Exp Econ (2016) 19:864–885
DOI 10.1007/s10683-015-9472-x

ORIGINAL PAPER

Patience auctions: the impact of time vs. money bidding on elicited discount rates

Christopher Y. Olivola[1] · Stephanie W. Wang[2]

Received: 13 April 2014 / Revised: 26 October 2015 / Accepted: 27 October 2015 /
Published online: 4 November 2015

Abstract We introduce, test, and compare two auction-based methods for eliciting discount rates. In these "patience auctions", participants bid the smallest future sum they would prefer -or- the longest time they would wait for a reward, rather than receive a smaller, immediate payoff. The most patient bidder receives the delayed reward; all others receive the immediate payoff. These auctions allow us to compare discounting when participants' attention is focused on the temporal versus monetary dimension of delayed rewards. We find that the estimated parameters in the three most commonly used discount functions (exponential, hyperbolic, and quasi-hyperbolic) differ across these two bidding methods (time-bids vs. money-bids). Specifically, our participants tend to show more impatience under time-bids. Furthermore, we find that people are more likely to exhibit exponential (as opposed to hyperbolic) discounting and exhibit less present bias under time-bids, compared to money-bids. To our knowledge, this paper is the first to directly compare time versus money preference elicitations, within the same subjects, using an incentive-compatible mechanism.

Electronic supplementary material The online version of this article (doi:10.1007/s10683-015-9472-x) contains supplementary material, which is available to authorized users.

✉ Christopher Y. Olivola
olivola@cmu.edu

Stephanie W. Wang
swwang@pitt.edu

[1] Tepper School of Business, Carnegie Mellon University, Posner Hall, 5000 Forbes Ave, Pittsburgh, PA 15213, USA

[2] Department of Economics, University of Pittsburgh, 4715 Wesley W. Posvar Hall, 230 South Bouquet Street, Pittsburgh, PA 15260, USA

Keywords Delay discounting · Preference elicitation · Intertemporal choice · Within-subject comparison

JEL Classifications C91 · D03 · D44 · D90

1 Introduction

Many of our most consequential choices involve tradeoffs between present and future rewards. Consider, for example, the decision to invest in a lucrative graduate degree (e.g., an M.B.A.). Doing so can boost one's future income but the increased salary only comes later (once the degree is completed) and typically entails forgoing one's current salary (while attending graduate school). A potential full-time graduate student therefore has to tradeoff a smaller, immediate payoff (her current income) against a larger, delayed payoff (her post-degree income). To resolve this tradeoff, she might ask herself how many years (of schooling) she is willing to wait in order to obtain a given salary increase (e.g., the salary increase typically associated with an M.B.A. degree). Alternatively, she might ask herself what size future salary increase would make her willing to forgo her current salary for a given period of time (e.g., 2 years of schooling). Another example would be a consumer deciding whether to wait to purchase a product at a discounted price (e.g., waiting in line for a popular sale or selecting the cheapest, and therefore slowest, shipping option) or to forgo waiting and pay the full price for it. The consumer can approach his decision either by determining the maximum length of time he would be willing to wait to receive a given discount or by determining the minimum discount he would be willing to receive to wait a given length of time.[1] As these examples illustrate, people sometimes approach intertemporal tradeoffs by focusing on the time dimension (e.g., "How long am I willing to wait for a given-sized reward if it means forgoing a sooner, smaller reward?"), and other times by focusing instead on the reward dimension (e.g., "What sized reward, receivable after a given delay, would make me forgo a sooner, smaller reward?"). An important question, then, is

[1] In fact, there is evidence that consumers may make inconsistent choices across these two normatively equivalent ways of approaching the decision to wait for a discount. Sussman and Olivola (2011; Experiment 2) asked a sample of consumers (N = 351) to indicate both their willingness to wait (in minutes) in line for a 9 % discount on a product, and the smallest discount (in percentage) on that same product for which they would be willing wait 15 min in line. Both sets of preferences were collected using multiple-binary-choice ladders (with waiting times ranging from 5 to 60 min and discounts ranging from 5 to 12 %, respectively), and all participants completed both ladders (with ladders presented in counterbalanced order and usually separated by other, unrelated questions). We reanalyzed the data from that study, focusing on the 339 participants who completed both ladders without multiple switching points and on the equivalent binary-choice question in each ladder (whether they were willing to wait 15 min for a 9 % discount on the product). Our analyses showed that Sussman and Olivola's participants were significantly more likely to accept waiting 15 min for a 9 % discount (i.e., when the ladder varied the waiting time) than they were to accept a 9 % discount in exchange for a 15 min wait (i.e., when the ladder varied the discount size): likelihood of accepting = 87 vs. 74 %, $z = 4.07$, $p < 0.0001$. A closer look at their responses showed that 19 % of the participants reported inconsistent preferences: 16 % accepted to wait 15 min for the 9 % discount but also refused a 9 % discount for the 15 min wait; 3 % exhibited the opposite inconsistency (i.e., refusing to wait 15 min for the 9 % discount but also accepting a 9 % discount for the 15 min wait).

whether decision makers exhibit similar levels of patience across these two dimensions of focus.

Surprisingly, this question has received very little attention. Despite decades of research and scholarship on intertemporal decision making—not only within economics, but also psychology (e.g., Killeen 2009; Pronin et al. 2008; Reimers et al. 2009; Scholten and Read 2010), neuroscience (e.g., Ersner-Hershfield et al. 2009; McClure et al. 2004), and philosophy (e.g., Elster 1984; Parfit 1984)—we know of no published research that directly compares time preferences expressed in terms of time versus reward.[2] In this paper, we attempt to fill this void. To do so, we introduce, test, and compare a pair of novel auction-based experimental methods for eliciting discount rates.[3] In these "patience auctions", participants bid the smallest sum they would prefer receiving in the future (money bid) -*or*- the longest time they would prefer waiting for a reward (time bid), rather than receive a smaller, immediate payoff. Only the most patient bidder receives the delayed reward; all other bidders receive the smaller, immediate payoff. These auctions allow us to compare how discounting varies depending on whether the bidding focuses participants' attention on the temporal or monetary dimension of delayed rewards. By having participants take part in both the money-bid and time-bid auctions we can draw within-subject comparisons, thereby controlling for bidder-specific fixed effects.

Only a few studies have used auctions to elicit discount rates (Horowitz 1991; Kirby 1997; Kirby and Marakovic 1995; Manzini, Mariotti, and Mittone unpublished manuscript), and their designs differed from our patience auctions. More importantly, none of these studies have compared discount rates obtained through money bidding and time bidding. By contrast, our study is designed to determine whether individuals are consistent in discounting the future when bidding with time versus money.

We find that the best-fitting discounting model is usually different in the money-bid versus time-bid auctions. Furthermore, we find that the estimated parameters in the three most commonly used discount functions (exponential, hyperbolic, and quasi-hyperbolic) are significantly different across the two bidding types. Taken together, our results suggest that compared to money bidding, time bidding leads to more exponential (as opposed to hyperbolic or quasi-hyperbolic) discounting and to greater impatience. To our knowledge, this paper is the first to directly demonstrate

[2] Research comparing decisions about spending (or earning) time versus money (as goods in their own right) has shown that people treat these two commodities very differently (Burmeister-Lamp and Schade unpublished manuscript; De Borger and Fosgerau 2008; Duxbury et al. 2005; Leclerc et al. 1995; Liu and Aaker 2008; Monga and Saini 2009; Okada and Hoch 2004; Saini and Monga 2008; Soman 2001; Soster et al. 2010; Zauberman and Lynch 2005). However, whereas these previous studies compared preferences for two different kinds of goods (money versus time), we are interested in examining how people value the *same class of goods* (delayed rewards) when they express their preferences along one of two different dimensions (money versus time). The fact that people approach expenditures (or earnings) of money and time differently tells us little about how their patience-levels might vary depending on whether they are focused on the monetary or temporal dimension of delayed rewards.

[3] Originally introduced by Samuelson (1937) in his seminal paper on the discounted utility model, discount rates describe how people trade off consumption in the present vs. future (Frederick et al. 2002). Discount rates can be thought of, essentially, as economic measures of impatience.

Table 1 Summary of experimental design: the six different conditions shown are made up of four auction sessions (two first-price auction sessions; two second-price auction sessions) and two different orderings of the hypothetical matching survey

Method	No. of choices and ordering	*N*
First-price auction	8 money bid, then 8 time bid	15
First-price auction	8 time bid, then 8 money bid	15
Second-price auction	8 money bid, then 8 time bid	15
Second-price auction	8 time bid, then 8 money bid	15
Hypothetical matching	8 money matching, then 8 time matching	15
Hypothetical matching	8 time matching, then 8 money matching	15

the effects of expressing preferences in terms of time versus money, using an incentive-compatible elicitation method (with real payoffs and real delays at stake).

Our paper is organized as follows. Section 2 describes the experimental design and procedures. We present the econometric specifications in Sect. 3 and the results of our estimations in Sect. 4. Section 5 discusses the implications of our experimental results and future directions.

2 Experimental design and procedures

We introduce an auction-based approach designed to elicit discount rates in an incentive-compatible way (i.e., giving participants an incentive to accurately report their discount rates). To empirically evaluate our patience auctions and compare the effects of bidding money versus time, we conducted 4 experimental auction sessions with 15 participants in each session and we also administered surveys to 30 additional participants, all in the fall of 2007. The participants were all registered undergraduate students at Princeton University who were recruited via email announcements. The auction sessions were conducted in a computer lab and all interactions took place through a *multistage* (http://multistage.ssel.caltech.edu) interface on networked computers. No one participated in more than one session. The treatment variables in our 3×2 design were *elicitation procedure*: first-price auction versus second-price auction versus hypothetical matching survey, and *bid type*: bidding money versus bidding time. Table 1 summarizes the experimental design.

We ran single-round sealed-bid auctions of both bid types in each auction session: the money-bid auction and the time-bid auction. In each of the four sessions, participants bid in eight money-bid auctions and eight time-bid auctions. In two of the sessions all the auctions were first-price auctions and in the other two sessions they were all second-price auctions. Furthermore, we varied the ordering of the bid type so as to control for possible order effects[4]: half the sessions involved eight money-bid auctions followed by eight time-bid auctions, while the other half

[4] Our analyses revealed no order effects within either type of auction (on either the discount rate or the β parameter representing present-bias).

involved eight time-bid auctions followed by eight money-bid auctions. Thus, the bid type (money vs. time) was a within-subject manipulation, whereas the auction mechanism (first-price vs. second-price) was a between-subjects manipulation. All other conditions were kept constant across the four auction sessions.

After entering the laboratory, participants were given an instruction packet and an auction agreement form before being seated at workstations with dividers (seating assignment was randomized). Unless otherwise specified, subjects were informed about all the features of the auctions in which they participated. Sample instructions, sample screenshots, and the auction agreement form can be found in the Online Appendix. Before the instruction period, participants first signed the agreement form, stating that they understood the possibility of receiving a delayed payoff in the future rather than a payoff at the end of the session, depending on the auction outcome. In order to make sure that subjects understood the rules and payoff structure of each auction, they were given both verbal and written instructions. Participants were first given a complete description of single round sealed bid auctions in general and told how the outcome is determined by the auction mechanism, either first-price or second-price depending on the session. They were also told exactly what information feedback they would receive and how their overall earnings would be determined.

They were then given the rules for the first block of eight auctions they would be taking part in—either money-bid auctions or time-bid auctions, depending on the session. They were also familiarized with the computer interface through which they would make all of their decisions by viewing sample screens. After the instructions for this first block of eight auctions were covered, participants had to complete a short comprehension quiz and they could not proceed to the actual auctions until all answers were correct. After the first set of eight auctions of one bid type concluded, participants were instructed on the rules of the next eight auctions with the other bid type and again familiarized with the corresponding computer interface. Participants again completed a short comprehension quiz before they could proceed to the second set of auctions. In each auction, a participant could either receive \$10 at the end of the session or a delayed payoff sometime in the future, as determined by the bidding process.

In each money-bid auction, the length of delay for the future payoff was pre-set (see Table 2) and participants simultaneously bid the monetary amount for that payoff. They did so by entering numbers into respective boxes for dollars and cents then clicking the "submit" button (see sample screenshot in Online Appendix). The lowest bidder obtained the bid-determined payoff at the end of the pre-set delay period and all other bidders received \$10 at the end of the session. In the first-price auction, the bid-determined payoff was the lowest bid whereas in the second-price auction, the bid-determined payoff was the second lowest bid.

In each time-bid auction, the monetary amount for the future payoff was pre-set (see Table 2) and participants simultaneously bid the length of delay for that payoff. They did so by entering numbers into the respective boxes for years, months, weeks, and days, then pressing the "submit" button (see sample screenshot in Online

Table 2 Experimental parameter values for money-bid and time-bid auctions

Money bid: pre-set delay d	Time bid: pre-set delayed payoff m_d
1 day	\$15
3 days	\$20
1 week	\$25
2 weeks	\$30
1 month	\$50
3 months	\$75
6 months	\$90
1 year	\$100

Appendix). The participants were allowed to use any combination of the boxes to place their bids. The only restriction was that they had to use at least one of the boxes to submit their bid. However, which box or boxes they chose to use and how many boxes they chose to fill-in was entirely up to them. We implemented this flexibility to minimize possible anchoring or framing effects (Kahneman and Tversky 1979; Tversky and Kahneman 1974). The "longest" bidder (i.e., who bid the longest delay) obtained the pre-set payoff at the end of the bid-determined delay period and all other bidders received \$10 at the end of the session. In the first-price auction, the bid-determined length of the delay was the longest delay bid, whereas in the second-price auction it was the second longest delay bid.

Table 2 shows all the pre-set values that were used in each session. We randomized the order of the eight pre-set delay lengths for the money-bid auctions and the order of the eight pre-set delayed payoff amounts for the time-bid auctions. This randomized order was held constant across the four sessions. The parameters are presented in order of ascending magnitude in Table 2, which does not correspond to the order used in the sessions.

Each participant submitted a single bid in each auction. They were given no feedback about other participants' bids or the outcome of each auction (neither the winning bid nor the identity of the winning bidder was revealed at any point during the session). In other words, none of the auction results were announced to participants until the end of the session (i.e., after all bids to all auctions were submitted). Moreover, at the end of the session, only the result of the randomly selected auction (that determined final payoffs) was *partially* revealed: participants only found out whether or not they had won that particular auction (the winning bidder also learned its winning bid). All subjects were informed of this strict anonymity policy at the very beginning of the experiment. This precaution was taken in order to minimize "auction fever" (i.e., the tendency to overbid in time-bid auctions or to underbid in money-bid auctions).[5] They were told that they would be paid based on the outcome of one of the sixteen auctions chosen at random, after all

[5] Overbidding has been documented in laboratory second-price auctions (Heyman et al. 2004; Kagel et al. 1987; Kagel and Levin 1993). This phenomenon could be due to a myriad of factors ranging from a failure to understand the dominant bidding strategy to "auction fever" over the possibility of "winning" (Delgado et al. 2008).

auctions were completed. Given this structure, they were told that the optimal strategy was to treat each auction as if (i) it was the only auction and (ii) its outcome would determine their real earnings. All participants were paid one at a time, in private, on their way out of the lab. They were each given a sealed envelope that contained $10, except the participant who received the delayed payoff. This person instead received a sheet of paper explaining the payoff amount he/she would receive and the length of the delay until he/she received that payoff, as well as instructions for receiving the delayed payoff (along with the experimenters' contact information). The participant would then contact the experimenters to specify how and where he/she would prefer to be paid.[6]

The payments in our auctions were not trivial, ranging from $10 to $100 (or potentially higher in the money-bid auctions). Given that the distributions of prices, monetary gains, and monetary losses that people experience are approximately power-law shaped (Stewart et al. 2006), this range represents the vast majority of payments that most people (especially undergraduates) experience most of the time. Furthermore, our auction design guards against a "house-money effect" (Thaler and Johnson 1990), since participants were neither paid a show-up fee nor given a starting budget with which to bid. This is important because the house-money effect has been shown to distort preferences (e.g., Thaler and Johnson 1990).

In addition to the auctions, we conducted surveys that were designed to be the hypothetical matching task[7] equivalents of the auctions: the parameter values used and their orderings were identical to the ones used in the auctions. For this matching task survey, we recruited another 30 Princeton University undergraduate students (no monetary compensation given) to fill-out a short paper–pencil questionnaire examining their willingness to trade off an imaginary $10 in the present with a hypothetical payoff in the future. These hypothetical value-matching tasks asked subjects to indicate the sizes or delays of future payoffs that would make them indifferent between receiving those later payments and receiving the $10 dollars now (see Online Appendix). In each money-matching task, they were asked to report what value of the delayed payoff would make them indifferent between receiving $10 now and that delayed payoff after a pre-set length of delay. In each time-matching task, they were asked to report the delay length that would make them indifferent between receiving $10 now and a pre-set payoff amount at the end of that delay. In order to further parallel the design of the first-price and second-price auctions, half the participants (N = 15) were assigned to a survey in which they first answered money-matching questions, followed by time-matching questions. For these participants, the first page contained 8 matching questions asking them to match on future payoff and the second page contained 8 matching questions asking them to match on delay. This ordering was reversed for the other

[6] Participants could choose from several options to receive their payments: cash payments in person, check payments by mail, or direct bank transfers. Doing so helped minimize the cost to them of receiving their payments and ensured that we could pay them even if they had left campus when their payment deadlines arrived.

[7] For a more extensive discussion of this procedure, see Frederick et al. (2002).

participants.[8] Participants completed the survey by themselves, usually as part of a larger set of unrelated studies. The hypothetical matching tasks offer baseline controls that allow us to compare the auction results with those obtained from a commonly used discount rate elicitation procedure (Frederick et al. 2002).

Patience auctions possess several desirable properties. First, second-price patience auctions provide incentive-compatible ways to measure individual discount rates. The subjects' bids in second-price patience auctions, combined with the values of the predetermined parameters, can be used to calculate their discount rates (as we explain in the next section). Second, patience auctions are both time- and cost-efficient. Using these auctions, we can simultaneously measure the discount rates of multiple participants in a single session without having to incur large costs. Although every person bids, and thus reveals his or her discount rate (or a close approximation), only one participant in each session receives the larger, delayed amount, with all other participants receiving the earlier smaller amount. Consequently, patience auctions allow us to pay *all* of our subjects according to their bids, rather than only a (randomly selected) fraction of them as is often done in these studies. The latter approach (only paying a fraction of participants according to their choices) is problematic as it introduces an additional element of uncertainty into the decision making process. The uncertainty of being paid according to one's choices may distort the way participants perceive their payoffs (Todorov et al. 2007) and thus alter their reported preferences. In the case of bidding on future payoffs, for example, participants who learn that they only have a small chance of receiving the delayed payoff, even if they "win" the auction, might be emboldened to bid too patiently since the risks associated with bidding too "long" (in the case of time-bids) or too low (in the case of money-bids) would be much lower under a probabilistic payoff scheme. By contrast, in our current experiment, *every* bidder is paid according to the outcomes of one randomly chosen auction, so they know that one of their decisions will be consequential (for certain). This is quite different (at least, psychologically) from an incentive scheme in which bidders face only a small probability of being paid according to *any* of their bids.[9] Finally, the participants in patience auctions do not bid their own money nor receive a fixed bidding budget, which offers at least two advantages: (i) the size of the delayed payoff is constrained only by the experimenter's budget, not the participants', so they can bid as much as

[8] In contrast to the auctions, we do observe order effects with the hypothetical matching task survey. Specifically, we find a marginally significant order effect on the discount rate under money biding (when we estimate exponential and hyperbolic models, but not when we estimate quasi-hyperbolic models), such that discount rates are higher under the time-then-money ordering (compared to the money-then-time ordering). And we find a significant order effect on the discount rate under time biding (for all three discounting models), such that discount rates are again higher under the time-then-money ordering. Finally, when we estimate the quasi-hyperbolic model, we find significant order effects on the β (present-bias) parameter: under money bidding, β is larger in the money-then-time ordering; under time bidding, β is larger in the time-then-money ordering.

[9] While the validity of the random-lottery incentive system (paying all participants according to one randomly selected choice) has received strong empirical support (Beattie and Loomes 1997; Cubitt et al. 1998; Hey and Lee 2005a, b; Lee 2008; Starmer and Sugden 1991), the validity of paying a random subset of participants according to their choices is less well established (see Baltussen et al. 2012). The few tests of this latter method have found that it biases preferences by lowering risk aversion (Baltussen et al. 2012).

they would like for the delayed payoff, even if its private value exceeds their personal budgets; (ii) since participants are explicitly trading-off a smaller immediate reward with a larger delayed one, neither of which they possess at the time of bidding, we avoid the potential effects of unobserved reference points (Carmon and Ariely 2000; Kahneman et al. 1990) and loss aversion (Kahneman and Tversky 1979; Tversky and Kahneman 1991).

Beyond their methodological advantages, patience auctions allow us to compare the level of discounting that occurs when bids are in terms of delay versus payoff, so that we can see whether people consistently trade off time and money. While eliciting discount rates through matching (i.e., measuring indifference points) along the payoff dimension is a relatively common procedure in the intertemporal choice literature, only a handful of studies have elicited discount rates through matching along the time dimension (e.g., Attema et al. 2010; Chesson and Viscusi 2000; Green et al. 1994; Mazur 1987; Takeuchi 2011). In fact, we only know of one previous attempt[10] to directly compare the discount rates obtained when subjects focus on the monetary versus temporal dimension of the delayed payoff: an unpublished dissertation by Roelofsma (cited in Roelofsma 1996 and in Frederick et al. 2002). Roelofsma's findings suggest that discount rates are higher with time matching than with payoff matching (a result that we replicate). However, as far as we know, he neither used an incentive-compatible elicitation mechanism, nor fitted models other than standard exponential discounting. More generally, to the best of our knowledge, no published research has directly compared (im)patience when people focus on the temporal versus monetary dimension, particularly under incentive-compatible conditions. Therefore, it is still an open question whether (and how) focusing on the time versus money dimension will impact discounting when real payoffs and delays are at stake.

3 Framework and econometric specification

Let m_0 be the immediate reward, which is \$10 in our experiment. Let d be the length of delay in days, which is pre-determined in the money-bid auctions and determined by the bids in the time-bid auctions. m_d is the delayed reward after d days, where the amount (m_d) is determined by the bids in the money-bid auctions and pre-determined in the time-bid auctions. Let r be the daily discount rate. Equation 1 expresses the relationship between the immediate reward and the discounted delayed reward where $D(r, \beta, d)$ is the discounting function:

[10] Recently, a working paper by David Eil has come to our attention. He utilizes a multiple-choice method (involving a series of binary choices between smaller payoffs today and larger payoffs in the future) to identify switching points in subject's preferences (in order to estimate their implied discounting parameters). He directly compares time preferences when the delay vs. amount of the future payoff is systematically varied. However, it should be noted that his experiment was conducted in May 2010, whereas ours was carried out in the fall of 2007, our results were first presented (at a major conference) in 2007, and the first complete draft of our manuscript had been written by 2009.

Table 3 Summary of discounting models

Discounting model	Functional form	Parameters
Exponential	$D = e^{-rd}$	d, r
Hyperbolic	$D = \frac{1}{1+rd}$	d, r
Quasi-hyperbolic	$D = \beta e^{-rd}$	d, r, β

$$m_0 = D(r, \beta, d)\, m_d \tag{1}$$

We compare the fits and estimated the parameter values of the three most commonly used discount functions: exponential, hyperbolic, and quasi-hyperbolic, as summarized in Table 3. The first specification corresponds to the canonical exponential discounting model, characterized by time-consistent preferences. The second specification corresponds to the hyperbolic discounting model, which, unlike exponential discounting, implies declining subjective discount rates with increasing delays. The third, and final, specification corresponds to the quasi-hyperbolic discounting model, which allows for present bias through the addition of a β parameter. Specifically, in this model, an individual demonstrates present-bias, or overweighs the current period, if $\beta < 1$ (and no present bias when $\beta = 1$). However, all future periods are still discounted exponentially at rate r in this model (Laibson 1997; Phelps and Pollak 1968).

We now consider the optimal bidding strategies for money- and time-bids, starting with a second-price auction. Vickrey (1961) showed that in a second-price private-value auction, it is a dominant strategy for each bidder to bid her true value for the object regardless of what the other bidders do. In our experiment, the assumption of a private-value auction is a valid one, especially given the lack of a secondary market for the future reward and the fact that subjects were not informed of others' bids. Each bidder's valuation of the future reward relative to the immediate gain is also assumed to depend on her discount rate only and not those of the other bidders. Therefore the bidders should, through their bids, provide accurate and exact measurements of their willingness to trade off smaller immediate gains with larger delayed ones. In the money-bid auction, the truthful bidding strategy would reveal the delayed payoff that would make the bidder indifferent between receiving that payoff after the set delay length and receiving \$10 now. In the time-bid auction, the truthful bidding strategy would reveal the delay length that would make the bidder indifferent between receiving the set payoff after that delay length and receiving \$10 now. In other words, in both auctions, the truthful bidding strategy equalizes the discounted future payoff and the immediate payoff (assuming risk neutrality[11]), as expressed in Eq. 1.

Consider a bidder in the money-bid auction: If anyone else bids lower than m_d then this bidder receives m_0. If m_d is the lowest bid, then this bidder receives the delayed payoff associated with the second-lowest bid after the set delay d. If this bidder ties with other bidders then one of them is randomly selected to receive the

[11] Note that our manipulation is within-bidder, and therefore controls for bidder-specific fixed effects (including risk attitudes and beliefs about the distribution of other bidders' risk and time preferences).

delayed payoff, so this bidder will either receive the delayed m_d or m_0 now, which are of equal value to her. If the bidder were to deviate from the truthful bidding strategy by bidding a lower delayed payoff than what would make her indifferent, she risks receiving a delayed payoff that is lower than the one that makes her just as well off as receiving $10 now (if the second lowest bid is also lower than her indifference point). If, instead, the bidder were to deviate by stating a higher delayed payoff than what would make her indifferent then she risks losing out on a delayed payoff that would make her better off than $10 now (if the first and second lowest bids are both higher than her indifference point and lower than her bid). The same logic applies to bidding one's true willingness to wait in the time-bid auctions.

Next, we verify the strategic equivalence of truthful bidding in the money-bid and time-bid auctions. Suppose the bidder's time preference is characterized by $D(r^*, \beta^*, d)$ where r^* is her discount rate and β^* is her present-bias parameter. For this bidder, the truthful bidding strategy is $m_d^* = (m_0/D(r^*, \beta^*, d))$ in the money-bid auction. If the bidder's preference is characterized by the exponential discount function, for example, then $m_d^* = m_0 e^{r^* d}$. In the time-bid auction, the truthful bidding strategy is d^* such that $m_d = (m_0/D(r^*, \beta^*, d^*))$. Again, under the exponential discounting model, $d^* = (ln(m_d/m_0))/r^*$. These strategically equivalent bids in the money- and time-bid auctions reflect the same underlying preference parameters (r^* and β^*). That is, $m_0 = D(r^*, \beta^*, d^*(m_d^*, m_0, r^*, \beta^*))m_d^*$.

The equilibrium bidding strategy for the first-price auctions, in contrast, is more complicated and depends on the bidders' beliefs about the other bidders' time preferences and subsequent valuations. Bidders are assumed to have common knowledge about the distribution of the private values. While this is indeed a strong assumption, it is a plausible one given the relative homogeneity of the study population. In our context, bidders believing that all other bidders have the same type of discounting function as well as bidders having common knowledge about the range of the parameters of that function within the bidding population would satisfy the assumption. To show strategic equivalence between the money- and time-bids in the first-price auction, we must maintain the assumption that the beliefs about other bidders' discount functions and parameters are the same across these two bid types.[12] The bidding strategy of the unique equilibrium of the first-price auction is to bid the expectation of the second-place bidder's value conditional on winning the auction (again, assuming risk neutrality—see Footnote 11). Given a reasonable distribution for the beliefs about other bidders' valuations, the greater the number of other bidders, the closer the bid should be to the true valuation. In our experiment, we had 15 bidders per auction. We estimate the discounting parameters from the first-price auctions as if the bids were truthful revelations because this does not require us to make any assumptions about the bidders' beliefs. While the equilibrium bidding strategy is not truthful revelation as it is in the second-price auctions, the difference is driven by beliefs about the other bidders' discounting parameters. If these beliefs are the same across the money- and time-bids (as they should be) then any difference in estimated discounting parameters cannot be attributed to different bidding strategies, but rather to differences in the underlying

[12] Testing this assumption is left to future work.

discount function and/or parameters. This is sufficient for our purpose, since our goal is not to compare estimated discounting parameters across first-price and second-price auctions, but rather to compare these parameters across money- and time-bids within each type of auction. Note that our money-bid versus time-bid manipulation is within-bidder, and therefore controls for bidder-specific beliefs about the distribution of other bidders' risk and time preferences (especially since they receive no feedback regarding the other bids in these auctions).

Following Benhabib et al. (2010), we do nonlinear least squares (NLS) estimation of the discount function parameters by individual using the econometric specification:

$$m_0 = D^i(m_d(m_0, d), d; r^i, \beta^i) m_d(m_0, d) \varepsilon^i(m_d, d) \quad (2)$$

where $\varepsilon^i(m_d, d)$ is distributed log-normally and i.i.d. with respect to individuals and the pre-determined parameters.

In sum, to account for heterogeneity in discounting across individuals and to test whether individuals discount differently in the money- versus time-bid auctions, we estimate r for each individual i (90 in total across first-price, second-price, and hypothetical survey treatments) using NLS where m_d is a function of m_0, d, and r in the money-bid auctions and d is a function of m_0, m_d, and r in the time-bid auctions.

4 Results

First, we examine how bids vary as a function of the elicitation mechanism and the parameter values. Bidding values are quite volatile, highly skewed, and the distribution of bids is necessarily censored at zero (no one could bid negative payoffs or negative delays). We therefore examine median bids, as these are more robust to outliers. Figures 1 and 2 plot median bids for each elicitation method, as a function of the auction parameter value. These figures suggest that the elicitation method had an impact on the level of patience that bidders revealed. In particular, we see that for both time-bid and money-bid auctions, bidders seemed to be most patient under a second-price auction mechanism, and least patient when responding to a hypothetical survey. There could be numerous reasons for these observed differences between elicitation procedures but this is outside the scope of the current paper.[13] More importantly, and as we explain below, we find that all three elicitation methods produce the same pattern of results when we compare patience expressed in terms of money versus time. These figures also reveal a monotonically increasing relationship between bid size and parameter value, as we would expect. After all,

[13] We can also check whether revenue equivalence holds between the first- and second-price money-bid auctions. Here we define "revenue" as the amount given to the "winner" of the money-bid auction at the delayed date instead of \$10 in the present. A Mann–Whitney test comparing revenues from the first-price money-bid auctions and revenues from the second-price money-bid auctions yields p value $= 0.057$. Note that numerous previous experiments have found lack of revenue equivalence between first- and second-price auctions (Kagel 1995). The predicted revenue equivalence hinges on the assumption that bidders are risk-neutral, so any lack of revenue equivalence found here could be attributed to non-risk-neutrality or to a number of other possible factors.

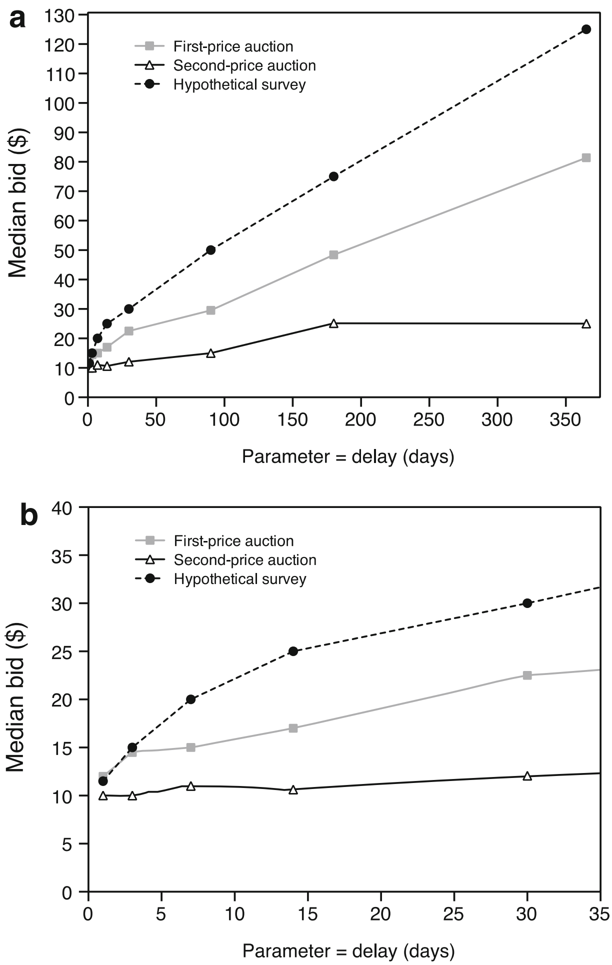

Fig. 1 **a** Median money bids in the money-bid auctions as a function of delay and elicitation method (full range of delays). **b** Median money bids in the money-bid auctions as a function of delay and elicitation method (lower range of delays). There were 30 bidders in each auction, so each data point is based on N = 30 observations

money bids should increase with the delay of the future payment and time bids should increase with the size of the future payment. The main exception to this monotonic relationship seems to be for the highest delayed payoff ($100) used in the second-price time-bid auctions, which saw a sizeable drop in the median delay bid. Otherwise, median bids are mostly coherent and rational in that they are sensitive to parameter values (and adjust in the correct direction).

Next we examine how well the three main discounting models found in the literature fit the bidding data. Figure 3 shows the proportion of subjects whose bidding data were best fit by a given model, as determined by the Bayesian

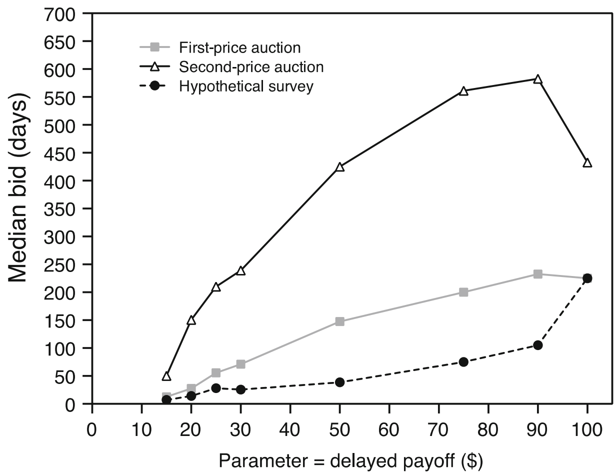

Fig. 2 Median time bids in the time-bid auctions as a function of delayed payoff and elicitation method (full range of delayed payoffs). There were 30 bidders in each auction, so each data point is based on N = 30 observations

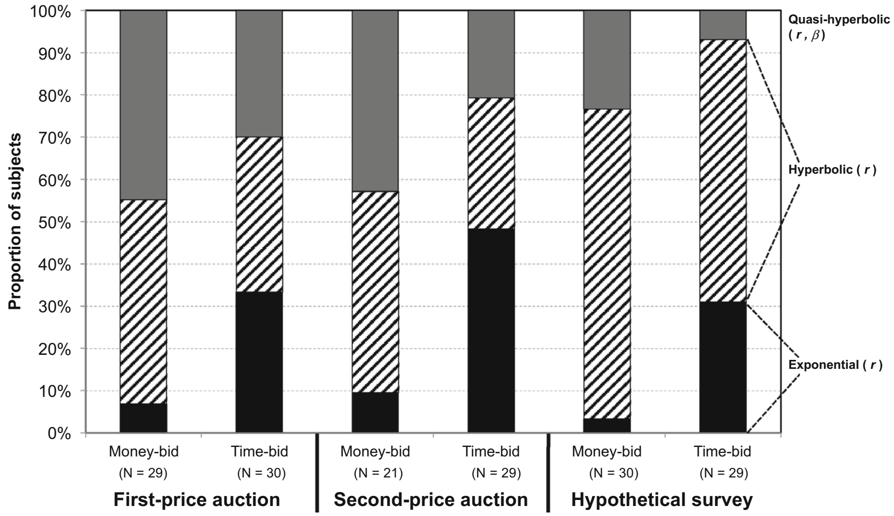

Fig. 3 The proportion of subjects (with estimable parameters) whose bidding data were best fit by a given model as a function of bid-type and elicitation method

Information Criterion (BIC—Schwarz 1978). We can see that, for money bids, most individuals' bidding patterns were best fit either by the hyperbolic or quasi-hyperbolic model, and only a small minority (<10 %) followed exponential discounting. With time bids, in contrast, between a third and one half of individuals followed exponential discounting, and they were always least likely to follow quasi-hyperbolic discounting. This difference, between time bids and money bids, in the proportion of subjects falling into the various models, was significant for the first-price auction ($\chi^2(2) = 6.41$, $p = 0.041$), the second-price auction ($\chi^2(2) = 8.59$, $p = 0.014$), and the hypothetical survey ($\chi^2(2) = 9.56$, $p = 0.008$). Furthermore, the model that a subject fell into when bidding money did not significantly predict which model that (same) person would fall into when bidding time: $\chi^2(4) = 3.17$, $p > 0.5$ (after collapsing across elicitation procedures to increase sample size). The weak categorical association between best-fitting models under money and time bids (Cramér's $V = 0.14$) suggests that the shape of the discount function may not be a stable individual characteristic (Meier and Sprenger 2015; Stewart et al. 2006).

We also calculated the median value of each model parameter,[14] for each bid-type, and for each elicitation procedure. These values are presented in Table 4. As this table shows, the model parameters consistently (and significantly) differed between money bids and time bids, for all discounting models and elicitation procedures. Consider the exponential model, which represents the standard rational model of discounting: the average median daily discount rate across the three elicitation procedures is $r_{avg} = 0.010$ for time-bids and $r_{avg} = 0.006$ for money-bids (i.e., bidders were 1.82 times as impatient when bidding time). Similarly, under the hyperbolic and quasi-hyperbolic functions, we find that bidders were more patient when they bid money than when they bid time (r is consistently lower for money bids than for time bids). Although bidders are more patient under the quasi-hyperbolic model when they bid money, they are also more present-biased. Specifically, the β parameter representing present-bias is closer to 1 ($\beta = 1$ indicates no present-bias) for time bids, and it is consistently lower for money bids, indicating more present bias[15] for money bidding than for time bidding. Nonparametric (Wilcoxon signed-rank) tests revealed that the β parameter was significantly less than 1 (indicating significant present bias) under the quasi-hyperbolic model in all cases (all p values < 0.01) except time bids elicited through second-price auctions ($p = 0.38$).

We take these estimation results, along with the best-fit discount function percentages, as suggestive evidence that people are more likely to behave like exponential discounters[16] when thinking about the temporal dimension of delayed rewards.

[14] These model parameters were estimated using data from *all* subjects with estimable parameters, not just those for whom the model represented the best fit according to the BIC.

[15] These results are consistent with those obtained by Eil (unpublished manuscript), who finds that eliciting time preferences through a multiple-choice method in which the delay of the future payoff is systematically varied (to identify a switching point) yields more impatience and less present bias than a multiple-choice method in which the amount of the future payoff is varied instead. However, whereas Eil finds evidence of future bias (i.e., $\beta > 1$, see also Takeuchi 2011) under the "time" multiple-choice method, we find lower or no present bias with time bids, but not future bias.

[16] We also fitted a model specification with a free θ (hyperbolicity) parameter (such that $\theta = 1$ produced the exponential and $\theta = 2$ produced the standard hyperbolic function; see Benhabib et al. 2010), and found that the estimated θ is lower and closer to 1 for time bids (compared to money bids) under all three elicitation procedures. Specifically, the θ parameter ranged from 1.1 to 1.8 for time bids, and from 2.5 to 3.4 for money bids. Thus, people seem to be less hyperbolic (i.e., more exponential) when they bid time, whereas they seem to be "supra"-hyperbolic ($\theta > 2$) when they bid money.

Table 4 Median daily discounting parameter estimates for each specification, elicitation mechanism, and bid-type (for subjects with estimable parameters)

Model (specification)	First-price auction (N_{money} = 29, N_{time} = 30)					
	Median r			Median β		
	Money-bid		Time-bid	Money-bid		Time-bid
Exponential (r)	0.006	**(p = 0.0001)**	0.010			
Hyperbolic (r)	0.020	**(p = 0.006)**	0.033			
Quasi-hyperbolic (r, β)	0.004	**(p < 0.0001)**	0.008	0.538	**(p = 0.0001)**	0.740
Model (specification)	Second-price auction (N_{money} = 21, N_{time} = 29)					
	Median r			Median β		
	Money-bid		Time-bid	Money-bid		Time-bid
Exponential (r)	0.0036	**(p = 0.046)**	0.0038			
Hyperbolic (r)	0.007	**(p = 0.01)**	0.013			
Quasi-hyperbolic (r, β)	0.003	**(p = 0.02)**	0.004	0.752	**(p = 0.135)**	0.979
Model (specification)	Hypothetical survey (N_{money} = 30, N_{time} = 29)					
	Median r			Median β		
	Money-bid		Time-bid	Money-bid		Time-bid
Exponential (r)	0.007	**(p < 0.0001)**	0.017			
Hyperbolic (r)	0.036	**(p = 0.0001)**	0.051			
Quasi-hyperbolic (r, β)	0.004	**(p = 0.0001)**	0.011	0.361	**(p = 0.002)**	0.631

The p values in bold represent the results of within-subject Wilcoxon signed-rank tests, which were conducted separately for each model parameter and for each elicitation procedure

These differences in discounting model parameter values between money bids and time bids seem, at first glance, to support the idea that individuals lack a stable underlying level of impatience and that time preferences may be largely "constructed" from the decision context (in this case, the method of elicitation) rather than revealed through it (Lichtenstein and Slovic 2006). However, it might be the case that individuals do possess stable underlying time preferences to begin with, but that the specific method of elicitation exerts an additional effect on these preferences. One way to examine this question is to look at within-individual correlations between discounting model parameters under money bidding and time bidding. To the extent that intertemporal choices are driven by stable underlying individual preferences, these correlations should be large and positive.

Table 5 presents the non-parametric (Spearman's rank) correlations between parameter values obtained from money bids and time bids, with each individual bidder representing a single data point (for a given discounting model and model parameter). As this table shows, there is significant within-individual consistency in

Table 5 Within-subject Spearman's rank correlations between parameter values obtained from money bids and time bids (for subjects with estimable parameters)

Model (specification)	First-price auction (N = 29)		Second-price auction (N = 21)		Hypothetical survey (N = 29)	
	r	*β*	*r*	*β*	*r*	*β*
Exponential (*r*)	**0.71**		**0.46**		**0.91**	
Hyperbolic (*r*)	**0.71**		**0.48**		**0.91**	
Quasi-hyperbolic (*r*, *β*)	**0.46**	**0.39**	***0.42***	0.28	**0.67**	−0.17

Results are presented for both the daily discount rate (*r*) and present-bias parameter (β)

Bold correlation coefficients are significant at the standard ($p < 0.05$) level. Bold italicized correlation coefficients are marginally significant ($p < 0.1$). All other correlation coefficients are not significantly different from 0 ($p > 0.1$)

r parameter values. This consistency is stronger for the simpler (1-parameter) models and weakest for the second-price auction. Although the correlation coefficient is larger for the hypothetical survey than the first-price auction, this could be due to the fact that the hypothetical survey presented participants with all 8 variations of the matching questions for each bid-type on the same page, thereby facilitating consistency. On the other hand, the hypothetical survey also produced a (non-significant) *negative* within-individual correlation for the β parameter.

These results suggest that, despite the main (within-subject) effect of bid type, individuals do possess time preferences that are stable to some extent (even if their functional form "type" may not be, as we saw above). Alas, these correlations are not direct evidence that preferences are revealed rather than constructed. It may be that individuals regularly construct their preferences from a fairly stable memory cache (Olivola and Sagara 2009; Stewart et al. 2006). This alternative account would also predict significant positive within-individual correlations.

5 Discussion

We find that expressing patience in terms of money versus time fundamentally alters the way people discount delayed payoffs. First, the best-fitting discounting model tends to differ (within-subject) between money- and time-bid auctions. Indeed, less than half (42 %) of our subjects were best fit by the same model across bid types. Second, even if the discounting model is held constant, its estimated parameter(s) differ(s) between money- and time-bid auctions (also within-subject). Overall, we find that under money bids, payoffs are discounted more heavily at initial time intervals and less so at later time intervals ($\beta < 1$ and $\theta > 2$—see Footnote 16). In contrast, under time bids, the rate of discounting seems to be spread more evenly across time intervals (β and θ are both closer to 1—see Footnote 16), thereby bringing people more closely in line with the standard exponential model. Together, these results suggest that intertemporal tradeoffs are governed by different factors when people are focused on the monetary versus temporal dimension of these tradeoffs. Exactly why this happens, however, is an as-yet-unanswered question that we leave for future research. Critically, since we compared the effects of money bidding and time bidding within the same subjects, our analyses control for any bidder-specific fixed effects. This means, among other things, that differences between money bidding and time bidding cannot be explained by reference to endogenous variables, including risk attitudes or other individual traits. Moreover, this pattern of differences is similar across the three elicitation procedures we employed. Given their robustness to the specific elicitation method used (first-price auction vs. second-price auction vs. hypothetical matching task), our results are unlikely to be entirely the product of the specific characteristics (e.g., incentive compatibility or complexity) of these procedures. Finally, it is worth noting that we perfectly counterbalanced bid-type order (whether money bids came before or after time bids), so the effects we observe cannot be attributed (in any simple way) to spillovers from one elicitation method to another, nor to whether one particular elicitation method came first or second. Any remaining concerns that learning may have occurred (and thus accounted for our results) are further minimized by the fact that our subjects received no feedback until the very end of each session (after *all* bids were submitted).

The finding that people seem more likely to be exponential discounters when they must bid time than when they must bid money is both theoretically interesting and practically important. Contrary to the other discount functions we estimated, exponential discounting leads to time-consistent preferences, which is generally considered a hallmark of rational decision making. Therefore, from a practical point of view, having people focus on the temporal dimension (rather than the monetary dimension) when they are making decisions involving delayed goods may promote more consistent choices across time. From a theoretical point of view, our results challenge the idea, popular within behavioral economics and psychology, that people are (consistently) hyperbolic (or quasi-hyperbolic) discounters. Rather, it seems that the tendency to be (quasi-)hyperbolic (as opposed to exponential) may depend critically on whether the intertemporal tradeoffs focus attention on the monetary or temporal dimension. Theoretical debates concerning whether intertemporal choices are better described by the standard, classic exponential model or by newer, behavioral (quasi-)hyperbolic models may miss the point, since the answer, it seems, is that it depends! Our results also contribute to the literature on preference reversals (Goldstein and Einhorn 1987; Grether and Plott 1979; Tversky, Slovic, and Kahneman 1990). Those previous studies showed that people's risk preferences depend on whether they are asked to choose or price gambles, and suggested that this occurs because each elicitation method focuses attention on a different dimension of the choice problem (either probability or payoff). Here, we show that time preferences similarly depend on whether the method of elicitation focuses attention on the time versus payoff dimension.

Furthermore, our time-bid auctions are analogous to the "marshmallow test" (Mischel et al. 1972) used to elicit the level of patience in young children. In both cases, the immediate reward and delayed reward are already set (one and two marshmallows, respectively, in their setting) but the delay is unspecified. Therefore, subjects are essentially choosing how long they are willing to wait for the larger, delayed reward. The difference is that in the “marshmallow test,” the young subjects could decide at every moment, before the unknown-to-them time arrived for the second marshmallow, whether to consume the single marshmallow or continue waiting for the additional marshmallow. Given the surprisingly strong correlations between responses on the “marshmallow test” and later life outcomes (Mischel et al. 1988, 1989), a modified version of our time-bid auction such as an ascending clock auction or a DOSE design (Wang, Filiba, and Camerer unpublished manuscript), in which subjects are asked to accept or reject different waiting times for the pre-determined delayed reward based on which time delays they have accepted or rejected thus far, might work well in the field.

Our patience auction results could also have implications for the design of auctions with a temporal component, such as those for highway procurement (Lewis and Bajari 2011) or treasury bills (Bartolini and Cottarelli 1997). Lewis and Bajari present evidence of efficiency gains when contractors bid along the time dimension, in addition to the standard cost dimension, in highway procurement auctions. Our results suggest that policymakers who are primarily concerned with the fast completion of projects (rather than their quality) should restrict bidding to the time dimension so that more attention, and perhaps more value, is given to time spent. Treasury bills always have a set maturity payment at a set maturity date, with an auction determining the current price at which these bills can be bought. Instead, a treasury bill could have a set price and payment, with the auction determining the maturity date. The bidders who are willing to wait the longest would receive the bills. Although this time-bid auction might elicit less revealed patience due to a focus on the temporal dimension, the auctioneer at least knows the revenue from the auction with certainty ex-ante because the price is set.

A future direction we are currently exploring is using our auction methods to elicit other kinds of preferences (Olivola and Wang unpublished manuscript). For example, instead of having participants trade off time and money in order to elicit their discount rates, we can just as easily have participants trade off probability and money in order to measure their risk preferences. Thus, by replacing time with probability, we can have participants trade off a smaller sure gain (e.g., \$10 with 100 % probability) and a larger risky gain (m_r dollars with probability p_r). In these auctions, participants either bid for the value of m_r, the payoff associated with the risky gamble, or for the value of p_r, the probability associated with the risky gamble. Here again, standard economic theory predicts that both methods should lead to identical risk preference. Our auctions would test this using an incentive-compatible procedure. James (2011) tested another type of incentive-compatible procedure (BDM) for eliciting probability equivalents instead of certainty equivalents for lotteries. Both studies can contribute insights into how patterns of preferences might vary with elicitation methods.

In conclusion, we find that people reveal markedly different discount functions and rates when they express their intertemporal preferences in terms of time versus money. When expressing preferences in terms of money, they are more “sensitive” (or “attentive”) to initial time intervals and less “sensitive” to later time intervals.

In contrast, when expressing preferences in terms of time, they seem to "spread" their sensitivity (or attention) more equally across time intervals.

Acknowledgments We would like to thank Abhijit Banerji, Roland Benabou, Colin Camerer, John Doyle, Jacob Goeree, Jens Großer, David Hardisty, Joel Huber, Eric J. Johnson, Nobuo Koida, Alexey Kushnir, Graham Loomes, J.V. Meenakshi, Daniel Mochon, C. Dan Myers, Thomas Palfrey, Uliana Popova Makarov, Ganna Pogrebna, Eldar Shafir, Rohini Somanathan, Dean Spears, Peter Wakker, and four anonymous reviewers (including one from the Russell Sage Foundation) for their helpful comments and suggestions. We also received helpful feedback from participants of SJDM 2007, BDRM 2008, Yale Whitebox 2008, ACR 2008, ESA 2008, SJDM 2008, SCP 2009, SPUDM 2009, the DR@W seminar at the University of Warwick, the ECO seminar at the University of East Anglia, and the Center for Development Economics seminar at the Delhi School of Economics. We are especially indebted to Chris Crabbe, who did a spectacular job programming the experiment. This research was generously funded by a grant from the Russell Sage Foundation's small grants program in Behavioral Economics and an award from the Society for Judgment and Decision Making.

References

Attema, A. E., Bleichrodt, H., Rohde, K. I. M., & Wakker, P. P. (2010). Time-tradeoff sequences for analyzing discounting and time inconsistency. *Management Science, 56*(11), 2015–2030.

Baltussen, G., Post, T., van den Assem, M. J., & Wakker, P. P. (2012). Random incentive systems in a dynamic choice experiment. *Experimental Economics, 15*(3), 418–443.

Bartolini, L., & Cottarelli, C. (1997). Treasury bill auctions: Issues and uses. In M. I. Blejer & T. Ter-Minassian (Eds.), *Macroeconomic dimensions of public finance: Essays in honour of Vito Tanzi* (pp. 267–336). London: Routledge.

Beattie, J., & Loomes, G. (1997). The impact of incentives upon risky choice experiments. *Journal of Risk and Uncertainty, 14*(2), 155–168.

Benhabib, J., Bisin, A., & Schotter, A. (2010). Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior, 69*(2), 205–223.

Carmon, Z., & Ariely, D. (2000). Focusing on the forgone: Why value can appear so different to buyers and sellers. *Journal of Consumer Research, 27*, 360–370.

Chesson, H., & Viscusi, W. K. (2000). The heterogeneity of time-risk tradeoffs. *Journal of Behavioral Decision Making, 13*, 251–258.

Cubitt, R. P., Starmer, C., & Sugden, R. (1998). On the validity of the random lottery incentive system. *Experimental Economics, 1*(2), 115–131.

De Borger, B., & Fosgerau, M. (2008). The trade-off between money and time: A test of the theory of reference-dependent preferences. *Journal of Urban Economics, 64*(1), 101–115.

Delgado, M. R., Schotter, A., Ozbay, E. Y., & Phelps, E. A. (2008). Understanding overbidding: Using the neural circuitry of reward to design economic auctions. *Science, 321*, 1849–1852.

Duxbury, D., Keasey, K., Zhang, H., & Chow, S. L. (2005). Mental accounting and decision making: Evidence under reverse conditions where money is spent for time saved. *Journal of Economic Psychology, 26*, 567–580.

Elster, J. (1984). *Ulysses and the sirens*. New York: Cambridge University Press.

Ersner-Hershfield, H., Wimmer, G. E., & Knutson, B. (2009). Saving for the future self: Neural measures of future self-continuity predict temporal discounting. *Social Cognitive and Affective Neuroscience, 4*(1), 85–92.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature, 40*, 351–401.

Goldstein, W. M., & Einhorn, H. J. (1987). Expression theory and the preference reversal phenomena. *Psychological Review, 94*(2), 236–254.

Green, L., Fristoe, N., & Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review, 1*, 383–389.

Grether, D. M., & Plott, C. R. (1979). Economic theory of choice and the preference reversal phenomenon. *American Economic Review, 69*(4), 623–638.

Hey, J. D., & Lee, J. (2005a). Do subjects remember the past? *Applied Economics, 37*(1), 9–18.

Hey, J. D., & Lee, J. (2005b). Do subjects separate (or are they sophisticated)? *Experimental Economics, 8*, 233–265.

Heyman, J. E., Orhun, Y., & Ariely, D. (2004). Auction fever: The effect of opponents and quasi-endowment on product valuations. *Journal of Interactive Marketing, 18*, 7–21.

Horowitz, J. K. (1991). Discounting money payoffs: An experimental analysis. In S. Kaish & B. Gilad (Eds.), *Handbook of behavioral economics* (pp. 309–324). Greenwich, CT: JAI Press.

James, D. (2011). Incentive compatible elicitation procedures. *Modsim Conference Proceedings, 19*, 1421–1427.

Kagel, J. H. (1995). Auctions: A survey of experimental research. In J. H. Kagel & A. E. Roth (Eds.), *Handbook of experimental economics* (pp. 501–585). Princeton, NJ: Princeton University Press.

Kagel, J. H., Harstad, R. M., & Levin, D. (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica, 55*, 1275–1304.

Kagel, J. H., & Levin, D. (1993). Independent private value auctions: Bidder behavior in first-, second- and third-price auctions with varying numbers of bidders. *Economic Journal, 103*, 868–879.

Kahneman, D., Knetsch J. L., & Thaler R. H. (1990). Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy, 98*, 1325–1348.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–292.

Killeen, P. R. (2009). An additive-utility model of delay discounting. *Psychological Review, 116*(3), 602–619.

Kirby, K. N. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General, 126*, 54–70.

Kirby, K. N., & Marakovic, N. N. (1995). Modeling myopic decisions: Evidence for hyperbolic delay-discounting within subjects and amounts. *Organizational Behavior and Human Decision Processes, 64*, 22–30.

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

Leclerc, F., Schmitt, B. H., & Dube, L. (1995). Waiting time and decision making: Is time like money? *Journal of Consumer Research, 22*, 110–119.

Lee, J. (2008). The effect of the background risk in a simple chance improving decision model. *Journal of Risk and Uncertainty, 36*, 19–41.

Lewis, G., & Bajari, P. (2011). Procurement contracting with time incentives: Theory and evidence. *Quarterly Journal of Economics, 126*(3), 1173–1211.

Lichtenstein, S., & Slovic, P. (Eds.). (2006). *The construction of preference*. New York: Cambridge University Press.

Liu, W., & Aaker, J. (2008). The happiness of giving: The time-ask effect. *Journal of Consumer Research, 35*, 543–557.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior, vol. 5: The effect of delay and of intervening events on reinforcement value* (pp. 55–73). Hillsdale, NJ: Lawrence Erlbaum.

McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science, 306*, 503–507.

Meier, S., & Sprenger, C. D. (2015). Temporal stability of time preferences. *Review of Economics and Statistics, 97*(2), 273–286.

Mischel, W., Ebbesen, E. B., & Zeiss, A. R. (1972). Cognitive and attentional mechanisms in delay of gratification. *Journal of Personality and Social Psychology, 35*, 204–218.

Mischel, W., Shoda, Y., & Peake, P. K. (1988). The nature of adolescent competencies predicted by preschool delay of gratification. *Journal of Personality and Social Psychology, 54*(4), 687–696.

Mischel, W., Shoda, Y., & Rodriguez, M. I. (1989). Delay of gratification in children. *Science, 244*, 933–938.

Monga, A., & Saini, R. (2009). Currency of search: How spending time on search is not the same as spending money. *Journal of Retailing, 85*(3), 245–257.

Okada, E. M., & Hoch, S. J. (2004). Spending time versus spending money. *Journal of Consumer Research, 31*, 313–323.

Olivola, C. Y., & Sagara, N. (2009). Distributions of observed death tolls govern sensitivity to human fatalities. *Proceedings of the National Academy of Sciences, 106*, 22151–22156.

Parfit, D. (1984). *Reasons and persons*. Oxford: Oxford University Press.

Phelps, E. S., & Pollak, R. A. (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies, 35*, 185–199.

Pronin, E., Olivola, C. Y., & Kennedy, K. A. (2008). Doing unto future selves as you would do unto others: Psychological distance and decision making. *Personality and Social Psychology Bulletin, 34*(2), 224–236.

Reimers, S., Maylor, E. A., Stewart, N., & Chater, N. (2009). Associations between a one-shot delay discounting measure and age, income, education and real-world impulsive behavior. *Personality and Individual Differences, 47*, 973–978.

Roelofsma, P. H. M. P. (1996). Modelling intertemporal choices: An anomaly approach. *Acta Psychologica, 93*, 5–22.

Saini, R., & Monga, A. (2008). How I decide depends on what I spend: Use of heuristics is greater for time than for money. *Journal of Consumer Research, 34*, 914–922.

Samuelson, P. (1937). A note on measurement of utility. *Review of Economic Studies, 4*, 155–161.

Scholten, M., & Read, D. (2010). The psychology of intertemporal tradeoffs. *Psychological Review, 117*(3), 925–944.

Schwarz, G. E. (1978). Estimating the dimension of a model. *Annals of Statistics, 6*, 461–464.

Soman, D. (2001). The mental accounting of sunk time: Why time is not like money. *Journal of Behavioral Decision Making, 14*, 169–185.

Soster, R. L., Monga, A., & Bearden, W. O. (2010). Tracking costs of time and money: How accounting periods affect mental accounting. *Journal of Consumer Research, 37*(4), 712–721.

Starmer, C., & Sugden, R. (1991). Does the random-lottery incentive system elicit true preferences? An experimental investigation. *American Economic Review, 81*(4), 971–978.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology, 53*, 1–26.

Sussman, A. B., & Olivola, C. Y. (2011). Axe the tax: Taxes are disliked more than equivalent costs. *Journal of Marketing Research, 48*(SPL), S91–S101.

Takeuchi, K. (2011). Non-parametric test of time consistency: Present bias and future bias. *Games and Economic Behavior, 71*(2), 456–478.

Thaler, R., & Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science, 36*(6), 643–660.

Todorov, A., Goren, A., & Trope, Y. (2007). Probability as a psychological distance: Construal and preferences. *Journal of Experimental Social Psychology, 43*(3), 473–482.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*, 1124–1131.

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference dependent model. *Quarterly Journal of Economics, 107*, 1039–1061.

Tversky, A., Slovic, P., & Kahneman, D. (1990). The causes of preference reversal. *American Economic Review, 80*(1), 204–217.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance, 16*, 8–37.

Zauberman, G., & Lynch, J. G. (2005). Resource slack and propensity to discount delayed investments of time versus money. *Journal of Experimental Psychology: General, 134*, 23–37.