Are Discount Rates Constant?

Reconciling Theory and Observation

by

Maribeth Coller, Glenn W. Harrison and E. Elisabet Rutström†

January 2005

Abstract. We conduct a series of laboratory experiments designed to test if individual discount rates vary with the time horizon for which the rate is elicited. We conclude that rates of time preference are not continuously declining, although observations generated in many prior experiments have appeared to reflect such preferences. Our observations are consistent with the intrinsically confounded notions of a "passion for the present" and other "costs of *any* delayed payment." The latter may simply be due to the subjective transactions costs of receiving payment at a later date, rather than reflecting time preferences. Each explanation results in a fixed premium being attached to payment options delayed beyond the present. We demonstrate the existence of this premium in choices between an immediate payment option and a delayed payment option. When both payment options are delayed, the premium does not vary between the options, and thus does not enter into the choice between them. The revealed rates then do not vary with the choice horizon. To the extent that many significant field choices or policy decisions do not involve choices with immediate payoffs, the use of constant discount rates remains appropriate. We also identify individual characteristics associated with higher or lower fixed premia when there is no delay in the earlier payoff. For some groups of individuals the premia is significantly higher than average, while for other groups it is virtually absent.

† School of Accounting, Moore School of Business, University of South Carolina, USA (Coller) and Department of Economics, College of Business Administration, University of Central Florida, USA (Harrison and Rutström). E-mail contacts: MBETH@MOORE.SC.EDU, GHARRISON@BUS.UCF.EDU and ERUTSTROM@BUS.UCF.EDU. Helpful comments have been received from Morten Lau, Melonie Sullivan, and seminar participants. We are grateful for funding from the Friends of the Accounting Department, School of Accounting, Moore School of Business, University of South Carolina. Rutström thanks the U.S. National Science Foundation for research support under grants NSF/IIS 9817518, NSF/MRI 9871019 and NSF/POWRE 9973669. Supporting data, statistical code and experimental instructions are stored in the ExLab Digital Library at http://exlab.bus.ucf.edu.

What generates seemingly declining discount rates? One hypothesis is that individuals do indeed have different rates of time preference for different time horizons, such as continuously hyperbolic preferences. An alternative hypothesis is referred to as quasi-hyperbolic preferences.[1] This hypothesis holds that the individual has a "fixed premium" that applies to money deferred for some initial time delay, but has constant discount rates for any monies deferred for further horizons beyond the initial delay.

Preferences involving this fixed premium could give rise to discount rates over time horizons that *appear* to be non-constant and "hyperbolicky." For example, assume that option A is $100 available immediately, and option B is some amount to be received later. If the individual demands $20 in compensation to wait *at all*, the subject would require $20, in addition to his interest requirement, in order to choose option B. If his annual discount rate is constant and equal to 10%, the subject would prefer option B over horizons of 1, 2 and 3 months only when option B is a future payoff of at least $120.83, $121.67 and $122.53, respectively. If the experimenter ignores the $20 fixed premium that the subject demands to wait for payment, he infers annual discount rates of 868%, 224% and 125%, respectively.[2] However, these declining rates are simply a result of the $20 premium being "spread" over longer time horizons. Alternatively, if the choice were between payments that both involved a delay, then the fixed premium would be associated with both options. For example, if option A is a payment of $100 in one month, and option B is a payment of more than $100 beyond one month, then this individual would prefer option B when it is $100.83, $101.67, and $102.53 over horizons of 1, 2 and 3 months, respectively. These choices correctly reveal a 10% discount rate.

Thus there are two alternative hypotheses with quite different testable implications for behavior. If declining discount rates are generated by the process underlying the notion of quasi-

[1] So-called "quasi-hyperbolic discount functions" were introduced by Phelps and Pollack [1968], and have been adopted by Laibson [1997], O'Donoghue and Rabin [1999], Angeletos et al. [2001] and others. Such specifications have been given various names. For example, O'Donoghue and Rabin [1999] call them "present-biased preferences" to emphasize that they reflect differences in pure rates of time preference.

[2] For example, 868% is the rate that equates $100 now with $120.83 in one months, and so on.

hyperbolic preferences, then the use of a front end delay (FED) should have a dramatic effect on behavior, assuming of course that the FED is longer than the shortest time horizon which generates the fixed premium. Specifically, a FED of any reasonable length should result in a marked reduction in elicited discount rates compared to a setting in which there was no FED. Moreover, the monetary premium required for delay should be a linear function of horizon.[3] However, if declining discount rates are generated by the process underlying the original notion of continuously hyperbolic preferences, then the FED should have no effect. That is, one should still expect to elicit higher discount rates for earlier horizons, despite the FED.[4]

We conduct a series of experiments that are designed to differentiate between the two hypotheses. We provide subjects with choices of two basic types: (I) money now or money later, and (ii) money later or money much later. A critical design choice is the length of the FED involved in choices of the second type. The goal is to provide a FED that is just long enough so that subjects do not perceive the earlier option as "immediate." We choose our FEDs with this in mind, and use 7 days as the minimum (non-zero) FED.[5] Our design also incorporates a FED of 30 days.

In our treatments without a FED, as has been employed in much of the empirical literature, we do indeed find individual discount rates that *appear* to decline over the choice horizon when the fixed premium is not included in the model. However, once a fixed premium is allowed for by using a FED, discount rates no longer decline over the horizon. Thus, although we reject the hypothesis of continuously declining discount rates, we can reconcile our findings with other experimental results that appear to support hyperbolic preferences. This leaves two possible explanations: either quasi-hyperbolic preferences, or constant preferences with some artefactual transactions costs

---

[3] This is equivalent to the claim that the elicited discount rates should exhibit no statistically significant pattern with respect to horizon.

[4] One could imagine a generalized hyperbolic specification in which the rates declined so rapidly with horizon that one could not tell that apart from the quasi-hyperbolic case empirically. We cannot rule such extreme cases out, but then such specifications would not be operationally meaningful in the sense of imposing any refutable restrictions on behavior.

[5] This choice is clearly subjective. "Very short" FEDs, such as less than one day were rejected since subjects may still attach premiums to the later option that they do not attach to the earlier option. We leave it to future research to determine the minimum FED to which such premia do not apply.

generating a fixed premium when there is no FED.

The implications of our findings depend upon the field context to which they are being applied. If the existence of a FED is consistent with the field context or policy decision being made, then our findings lead to the conclusion that preferences are constant over varying horizons.[6] We believe that this is often the case. Because significant policy decisions involving trade-offs between the "present" and the future do not typically involve *immediate* payoffs, which in this instance refer to payments being made on the day of the decision, most policy decisions inherently do involve a FED. However, if the FED is inappropriate for a specific field application, then our findings lead to the conclusion that subject choices might *appear* to reflect hyperbolic discount rates, albeit only over very short horizons. As our experimental design suggests, however, such preferences could be the result of subjective factors resulting in a premium attached to payments involving a delay of any length, rather than reflecting pure time preference.

Our results also demonstrate considerable individual heterogeneity in the fixed premium required to delay payouts. We do find, however, that the fixed premium is significantly correlated with several observable individual characteristics. For example, Asian-Americans and those with higher household incomes appear to require no fixed premium for any delay. On the other hand, older undergraduates, those with higher GPAs, those majoring in business, and those with richer parents tend to have relatively higher fixed premia. Thus one begins to see a richer characterization of the domain of violations of constant discount rates when there is no front end delay on payouts.

[6] This position is perfectly consistent with O'Donoghue and Rabin [1999; p.107, fn.10], although they do not say how long the FED needs to be.

1. Experimental Design

Our experimental design tests for constancy of individual discount rates (IDRs) using two basic treatments:

- DELAYED PAYOFF EXPERIMENTS. In these experiments we elicit individual discount rates for monetary streams using a positive FED, such that both the "shorter horizon" and "longer horizon" payoffs occur in the future. That is, *both* payoffs occur with a delay. We vary the length of time between the shorter and longer horizon payoffs that each individual faces. Each individual faces just one such time horizon.
- INSTANT PAYOFF EXPERIMENTS. In these experiments, which otherwise mimic the delayed payoff experiments, we have no FED. In other words, the shorter horizon payoffs occur immediately and the longer horizon payoffs occur in the future. This feature is widely used in much of the empirical literature associated with extremely high discount rates.

These treatments represent simple extensions of the experimental design employed by Coller and Williams [1999] and Harrison, Lau and Williams [2002].

The basic question used to elicit individual discount rates is simple, and is of this general form: do you prefer \$100 today or \$100+x in the future, where x is some positive amount? If the subject prefers the \$100 today then we can infer that the discount rate is higher than x% over the period of the delay; otherwise, we can infer that it is x% or less over the period of the delay. The format of our experiment modifies and extends this basic question in five ways.

First, we pose a number of such questions to each individual, each question varying x by some amount. When x is zero we would expect the individual to reject the option of waiting for a zero rate of return. As we increase x we would expect more individuals to take the future income option. For any given individual, the point at which he switches from choosing the current income option to taking the future income option provides a bound on his discount rate. That is, if an individual takes the current income option for all x from 0 to 10, then takes the future income option for all x from 11 to 100, we can infer that his discount rate lies between 10% and 11% for

this time interval. The finer the increments in x, the finer is the discount rate of the individual that we will be able to bracket.[7] We then select one question at random for actual payment after all responses have been completed by the individual. In this way the results from one question do not generate income effects which might influence the answers to other questions.

Second, in the delayed payoff experiments we provide two future income options rather than one "instant income" option and one future income option. For example, we offer \$500 in 30 days and \$500+x in 90 days, interpreting the revealed discount rate as applying to a time horizon of 60 days. [8] This is intended to neutralize the effect of any fixed, horizon-independent, premia which subjects may attach to future income options. For example, if delayed options involve greater transactions costs than immediate options, then the compensation demanded to wait for the delayed option would include these subjective transactions costs. Similarly, if preferences include a "passion for the present," then the compensation demanded to accept a payment with any delay would include an amount to compensate for this. In either case the revealed discount rate would be biased upwards, unless we account for a fixed premium. By having both options entail future income, we hold these subjective costs constant across the payoff options and hence gain greater control over the comparability of the two monetary payoffs.[9]

Third, we consider many different horizons, ranging from 1 day to 60 days after the FED. We assign one of these time horizons to each subject at random, and elicit discount rates pertaining only to that horizon. We are thus able to plot discount rates over a variety of horizons and explicitly test whether rates are related to the time horizon.

Fourth, we provide respondents with information on the implied interest rates associated

---

[7] The same procedure, which can be referred to as a Multiple Price List (MPL) auction, is used to elicit discount rates by Kirby and Maraković [1996] and Kirby, Petry and Bickel [1999]. They randomize the order in which the options are presented to subjects, whereas we deliberately order them in terms of increasing rates of return. Because we seek the cleanest test of the hypothesis of interest, we do not address possible confounds such as task complexity.

[8] Because our subjects are students, we take some care in our experiments to avoid payouts that fall outside of the regular academic year, during recognized holidays, or on weekends.

[9] Although the logic of this point is intuitive enough, and appears to have been first stated by Roberts [1991], it was first tested by design in Coller and Williams [1999].

with the delayed payment option. This is an important control feature if field investments are priced in terms of interest rates. To the extent that some subjects are attempting to compare the lab investment to their field options, this feature may serve to reduce comparison errors since both the lab and field options are now priced using the same metric.[10]

Table 1 shows a payoff table for the case of a subject facing a 60 day horizon and a 30-day FED. In other words, the later payment, option B, occurs 60 days beyond the earlier payment, option A. The subject circled which option, A or B, he would prefer in each row. For subjects facing different time horizons, the annual and effective rates were the same as those in Table 1, but the payment amounts for option B were appropriately different.

2. Previous Literature

Our goal is to place the existing empirical findings in context, so that we can see when and why non-constant discount rates emerge. Our experimental design therefore differs from much of the previous literature in some respects, although each of the key features has been employed previously in other experiments.[11] The most important design feature is that we employ a FED on the choices presented to subjects in order to control for any confounding effects from fixed premia as explained above.

O'Donoghue and Rabin [1999; p.103] motivate their analysis of time-inconsistent preferences with the following passage:

People are impatient — they like to experience rewards soon and to delay costs until

---

[10] Coller and Williams [1999] show that providing these rates has a significant and negative effect on the mean and variance of responses. There is an important field counterpart to providing the rates associated with the payment options: truth in lending and savings laws require disclosure of both the annual and effective interest rates associated with credit market instruments. Hence, the provision of the implied rates serves to provide information in the laboratory that is consistent with the information available to individuals in formal credit markets.

[11] Two design features are discussed further in an appendix, since they may be of more technical interest. One is the fact that we ask our participants to make choices over real rather than hypothetical assets. Although we are not the first to do so, the bulk of the literature has relied on the use of hypothetical assets. The other design feature is that we employ an elicitation method with known incentive compatibility properties.

> later. Economists almost always capture impatience by assuming that people discount streams of utility over time exponentially. Such preferences are time-consistent: A person's relative preference for well-being at an earlier date over a later date is the same no matter when she is asked.
>
> Casual observation, introspection, and psychological research all suggest that the assumption of time consistency is importantly wrong. It ignores the human tendency to grab immediate rewards and to avoid immediate costs in a way that our "long-run selves" do not appreciate. For example, when presented a choice between doing seven hours of an unpleasant activity on April 1 versus eight hours on April 15, if asked on February 1 virtually everyone would prefer the seven hours on April 1. But come April 1, given the same choice, most of us are apt to put off the work until April 15. We call such tendencies present-biased preferences: When considering trade-offs between two future moments, present-biased preferences give stronger relative weight to the earlier moment as it gets closer. (Footnotes omitted)

This passage seems to confound two things, each of which are important.[12] The first is whether the length of the time horizon over which the discount rate is being elicited affects the discount rate, such as it does with continuously hyperbolic preferences. The second is whether the discount rate elicited with a FED is different than a discount rate elicited with no FED, i.e. whether preferences are quasi-hyperbolic. For an experimenter, and for subjects evaluating the credibility of being paid, these are very different questions.

Our design is intended to separate the effects of the FED from the pure effects of the length of the time horizon. The potential importance of this distinction seems to have been first noticed by Benzion, Rapoport and Yagil [1989].[13] It was also highlighted by Roberts [1991; p.344], in the context of comments on Ainslie and Haendel [1983] and Winston and Woodbury [1991]:

> There is a bias toward choosing the small-early reward, particularly in the study by Ainslie and Haendel [1983] where real money changed hands. Ainslie and Haendel find ridiculously high implicit interest rates. An individual who prefers $10 today to $12.50 three days from now is turning down a rate of return that is a compounded annual rate of well over a trillion percent. Ainslie and Haendel, in a passage quoted by Winston and Woodbury, suggest that this is how people behave in one-time events and is not inconsistent with people willing to put money in a savings account at 5 percent - a repetitive activity.

---

[12] Their formal analysis employs quasi-hyperbolic preferences, which do not suffer from these confounds.

[13] Holcomb and Nelson [1992] re-examine the role of a FED with monetary payoffs, motivated by a concern that Benzion, Rapaport and Yagil [1989] only studied hypothetical choices. Their FED was only one day, so it is not obvious that the subjects viewed this as substantially different from there being no FED. They observed no apparent effect of the one-day FED on behavior.

> An alternative explanation is that Ainslie and Haendel have revealed the role of uncertainty in experiments lasting more than a single day. A student is likely to prefer $10 today to $12.50 in three days if there is some uncertainty about whether the experimenters will return in three days. I suspect that if respondents saw the experimenters consistently handing out money three days later at 25 percent interest rates so that they could be convinced that the experimenters were a going concern, they would take the larger return and not act as if they required scientific notation to express their discount rate.

One important reason that the FED design was introduced into discount rate experiments is the concern about differential credibility. While it may not completely solve the potential credibility problem,[14] it arguably mitigates it. For example, consider a FED of 30 days such that subjects choose between receiving different payment amounts in 30 days and in 60 days. While the subject may have some doubt about actually receiving payment in 30 days, this doubt is not likely to differ much from the doubt about receiving payment in 60 days. Similarly, the FED equalizes transactions costs between the two payment options. While there are some costs to returning at a later date to receive payment, these costs are not likely to differ between returning in 30 days vs. returning in 60 days.[15] The FED also serves to equalize any other unspecified differences subjects may perceive between the two payment options. For example, if subjects have a "passion for the present," they demand a premium in order to accept a delay of any length. In a choice between immediate payment and delayed payment, this premium is attached only to the delayed payment. However, if both payments are delayed, the premium applies to both choices and thus becomes irrelevant in choosing between them.

Having said this, there are many field settings in which the relevant issue is what the discount rate is for "money today" versus "money in the future."[16] Even if the experimenter faces the

---

[14] It could be that credibility in the mind of the subject is a decreasing function of the time between the experiment and the next payoff. Hence an experiment with a FED of 10 years would be tantamount to a hypothetical experiment in most settings. Conversely, a FED of 5 minutes would be the same as having no FED.

[15] This assumes no dramatic change in circumstances that makes collecting the payment more difficult later. For this reason, we are careful to only offer choices that pay off on a week day during the current teaching semester.

[16] Such settings might include individual decisions of whether to consume now or save for future consumption, or to purchase a more expensive but energy efficient appliance. We believe that individual

inferential problem of having to then tease apart the effects of time horizon from credibility, transactions, or other subjective costs, it is entirely appropriate that experiments with no FED be considered. If there is a finding that discount rates are not constant when there is no FED, then it is a matter for interpretation as to whether this is a subjective differential cost effect or a time-inconsistency effect (or both).

Evidence for the behavioral importance of a 30-day FED was provided by Coller and Williams [1999]. In one of their experimental treatments they had no such delay, and the results from those experiments can be directly compared to their other experiments. After some minor modifications to their statistical analysis, their results provide evidence that the use of a FED decreases elicited rates by a large amount. The average effect of having no FED is to increase elicited rates by 28 percentage points, with a 95% confidence interval between 52 percentage points and 3 percentage points; the coefficient on the dummy variable denoting this treatment is statistically significant at the 4.8% level.[17]

Kirby and Santiesteban [2003] use a FED of 1 day in an experiment in which they elicited the individual discount rates of financially-motivated subjects for horizons of between 1 and 43 days. They compare results with comparable experiments using no FED, and find that there is essentially no difference in the pattern: discount rates decline with time horizon. These results are valuable, and accord with our priors. They show that 1 day appears to be insufficient to overcome the subjective costs and resulting fixed premium we hypothesize.

---

decisions involving more significant sums of money or public policy decisions are better characterized as having a FED.

[17] Read [2001] develops an experimental design which includes a FED, and finds that it *increases* the elicited discount rate (reported as a decline in the discount factor) in the one experiment in which he used salient monetary rewards. Unfortunately, the method he uses to elicit discount rates is not incentive compatible, and he was obliged to drop some subjects that appeared to have exploited this flaw by claiming that they always preferred the shorter horizon option. Even if some subjects did not exploit the flaw to the point where they were dropped, it is possible that they were aware of it and inflated their responses to a point that they thought would not be detected. Either way, there are obvious problems of control over incentives to truthfully reveal discount rates. His design also examined only one horizon (8 months for the non-salient experiments, and 6 months for the salient experiment), making it difficult to tell if the effect of the FED persisted across other time horizons, which is the focus of our design; the same limitation applies to the design of Coller and Williams [1999].

3. Experimental Results

We recruited 217 subjects from the general undergraduate population of the University of South Carolina to participate in six sessions. There were 87 subjects (two sessions) in the experiments with no FED, 90 subjects (three sessions) in the experiments with a FED of 30 days, and 40 subjects (one session) in the experiments with a FED of 7 days. Subjects generally faced horizons between 1 day and 60 days after the FED.[18] Subjects were randomly assigned to sessions, and then to horizon.

The subjects were initially presented with an informed consent form and then with the written instructions. These instructions were read out loud by an experimenter, who then ran the subjects through a "hands on" trainer in which they made choices over receiving Hershey's Kisses immediately or more Hershey's Kisses at the end of the experiment. After the trainer was completed, and all subjects had an opportunity to ask any questions, the main experiment was implemented. When all subjects had made their choices on record sheets similar to the one appearing in Table 1, the record sheets were collected and subjects were asked to complete a questionnaire eliciting socio-demographic information as well as some information about their personal finances. One payoff alternative was selected at random, and then one subject was selected at random for payment. All random choices were implemented with a visible bingo cage. Each experiment lasted an average of 45 minutes.

Focusing directly on the dollar difference between options A and B, we can directly investigate the hypothesis that the absence of a FED results in a fixed dollar premium for delay.[19] Figure 1 displays the observed data for the experiments with no FED. Because the vertical axis is in

---

[18] We employed a horizon of only 30 days in the session with a FED of 7 days because, given the date on which it was conducted, a longer horizon would mean that subjects would be evaluating options that were to be paid after the end of the regular teaching semester. This would have added possible transactions costs to the deferred option that did not apply to the shorter-term option.

[19] Keller and Strazzera [2002] evaluate alternative discounting models in terms of the monetary premia demanded by subjects to delay payouts. One can also equivalently examine the discount rates implicitly chosen by the subject, as in Harrison, Lau and Williams [2002]. In an appendix we evaluate some survey data collected from our subjects to suggest that most of our subjects did focus on the monetary premium, *inter alia*.

terms of dollars required to choose the later payment option, a positive relation is expected between the delay premium and the time horizon. The slope of the line describing the data reflects the interest rate, while a significant intercept reflects the fixed premium. The light, dashed line shows the maximal delay premium available for that horizon; this amount corresponds to a 100% annual rate (or a 171.45% annual effective rate) of return. Subjects that reject the delayed option at all interest rates are shown on that line.[20] We observe many subjects on that maximal premium line for short horizons. Under the assumption that a fixed delay premium exists, this is not surprising. For example, the dollar premium equating a 100% annual return with a delay of 1 day is only $1.39. Thus, a subject who is unwilling to wait at all for $1.39 may simply have a fixed premium greater than $1.39; yet he appears to demand more than 100% interest. In fact, the maximal dollar premium exceeds $10 only for a horizon of 8 days or more, so subjects with horizons of a week or less that required a fixed premium of, say, $10 for *any* delay would simply never delay.

Employing appropriate statistical models for censored responses in the experiments with no FED, we estimate the relation between delay premium and horizon shown by the dark, solid line in Figure 1. This relation clearly has a positive intercept, representing the fixed premium, at a horizon of 0. The fixed premium is estimated to be $10.28 with a p-value of 0.010. The slope of the line is estimated to be $0.49 per day of delay (or 0.098% of the $500 principle), and has a p-value of 0.007. Figures 2 and 3 show comparable results for the two treatments with a FED. No significant fixed premium is found in the case of the 7-day FED ($0.31, $p = 0.91$), and only a marginally significant fixed premium is estimated in the case of the 30-day FED ($4.05, $p = 0.16$).

Considerable heterogeneity across subjects is reflected in the fixed monetary premium required. Table 2 includes the results of a statistical model designed to identify the significance of these differences, based on observable individual characteristics. This model is estimated only on the responses in the treatment with no FED. Controlling for individual characteristics results in a fixed

[20] More accurately, such subjects never choose the delayed premium and hence have a delay premium that is *greater* than the maximal amount. We treat these observations as censored at the maximal amount.

premium of $15.49 for the individuals included in the implicit categories that are omitted (e.g., female, aged 21, neither black nor Asian, etc.). Restricting attention to estimates that are statistically significant at the 5% level, older students have a much higher fixed monetary premium ($38.95 higher). Because graduate students, a group which is entirely composed of older students, have a fixed premium that is $40.13 lower, the age effect is entirely due to older undergraduates. Asians have a fixed premium that is $11.53 lower than the overall group. Those with the best GPA levels have a fixed premium that is $6.94 higher on average. Those subjects with a household income above $15,000 have virtually a fixed premium that is $9.21 lower, and those students with richer parents tend to have a fixed premium that is $9.43 higher than the overall group.

Pooling responses across all FED treatments, we can undertake a more controlled statistical comparison of the main treatment. We condition on the individual characteristics of each sample, as listed in Table 2, and also correct for any heteroskedasticity that depends directly on the absence of the FED and time horizon. The estimated intercept in this pooled analysis for the responses observed with no FED is strongly significant, although slightly smaller in size compared to the estimates that did not condition on differences in the characteristics of the sample ($6.23, $p = 0.007$, with 95% confidence intervals of $1.74 and $10.72). The comparable intercepts for the treatments with a FED are not statistically significant.[21] Thus we have clear evidence that *the fixed premium effect is significant in the absence of a FED, and that it disappears with a FED.*

4. Conclusions

We have shown that if the FED is 30 days, subjects behave as if they have constant discount rates over time. We find similar behavior when the FED is only 7 days. On the other hand, we have also shown that if the FED is 0 days, subjects behave as if they have extremely high discount rates

[21] It is easy to test for possibility that discount rates continuously decline over horizon, by directly testing if the monetary premium demanded is a non-linear function of horizon. Adding quadratic and cubic terms to the above statistical specifications leads to a convincing rejection of the hypothesis of any non-linearity.

for shorter time horizons that decline during the next few weeks.

Our results focus attention on the interpretation of the FED feature of intertemporal choice experiments. This feature controls for plausible contaminants of behavior in choice settings such as the asymmetric credibility of a payment by the experimenter "today" rather than "in the future" and asymmetric transactions costs to the subject of collecting the payoffs, as well as any other asymmetric subjective costs that the subject may attach to having to face a delay of any length. Any such subjective costs would result in the demand for additional compensation in the form of a fixed premium.

Although we do not provide evidence as to why subjects attach a premium (in addition to their required rate of interest) to deferred payment options, we provide evidence that this fixed premium affects the compensation demanded only in choices made in the presence of no FED. Furthermore, the amounts required by subjects to compensate them for such a fixed premium generate observations appearing to reflect hyperbolic discounting. When subjects choose between payments that both involve a delay, there is no differential fixed premium and no evidence of rates declining as the choice horizon lengthens.

If the FED is viewed as a necessary device to control for the subjective cost confounds of earlier experiments, then we have shown that one cannot reject the hypothesis that discount rates are constant over the time horizons considered. We view the FED in this way, hence we reject the alternative hypothesis of continuously declining discount rates.

Whether our results are due to the elimination of differential transactions costs and risk premia between the choices offered, or to subjects simply having some other "passion for the present" which does not come into play in the presence of a FED, remains unknown. The difficulty in designing an experiment to separate these possibilities is obvious. To do so would require a decision task that involves "now vs. later," allowing for present-biased preferences to manifest themselves, while keeping other costs constant between the two choices. Clearly any choice that involves "now vs. later" must involve at least some transactions cost (if nothing more than the effort

of keeping track of the arrangements needed to receive payment) and some (however minimal) risk of default in payment. These difficulties call into question the relevance of drawing such a distinction in the first place. Field choices suffer from the same intrinsic confounds as lab choices. For example, the simple choice of whether to consume presently or to invest in a certificate of deposit involves not only a possible present-biased desire to "spend now," but also requires consideration of any default risk of the investment instrument and some effort in keeping track of the investment and redeeming it when it necessary.

We conclude that rates of time preference are not continuously declining, but that it is simply not possible to test the elements of the quasi-hyperbolic specification in an operationally meaningful manner. The *appearance* of continuously declining rates of time preference is consistent with either a "passion for the present" or with other "subjective costs of *any* delayed payment," and these appear to be intrinsically confounded. To the extent that many significant field choices or policy decisions do not involve choices with immediate payoffs, the use of constant discount rates remains appropriate. Furthermore, we can identify significant groups of individuals that tend to suffer more or less from the need for a fixed premium for any time delay.

Table 1: Illustrative Payoff Matrix

Payoff Alternative	**Payment Option A** (pays amount below in 30 days)	**Payment Option B** (pays amount below in 90 days)	**Annual Interest Rate** (AR)	**Annual Effective Interest Rate** (AER)	**Preferred Payment Option** (Circle A or B)
1	$500	$501.67	2.00%	2.02%	A B
2	$500	$502.51	3.00%	3.05%	A B
3	$500	$503.34	4.00%	4.08%	A B
4	$500	$504.18	5.00%	5.13%	A B
5	$500	$506.29	7.50%	7.79%	A B
6	$500	$508.40	10.00%	10.52%	A B
7	$500	$510.52	12.50%	13.31%	A B
8	$500	$512.65	15.00%	16.18%	A B
9	$500	$514.79	17.50%	19.12%	A B
10	$500	$516.94	20.00%	22.13%	A B
11	$500	$521.27	25.00%	28.39%	A B
12	$500	$530.02	35.00%	41.88%	A B
13	$500	$543.42	50.00%	64.81%	A B
14	$500	$566.50	75.00%	111.53%	A B
15	$500	$590.54	100.00%	171.45%	A B

Figure 1: Choices with No FED

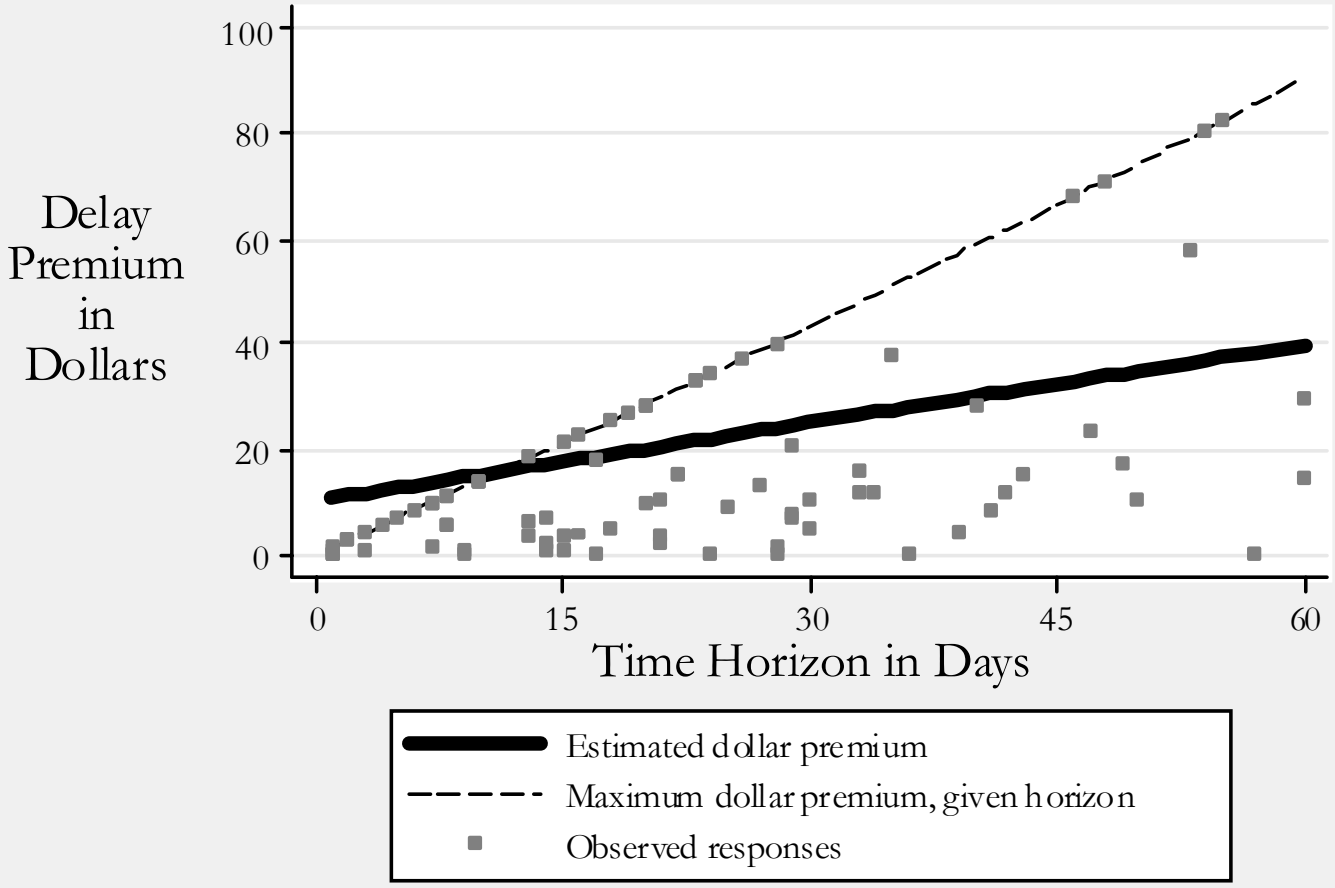

Figure 2: Choices with FED of 30 Days

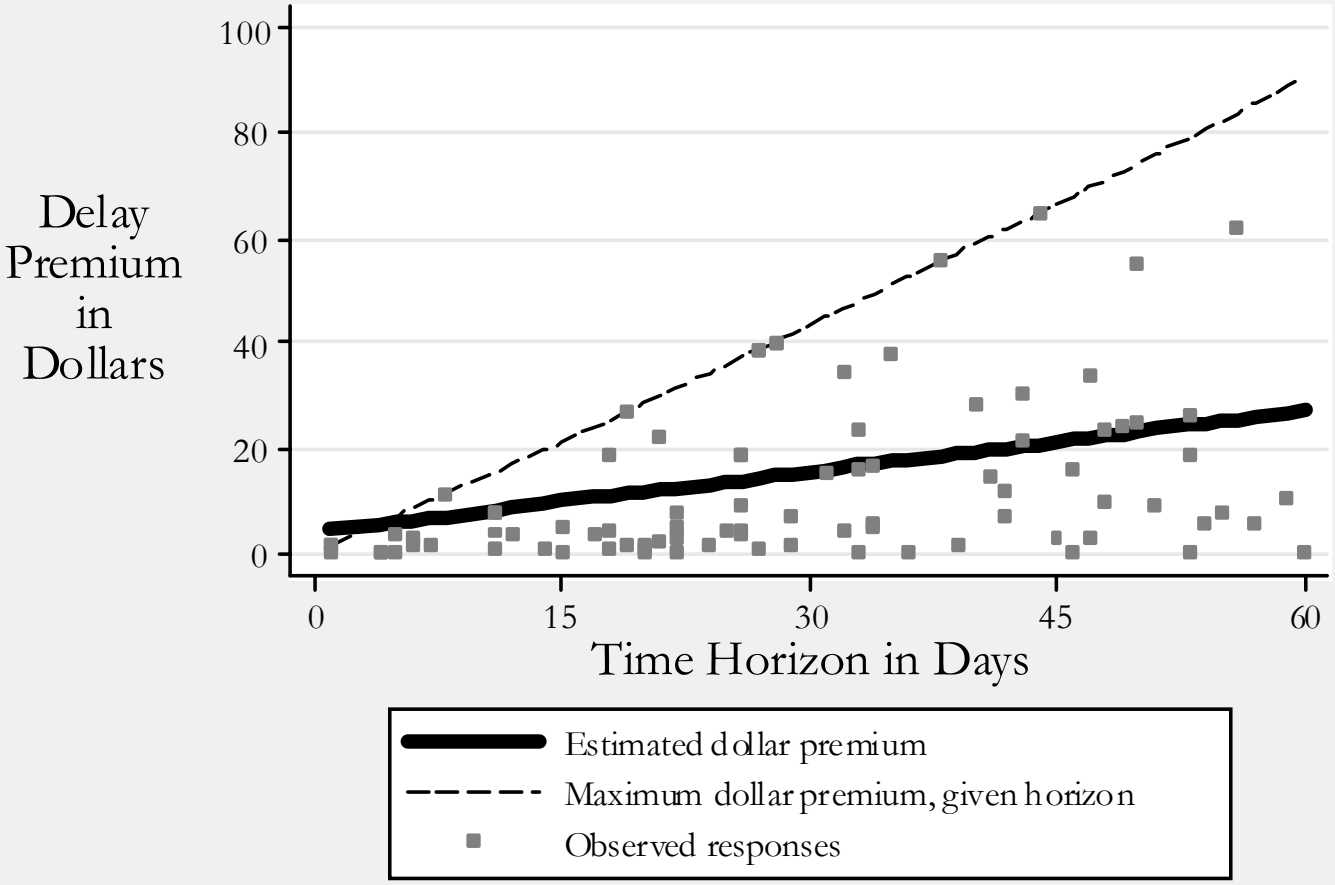

Figure 3: Choices with FED of 7 Days

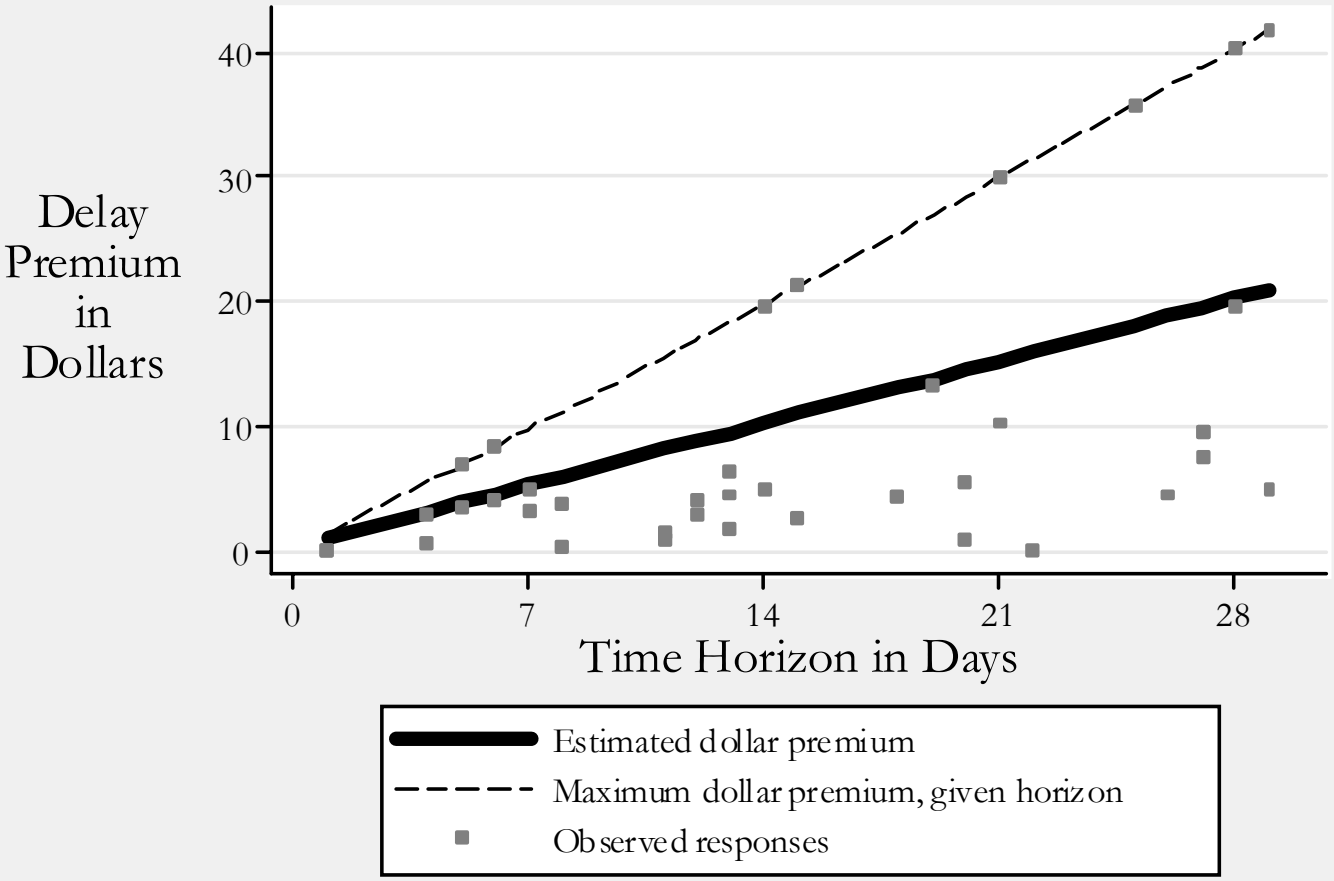

Table 2: Estimates of Determinants of Monetary Premium

Interval regression estimates of mean effects, with robust Huber-White standard errors
Dependent variable is the monetary premium chosen by the subject, expressed as an interval
Responses only for zero FED treatment; N=85
Multiplicative heteroskedasticity assumed, but estimates not reported

Variable	Description	Estimate	Standard Error	z	*p*
	Constant	15.49	7.441	2.08	0.037
HORIZON	Time horizon for option B in days	0.07	0.102	0.73	0.465
MALE	Male	1.65	4.068	0.41	0.685
YOUNGER	Younger than 21	3.42	3.180	1.07	0.283
OLDER	Older than 21	38.95	9.212	4.23	0.000
BLACK	Black	6.05	4.746	1.28	0.202
ASIAN	Asian or Asian-American	-11.53	2.247	-5.13	0.000
BUSINESS	Business major	-4.68	3.755	-1.25	0.213
STANDING	Freshman or Sophomore	-3.12	1.899	-1.64	0.101
GRADUATE	Graduate Student	-40.13	8.290	-4.84	0.000
GPAAA	GPA mainly has A's	6.94	2.119	3.28	0.001
GPAAB	GPA has A's and B's	2.37	3.856	0.61	0.540
HHY2	Household income above $15k	-9.21	1.890	-4.87	0.000
PARY2	Parental income between $65k and $80k	2.37	4.723	0.50	0.616
PARY3	Parental income above $80k	9.43	3.494	2.70	0.007
BALANCE	Carries a loan or credit card balance	0.39	3.624	0.11	0.915

Appendix A: Additional Discussion of Experimental Design Choices

1. Hypothetical Responses

One major difference in experimental design between the experiments reported here and most of the existing literature is our use of real, rather than hypothetical, payments. One of the hallmarks of experimental economics is the use of instructions and payoffs designed to ensure control over the incentives faced by individuals. This control facilitates interpretation in terms of existing theory. One aspect of control is the use of payoffs to the subject that vary with the responses made by the subject, and in a way that the subject understands. This is called "salience" in the terminology of Smith [1982]. It is also important that these payoffs be observable and measurable by the experimenter, thus ruling out "intrinsic motivation" as a candidate explanation for observed behavior. We cannot rely on subjects' personal desires to do the task correctly for their own satisfaction, since that is not observable and subjects may differ in this regard. Hence monetary rewards are customarily used by experimental economists to motivate subjects.

The existing literature relies heavily on responses to hypothetical choice situations which do not employ salient monetary incentives.[22] While hypothetical scenarios may be entirely appropriate in some contexts, the argument offered in favor of using hypothetical scenarios in many studies is simply convenience. The notion here is that there are certain "realistic" questions which one cannot feasibly ask within the usual constraints of budgets and ethical review boards.

Loewenstein and Thaler [1989; p. 184] offer the following argument in favor of using hypothetical choice scenarios:

> In this study, and some others described here, the questions asked were hypothetical. Of course, all things being equal it would be better to study actual choices. However, there are serious trade-offs between hypothetical and real money methods. Using hypothetical questions one can ask subjects to consider options that incorporate large amounts of money, both gains and losses, and delays of a year or more. In studies using real choices, the experimenter must reduce the size of the stakes and the length of the delay, and it is difficult to investigate actual losses. Also, in a hypothetical question, one can ask the subject to assume that there is no risk associated with future payments, while in experiments using real stakes, subjects must assess the experimenter's credibility.

[22] Frederick, Loewenstein and O'Donoghue [2002] review the literature.

The flexibility made possible by using hypothetical payments prompt Loewenstein and Thaler [1989] to utilize a hypothetical scenario in their study. However, they acknowledge that real choices made in the presence of economic incentives are more credible. It is only when the research question necessitates the use of unaffordably large prizes, or runs counter to ethical constraints which prohibit the imposition of losses, that hypothetical responses are reasonable substitutes for real ones. Because sufficient evidence[23] exists that hypothetical responses *can* be misleading in valuation and choice settings of interest to economists, we choose not to rely on inferences made from hypothetical choices.

B. Eliciting Truthful Responses

An important aspect of salience is the understanding that subjects have for the way in which payoffs are affected by their decisions. Furthermore, for payoffs to be salient, subjects must feel confident that the earned payoffs will actually be received. A second important design consideration that distinguishes the experiments reported here from many others is the effort to ensure credibility and to provide a simple, incentive compatible payment mechanism.

Problems with credibility and subject understanding can greatly limit the inferences drawn from any experiment. For example, Horowitz [1991; p.320] honestly reports problems with experimental control that should cause some pause about accepting his results:

> If anything, the results reported in this paper understate how unusual the behavior was in our experiment. [...] First, the winners of Auctions 1 and 2 were extremely reluctant to pay for their bonds. Two out of six of the bond purchasers refused to pay for the bonds they were supposed to buy; the four other winners all had elaborate excuses for their inability to pay for the bonds immediately and took as long as a week to pay eventually. This reaction was *especially* pronounced in Auction 2.
>
> Second, a dramatic change in bidding behavior was observed between Auction 1 and Auction 2. Over 55 percent (40 out of 70) of the participants were willing to pay less for the second bond than they were for the first one, despite the fact that the payoff date was 30 days closer. The four winners of the first auction bid especially low in the second. Two of them submitted bids of $0 and one submitted a bid of $5.00; the

[23] See the review in Harrison and Rutström [2005] of valuation tasks designed to test for hypothetical bias.

> fourth winner dropped the class between the two auction dates. In scrutinizing the reliability of our data, one might argue that a single auction (Auction 1) is insufficient to familiarize individuals with the auction mechanism or the tradeoffs involved with our bonds. But the striking change in behavior between Auctions 1 and 2 suggests that individuals learned a lot from Auction 1.

The credibility of the payoffs is an issue that was also recognized by Neill, Cummings, Ganderton, Harrison and McGucken [1994] in experiments designed to elicit real payments for a physical good (an art object). They addressed it by modifying the experimental procedures for their Vickrey auction by requiring subjects to include the cash or a check for their bid in an envelope, to make their bid effective. The losers would have their envelopes returned unopened, and the winner would simply be refunded the difference between their bid and the second-highest bid. Quite apart from ensuring that the experimenter received some real payment from the winner for the object sold, this procedural device undoubtedly clarified for all subjects that there were real consequences of their bid. In this manner, control over incentives was regained.

Kirby and Maraković [1995] used real rewards for their subjects in discount rate experiments in which the subjects were asked to state an amount of money they would accept immediately in exchange for a deferred payment. The auction institution used was a first-price sealed-offer auction, in which the subjects were told that the winner would be the person that offered the smallest amount and that the winner would receive that amount instead of the deferred payment. The first-price auction is not one in which subjects have a rational incentive to reveal their true value, although it could be inferred by the experimenter with some strong auxiliary assumptions.[24] Another problem was that the actual procedures for payments were quite different than the procedures explained to subjects: "In fact, the auction was entirely simulated and subjects' bids were never really compared. The computer randomly determined whether the subject "won" the bid with a

[24] Specifically, the experimenter would have to assume that the subject followed a symmetric Nash Equilibrium offer function, relating his offer to the number of active subjects in the auction and the range of possible valuations that individuals could have drawn. In addition, some assumptions about risk attitudes would need to be made. In the simplest case of a first-price sealed-*bid* auction with risk neutral bidders and valuations drawn from the unit interval, the Nash Equilibrium bidding rule is for the subject to bid a fraction of his valuation, equal to the number of bidders minus one, divided by the number of bidders. Hence one could just invert this fraction, multiply it by the observed bid, and infer the true valuation. Quite apart from the plausibility of these auxiliary assumptions being true for these subjects, the first-price auction is known to provide weak incentives to the subject to reveal the optimal bid accurately.

probability inversely related to the size of the bid." (p. 24). One might say that this is immaterial since the subjects did not know this, but such deception calls into question the credibility of anything that the experimenter has to say to the subject.

Both of these concerns were addressed in an important follow-up study by Kirby [1997]. This study does use procedures that meet the salience standards of experimental economics. The subjects were provided real rewards, and were required to come to the experiment with $20 in cash to use for bidding (p.59). He also employed an example, patterned after Neill et al. [1994], in which an auction for a used car was used to explain the notion that truthful bidding is a dominant strategy. Finally, he studied the effect of not providing feedback on winning bids during repetitive trials.[25]

Although Vickrey auctions are in general demand-revealing, it is possible that we may not elicit truthful responses with this institution if subjects do not understand the dominant strategy logic.[26] An alternative procedure, known as the Multiple Price List (MPL) auction and shown in Table 1, has been employed by Kirby and Maraković [1995], Kirby, Petry and Bickel [1999], Coller and Williams [1999] and Harrison, Lau and Williams [2002]. The idea is to offer subjects a series of choices between a short-term reward and a longer-term reward, and to vary the parameters of the reward in the series offered to any one subject. One could vary the principal amount, the time horizon between the two rewards, the FED, the rate of return implied by the choices, the way in which the alternatives are ordered when presented to subjects, or any combination of these. The logic of telling the truth is even more transparent than in the bidding setting, since the subject literally gets the binary choice he or she makes.

One difference between the implementation of the MPL procedures across these studies is the ordering of the payoff choices. Kirby and Maraković [1995] and Kirby, Petry and Bickel [1999] presented the options to each subject in a random order, whereas Coller and Williams [1999] and

---

[25] His experiment 3 provided this information to 2 of the 4 bidders in each trial; neither of his previous experiments provided any feedback. The concern with providing feedback is that it *may* lead to affiliation of beliefs about the valuation of the object for sale, as explained by Harrison, Harstad and Rutström [2004].

[26] See Rutström [1998] and Kagel, Harstad and Levin [1987] for discussions of this problem with Vickrey auctions.

Harrison, Lau and Williams [2002] presented them in increasing order of the implied rate of return. The latter ordering was employed to make the trade-off to subjects transparent, and hence to focus on the trade-off rather than on the computational difficulty of comparing the alternatives. Moreover, Coller and Williams [1999] and Harrison, Lau and Williams [2002] presented the annual rate of return and the annual effective rate of return for each option, whereas the other studies did not. In addition to achieving consistency with the informational requirements of "fair lending laws" in most developed countries, these procedures likely minimize the informational burden of the task and focus the subject on the trade-offs entailed by the alternatives.

Another difference across studies is whether the principal amount was varied. Coller and Williams [1999] and Harrison, Lau and Williams [2002] maintained the same principal amount for all of the short-term options, so that the only thing that varied for each subject was the rate of return in the longer-term reward. Again, the objective was to make the trade-off transparent, rather than additionally test if the subject was able to make these calculations. To illustrate the potential importance of this simple design feature, Table 3 lists the options offered in Kirby and Maraković [1995] and Kirby, Petry and Bickel [1999] in order of the implied annual effective rate of return on each choice: the actual order presented to the subjects is listed in the first column. Table 3 also shows the percentage of subjects that chose the deferred payment option.[27] It is obvious from Table 3 that the change in the percentage of subjects choosing the deferred option is not monotonic, a fact which may reflect the additional computational burden of the task.

Yet another difference is whether individual subjects faced differing choice horizons. Coller and Williams [1999] and Harrison, Lau and Williams [2002] provided the subjects with only one time horizon in each payoff table, while Kirby and Maraković [1995] and Kirby, Petry and Bickel [1999]

[27] Kirby and Maraković [1996] deleted some subjects due to inconsistent responses, and we summarize their final sample of 621. Kirby, Petry and Bickel [1999] used heroin addicts as well as 60 control subjects that were not heroin addicts, and we report results for the latter sample only. The financial incentives for the Kirby and Maraković [1996] study were also relatively weak: 2,000 students were asked to return the questionnaire, knowing that one of those that returned the survey would be eligible to receive *up to* $85 based on their responses. If the students correctly guessed the potential sample size of 2000, and assumed optimistically that they were awarded the $85 prize, the expected payoff from returning the questionnaire was only 4.25 cents. If they rationally estimated that only 672 would return the questionnaire, and continued to be optimistic with respect to the $85 prize, the expected payoff changes to 13 cents.

had many different time horizons in the same payoff table. Although Harrison, Lau and Williams [2002] vary the time horizon across subjects and ask some of their subjects to fill in responses for all four horizons, within each MPL payoff table the subject had only one time horizon to consider.

There may be legitimate reasons for using these "jumbled up" tasks, such as when there are strong reasons to suspect framing effects,[28] but such tests make any inferences about discount rates *joint hypothesis tests*, that the subject is able to make these comparisons despite their computational difficulty, and that the subject chooses the preferred alternative in accord with his true discount rate.[29] In the experiments reported here, we choose to simplify the computational aspect of the choice we ask subjects to make. We employ the MPL mechanism using a constant principal amount with choices ordered in terms of increasing rates of return. We provide subjects with both the annual interest rate and the annual effective rate implied by each choice. We also implement the different time horizons on a between subject basis.

[28] One of the intentions behind such "jumbled up" payoff matrices is to avoid framing effects which may occur when subjects evaluate options in an ordered sequence. These framing effects emerge when subjects anchor on the first option and evaluate all subsequent offers in relation to the first. Since the MPL presents all options simultaneously to subjects, in one payoff matrix, we do not believe that the choice between jumbling or ordering will significantly affect framing in our setting. One concern with the MPL procedures used in Coller and Williams [1999] and Harrison, Lau and Williams [2002] is that the use of a fixed array of alternatives might "cue" subjects to switch their payment preference somewhere in the middle of the table. This hypothesis would be easy to test by varying the number of rows or changing the numerical values in the same number of rows. However, our design is focused on a comparison across FED treatments that all use the MPL method, so any effect from that method would be applicable to all of our FED treatments.

[29] Such joint tests are quite interesting, and of some importance for policy settings in which there are no regulatory constraints on the information that must be provided in financial transactions, or where decision-makers might not understand the import of the information provided, but they are tests that are best undertaken after we have resolved how to elicit discount rates in the most transparent setting possible. We would add to this list tests of the effects of elicitation mode: for reasons described in Harrison, Harstad and Rutström [2004], we are not surprised that there is considerable variation in elicited responses to discount rate questions in the psychology literature when the choices are framed differently.

Table 3: Aggregate Results from Comparable Psychology Experiments

Question Order	Amount Today	Amount in the Future	Difference in Amounts	Delay in Days	Annual Effective Rate	Percent of Subjects Choosing Future Payoff
A. KIRBY AND MARAKOVIĆ [1996]						
15	$53	$55	$2	55	27%	12
4	$34	$35	$1	43	27%	12
7	$83	$85	$2	35	28%	12
12	$65	$75	$10	50	180%	44
20	$27	$30	$3	35	196%	17
9	$48	$55	$7	45	197%	34
8	$21	$30	$9	75	454%	36
16	$47	$60	$13	50	480%	57
18	$50	$80	$30	70	1021%	74
3	$67	$85	$18	35	1056%	70
10	$40	$65	$25	70	1114%	67
14	$30	$35	$5	20	1503%	44
19	$45	$70	$25	35	9312%	90
2	$40	$55	$15	25	9708%	71
11	$25	$35	$10	25	12613%	68
21	$16	$30	$14	35	64171%	86
6	$32	$55	$23	20	1713182%	94
17	$40	$70	$30	20	2369554%	97
1	$30	$85	$55	14	42708226751144%	99
13	$24	$55	$31	10	923532365791074%	99
5	$15	$35	$20	10	1766705945627180%	99
B. KIRBY, PETRY AND BICKEL [1999]						
13	$34	$35	$1	186	6%	0
9	$78	$80	$2	162	6%	2
1	$54	$55	$1	117	6%	2
20	$28	$30	$2	179	15%	0
17	$80	$85	$5	157	15%	0
6	$47	$50	$3	160	15%	2
26	$22	$25	$3	136	40%	0
12	$67	$75	$8	119	41%	3
24	$54	$60	$6	111	41%	5
16	$49	$60	$11	89	127%	10
22	$25	$30	$5	80	127%	3
15	$69	$85	$16	91	128%	20
2	$55	$75	$20	61	524%	33
10	$40	$55	$15	62	535%	20
3	$19	$25	$6	53	545%	22
21	$34	$50	$16	30	10130%	52
18	$24	$35	$11	29	10716%	30
25	$54	$80	$26	30	11078%	67
14	$27	$50	$23	21	3868337%	85
23	$41	$75	$34	20	5259992%	90
5	$14	$25	$11	19	5904279%	72
8	$25	$60	$35	14	598219411260%	93
19	$33	$80	$47	14	774637209375%	100
7	$15	$35	$20	13	1549267326926%	90
11	$11	$30	$19	7	256388099524044320000000%	100
4	$31	$85	$54	7	337798439950498720000000%	100
27	$20	$55	$35	7	392873581121155360000000%	98

Appendix B: What Did the Subjects Say They Were Focusing On?

In an effort to get some insight from the subjects as to what they were focusing on when making their choices, interest rates or money amounts, we asked them some questions directly on this after their choices were made and collected by the experimenter. One question was: "Which one of the variables below (interest rate or dollar amount of option B) did you focus on more when making your decision in today's experiment?" The subjects were asked to check one of "Interest rate," the "Dollar amount of option B," or "Equal focus on both." This question was not obviously salient, and could have been interpreted in an ambiguous manner, but the responses are nonetheless worth looking at for exploratory purposes.

Figure 4 reports the results, broken down by FED. The results point to the subjects focusing much more on the dollar amount when there was no FED, compared to both the 7 day FED and 30 day FED experiments.[30] This finding is even stronger for the shorter horizons: for horizons of 15 days or less, *no* subject in the experiments with no FED claimed to focus solely on the rate, and 73% claimed to focus solely on the dollar amount. These findings *suggest* that one reason that the use of a FED leads to less extreme responses is that it serves to encourage the subject to focus more on the implied interest rate. This argument is consistent with subjects attaching a fixed premium to the delayed payment when they are faced with a choice between immediate payment and delayed payment (i.e., when there is no FED).

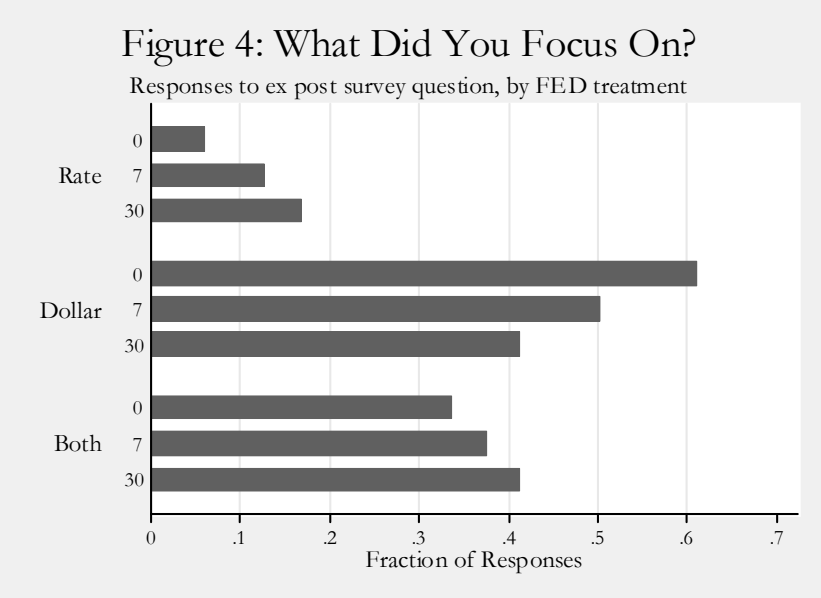

[30] This conclusion is supported by the results of Fisher's exact test applied to a contingency table showing the fraction of responses in the 0 FED case with those in the positive FED cases. The test statistic rejects the null of no association at the 1.9% level.

References

Ainslie, George, and Haendel, Vardm, "The Motives of Will," in E. Gottheil, K. Druley, T. Skolda and H. Waxman (eds.), *Etiologic Aspects of Alcohol and Drug Abuse* (Springfield, IL: Charles C. Thomas, 1983).

Angeletos, George-Marios; Laibson, David; Repetto, Andrea; Tobacman, Jeremy, and Weinberg, Stephen, "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 15(3), Summer 2001, 47-68.

Benzion, Uri; Rapoport, Amnon, and Yagil, Joseph, "Discount Rates Inferred from Decisions: An Experimental Study," *Management Science*, 35, March 1989, 270-284.

Coller, Maribeth, and Williams, Melonie B., "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 1999, 107-127.

Frederick, Shane; Loewenstein, George; and O'Donoghue, Ted, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40(2), June 2002, 351-401.

Harrison, Glenn W., Harstad, Ronald M., and Rutström, E. Elisabet, "Experimental Methods and Elicitation of Values," *Experimental Economics*, 7, June 2004, 123-140.

Harrison, Glenn W.; Lau, Morten Igel, and Williams, Melonie B., "Estimating Individual Discount Rates for Denmark: A Field Experiment," *American Economic Review*, 92(5), December 2002, 1606-1617.

Harrison, Glenn W., and Rutström, E. Elisabet, "Experimental Evidence on the Existence of Hypothetical Bias in Value Elicitation Methods," in C.R. Plott and V.L. Smith (eds.), *Handbook of Experimental Economics Results*, North-Holland: Amsterdam, 2005.

Holcomb, James H., and Nelson, Paul S., "Another Experimental Look at Individual Time Preference," *Rationality and Society*, 4(2), April 1992, 199-220.

Kagel, John H., Harstad, Ronald M., and Levin, Dan, "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," *Econometrica*, 55, 1987, 1275-1304.

Horowitz, John K, "Discounting Money Payoffs: An Experimental Analysis," *Handbook of Behavioral Economics* (Greenwich, CT: JAI Press, Inc., v. 2B, 1991, 309-324).

Keller, L. Robin, and Strazzera, Elisabetta, "Examining Predictive Accuracy Among Discounting Models," *Journal of Risk and Uncertainty*, 24(2), March 2002, 143-160.

Kirby, Kris N., "Bidding on the future: Evidence against normative discounting of delayed rewards." *Journal of Experimental Psychology: General*, 1997, 126, 54-70.

Kirby, Kris N. and Maraković, Nino N., "Modeling myopic decisions: Evidence for hyperbolic delay-discounting with subjects and amounts," *Organizational Behavior & Human Decision Processes*, 1995, 64, 22-30.

Kirby, Kris N. and Maraković, Nino N., "Delay-discounting probabilistic rewards: Rates decrease as amounts increase," *Psychonomic Bulletin & Review*, 1996, 3(1), 100-104.

Kirby, Kris N.; Petry, Nancy M., and Bickel, Warren K., "Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls," *Journal of Experimental Psychology: General*, 1999, 128(1), 78-87.

Kirby, Kris N., and Santiesteban, Mariana, "Concave utility, transaction costs, and risk in measuring discounting of delayed rewards," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 2003, 29(1), 66-79.

Laibson, David I., "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 62(2), 1997, 443-478.

Loewenstein, George, and Thaler, Richard H., "Anomalies: Intertemporal Choice," *Journal of Economic Perspectives*, 3(4), Fall 1989, 181-193.

Neill, Helen R.; Cummings, Ronald G.; Ganderton, Philip T.; Harrison, Glenn W., and McGuckin, Thomas, "Hypothetical Surveys and Real Economic Commitments," *Land Economics*, 70(2), May 1994, 145-154.

O'Donoghue, Ted, and Rabin, Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 89(1), 1999, 103-124.

Phelps, Edmund S., and Pollak, Robert A., "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, 35, 1968, 185-199.

Read, Daniel, "Is Time-Discounting Hyperbolic or Subadditive?" *Journal of Risk and Uncertainty*, 23(1), July 2001, 5-32.

Roberts, Russell D., "Myopic Discounting: Empirical Evidence – Comment," *Handbook of Behavioral Economics* (Greenwich, CT: JAI Press, Inc., v. 2B, 1991, 342-345).

Rutström, E. Elisabet, "Home-grown Values and Incentive Compatible Auction Design," *International Journal of Game Theory,* 27, 1998, 427-441.

Smith, Vernon L., "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72, December 1982, 923-955.

Winston, Gordon C., and Woodbury, Richard G., "Myopic Discounting: Empirical Evidence," *Handbook of Behavioral Economics* (Greenwich, CT: JAI Press, Inc., v. 2B, 1991, 325-342).