$ indicates self-reported average daily expenditure, which varies across subjects. λ is included in the discount factor when the front-end delay is 28 days, the longest front-end delay considered in our experiment.

Table A2: Parameter Estimates: Subjects with No GARP Violations

Parameter	NLS (1)	NLS (2)	NLS (3)	NLS (4)	Tobit (5)	Tobit (6)
β	0.887*** (0.022)	0.924*** (0.018)	0.940*** (0.017)	0.919*** (0.022)	0.858*** (0.044)	0.885*** (0.052)
δ	0.992*** (0.002)	0.991*** (0.001)	0.994*** (0.001)	0.991*** (0.001)	1.000*** (0.004)	1.004*** (0.003)
ρ	0.529*** (0.029)	1.277*** (0.129)	1.366*** (0.157)	0.895*** (0.049)	0.715*** (0.004)	1.545*** (0.003)
$\omega_t = \omega_{t+k}$	0 —	569.025*** (104.332)		$\bar{\omega}_i$ —	0.01 —	$\bar{\omega}_i$ —
ω_t			619.666*** (123.881)			
ω_{t+k}			682.152*** (134.554)			
$H_0: \beta = 1$	0.000	0.000	0.001	0.000	0.001	0.027
$H_0: \delta = 1$	0.000	0.000	0.000	0.000	0.926	0.248
$H_0: \omega_t = \omega_{t+k}$			0.014			
Observations	8832	8832	8832	8832	8832	8832
Subjects	184	184	184	184	184	184

Robust standard errors clustered at the subject level. $\bar{\omega}_i$ indicates self-reported average daily expenditure, which varies across subjects. Sample restricted to subjects with no violations of the Generalized Axiom of Revealed Preferences.

Table A3: Parameter Estimates: Subjects with Basic Consistency Indices ≥ 0.85

Parameter	NLS (1)	NLS (2)	NLS (3)	NLS (4)	Tobit (5)	Tobit (6)
β	0.913*** (0.021)	0.902*** (0.018)	0.910*** (0.018)	0.928*** (0.023)	0.907*** (0.041)	0.938*** (0.044)
δ	0.993*** (0.002)	0.991*** (0.002)	0.996*** (0.001)	0.992*** (0.001)	1.005*** (0.003)	1.005*** (0.003)
ρ	0.467*** (0.024)	0.185*** (0.012)	0.232*** (0.014)	0.890*** (0.097)	0.617*** (0.003)	1.253*** (0.003)
$\omega_t = \omega_{t+k}$	0 —	-119.026*** (8.621)		$\bar{\omega}_i$ —	0.01 —	$\bar{\omega}_i$ —
ω_t			-112.477*** (8.268)			
ω_{t+k}			-49.975** (21.158)			
$H_0: \beta = 1$	0.000	0.000	0.000	0.002	0.022	0.157
$H_0: \delta = 1$	0.000	0.000	0.003	0.000	0.149	0.092
$H_0: \omega_t = \omega_{t+k}$			0.004			
Observations	9264	9264	9264	9264	9264	9264
Subjects	193	193	193	193	193	193

Robust standard errors clustered at the subject level. $\bar{\omega}_i$ indicates self-reported average daily expenditure, which varies across subjects. Sample restricted to subjects with basic consistency indices above 0.85 — high enough that they almost certainly could not have occurred at random.

Table A4: Parameter Estimates: Subjects with Fewer than 50 Percent Corner Solutions

Parameter	NLS (1)	NLS (2)	NLS (3)	NLS (4)	Tobit (5)	Tobit (6)
β	0.872*** (0.026)	0.889*** (0.024)	0.908*** (0.024)	0.900*** (0.023)	0.865*** (0.038)	0.895*** (0.037)
δ	0.997*** (0.002)	0.996*** (0.002)	1.002*** (0.001)	0.994*** (0.002)	1.001*** (0.003)	1.001*** (0.002)
ρ	0.829*** (0.035)	1.070*** (0.114)	1.119*** (0.130)	1.277*** (0.064)	0.827*** (0.003)	1.351*** (0.002)
$\omega_t = \omega_{t+k}$	0 —	127.045** (53.807)		$\bar{\omega}_i$ —	0.01 —	$\bar{\omega}_i$ —
ω_t			127.511** (59.528)			
ω_{t+k}			181.677*** (66.849)			
$H_0: \beta = 1$	0.000	0.000	0.000	0.000	0.000	0.004
$H_0: \delta = 1$	0.166	0.011	0.266	0.000	0.635	0.767
$H_0: \omega_t = \omega_{t+k}$			0.006			
Observations	8400	8400	8400	8400	8400	8400
Subjects	175	175	175	175	175	175

Robust standard errors clustered at the subject level. $\bar{\omega}_i$ indicates self-reported average daily expenditure, which varies across subjects. Sample restricted to subjects who chose corner solutions less than half the time.

Table A5: Convex Time Budget Decision Problems

Set	Decision	Front-End Delay (t)	Early Later (k)	vs. Delay	Early Max	Later Max	$1 + r$
1	1	14	14		400	440	1.1
1	2	14	14		400	500	1.25
1	3	14	14		400	700	1.75
1	4	14	14		400	800	2
1	5	14	14		400	1200	3
1	6	14	14		400	1600	4
2	7	0	28		400	440	1.1
2	8	0	28		400	500	1.25
2	9	0	28		400	700	1.75
2	10	0	28		400	800	2
2	11	0	28		400	1200	3
2	12	0	28		400	1600	4
3	13	0	14		400	440	1.1
3	14	0	14		400	500	1.25
3	15	0	14		400	700	1.75
3	16	0	14		400	800	2
3	17	0	14		400	1200	3
3	18	0	14		400	1600	4
4	19	14	14		600	660	1.1
4	20	14	14		600	750	1.25
4	21	14	14		600	1050	1.75
4	22	14	14		600	1200	2
4	23	14	14		600	1800	3
4	24	14	14		600	2400	4
5	25	28	14		400	440	1.1
5	26	28	14		400	500	1.25
5	27	28	14		400	700	1.75
5	28	28	14		400	800	2
5	29	28	14		400	1200	3
5	30	28	14		400	1600	4
6	31	28	28		400	440	1.1
6	32	28	28		400	500	1.25
6	33	28	28		400	700	1.75
6	34	28	28		400	800	2
6	35	28	28		400	1200	3
6	36	28	28		400	1600	4
7	37	0	14		600	660	1.1
7	38	0	14		600	750	1.25
7	39	0	14		600	1050	1.75
7	40	0	14		600	1200	2
7	41	0	14		600	1800	3
7	42	0	14		600	2400	4
8	43	14	28		400	440	1.1
8	44	14	28		400	500	1.25
8	45	14	28		400	700	1.75
8	46	14	28		400	800	2
8	47	14	28		400	1200	3
8	48	14	28		400	1600	4

Table A6: Multiple Price List Decision Problems

Set	Decision	Front-End Delay (t)	Early Later Delay (k)	vs.	Early Max	Later Max	$1 + r$
1	1	14	14		400	440	1.1
1	2	14	14		400	500	1.25
1	3	14	14		400	700	1.75
1	4	14	14		400	800	2
1	5	14	14		400	1200	3
1	6	14	14		400	1600	4
2	7	0	28		400	440	1.1
2	8	0	28		400	500	1.25
2	9	0	28		400	700	1.75
2	10	0	28		400	800	2
2	11	0	28		400	1200	3
2	12	0	28		400	1600	4
3	13	0	14		400	440	1.1
3	14	0	14		400	500	1.25
3	15	0	14		400	700	1.75
3	16	0	14		400	800	2
3	17	0	14		400	1200	3
3	18	0	14		400	1600	4
4	19	14	28		400	440	1.1
4	20	14	28		400	500	1.25
4	21	14	28		400	700	1.75
4	22	14	28		400	800	2
4	23	14	28		400	1200	3
4	24	14	28		400	1600	4

Figure A1: Histograms of Individual-Level ρ_i Parameter Estimates




Figure A2: Screenshot of a CTB Decision

The screenshot displays a decision interface for a CTB (Continuous Time Budget) choice. At the top, a horizontal bar with a color gradient from green to blue is labeled 'V'. Below the bar, the words 'EARLIER' and 'LATER' are positioned at the left and right ends, respectively. Two boxes provide details for each option:

- EARLIER:** Paid on July 27, 2015; Maximum: 400; **207 Ksh**
- LATER:** Paid on August 10, 2015; Maximum: 600; **290 Ksh**

Below these boxes, a note states: "August 10, 2015 is 2 weeks after July 27, 2015". At the bottom, there are two buttons: a red "Clear" button and a green "OK" button.

Figure A3: Screenshot of an MPL Decision

The screenshot displays a decision interface for an MPL (Multi-Period Lottery) choice, titled "Choice 1". It features two boxes with payment details:

- EARLIER:** Paid on August 10, 2015; **400 KSH**
- LATER:** Paid on August 24, 2015; **440 KSH**

Below these boxes, a note states: "August 24, 2015 is 2 weeks after August 10, 2015". At the bottom, there are two buttons: a green "CHOOSE THE EARLIER PAYMENT" button and a yellow "CHOOSE THE LATER PAYMENT" button.