Measuring time preferences: A comparison of experimental methods[☆]

James Andreoni[a], Michael A. Kuhn[b], Charles Sprenger[c,*]

[a] *University of California, San Diego, Department of Economics, 9500 Gilman Drive #0508, La Jolla, CA 92093, United States*
[b] *University of Oregon, Department of Economics, 1285 University of Oregon, Eugene, OR 97403, United States*
[c] *University of California, San Diego, Rady School of Management, 9500 Gilman Drive #0553, La Jolla, CA 92093, United States*

A B S T R A C T

Eliciting time preferences has become an important component of both laboratory and field experiments, yet there is no consensus as how to best measure discounting. We examine the predictive validity of two recent, simple, easily administered, and individually successful elicitation tools: convex time budgets (CTB) and double multiple price lists (DMPL). Using similar methods, the CTB and DMPL are compared using within- and out-of-sample predictions. While each perform equally well within sample, the CTB significantly outperforms the DMPL on out-of-sample measures.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Time preferences are fundamental to theoretical and applied studies of decision-making, and are a critical element of much of economic analysis. At both the aggregate and individual level, accurate measures of discounting parameters can provide helpful guidance on the potential impacts of policy and provide useful diagnostics for effective policy targeting.

Though efforts have been made to identify time preferences from naturally occurring field data,[1] the majority of research has relied on laboratory samples using variation in monetary payments.[2] Despite many attempts, however, the experimental

---

[1] These methods investigate time preferences by examining durable goods purchases, consumption profiles or annuity choices (Hausman, 1979; Lawrance, 1991; Warner and Pleeter, 2001; Gourinchas and Parker, 2002; Cagetti, 2003; Laibson et al., 2003, 2007). While there is clear value to these methods they may not be practical for field settings with limited data sources or where subjects make few comparable choices.

[2] Chabris et al. (2008b) identify several important issues related to this research agenda, calling into question the mapping from experimental choice to corresponding model parameters in monetary discounting experiments. Paramount among these issues are clear arbitrage arguments such that responses in monetary experiments should reveal only the interval of borrowing and lending rates, and thus limited heterogeneity in behavior if subjects face similar

community lacks a clear consensus on how best to measure time preferences; a point made clear by Frederick et al. (2002). One natural challenge which has gained recent attention is the confounding effect of utility function curvature. Typically, linear utility is assumed for identification, invoking expected utility's necessity of risk neutrality for small stakes decisions (Rabin, 2000). However, in an important contribution, Andersen et al. (2008) show that if utility is assumed to be linear in experimental payoffs (risk neutrality) when it is truly concave (risk aversion), estimated discount rates will be biased upwards.[3] This observation has reset the investigation of new elicitation tools.

Andersen et al. (2008) (henceforth AHLR) use of measures of risk taking to incorporate utility function curvature, which we refer to as a double multiple price list (DMPL: one multiple price list for time and one for risk). Andreoni and Sprenger (2012a) (henceforth AS) used variation in linear budget constraints over early and later income to identify convexity of preferences, a device they call a convex time budget (CTB). This technique is motivated by early developments in risk preference elicitation such as Gneezy and Potters (1997) and is already being used in field settings (e.g. Giné et al., 2012). The objective of this study is to work toward a consensus by comparing these two methods.

Our comparison criteria are both experimental and empirical. The key experimental criterion is simplicity. In particular, researchers eliciting preferences put a premium on devices that are simple for subjects, easy to administer, transportable to the field, and can be easily folded into a larger research design. Both methods seem to succeed equally well on this dimension.

More central to our analysis, we propose empirical predictive validity as the second and most relevant criterion. In particular, parameter estimates generated from a specific data set should yield good in-sample fit, have out-of-sample predictive power, and predict relevant, genuine economic activity.[4]

We document two main findings when examining predictive validity. First, we reproduce the broad conclusions of both AHLR and AS: there are clear confounding effects of utility function curvature that need to be controlled for in estimating discounting. Second, when taking these estimates out-of-sample we find that the CTB-based estimates markedly outperform the DMPL-based estimates when predicting intertemporal choice.

Determining why the CTB outperforms the DMPL is not the main focus of this paper. However, we suggest that there are three important theoretical distinctions that can guide the design of future preference elicitation techniques:

- using only domain-specific data to identify preferences,
- designing the elicitation to permit preferred estimation strategies,
- increasing the preference-identifying informational content of each choice.

In the context of our comparison exercise, the CTB lets the researcher avoid the worry that the time and risk domain may not be perfectly related, use demand theory rather than a probabilistic choice model for identification[5] and represent each choice as defining an equality rather than an inequality constraint.[6] Section 2 discusses each of these issues in more detail.

The issue of informational content is closely related to a criticism of the CTB: that the high frequency of observed corner choices is a shortcoming of the technique (Harrison et al., 2013). In fact, the frequency of corner solutions in the CTB is precisely its feature that generates the empirical improvements in predictive validity. This is because a corner solution from a CTB carries more information about preferences than the exact same choice from a DMPL. Specifically, it implies that preferences over time-dated experimental payments may be close to linear. In contrast, curvature in DMPL-elicited preferences is informed primarily from risky choices, not from choices over time. Hence, CTB and DMPL estimates differ largely in their identified degree of utility function curvature. Indeed, it is the near linearity in CTB estimated preferences that generate the improved predictive performance when compared to DMPL estimates.

Section 2 describes our preference elicitation techniques and experimental protocol. Section 3 presents estimation results and evaluates the success of the CTB and DMPL at predicting choice both in- and out-of-sample. Section 4 concludes.

---

credit markets (Cubitt and Read, 2007; Andreoni and Sprenger, 2012a,b). This last concern may be beyond the reach of most experimental samples. Evidence from Coller and Williams (1999) suggests that even when the entire arbitrage argument is explained to subjects, heterogeneity remains and responses do not collapse to reasonable intervals of borrowing and lending rates. Following most of the literature, the experiments we conduct will focus on monetary choices, taking the laboratory offered rates as the relevant ones for choice. Importantly, the methods we describe are easily portable to other domains with less prominent fungibility problems. One recent example using the convex time budget described below with choices over effort is Augenblick et al. (2013).

[3] Frederick et al. (2002) also provide discussion of this confound and present three strategies for disentangling utility function curvature from time discounting: (1) eliciting utility judgments such as attractiveness ratings at two points in time; (2) eliciting preferences over temporally separated probabilistic prospects to exploit the linearity-in-probability property of expected utility; and (3) "separately elicit the utility function for the good in question, and then use that function transform outcome amounts to utility amounts, from which utility discount rates could be computed" (p. 382). The third of these techniques is close in spirit to the double multiple price list implemented by Andersen et al. (2008) described below.

[4] Though this seems a natural objective, there are relatively few examples of research linking laboratory measures of time preference to other behaviors or characteristics (Ashraf et al., 2006; Dohmen et al., 2010; Meier and Sprenger, 2010, 2012; Mischel et al., 1989). These exercises at times demonstrate the lack of explanatory power for prior time preference estimates (Chabris et al., 2008a).

[5] In this sense, the CTB elicitation and estimation techniques are not separate advances: they go hand in hand.

[6] Or, depending on econometric approach, multiple inequality constraints simultaneously rather than a single inequality constraint.

2. Techniques and protocol

Before introducing the two considered elicitation techniques, we first outline the nature of preferences. Consider allocations of experimental payments, x_t and x_{t+k} between two periods, t and $t+k$. Preferences over these experimental payments are assumed to be captured by a stationary, time-independent constant relative risk averse utility function $u(x_t) = x_t^\alpha$. We assume a quasi-hyperbolic structure for discounting (Laibson, 1997; O'Donoghue and Rabin, 1999), such that preferences over bundles are described by

$$U(x_t, x_{t+k}) = \begin{cases} x_t^\alpha + \beta\delta^k x_{t+k}^\alpha & \text{if} \quad t = 0 \\ x_t^\alpha + \delta^k x_{t+k}^\alpha & \text{if} \quad t > 0. \end{cases} \tag{1}$$

The parameter δ captures standard long-run exponential discounting, while the parameter β captures a specific preference toward payments in the present, $t=0$. The one period discount factor between the present and a future period is $\beta\delta$, while the one period discount factor between two future periods is δ. Present bias is associated with $\beta < 1$ and $\beta = 1$ corresponds to the case of standard exponential discounting.[7]

We consider two elicitation techniques, the DMPL and the CTB, designed to provide identification of the three parameters of interest, α, δ, and β, corresponding to utility function curvature, long-run discounting, and present bias, respectively. Given that any functional form of utility one estimates will be misspecified to some degree, different methods are likely to yield different parameter estimates. While these differences are important, our view us that the first concern is to have a method that is useful as a predictive tool for the research community.[8]

2.1. Elicitation techniques

We begin by presenting the DMPL, which consists of two stages. The first identifies discounting, potentially confounded by utility function curvature. The second is designed to un-confound the first stage by providing information on utility function curvature through decisions on risky choice. In the first stage, individuals make a series of binary choices between smaller sooner payments and larger later payments. Such binary choices are organized into multiple price lists (MPLs) in order of increasing gross interest rate (Coller and Williams, 1999; Harrison et al., 2002). Where an individual switches from preferring the smaller sooner payment to the larger later payment carries interval information on discounting. Fig. 1, Panel A, presents a sample intertemporal MPL.[9]

Importantly, one cannot make un-confounded inference for time preferences based on these intertemporal responses alone. Consider an individual who prefers \$X at time t over \$Y at time $t+k$, but prefers \$Y at time $t+k$ over \$X′<\$X at time t. If $t \neq 0$ then one can infer the bounds on δ to be $\delta \in (X'^\alpha/Y^\alpha, X^\alpha/Y^\alpha)$. Though standard practice for identifying δ often (at times implicitly) assumes linear utility, $\alpha = 1$, it is clear that a concave utility function, $\alpha < 1$, will bias discount factor estimates downwards, understating the true bounds.[10] Further, without some notion of the extent of curvature, one cannot un-confound the measure. This motivates the second stage.

The second stage of the DMPL is designed to account for utility function curvature by introducing a second experimental measure. In particular, a Holt and Laury (2002, henceforth HL) risk preference task is conducted alongside the intertemporal decisions. Subjects face a series of decisions between a safe and a risky binary gamble. The probability of the high outcome in each gamble increases as one proceeds through the task, such that where a subject switches from the safe to the risky gamble carries information on risk attitudes. Fig. 1, Panel B, presents a sample HL task. The risk attitude elicited in the HL task identifies the degree of utility function curvature, α, which is then applied to the intertemporal choices to un-confound the discounting bounds. In effect, α is identified from risky choice data, and δ and β are identified from intertemporal choice data.

The CTB takes a different approach to identification. Instead of incorporating a second experimental elicitation, the CTB recognizes a key restriction of the standard multiple price list approach. When making a binary choice between a smaller sooner payment, \$X, and a larger later payment, \$Y, subjects are effectively restricted to the corner solutions in (sooner, later) space, (\$X, \$0) and (\$0, \$Y). That is, they maximize the utility function in (1) subject to the discrete budget constraint

---

[7] We abstract away from any discussion of sophistication or naiveté wherein individuals are potentially aware of their predilection of being more impatient in the present than they are in the future. Our implemented experimental techniques will be unable to distinguish between the two.

[8] While the quasi-hyperbolic utility function offers the intuitive appeal of separate present-bias and exponential discounting estimates and thus nests the exponential discounting model, the hyperbolic utility function, $U(x_t) = \frac{1}{1+\gamma t} \cdot x_t^\alpha$, is an alternative specification widely used in other literatures. Appendix A.9 presents estimates of discounting using this functional form. This approach leads to lower estimates of annual discounting for both the CTB and DMPL methods, while estimates of utility curvature are unaffected.

[9] This implementation appears slightly different from others for coherence with our implementation of the CTB. In effect, individuals choose between smaller sooner payments and larger later payments. However, we clarify that choosing the smaller sooner payment implies a subject will receive zero at the later date, and vice versa.

[10] Correspondingly, a convex utility function biases discount factors upwards. A similar issues exists for identifying β when $t=0$.

TODAY _and_ 5 WEEKS from today

For each decision number (1 to 6) below, decide the **AMOUNTS** you would like for sure **today** AND **in 5 weeks** by checking the corresponding box.

Example: In Decision 1, if you wanted $19.00 today and $0 in five weeks you would check the left-most box. Remember to check only one box per decision!

1.	payment TODAY …	$19.00		$0
	and payment in 5 WEEKS	$0		$20.00
		☐		☐
2.	payment TODAY …	$18.00		$0
	and payment in 5 WEEKS	$0		$20.00
		☐		☐
3.	payment TODAY …	$17.00		$0
	and payment in 5 WEEKS	$0		$20.00
		☐		☐
4.	payment TODAY …	$16.00		$0
	and payment in 5 WEEKS	$0		$20.00
		☐		☐
5.	payment TODAY …	$14.00		$0
	and payment in 5 WEEKS	$0		$20.00
		☐		☐
6.	payment TODAY …	$11.00		$0
	and payment in 5 WEEKS	$0		$20.00
		☐		☐

Panel A: Intertemporal Multiple Price List

Decision		Option A						Option B				
		If the die reads	you receive	and	If the die reads	you receive		If the die reads	you receive	and	If the die reads	you receive
1	☐	1	15		2-10	8.31	☐	1	20		2-10	0.52
2	☐	1-2	15		3-10	8.31	☐	1-2	20		3-10	0.52
3	☐	1-3	15		4-10	8.31	☐	1-3	20		4-10	0.52
4	☐	1-4	15		5-10	8.31	☐	1-4	20		5-10	0.52
5	☐	1-5	15		6-10	8.31	☐	1-5	20		6-10	0.52
6	☐	1-6	15		7-10	8.31	☐	1-6	20		7-10	0.52
7	☐	1-7	15		8-10	8.31	☐	1-7	20		8-10	0.52
8	☐	1-8	15		9-10	8.31	☐	1-8	20		9-10	0.52
9	☐	1-9	15		10	8.31	☐	1-9	20		10	0.52
10	☐	1-10	15		-	8.31	☐	1-10	20		-	0.52

Panel B: Holt-Laury Risk Elicitation

Fig. 1. Sample DMPL decision sheets.

$(x_t, x_{t+k}) \in \{(X, 0), (0, Y)\}$. If the utility function is indeed linear, such that $\alpha = 1$, the restriction to corners is non-binding.[11] However, if $\alpha < 1$, individuals have convex preferences in (sooner, later) space, preferring interior solutions, and leading the restriction to corners to meaningfully restrict behavior.

---

[11] While the restriction of the data to corner solutions in non-binding in the case of linear utility, it does not mean that the same set of choices on restricted and unrestricted data will yield the same parameter estimates. Corner choices on unrestricted data have very different implications for utility curvature than corner choices on restricted data.

	TODAY _and_ 5 WEEKS from today					
	For each decision number (1 to 6) below, decide the **AMOUNTS** you would like for sure **today** **AND** **in 5 weeks** by checking the corresponding box. _Example:_ In Decision 1, if you wanted $19.00 today and $0 in five weeks you would check the left-most box. Remember to check only one box per decision!					
1.	payment TODAY …	$19.00	$15.20	$11.40	$7.60	$3.80	$0
	and payment in 5 WEEKS	$0	$4.00	$8.00	$12.00	$16.00	$20.00
		☐	☐	☐	☐	☐	☐
2.	payment TODAY …	$18.00	$14.40	$10.80	$7.20	$3.60	$0
	and payment in 5 WEEKS	$0	$4.00	$8.00	$12.00	$16.00	$20.00
		☐	☐	☐	☐	☐	☐
3.	payment TODAY …	$17.00	$13.60	$10.20	$6.80	$3.40	$0
	and payment in 5 WEEKS	$0	$4.00	$8.00	$12.00	$16.00	$20.00
		☐	☐	☐	☐	☐	☐
4.	payment TODAY …	$16.00	$12.80	$9.60	$6.40	$3.20	$0
	and payment in 5 WEEKS	$0	$4.00	$8.00	$12.00	$16.00	$20.00
		☐	☐	☐	☐	☐	☐
5.	payment TODAY …	$14.00	$11.20	$8.40	$5.60	$2.80	$0
	and payment in 5 WEEKS	$0	$4.00	$8.00	$12.00	$16.00	$20.00
		☐	☐	☐	☐	☐	☐
6.	payment TODAY …	$11.00	$8.80	$6.60	$4.40	$2.20	$0
	and payment in 5 WEEKS	$0	$4.00	$8.00	$12.00	$16.00	$20.00
		☐	☐	☐	☐	☐	☐

Fig. 2. Sample CTB decision sheet.

This observation leads to a natural solution. If one wishes to identify preferences in (sooner, later) space, one can convexify the decision environment. In a CTB, subjects are given the choice of (X, $0), ($0, Y) or anywhere along the intertemporal budget constraint connecting these points such that $Px_t + x_{t+k} = Y$, where $P = \frac{Y}{X}$ represents the gross interest rate. Fig. 2 presents a sample CTB allowing for interior solutions between the two corners.[12] The most important distinction between the two methods is the source of identification of curvature. The DMPL identifies utility function curvature based on the degree of risk aversion elicited in the HL risky choice. In contrast, the CTB identifies curvature based on the degree of price sensitivity in intertemporal choice. These varying sources of information for the shape of the utility function should be equivalent under the utility formulation in (1). The parameter α determines both the extent of intertemporal substitution and the extent of risk aversion.[13] However, there may be reason to expect differences in the extent of measured utility function curvature and hence discounting estimates across the two methods. AHLR document substantial utility function curvature in HL tasks, leading to substantial changes in discounting estimates when accounted for in the DMPL. In contrast, AS document substantially less utility function curvature from CTB choices.[14]

2.2. Experimental design

In order to assess the predictive validity of the DMPL and CTB elicitation methods, we designed a simple within-subject experiment. Subjects faced 4 intertemporal MPLs, 2 HL risk tasks, and 4 CTBs of the form presented in Figs. 1 and 2. For the intertemporal decisions the CTBs and MPLs took the exact same start dates, t, delay lengths, k, and gross interest rates, P. The experimental budget was always $20 such that the intertemporal budget constraint in each decision was $Px_t + x_{t+k} = 20$. Hence, as presented in Figs. 1 and 2, the only difference between the implemented CTBs and MPLs was the presence of interior allocations. Table 1 summarizes the parameters of the intertemporal choice portion of the experiment. The interest rates, experimental budgets and delay lengths are chosen to be comparable to those of AS. As presented in Fig. 1, Panel B, in the two HL tasks subjects faced a series of decisions between a safe and a risky gamble. In the first HL task, HL_1, the safe gamble outcomes were $10.39 and $8.31, while the risky gamble outcomes were $20 and $0.52. In the second HL task, HL_2, the safe gamble outcomes were $13.89 and $5.56, while the risky gamble outcomes were $25 and $0.28. These values were

---

[12] Notably, the version of the CTB we use is different than that of AS. AS used a computer interface to offer individuals 100 tokens that could be allocated to the sooner or later payoffs in any proportion. By condensing the budget to 6 options, we can represent the choice in a check-the-box format that fits onto a sheet of paper. While information is lost in this discretization, it puts the CTB on the same footing as the DMPL in terms of ease-of-administration and portability.

[13] Provided α is the sole source of curvature and expected utility maintains in atemporal choice.

[14] However, the AS estimates do differ significantly from linear utility. Further, AS show that the extent of CTB utility function curvature is correlated with the distance between standard price list discount factor estimates and CTB discount factor estimates. Individuals with more concave CTB-measured utility functions have more downwards-biased discount factor price list estimates.

Table 1
Intertemporal experimental parameters.

Choice set	t (days until first payment)	k (delay)	P (price ratios): $Px_t + x_{t+k} = 20$
CTB_1, MPL_1	0	35	1.05, 1.11, 1.18, 1.25, 1.43, 1.82
CTB_2, MPL_2	0	63	1.00, 1.05, 1.18, 1.33, 1.67, 2.22
CTB_3, MPL_3	35	35	1.05, 1.11, 1.18, 1.25, 1.43, 1.82
CTB_4, MPL_4	35	63	1.00, 1.05, 1.18, 1.33, 1.67, 2.22

Note: The price ratios for $k = 35$ correspond to yearly (compounded quarterly) interest rates of 65%, 164%, 312%, 529%, 1301% and 4276%. The price ratios for $k = 63$ correspond to rates of 0%, 33%, 133%, 304%, 823% and 2093%.

chosen to provide a measure of curvature at monetary payment values close to those implemented in the intertemporal choices and are scaled versions of those used in the original HL tasks.[15]

Our sample consists of 64 undergraduates, evenly divided into 4 sessions, conducted in February of 2009. Upon arriving in the laboratory, subjects were told they would be participating in an experiment about decision-making over time. Subjects were told that based on the decisions they made, and chance, they could receive payment as early as the day of the experiment, as late as 14 weeks from the experiment, or other dates in between. All of the payments dates were selected to avoid holidays or school breaks, and all payments were designed to arrive on the same day of the week. All choices were made with paper and pencil and the order in which subjects completed the tasks was randomized. Two orders were implemented with the HL tasks acting as a buffer between the more similar time discounting choices: (1) MPL, HL, CTB; (2) CTB, HL, MPL.[16] Subjects were told that in total they would make 49 decisions. One of these decisions would be chosen as the 'decision-that-counts' and their choice would be implemented.[17] The full instructions are provided in Appendix A.10.

A primary concern in the design of discounting experiments is to equalize all transaction costs between different dates of payment. Eliminating any uncertainty over delayed payments and convenience of immediate payments is key to obtaining accurate results. We follow the techniques used in AS and take six specific measures to equate transaction costs and ensure payment reliability.[18] Subjects were surveyed extensively after the completion of the experiment. Importantly, 100% of subjects said that they believed that their earnings would be paid out on the appropriate dates.

Once the decision-that-counts was chosen, subjects participated in a Becker et al. (1964, henceforth BDM) auction eliciting their lowest willingness-to-accept amount in their sooner payment to forgo a claim to an additional $25 in their later payment with a uniform distribution of random prices drawn from [$15.00, $24.99]. This was presented as a bonus that would build on the previous earnings of one individual in the session, drawn at random at the end of the study.[19] The instructions outlined the procedure and explicitly informed subjects that "the best idea is to write down your true value . . .".[20] Subsequently, subjects completed a survey including demographic details as well as two hypothetical measures of patience. The first hypothetical measure asked subjects to state the dollar amount of money today that would make them indifferent to $20 in one month. The second hypothetical measure asked subjects to state the mount of money in one month that would make them indifferent to $20 today.[21]

While there were 64 subjects in total, our estimation sample for the remainder of the paper consists of 58 individuals. Five individuals exhibited multiple switching at some point in the HL task. One individual never altered their decision from a specific corner solution in all 4 CTBs and thus provided insufficient variation for the calculation of utility parameters. These

---

[15] See Appendix A.10 for the full instructions. In the HL baseline task, the safe gamble outcomes were $2.00 and $1.60 and the risky gamble outcomes were $3.85 and $0.10. Our HL_1 scales the largest payment to $20 and keeps all ratios the same. The second task, HL_2, increases the highest payment to $25 and increases the variance.

[16] No order effects were observed.

[17] Our randomization device for implementing the decision-that-counts favored the intertemporal choices over the HL choices. Whereas each time preference allocation was viewed as a choice (48 in total), the HL tasks were viewed as a single choice. When the HL tasks were explained, subjects were told that if these were chosen as the decision-that-counts, then a specific HL choice would be picked at random (with equal likelihood) and a 10-sided dice would be rolled to determine lottery outcomes. Payment would be made in cash immediately in the lab, and subjects would receive a show-up fee of $10 immediately as well. We recognize that this favored randomization may limit the attention subjects pay to the HL tasks. Our results, however, are comparable to other findings of risk aversion in Holt and Laury (2002) and to other implementations of the DMPL (Andreoni and Sprenger, 2012b).

[18] As in AS, all participants lived on campus at UC San Diego, which meant that they had 24 h access to a locked personal mailbox. Our first measure was to use these mailboxes for intertemporal payments. Second, intertemporal payments were made by personal check from Professor James Andreoni. Although this introduces a transaction cost, it ensures an equal cost in all potential periods of distribution. In addition, these checks were drawn on an account at the on-campus credit union. Third, for intertemporal payments the $10 show-up fee was split into two $5 minimum payments avoiding subjects loading on one experimental payment date to avoid cashing multiple checks. Fourth, the payment envelopes were self-addressed, reducing risk of clerical error. Fifth, subjects noted payment amounts and dates from the decision-that-counts on their payment envelopes, eliminating the need to recall payment values and reducing the risk of mistaken payment. Sixth, all subjects received a business card with telephone and e-mail contacts they could use in case a payment did not arrive. Subjects were made aware of all of these measures prior to the choice tasks.

[19] Subjects were potentially aware of their payment amounts at this point if they remembered their choice exactly.

[20] This follows the protocol of Ariely et al. (2003). A copy of the elicitation and instructions can be found in Appendix A.10.

[21] The exact wording of the first question was 'What amount of money, $X, if paid to you today would make you indifferent to $20 paid to you in one month?' The exact wording of the first question was 'What amount of money, $Y, would make you indifferent between $20 today and $Y one month from now?'

6 subjects are dropped to maintain a consistent number of observations across estimates. This choice does not alter the conclusions of our paper.[22]

2.3. Parameter estimation strategies

The data collected in the experiment are used to separately identify the key parameters of utility function curvature, α, discounting, δ, and present bias, β for both the CTB and the DMPL. Preferred estimation strategies for recovering these parameters differ between the two elicitation techniques. The CTB is akin to maximizing discounted utility subject to a future value budget constraint. Hence, a standard intertemporal Euler equation maintains,

$$MRS = \frac{x_t^{\alpha-1}}{\beta^{t_0}\delta^k x_{t+k}^{\alpha-1}} = P,$$

where t_0 is an indicator for whether $t = 0$. This can be rearranged to be linear in our experimental variations, t, k, and P,

$$\ln\left(\frac{x_t}{x_{t+k}}\right) = \frac{\ln(\beta)}{\alpha-1}t_0 + \frac{\ln(\delta)}{\alpha-1}k + \frac{1}{\alpha-1}\ln(P). \tag{2}$$

Assuming an additive error structure, this is estimable at either the group or individual level, with parameters of interest recovered via non-linear combinations of regression coefficients and standard errors calculated via the delta method. Eq. (2) makes clear the mapping from the variation of experimental parameters to structural parameter estimates. Variation in the gross interest rate, P, delivers the utility function curvature, α. For a fixed interest rate, variation in delay length, k, delivers δ, and variation in whether the present, $t = 0$, is considered delivers β.

Three natural issues arise with the estimation strategy described above. First, the allocation ratio $\ln\left(\frac{x_t}{x_{t+k}}\right)$ is not well defined at corner solutions.[23] Second, even if the optimality condition were defined at corner solutions, the preferences we assume cannot generate such choices in the form of point-identified maxima.[24] Indeed, this issue is a common point of criticism of CTB approaches (Harrison et al., 2013). Third, this strategy effectively ignores the interval nature of the data, created by the discretization of the budget constraint.

To address the first issue, one can use the demand function to generate a non-linear regression equation based upon

$$x_t = \frac{20(\beta^{t_0}\delta^k P)^{\frac{1}{\alpha-1}}}{1 + P(\beta^{t_0}\delta^k P)^{\frac{1}{\alpha-1}}}, \tag{3}$$

which avoids the problem of the logarithmic transformation in (2). However, this demand function is only defined for $\alpha \in (0, 1)$, so the use of either of these techniques is still subject to bias incurred by the second issue above.[25] While this issue is minimized by the fact that our metric for success is predictive validity, we propose a third technique, interval censored tobit (ICT) regression, that is robust to all three issues mentioned above. While this technique is less transparent and more complicated to perform, it serves as a robustness check for approaches (2) and (3). The details are discussed in Appendix A.1.[26]

Preferred methodology for estimating intertemporal preference parameters from DMPL data, as per AHLR, relies on maximum likelihood methods. Binary choices between $X sooner and $Y later are assumed to be guided by the utilities $U_X = \delta^t X^\alpha$ and $U_Y = \beta^{t_0}\delta^{t+k}Y^\alpha$. AHLR assign choice probabilities using Luce's (1959) formulation based on these utility values:

$$Pr(Choice = X) = \frac{U_X^{\frac{1}{\nu}}}{U_X^{\frac{1}{\nu}} + U_Y^{\frac{1}{\nu}}}, \tag{4}$$

---

[22] Excluding multiple switchers on MPL and HL tasks is common practice in the field. However, given the probabilistic models of choice used to estimate preferences, multiple switching should occasionally occur. Therefore, Appendix A.8 reproduces our main estimates for the sample that includes the multiple switchers.

[23] In our application we solve this issue operationally, by transforming the $0 payment in a corner solution to $0.01 such that the log allocation ratio is always well-defined. Additionally, we consider exercises adding in the fixed $5 minimum payments to each payment date and qualitatively similar results. See Appendix Table A2.

[24] Related to these two points is the issue of whether background income consumption outside of the laboratory need to be integrated into the estimation. Past work, including Andersen et al. (2008) includes an estimate of daily consumption as a baseline for experimental income. As we are concerned about the use of consumption and income jointly, we avoid this approach and for the purposes of our main results treat the experimental income as a positive prospect viewed in isolation. In Appendix A.5, we integrate the $5 show-up fee payments into the estimation. This has the mechanical effect of making individuals appear less willing to accept income receipts of zero in some periods, thus generating estimates of a marginal utility of income that diminishes at an unbelievable rate and risk aversion of an impossible degree. Despite this, the qualitative differences between the DMPL and CTB estimates are the same, implying similar relative predictive abilities of the techniques.

[25] Assuming that the degree of misspecification depends on the experimentally varied parameters to some degree, this will be problematic.

[26] AS provide a variety of estimates using both demand functions and Euler equations and several utility formulations such as CARA and CRRA. Broadly consistent estimates are found across techniques.

Table 2

Aggregate utility parameter estimates.

	Discounting	Curvature	Discounting and curvature			
Elicitation method:	MPL	HL	DMPL	CTB		
Estimation method:	ML	ML	ML	OLS	NLS	ICT
	(1)	(2)	(3)	(4)	(5)	(6)
Utility parameters						
r	1.022	–	0.472	0.741	0.679	0.630
	(0.223)	–	(0.103)	(0.390)	(0.148)	(0.230)
β	0.986	–	0.992	1.010	0.988	0.997
	(0.010)	–	(0.006)	(0.022)	(0.009)	(0.016)
α	–	0.549	0.549	0.947	0.928	0.867[a]
	–	(0.044)	(0.044)	(0.003)	(0.007)	(0.017)
Error parameters						
ν	0.085	–	0.046	–	–	–
	(0.010)	–	(0.007)	–	–	–
μ	–	0.096	0.096	–	–	–
	–	(0.010)	(0.010)	–	–	–
Clustered SE's	Yes	Yes	Yes	Yes	Yes	Yes
# Clusters	58	58	58	58	58	58
N	1392	1160	2552	1392	1392	1392
Log likelihood	−546	−327	−873	–	–	−2102
R^2	–	–	–	0.401	0.591	–

Note: Standard errors clustered at the individual level in parentheses. Each individual made 20 decisions on the HL, 24 decision on the MPL (and therefore 44 decisions on the DMPL) and 24 decisions on the CTB. In columns (1) through (3) HL, MPL and DMPL estimates are obtained via maximum likelihood using Luce's (1959) stochastic error probabilistic choice model. The CTB is estimated in three different ways: ordinary least squares (OLS) using the Euler Eq. (2), non-linear least squares (NLS) using the demand function (3) and interval-censored tobit (ICT) maximum likelihood using the Euler equation (2). All maximum likelihood models are estimated using the Broyden–Fletcher–Goldfarb–Shanno (BFGS) optimization algorithm in Stata 10. Elements of this algorithm have changed in subsequent versions of Stata.

[a] The ICT estimate for α is only identified up to a constant of proportionality. See Appendix A.1 for details.

where ν represents stochastic decision error. As ν tends to infinity all decisions become random and as ν tends to zero, all decisions are deterministic based on the assigned utilities. The log of this choice probability represents the likelihood contribution of a given observation.

In order to simultaneously estimate utility function curvature and discounting parameters, AHLR also define a similar likelihood contribution for a HL risk task observation, constructed under expected utility. An alternate stochastic decision error parameter, μ, is estimated for risky choice. As in AHLR, we provide estimates based on only the intertemporal decisions, assuming $\alpha = 1$, and on the combination of time and risk choices. We additionally provide estimates using only the risky data to demonstrate the extent to which estimated utility function curvature is informed by the HL choices. Appendix A.2 provides full detail of the maximum likelihood strategies for DMPL data.

A subtle, but critical difference between these estimation strategies is how choice 'errors', instances in which the option with the highest utility conditional on the estimated parameters is not selected, occur. Errors enter the CTB specification nested in the context of optimality: unobserved mean-zero shocks specific to one decision that perturb the tangency condition from what would be expected based on estimated parameters. In the DMPL framework, 'errors' come from estimated parameters, ν and μ, that are constant across the estimation sample, and represent how deterministic the relationship is between utility, conditional on estimated parameters, and choice. An econometric model of probabilistic choice cannot be derived from a model of economic optimization without the use of a specialized distributional assumption on the unobservables. If one is concerned about the applicability of the estimates to a more general choice space, it is worth carefully evaluating the preferred source of the structural assumptions that provide identification. We return to this issue in Section 3.2.3.

3. Results

We present the results in two stages. First, we provide estimation results based on the DMPL and CTB elicitation techniques, drawing some contrasts between the parameter estimates across the two methods. Second, we move to choice prediction and conduct two complementary analyses, attempting to predict choice across methods and attempting to predict choice out-of-sample to our BDM and hypothetical choice data.

3.1. Parameter estimates

Our main estimation results are presented in Table 2, providing aggregate estimates of α, β, and an annualized discount rate $r = \delta^{-365} - 1$ for both elicitation techniques and the variety of estimation strategies described in Section 2.3.[27] Standard errors are clustered on the individual level. We also estimate the parameters of interest on an individual level. These estimates will be used for the prediction exercises in the following section and the median individual estimates correspond generally to those in Table 2. These results and additional discussion are found in Appendix A.4. To begin, in columns (1) and (2) we separately analyze the two components of the DMPL. In column (1), we assume linear utility and use the intertemporal choice data to estimate β and r. When assuming linear utility, we estimate an annual discount rate of 102.2 percent (s.e. 22.3 percent). In column (2), we use only the HL data to estimate utility function curvature, estimating α of 0.549 (0.044), comparable to other experimental findings on the extent of small stakes risk aversion (e.g., Holt and Laury, 2002). Based on this curvature estimate, an individual would be indifferent between a 50–50 gamble over $20 and $0 and $5.67 for sure, implying a risk premium of $4.33. The extent of concavity found in column (2) suggests that the estimated annual discount rate of 102% in column (1) is dramatically upwards-biased. In column (3) we use both elements of the DMPL to simultaneously estimate utility function curvature and discounting. Indeed, we find that the estimated annual discount rate falls dramatically to 47.2% (s.e. 10.3%). The difference in discounting with and without accounting for curvature is significant at all conventional levels, ($\chi^2(1) = 15.71$, $p < 0.01$). This finding echoes those of AHLR, though our estimated discount rates are higher in general. Note that the curvature estimate is virtually identical across columns (2) and (3), indicating the extent to which the measure is informed by risky choice responses.

Next, we consider the CTB estimates. Table 2, columns (4)–(6) contain estimates based on the three methods described in Section 2.3. In column (4), ordinary least squares estimates based on the Euler equation (2) are presented.[28] The annual discount rate is estimated to be 74.1% (s.e. 39%), generating wide intervals for the extent of discounting. Hence, the discounting estimate from the DMPL method would lie in the 95 percent confidence interval of the CTB estimate. Importantly, the estimates of utility function curvature in column (4) are far closer to linear utility than that obtained from the DMPL. Based on CTB methods, we estimate α of 0.947 (s.e. 0.003). With this level of curvature, an individual would be indifferent between a 50–50 gamble over $20 and $0 and $9.62 for sure, implying a risk premium of $0.38. Column (5) provides non-linear least squares estimates based on the demand function (3). Broadly similar findings are obtained. Column (6) presents interval censored tobit estimates based on the Euler equation (2), accounting for the interval nature of the response data. We draw attention to the estimate of α, which is not directly comparable to our other estimates as this parameter is only identified up to a constant of proportionality (see Appendix A.1 for detail). Beyond this difference, similar estimates for discounting parameters are obtained. Though our estimated discount rates are higher than those of AS, broad consistency in discounting and curvature estimates are obtained across techniques with CTB data.

One point of interest in all of the estimates from Table 2, is the extent of dynamic consistency. Confirming recent findings with monetary payments when transaction costs and payment risk are closely controlled, we find minimal evidence of present bias (Andreoni and Sprenger, 2012a; Giné et al., 2012; Andersen et al., 2014; Augenblick et al., 2013).[29] Across elicitation techniques and estimation strategies, the present bias parameter, β, is estimated close to one.

3.2. Predictive validity

We consider predictive validity in two steps, using individual-specific parameter estimates for both. First, we test within and between methods. That is, we examine the in- and out-of-sample fit for CTB and DMPL estimates on the CTB data. Correspondingly we examine the in- and out-of-sample fit for CTB and DMPL estimates on the DMPL data. Though one would expect the in-sample estimates to outperform the out-of-sample estimates, this exercise does yield one critical finding: the CTB estimates perform about as well out-of-sample as the DMPL estimates perform in-sample for intertemporal choices.

Second, we test strictly out-of-sample for both methods. We examine behavior in a BDM mechanism eliciting willingness to accept to relinquish a claim for $25 at a later date and two hypothetical measures for patience. These three out-of-sample environments are constructed such that model estimates generate point predictions for behavior. Hence, one can analyze differences between predicted and actual behavior and the correlation between the two. Importantly, in both exercises we account for individual heterogeneity by estimating discounting parameters for each individual separately (see Appendix A.4 for details). For the CTB, individual level estimates are constructed based upon the estimation strategy of Table 2, Column

---

[27] For a summary of the raw results, please see Appendix Fig. A1 in Section A.3, which presents the choice proportions for the binary intertemporal MPL and HL data and the average allocations for the CTB data. All analysis was conducted using Stata 10. Elements of maximum likelihood optimization routines have been altered in subsequent versions of Stata. These changes influence the quantitative results in Table 2, column (6) and elsewhere, but do not alter any qualitative conclusions. Please contact the authors for details.

[28] The dependent variable is taken to be the chosen option in all interior allocations. For corner solutions in order for the log allocation ratio to be well defined we transform the value $0 to $0.01.

[29] For a general discussion of this issue and a direct demonstration of more severe present bias using estimates from the labor effort domain, see Augenblick et al. (2013). Also see Andreoni and Sprenger (2015) for a discussion of applications of these methods to different subject populations, with money and with goods, that have found present bias.

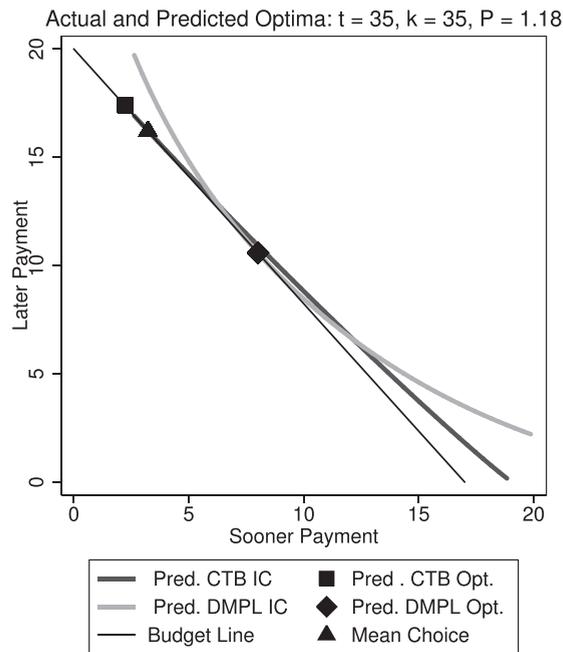**Fig. 3.** CTB and DMPL prediction of CTB data.

(4). Individual level estimates of α, β and r are obtained for all 58 subjects.[30] For the DMPL, individual level estimates are constructed based upon the estimation strategy of Table 2, Column (3). Individual level estimates of α, β and r are obtained for all 58 subjects. These analyses demonstrate that CTB-based estimates outperform DMPL-based estimates in all three out-of-sample environments.[31]

3.2.1. Within and between methods

We begin by analyzing the CTB data. First, consider the in-sample fit for the CTB estimates. We use the individual parameter estimates to construct utilities for each option within a budget and compare the predicted utility-maximizing option to the chosen option. Using the CTB estimates, the predicted utility maximizing choice was chosen 75% of the time.[32] Next, consider the out-of-sample fit for the individual DMPL estimates. They predict 16% of CTB choices correctly.[33]

The key out-of-sample failure for the DMPL estimates on the CTB data is generated by the high degree of estimated utility function curvature. Indeed, the majority of CTB choices are close to budget corners.[34] Fig. 3 presents an example budget with corresponding predicted indifference curves and choices based on CTB and DMPL estimates in which the DMPL prediction is far too close to the middle. The high degree of curvature prevents the DMPL estimates from making corner predictions and hence leaves the estimates unable to match many data points.[35]

We perform an identical exercise for the DMPL data. We focus specifically on the intertemporal MPL choices in this section. The HL data are considered in Appendix A.7 and demonstrate, in accordance with the idea that risk is a separate domain, that the DMPL estimates vastly outperform the CTB estimates on the HL data. In-sample individual DMPL estimates

---

[30] We opt to use the OLS estimates from Table 2, column (4), because individual level estimates are obtained for all 58 subjects. Using the NLS estimates of Table 2, column (5) very similar results are obtained, though the individual-level estimator converges for only 56 of 58 subjects.

[31] To account for estimation error, we also used the standard errors of the estimation to bootstrap the CTB and DMPL estimates for each person-choice combination. Since the results are quantitatively and qualitatively similar to those using the estimates alone, we do not report them here. One important dissimilarity, however, should be noted. When making DMPL predictions the bootstrapping procedure generates negative estimates of α in about 40% of the cases. If we exclude these, the predictive success of the bootstrapped individual level DMPL estimates is modestly better than the estimates alone. However, if we count these as incorrect predictions, the predictive success of the individual level DMPL estimates is reduced dramatically. Excluding negative α's skews the remaining distribution toward 1, which we demonstrate below favors more accurate predictions.

[32] Using the aggregate CTB estimates to construct utilities reduces success to 45%.

[33] Using the aggregate DMPL estimates to construct utilities reduces success to 3%.

[34] To be specific 88 percent of CTB allocations are at one of the two budget corners. Additionally, 35 of 58 subjects have zero interior allocations, consistent with linear utility.

[35] See Appendix A.6 for the exercise conducted on all experimental budgets.
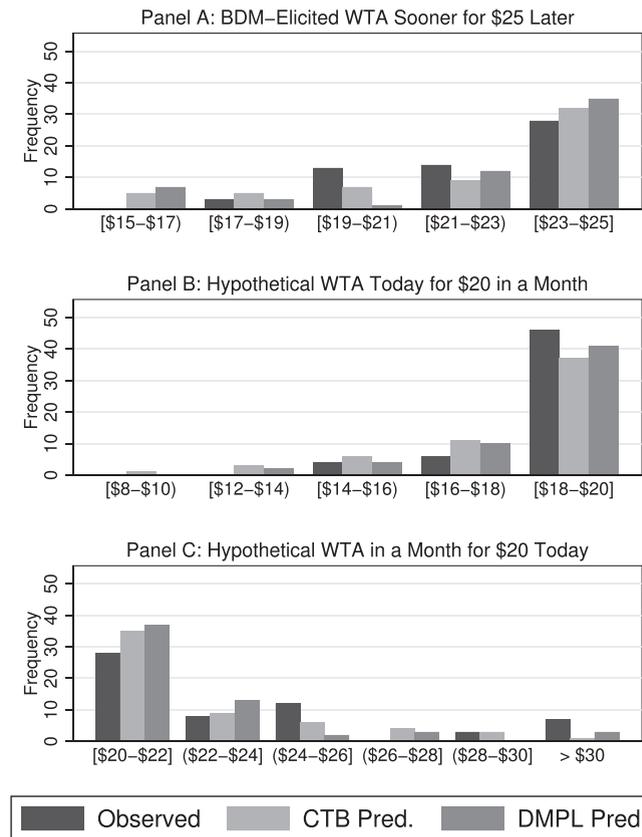
Fig. 4. Out-of-sample distributions.

predict 89% of MPL choices correctly.[36] Interestingly, the CTB estimates perform almost as well out-of-sample as the DMPL estimates perform in-sample. Aggregate CTB estimates predict 86% of MPL choices correctly.[37]

From this exercise we note that using individual level estimates both estimation techniques perform well in-sample. The CTB estimates predict out-of-sample with greater accuracy than the DMPL estimates. We next consider the predictive ability of the techniques in environments where both sets of estimates are out-of-sample.

3.2.2. Pure out-of-sample

Following the experimental implementation of the CTB and DMPL, subjects were notified of their two payment dates, based on a randomly chosen experimental decision. We then elicited the amount they would be willing to accept in their sooner check instead of $25 in the later check using a BDM technique with a uniform distribution of random prices drawn from [$ 15.00, $ 24.99].[38] All 58 subjects from our estimation exercise provided a BDM bid. The mean willingness to accept was $22.36 (s.d. $2.18). Fig. 4, Panel A presents the distribution of willingness to accept BDM responses.

Based on the payment dates, we use the individual parameter estimates from the CTB and DMPL to predict subject responses. These predictions account for the fact that relevant payment dates may involve different values of t and k. Responses that are predicted to fall outside of the price bounds described above are top and bottom-coded, accordingly. The mean CTB based prediction is $22.47 (s.d. $3.09), while the mean DMPL prediction is $22.48 ($2.95). Tests of equality demonstrate that we fail to reject the null hypothesis of equal means between the true data and both our CTB and DMPL estimates, ($t_{57} = -0.247$, $p = 0.81$), ($t_{57} = -0.251$, $p = 0.80$), respectively. The predicted distributions from the CTB and DMPL estimates are also presented in Fig. 4, Panel A. Though similar patterns to the true data emerge, Panel A does demonstrate some distributional differences, particularly at extreme values. Indeed, Kolmogorov-Smirnov tests of distributional equality

---

[36] Success falls to 81% with aggregate estimates.

[37] Success falls to 81% with aggregate estimates. Switching from the OLS to the NLS estimation technique for the individual CTB parameters generates an out-of-sample success rate above 90%.

[38] Hence, stating a willingness to accept greater than or equal to $25 implied a preference for the later payment in all states. Four subjects provided BDM bids of exactly $25 and no subjects provided a BDM bid greater than $25. Stating a willingness to accept lower than $15 implied a preference for any sooner payment. No subjects provided a BDM bid less than $15.

Table 3
Out-of-sample prediction

	(1) CTB predictions only	(2) DMPL predictions only	(3) CTB and DMPL predictions
Panel A: BDM-elicited WTA sooner for $25 later			
CTB prediction	0.230**	–	0.292**
	(0.094)	–	(0.118)
DMPL prediction	–	0.079	-0.107
	–	(0.103)	(0.125)
Constant	17.273	20.658	18.310
	(2.124)	(2.339)	(2.433)
Pseudo R^2	0.023	0.002	0.026
N	58	58	58
Panel B: Hypothetical WTA today for $20 in one month, X_{today}			
CTB prediction	0.545***	–	0.465***
	(0.092)	–	(0.121)
DMPL Prediction	–	0.600***	0.158
	–	(0.129)	(0.164)
Constant	9.268	8.217	7.805
	(1.672)	(2.633)	(2.267)
Pseudo R^2	0.152	0.084	0.157
N	55	55	55
Panel C: Hypothetical WTA in one month for $20 today, Y_{month}			
CTB prediction	0.541*	–	0.956**
	(0.322)	–	(0.448)
DMPL prediction	–	0.102	-0.987
	–	(0.535)	(0.736)
Constant	9.931	19.798	22.264
	(7.409)	(11.829)	(11.596)
Pseudo R^2	0.010	0.000	0.016
N	55	55	55

Note: All correlation estimates are from tobit regressions of actual choices on individual-specific choice estimates generated from utility function parameters. The predicted choices are top and bottom-coded in the following way: Panel A top and bottom-coded at BDM price distribution bounds. Panel B top-coded at $20. Panel C bottom-coded at $20. Of the 58 subjects for whom we have parameter estimates and BDM bids, 3 are dropped from the hypothetical choice analysis. 2 of these 3 failed to provide survey responses for either hypothetical question and another is excluded due to extreme outlying DMPL predictions.

* $p < 0.10$.
** $p < 0.05$.
*** $p < 0.01$.

reject the null hypothesis of equal distributions between observed and both CTB and the DMPL predictions, ($D = 0.414$, $p < 0.01$), ($D = 0.241$, $p = 0.06$), respectively.

More important to us than distributional accuracy is whether the predictions order individuals successfully, as judged by their actual BDM choices. Table 3, Panel A, columns (1) through (3) present tobit regressions analyzing the correlation between predicted and actual BDM behavior. In column (1) we show the CTB prediction to be significantly and positively correlated with BDM bids. In contrast, an insignificant correlation is obtained in column (2) where the independent variable is the DMPL predicted bid. Further, in column (3) when both predictions are used in estimation, we find that DMPL predictions carry little explanatory power beyond that of the CTB. This indicates predictive validity of the CTB estimates, though not the DMPL estimates, at the individual level.

Our final two prediction exercises involve hypothetical data collected during the post-experiment survey. First, we asked subjects what amount of money, X_{today}, today would make them indifferent to $20 in a month. Second, we asked subjects what amount of money, Y_{month}, in a month would make them indifferent to $20 today. Both measures are noisy with subjects at times answering free-form.[39] 56 of 58 subjects from our estimation exercise provided values for X_{today} and Y_{month}. Fig. 4, Panels B and C present these data. The data for X_{today} are top-coded at $20 while the data for Y_{month} are bottom-coded at $20. Following an identical strategy to that above, Panels B and C also present the distribution of responses predicted from CTB and DMPL individual estimates, top and bottom-coded accordingly. One subject's DMPL estimates produced a predicted value of Y_{month} in excess of $1000 and a X_{today} value of approximately $0. Excluding this outlier, our analysis focuses on 55 subjects. In nearly all cases, we reject the null hypothesis of equal means between predicted values and actual values.[40]

---

[39] In the first question, one subject responded 'Any amount over $20'. This response was coded as $20. This subject gave the same response in the second question and was again coded as $20. In the second question, one subject responded, '$19.05 plus one dollar in a month'. This was coded as $20.05.

[40] The mean actual value of X_{today} is $18.79 (s.d. $1.50). The CTB-based prediction for X_{today} is $18.29 (s.d. $2.36). The DMPL-based prediction for X_{today} is $18.44 (s.d. $1.76). We reject the null hypothesis of equal means between the true data and our CTB estimates, though not our DMPL estimates, ($t_{54} = 2.13$, $p = 0.04$), ($t_{54} = 1.63$, $p = 0.11$), respectively. The mean actual value of Y_{month} is $24.27 (s.d. $6.62). The CTB-based prediction for Y_{month} is $22.35 (s.d. $3.86). The DMPL-based prediction for Y_{month} is $21.92 (s.d. $2.46). We reject the null hypothesis of equal means between the true data and both our CTB and DMPL estimates, ($t_{54} = 2.04$, $p = 0.05$), ($t_{54} = 2.48$, $p = 0.02$), respectively.

Further, distributional tests frequently reject the null hypothesis of equality suggesting limited predictive validity at the distributional level.[41]

When considering the extent of correlations at the individual level, a different conclusion is drawn. Table 3, Panels B and C present tobit regressions similar to Panel A, where the dependent variable is either X_{today} or Y_{month}. Again we find the CTB predictions to carry significant correlations with the true measures. Though in Panel B, the DMPL prediction does significantly correlate with observed behavior, the DMPL predictions provide limited added predictive power beyond the CTB predictions. This again indicates predictive validity of the CTB estimates at the individual level.

Across our three out-of-sample exercises we find that both the CTB and DMPL can mis-predict, at times importantly, the distribution of behavior. However, at the individual level predictive validity is apparent, particularly for CTB-based estimates. DMPL-based estimates at times provide little independent and additional predictive power in our out-of-sample environments.

3.2.3. Probabilistic choice and multiple switching

While all of the predictions discussed above were generated via utility maximization, conditional on parameter values, the Luce model strategy suggests that another way of doing so would be to use a utility index with a decision error parameter to construct choice probabilities. This decision error allows one to connect preferences to choice probabilities via some functional form.[42] The aggregate in-sample fit of these models (estimated via maximum likelihood) may be very good but the out-of-sample prediction may falter. This may be for reasons of the parameter estimates being inapplicable or due to the assumption of probabilistic choice itself. In the case of price lists, a lot of decision error means a lot of multiple switching.

Another way of asking whether the estimates faithfully describe the data is to consider the degree of randomness in choice exhibited and the degree predicted. Of the 64 subjects who took part in the experiment, none exhibited multiple-switching behavior in the MPL task.[43] However, the Luce probabilistic choice model used to estimate the DMPL parameters (and probabilistic choice models generally) predicts choice probabilities that necessarily allow for switching more than once with some non-zero probability. We simulate 1000 sets of our MPL data using these predicted choice probabilities and find that the DMPL parameters and Luce model predicts that 86% of subjects should exhibit at least one "irrational" switch.

How much of this gap is due to the model and how much is due to the parameters themselves? To determine this, we run the CTB parameters through the Luce model, borrowing the DMPL estimate of ν, to again simulate 1000 sets of our data. This exercise predicts that 57% of subjects should exhibit at least one "irrational" switch. Given that a curvature parameter away from 1 directly attenuates utility differences between options, the CTB-DMPL gap make sense. The remaining Data-CTB gap is due to the Luce model itself; there are no hallmarks of probabilistic choice in the data.

4. Conclusion

We compare two recent innovations for the experimental identification and estimation of time preferences, the convex time budget (CTB) of Andreoni and Sprenger (2012a) and the double multiple price list (DMPL) of Andersen et al. (2008). Both innovations focus on generating measures of discounting which are not confounded by utility function curvature. The primary avenue along which the methods are compared is predictive validity. We examine the extent to which estimated utility parameters can predict behavior across experimental methods and in out-of-sample environments. At the distributional level, we find that both methods make predictions close to average behavior, though they often miss key elements of the distribution. At the individual level, we find CTB-based estimates to have increased predictive power relative to DMPL estimates.

We suggest three explanations for the observed differences between CTB and DMPL-based estimates: domain specificity, a failure of probabilistic choice and informational efficiency. All three of these explanations are linked together by a common thread, the identification of utility curvature. When it is estimated from a different domain, that of risky choice, choice predictions are inconsistent with the CTB data. This excess curvature mitigates the utility consequences of choosing one corner over the other in the MPL data. In the context of a probabilistic choice model, near-indifference between corners predicts common multiple switching, which is inconsistent with the MPL data. The greater informational content of choices, specifically corner choices in the CTB, rules out the excess curvature that causes these problems.

In motivating our study we suggested predictive power as a primary metric of success. We take the first step in this direction by exploring out-of-sample choices of our subjects made in the experiment. An essential test that remains, is to use these measurements of time preference to predict behavior outside of the experiment. In addition to laboratory refinement of the techniques presented here, a key next step is expanding to target populations for whom extra-lab choices are observable. Linking precisely measured discounting parameters to important intertemporal decisions is a promising avenue of future research.

---

[41] The KS statistic for the comparison of X_{today} across the true data and the CTB prediction is $D = 0.182$, ($p = 0.33$). For the comparison of X_{today} across the true data and the DMPL prediction is $D = 0.222$, ($p = 0.10$). The KS statistic for the comparison of Y_{month} across the true data and the CTB prediction is $D = 0.218$, ($p = 0.15$). For the comparison of Y_{month} across the true data and the DMPL prediction is $D = 0.259$, ($p = 0.04$).

[42] The specific functional form comes from the assumed random utility model and error distribution.

[43] All exclusions for multiple-switching were from violations on the HL task.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.jebo.2015.05.018.

References

Andersen, S., Harrison, G.W., Lau, M.I., Elisabet Rutström, E., 2008. Eliciting risk and time preferences. Econometrica 76 (3), 583–618.
Andersen, S., Harrison, G.W., Lau, M.I., Elisabet Rutström, E., 2014. Discounting behavior: a reconsideration. Eur. Econ. Rev. 71, 15–33.
Andreoni, J., Sprenger, C., 2012a. Estimating time preferences from convex budgets. Am. Econ. Rev. 102 (7), 3333–3356.
Andreoni, J., Sprenger, C., 2012b. Risk preferences are not time preferences. Am. Econ. Rev. 102 (7), 3357–3376.
Andreoni, J., Sprenger, C., 2015. Risk preferences are not time preferences: reply. Am. Econ. Rev. 107 (7), 1–8.
Ariely, D., Loewenstein, G., Prelec, D., 2003. Coherent arbitrariness: stable demand curves without stable preferences. Q. J. Econ. 118 (1), 73–106.
Ashraf, N., Karlan, D., Yin, W., 2006. Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines. Q. J. Econ. 121 (2), 635–672.
Augenblick, N., Niederle, M., Sprenger, C., 2013. Working Over Time: Dynamic Inconsistency in Real Effort Tasks. NBER Working Paper No. 18734.
Becker, G.M., DeGroot, M.H., Marschak, J., 1964. Measuring utility by a single response sequential method. Behav. Sci. 9 (3), 226–232.
Cagetti, M., 2003. Wealth accumulation over the life cycle and precautionary savings. J. Bus. Econ. Stat. 21 (3), 339–353.
Chabris, C.F., Laibson, D., Morris, C.L., Schuldt, J.P., Taubinsky, D., 2008a. Individual laboratory-measured discount rates predict field behavior. J. Risk Uncertain. 37 (2), 237–269.
Chabris, C.F., Laibson, D., Schuldt, J.P., 2008b. Intertemporal choice. In: Durlauf, S., Blume, L. (Eds.), New Palgrave Dictionary of Economics. Palgrave Macmillan.
Coller, M., Williams, M.B., 1999. Eliciting individual discount rater. Exp. Econ. 2, 107–127.
Cubitt, R.P., Read, D., 2007. Can intertemporal choice experiments elicit preferences for consumption? Exp. Econ. 10 (4), 369–389.
Dohmen, T., Falk, A., Huffman, D., Sunde, U., 2010. Are risk aversion and impatience related to cognitive ability? Am. Econ. Rev. 100 (3), 1238–1260.
Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference: a critical review. J. Econ. Lit. 40 (2), 351–401.
Giné, X., Goldberg, J., Silverman, D., Yang, D., 2012. Revising Commitments: Field Evidence on the Adjustment of Prior Choices. NBER Working Paper No. 18065.
Gneezy, U., Potters, J., 1997. An experiment on risk taking and evaluation periods. Q. J. Econ. 112 (2), 631–645.
Gourinchas, P.-O., Parker, J.A., 2002. Consumption over the life cycle. Econometrica 70 (1), 47–89.
Harrison, G.W., Lau, M.I., Williams, M.B., 2002. Estimating individual discount rates in Denmark: a field experiment. Am. Econ. Rev. 92 (5), 1606–1617.
Harrison, G.W., Lau, M.I., Elisabet Rutström, E., 2013. Identifying Time Preferences with Experiments: Comment. Working Paper.
Hausman, J.A., 1979. Individual discount rates and the purchase and utilization of energy using durables. Bell J. Econ. 10 (1), 33–54.
Holt, C.A., Laury, S.K., 2002. Risk aversion and incentive effects. Am. Econ. Rev. 92 (5), 1644–1655.
Laibson, D., 1997. Golden eggs and hyperbolic discounting. Q. J. Econ. 112 (2), 443–477.
Laibson, D., Repetto, A., Tobacman, J., 2003. A debt puzzle. In: Aghion, P., Frydma, R., Stiglitz, J., Woodford, M. (Eds.), Knowledge, Information and Expectations in Modern Economics: In Honor of Edmund S. Phelps. Princeton University Press, pp. 228–266.
Laibson, D., Repetto, A., Tobacman, J., 2007. Estimating Discount Functions with Consumption Choices over the Lifecycle. NBER Working Paper No. 13314.
Lawrance, E.C., 1991. Poverty and the rate of time preference: evidence from panel data. J. Polit. Econ. 99 (1), 54–77.
Luce, R.D., 1959. Individual Choice Behavior. John Wiley & Sons, New York.
Meier, S., Sprenger, C., 2010. Present-biased preferences and credit card borrowing. Am. Econ. J.: Appl. Econ. 2 (1), 193–210.
Meier, S., Sprenger, C., 2012. Time discounting predicts creditworthiness. Psychol. Sci. 23, 56–58.
Mischel, W., Shoda, Y., Rodriguez, M.L., 1989. Delay of gratification in children. Science 244, 933–938.
O'Donoghue, T., Rabin, M., 1999. Doing it now or later. Am. Econ. Rev. 89 (1), 103–124.
Rabin, M., 2000. Risk aversion and expected-utility theory: a calibration theorem. Econometrica 68 (5), 1281–1292.
Warner, J.T., Pleeter, S., 2001. The personal discount rate: evidence from military downsizing programs. Am. Econ. Rev. 91 (1), 33–53.