2009s-37

Willingness to Pay to Reduce Future Risk

*Jim Engle-Warnick, Julie Héroux,
Claude Montmarquette*

Série Scientifique
Scientific Series

Montréal
Août 2009

CIRANO
Allier savoir et décision

Centre interuniversitaire de recherche en analyse des organisations

Willingness to Pay to Reduce Future Risk[*]

Jim Engle-Warnick[†], Julie Héroux[‡],Claude Montmarquette[§]

Résumé / *Abstract*

Nous mesurons la volonté des participants de payer pour réduire les risques futurs. Au cours de nos séances expérimentales, les participants reçoivent une dotation en espèces et une loterie risquée. Ils signalent leur volonté de payer pour échanger la loterie risquée pour une loterie moins risquée. Les participants jouent à la loterie soit immédiatement, ou huit semaines plus tard, ou vingt-cinq semaines plus tard. Ainsi, dans ces expériences, la loterie et le futur forment deux sources d'incertitude. Lors de deux traitements additionnels, nous contrôlons l'aspect incertain de l'avenir avec une probabilité de continuation, constante et indépendante à travers les périodes, qui simule les chances de ne pas revenir jouer à la loterie après huit et vingt-cinq périodes. Nous avons trouvé des preuves d'un biais pour le présent à la fois dans les séances avec un délai temporel, que dans les séances avec une probabilité de continuation, ce qui suggère que cette tendance persiste avec vigueur dans les environnements comprenant de l'incertitude provenant à la fois du risque et du futur. Ceci suggère que cette règle d'arrêt peut constituer un outil efficace pour étudier ce domaine sans la nécessité de retarder les paiements dans le futur.

Mots clés : escompte hyperbolique, incertitude, risque, expériences

We elicit subjects' willingness to pay to reduce future risk. In our experiments, subjects are given a cash endowment and a risky lottery. They report their willingness to pay to exchange the risky lottery for a safe one. Subjects play the lottery either immediately, eight weeks later, or twenty-five weeks later. Thus, both the lottery and the future are sources of uncertainty in our experiments. In two additional treatments, we control for future uncertainty with a continuation probability, constant and independent across periods, that simulates the chances of not returning to play the lottery after eight and twenty-five periods. We find evidence for present bias in both the time-delay sessions and the continuation probability sessions, suggesting that this bias robustly persists in environments including both risk and future uncertainty, and suggesting that the stopping rule may be a tool to continue study in this area without the need to delay payments into the future.

Keywords: *Hyperbolic discounting, uncertainty, risk, experiments*

Codes JEL : *C91, D81*

[†] McGill University and CIRANO. Email: Jim.Engle-Warnick@cirano.qc.ca.
[‡] CIRANO.
[§] Université de Montréal and CIRANO.

Introduction

When people make decisions about the future, there are often two types of uncertainty involved: the uncertainty associated with possible outcomes, and the uncertainty that is associated with the future itself. For example, investing in human capital, in precautionary health, for retirement, or for a cleaner environment, by involving a cost in the present to reduce risk in the future, all combine both types of uncertainty. Failing to make such an investment could be the result of not only standard risk and time preferences, but also of uncertainty regarding the ability to be present in the future to receive an outcome, or uncertainty about whether the payment of the outcome will be realized. The decision to invest or not to invest could also be influenced by present bias or self-control issues.

The combination of standard risk and uncertainty about the future is of significant importance in the field, but difficult to study there. In the laboratory, where both types of risk can be better controlled, researchers have begun to make an effort to study both together. These studies tend to focus on whether or not adding uncertainty to future outcomes reduces immediacy, or present bias.

One well-known and intuitive way to model present bias is with quasi-hyperbolic discounting (Laibson (1997) and Phelps and Pollack (1968)). This method of discounting includes the standard geometric discounting and adds a discount factor that uniformly discounts all future periods from the present. Evidence for the presence of this type of discounting comes from preference reversals, whereby subjects prefer, say, $30 today to $31 tomorrow, but also prefer $31 in 31 days to $30 in 30 days.

There is a lot of this type of evidence, for which Frederick, Loewenstein and O`Donoghue (2002) provide a review, as well as an evolution of the notion of time-varying discount factors (note also Prelec (2004)). More recent evidence involves a wealth of creative investigations that

document the behavior, design mechanisms to account for it, and provide biological basis for the two-part discounting function.

For example, Della Vigna and Malmandier (2005) document the purchase of $75 gym memberships, resulting in a cost of $19/visit, when the cost per visit was $10. Such behavior can be explained with time-varying discount factors. Benartzi and Thaler (2004) present a mechanism with a commitment device for retirement savings, and show that people sign up for the plan, and that their saving increases, proving that institutions can be designed to take account of the behavior. McClure, Laibson, Loewenstein, and Cohen (2004) provide neurological evidence that time discounting results from the combined influence of the mesolimbic dopamine system and the fronto-parietal system. The former, which is emotional, is impatient, while the latter, which is analytic, is patient. Thus, they provide biological foundations for the two-parameter discount function.

Experiments are also used to discriminate between competing functional forms for discounting. Benhabib, Bisin and Schotter (2007) use a nested model and select, using experimental data where discount rates decline with both time and delay, a model that contains a fixed cost for present bias associated with a magnitude effect. They reject standard exponential discounting, but find little support for the quasi-hyperbolic discounting functional form. For his part, Rubinstein (2003) rejects both the exponential and hyperbolic discounting in favor of a procedural model that avoids the maximisation of a utility function, and thus is more in line with the psychological decision-making process.

Which brings us back to combining uncertainty about outcomes with uncertainty about the future, where the emphasis in the laboratory has been on making future outcomes uncertain to test whether there is an effect on present bias. Keran and Roelofsma (1995) exploit the similarity in explanation of preference reversals with future payments and those with the well-known common

ratio effect, and find that making the future outcome uncertain brings choices more in line with standard discounting. Weber and Chapman (2005) in a replication exercise concluded that these results depended on the presentation of the choice problem. Halevy (2008) discusses a theoretical model that models future risk, i.e., the risk of mortality or disappearance, with a stopping rule characterized by a continuation probability drawn from a constant and independent distribution after each period. If the continuation probability enters with a convex function, the type of impatient decisions associated with hyperbolic discounting can occur.

Our goal in this paper is to extend our knowledge with regard to behavior that involves the two types of uncertainty, using a decision problem that is a metaphor for many types of different field decisions. Our contribution is to specifically control for the uncertainty about the future so that we can begin to disentangle the effects of the two types of uncertainty. We do this by eliciting from the participants their willingness to pay to play a less risky lottery in the future than the one that is initially chosen for them. We delay payments to the subjects and we simulate the delay of payments in the laboratory.

In our experiments, subjects are given a risky lottery and $30. They are asked how much money they would pay to trade the risky lottery for a safe lottery. In three separate treatments subjects played their lottery immediately, eight weeks, and twenty-five weeks in the future. In two additional treatments we controlled specifically for future uncertainty by simulating a stopping rule for eight and twenty-five periods: we drew an integer from one to one-hundred from the uniform distribution, and if ever we drew the number one-hundred, the draws ended and the subject did not play the lottery; otherwise, the subject played the lottery at the end of the session.

Our results with time-delayed payments support hyperbolic discounting with, however, a dominant present bias: subjects discount the future at a rate that is almost the same whether their

future payoff occurs in eight weeks or in twenty-five weeks. We also find that the participants are moderately risk-averse, which is consistent with existing findings without the time delay.

Our results with the stopping rule are surprising. These results suggest that when the future risk is controlled within the laboratory, subjects continue to make decisions consistent with present bias, while not changing their risk preferences. In other words, their behavior is qualitatively very similar whether the payment is delayed or whether it is determined immediately with the stopping rule. This suggests a role for the stopping rule in explaining behaviour in this environment. It also suggests the stopping rule as a tool to continue study in this area without the need to delay payments into the future.

We close this section by noting that estimation of time discounting parameters has typically not taken into account risk preferences, i.e., that the vast majority of the work on time discounting has estimated time discount rates without taking into account subject risk preferences. This upward-biases the estimates of the discount factors. The notable exception is Andersen, Harrison, Lau and Rutstrom (2008), who jointly estimate risk and time preferences using behavioural instruments to directly measure both preferences, and a structural, maximum likelihood procedure to perform the joint estimation. They find that discount rates are lower than typically found in previous studies. While our methodology is different, namely, we do not directly estimate each preference separately, the point is that time and risk preferences must not be handled separately when analyzing decision making over time.

The Models

An agent must determine her willingness to pay to reduce future risk. Let s represent initial wealth, g_s a safe gamble with outcome x_l that occurs with probability p and x_h that occurs with probability $1\text{-}p$. Let g_r be a risky gamble with outcome y_l that occurs with probability q and y_h that

occurs with probability *1-q*. Let *D(t)* be a general discount factor for time *t*, let β be a standard (discrete) time preference parameter, and δ a continuation probability. Let *T* represent a terminal period, and *g( )* be a convex function. The agent's problem is to state her willingness to pay to exchange g_r for g_s where both gambles will pay in the future for two treatments, but all else pays in the present.

Let *W* represent the willingness to pay, and begin with the setup where all earnings are realized in the present. The expected utility maximizer will report *W* such that the following equality holds:

Now, let the outcomes of the gambles be realized in the future, but the payment to reduce the risk be paid in the present. The standard approach in discrete time is:

.

Letting $D(T) = \beta^T$ gives the standard geometric time-invariant discounting. There are at least two ways to admit hyperbolic-like behaviour. One method is to apply the following modification to standard discounting:

$$D(t) = \left\{ \begin{array}{ll} 1 & \text{if } t = 0 \\ \alpha\beta^t & \text{otherwise} \end{array} \right\}.$$

This specification allows time discounting between *t* = 0 and all other times in the future to be different than time discounting between any two times in the future. This quasi-hyperbolic function is tractable and captures the qualitative implication of the hyperbolic discounting.

A second way to model observed behavior is to assume that there are two components to time discounting, as in Halevy (2008). The first is the standard time-invariant discount factor. The second is a stopping rule, which takes into account the belief that the game proceeds each period

5

into the future with a probability less than one. This models, for example, the belief that there is a chance that the subject will not be around to collect the gamble's outcome in the future, or that the person who is supposed to pay will not. A convex function $g( )$ of this stopping rule gives the desired form, again separating the present from the future payoffs:

$$D(t) = g(\delta^t)\beta^t$$

While the utility function with parameter r measures attitude toward risk, the function $g( )$ can simply be of the CRRA type, as we assume for risk preferences. The CRRA estimated parameter, $v,$ will have to be less than zero to obtain a convex $g( )$ function. Negative values would reflect diminishing impatience.

Notice that for the $T = 0$ case to be consistent with the future payoff cases, we need to assume a separability condition between the lottery outcomes and the initial endowment s to get the same functional form as in the future payments. This hypothesis, we feel, is justifiable in the context of prospect theory, which assumes that subjects care about changes in wealth rather than levels (i.e., narrow bracketing). In other words, in our subsequent analysis, we will assume that subjects consider the effect of the endowment and the effect of the lotteries separately, even when they are paid both in the same session. Regardless, our conclusion of present bias is not affected by this assumption.

Experimental Design

In our experiment, subjects are given a cash endowment of $30 and a risky lottery. They report their willingness to pay to exchange the risky lottery for a safer one, as shown in Figure 1, which replicates the actual decision sheet used in the experiments. For example, the first row of the table represents a situation where the subject, upon sitting down to play the game, receives a 10/90 chance of receiving $80 or $2, plus $30 in cash for sure. They then report the minimum amount of

6

money that they are willing to pay to exchange the lottery for a 50/50 chance of receiving either $40 or $30. The money they pay, if they do exchange lotteries, comes from the $30 cash endowment. They are paid what is left of the endowment immediately.

Moving down the five rows, Figure 1 reveals that the probability of receiving the better outcome of the risky gamble increases by ten chances out of one-hundred in each subsequent row, up to a maximum of 50 chances out of 100. The safe and risky lotteries are based on the risk preference measuring instrument of Eckel and Grossman (2008). The subjects make these five decisions shown in Figure 1 in the experiment. The subjects are paid for a randomly chosen one of the five decisions.

To illustrate the incentives present in our experimental design, we present numerical solutions of the theoretical model in Figure 2. These predictions are for standard discounting, with a discount rate set at $\beta = 0.99$. For each of the five safe lotteries, we present predictions for a risk neutral expected utility maximizer, and for a risk averse agent with CRRA risk parameter *r = 0.5.* For each lottery choice, the first three columns indicate the willingness to pay to exchange lotteries for *T = 0, T = 8,* and *T = 25.*

Figure 2 shows that the level of willingness to pay is sensitive to risk aversion. For example, for lottery choice 1, a risk neutral expected utility maximizer would pay about $25 to exchange lotteries if they paid in the present, while a risk averse expected utility maximizer with a CRRA parameter of 0.5 (which has been reported in laboratory studies) would pay almost $27. Furthermore, these amounts are -$6 and $7 for lottery choice 5.

Predicted willingness to pay decreases with discounting through time in uniform proportions. For example, for lottery choice 1 and a risk neutral subject, the amount that makes the exchange worthwhile is $25, $23, and less than $20 if the lottery pays in the present, in eight weeks, or in

twenty-five weeks respectively. The figure shows that for reasonable parameter values, and for agents with different degrees of risk aversion, optimal decisions for an expected utility maximizer vary significantly across decision tasks. As the discount factor tends toward one, the decision making looks more and more as in the case we call $T = 0$.

The subjects reported their minimum willingness to pay to exchange the risky lottery for the safe gamble. A BDM procedure (Becker, Degroot and Marschack (1960)) then determined whether the subjects played the lottery or paid the experimenter to exchange it for the safe lottery. In this procedure, the computer draws a number in dollars and cents from $0.00 to $30.00. Every number from $0.00 to $30.00 has the same chance of being drawn. If the number drawn is less than or equal to the number reported by the participant, then the participant pays the amount drawn by the computer and switches lotteries. If the number drawn is greater than the amount reported by the subject, then the subject pays nothing and plays the risky gamble. We chose $30 for the upper bound because it is beyond anything any reasonable subject should wish to pay (note that it is also the amount of money they were given to spend on exchanging lotteries).

The BDM procedure has its costs and benefits. Costs involve complexity and theoretical justification (see Halevy (2007) for a complete discussion, and an example of the alternative procedure called "multiple price list"), while benefits involve incentivization of the decision. The procedure amounts to a second-price auction, and it is required to make the valuation for the exchange of gambles incentive compatible. The subjects were told that it was in their best interest to report their true willingness to pay to make the exchange, and they were given examples of how misreporting could lead to an outcome that they did not prefer, but could never be advantageous.

The subjects practiced the BDM procedure with a single lottery that was not part of the paid portion of the experiment. They entered a willingness to pay to switch fifteen times in order to

become acquainted with how the procedure works. We chose the willingness to pay design, and a range for the computer draw that exceeds what any subject should reasonably report as their true valuation, because it has been shown that this design elicits the most stable valuations among willingness to pay/willingness to accept procedures (James, (2007)).

Each participant received the show-up fee and the $30 to play the game minus the amount, if any, they paid to exchange lotteries at the end of the initial session. In the base treatment the designed lottery was played immediately by tossing a ten-sided die. For the time delay treatments the outcome and payment of the designed lottery occurred respectively eight and twenty-five weeks after the session.

We then ran two treatments in which the subjects were subjected to a stopping rule after making their choices and discovering which decision would be paid. After subjects made their decisions, and after one decision was chosen for pay, with probability 0.01, the computer stopped the session for the subject after each of eight simulated periods in one treatment, and after each of twenty-five simulated periods in a second treatment. If ever the computer stopped the session before the final period, the subject was not paid for the lottery choice.

The stopping rule simulates a continuation probability of 0.99 with the same number of periods as there were weeks in our two time-delayed treatments. In these sessions, subjects were shown thirty simulations of the stopping rule, as well as given detailed information as to the probability the session would continue after each period in the instructions. To implement the stopping rule, subjects clicked on the screen each period and were shown the random number, from $1 - 100$, drawn by the computer. If ever the number 100 game up, the draws ended, and the lottery was not played.

The sessions were run at the CIRANO experimental laboratory in Montreal. Between forty-three and forty-seven subjects participated in each of the five experimental treatments. The average payoff, which was influenced by the $30 cash endowment, was approximately $64.

Experimental Results

Delayed Payment Treatments

In Figure 3 we present the observed average willingness to pay to exchange lotteries for all lottery choices in the $T=0$ and delayed payments sessions. All values are positive suggesting that the average participant is risk averse. The average willingness to pay is clearly higher for the present than for the future, but in contrast to the numerical solutions presented in Figure 2, we observe little differences in discounting between $T = 8$ weeks and $T = 25$ weeks. In Table 1, a t-test confirms that there is a statistical difference between willingness to pay to reduce risk between the present and both future dates, but that there is no statistical difference between both future dates (this evidence is confirmed with a Mann-Whitney-Wilcoxon and a Robust Rank Order non-parametric test, which are not reported). This visual evidence is consistent with present bias.

In Table 2 we present results from regressions using nonlinear least squares on various specifications. The first column estimates the risk preference parameter for the sessions in which the lottery was paid immediately. The second column presents estimates for the hyperbolic specification and includes all three time framework sessions. The third and fourth columns estimate the stopping rule model.

The results are robust for the risk parameter around 0.45. As found elsewhere, (see Eckel and Grossman (2008) for a review), the negative coefficient on the *r x male* variable indicates that men are less risk averse than women. The present bias coefficient estimate of α falls far short of one, clearly rejecting the exponential specification. The present bias is very important to compare to the

¨hyperbolic¨ coefficient that seems to play a minor role as was shown in Figure 3. In fact, our point estimates for this model are very similar to those found in Laibson, Repetto, and Tobacman (2003), who estimate a life-cycle model with a similar specification for discounting.

We next estimate the parameter for the stopping rule function, v. To do this, we make the assumption that the risk preference is modeled by expected utility. Thus, the special case of $v=0$ corresponds to expected utility maximization with the standard geometric discount factor. Our purpose in running these regressions is simply to demonstrate that the present bias we observe in the data is consistent with a convex function.

Note that we cannot jointly identify both the geometric discount factor, and the continuation probability with its function parameter. Thus to estimate the parameter of the stopping rule function, we assume values for the geometric discount factor. If we set $\beta=1$, presented in the third column of Table 2, we estimate $v=-0.26$, which is a convex function. However, a 95% confidence interval includes $\delta=1$. With an assumption of the geometric discount factor of 0.99, the parameters of the stopping rule equation simply adjust. In this case we obtain a continuation probability not significantly different from 1, and $v=-0.26$. The impossibility to distinguish the discount rate from the stopping rule raises concern about the ability to identify this notion of discounting from empirical data. We can say, however, that the present bias we observe in our data is consistent with the convex function estimate.

We have other evidence in the data with regard to the stopping rule effect. Approximately 14% of our participants did not show up for the draw of the selected lottery, eight or twenty-five weeks later. If the stopping rule is a real behavioral issue, many participants may have chosen ex ante not to return for payments, and thus offer little or no money to trade for a less risky lottery.

We ran regressions with the independent variable a dummy taking on the value of 1 if the subject did not return to play her lottery. Explanatory variables are average WTP over all five decisions, dummies for the lottery choice selected by the computer for payment being a risky one, a dummy for whether the subject reported they understood the experiment, and a identifier for male gender. A negative coefficient on WTP would suggest that subjects underbid, anticipating a low probability of returning. Coefficients on the lotteries chosen for pay would suggest that returning is conditional on the lottery awaiting the subject upon her return to the laboratory.

Table 3 reports the results, where the omitted variable is the indicator for returning to play the riskiest lottery with the lowest expected value, in choice 1. Notice that relative to this lottery, the probability of not returning is lower for all other four lotteries, and that none of the other variables have significant coefficient estimates.[1] Note that all subjects returned to play lottery 4, which explains its exclusion from the regression. We conclude that subjects are failing to return because they are facing the worst lottery, i.e., the one with a 10% chance of winning the high amount in the risky lottery.

Stopping Rule Treatments

Figure 4 presents theoretical willingness to pay with the stopping rule model. For these computations we set the geometric discount factor to 1 (subjects are always paid at the end of the session), the continuation probability to 0.99, and the convex function parameter to -0.25, which is roughly what we estimated from our data (shown below). The figure shows the same pattern of behavior presented in Figure 2, and confirms that for this model the economic incentives were significant.

We first compare results from each stopping rule treatment. Figure 5 presents average willingness to pay from each of the two treatments, and compares them with the *T=0* treatment in

---

[1] The difference for lottery 3 is statistically insignificant.

which subjects played their lotteries for sure. Figure 5 is identical in structure to Figure 3, which provided the same information for the delayed payment sessions. The figure reveals that, conditional on initially holding a relatively risky lottery, subjects discount for uncertainty of the stopping rule. Statistical tests conducted as in Table 2 (not reported) confirm that discounting is different between $T=0$ and the two stopping rule treatments for lotteries $1 - 4$. However, as the endowed lottery becomes less risky, as in choice five, the subjects' discounting is indistinguishable between the two treatments. Indeed, with lottery choice five, willingness to pay was statistically the same whether there existed a stopping rule or not. Note, however, that the average willingness to pay is less for the stopping rule sessions than for the $T=0$ sessions.

We repeated the regressions reported in Table 2 with the data from the stopping rule treatments. For the stopping rule specification, we used the laboratory continuation probability of 0.99, and a geometric discount factor of 1, since continuation probability was essentially the only future risk present. The results are presented in Table 4. The first column of the table is identical to the first column of Table 2, and the second two columns estimate the two models respectively.

The hyperbolic model again reveals a present bias, slightly greater than the time delay treatments, but statistically different from one. In the stopping rule model, again the estimate for the stopping rule function parameter implies a concave function. In both specifications, qualitative results are similar as they were for the time delay sessions. Specifically, estimates for the risk parameter are almost the same, and the men are less risk averse, though again not statistically significantly so.

We ran a set of nonlinear regressions on willingness to pay per lottery in which we used the point estimates for the risk preference parameter in the $T = 0$ sessions to estimate the implied continuation probability in all other treatments. Conditional on the assumption that the risk

preference is stable across treatments, we can test how close the implied continuation probability from the data matches the actual probability. Such a test can only be performed in the laboratory where the continuation probability is controlled.

The results are shown in Table 5. The table displays two interesting results. First, the average estimate of the continuation probabilities for the eight-period stopping rule and the twenty-five-period stopping rule are different: 0.962 and 0.982 respectively, suggesting more impatience in one treatment than in the other. Second, the difference between these estimates for the time-delay sessions and the stopping rule sessions are small: 0.948 vs. 0.962 for the eight-period sessions, and 0.981 vs. 0.982 for the twenty-five-period sessions.

Conclusion

We report results from a series of experiments in which subjects reported their willingness to pay to reduce future risk. Our experiment involved a decision task which is a metaphor for many important real-life decisions: it includes uncertainty about outcomes, and uncertainty about the future. We conducted sessions in which subjects waited to play their lottery in the future, and in which we simulated the risk of waiting until the future to play the lottery.

When the timing of playing the lottery was delayed eight and twenty-five weeks, the subjects made decisions consistent with risk aversion and present bias. When we controlled for the uncertainty of the future by subjecting subjects to a stopping rule, their risk aversion was qualitatively similar, present bias persisted, and the stopping rule function was concave. This suggests that the present bias finding is rather robust in this framework. It also suggests that the stopping rule methodology may be useful in some cases in place of delaying payment in similar experiments.

References

Anderson, S., G. Harrison, M. Lau, E. Rutstrom (2008). "Eliciting Time and Risk Preferences", *Econometrica*, 76(3): 583-618.

Benartzi, S. and R.*Thaler*. (2004). "Save More Tomorrow: Using Behavioral Economics to Increase Employee Savings", *Journal of Political Economy*, 112:1:S164-S187.

Benhabib J., A. Bisin and A. Schotter. (2007). ¨Present-Bias, Quasi-Hyperbolic Discounting, and Fixed Costs¨, working paper, NYU.

Della Vigna, S. and U. Malmendier. (2005), "Paying Not to Go to the Gym," *American Economic Review*, forthcoming.

Eckel C.C. and P.J. Grossman. (2008). ¨Forecasting risk attitudes: An experimental study using actual and forecast gamble choices¨, *Journal of Economic Behavior & Organisation*, 68, 1-17.

Frederick S., G. Loewenstein and T. O`Donoghue. (2002). ¨Time Discounting and Time Preference: A Critical Review¨, *Journal of Economic Literature. 40*, 351-401.

Halevy Y. (2008). ¨Stroz Meets Allais: Diminishing Impatience and the Certainty Effect¨, *American Economic Review* 98(3), 1145-62.

Halevy, Y. (2007). "Ellsberg Revisited: An Experimental Study", *Econometrica* 75:503-536.

James, D. (2007). "Stability of Risk Parameter Estimates within the Becker-DeGroot-Marschak Procedure", *Experimental Economics,* 10:123-141.

Keren, G. And P. Roelofsma (1995). "Immediacy and Certainty in Intertemporal Choice", *Organizational Behavior and Human Decision Making*, 63(3), 287-297.

Laibson, D., A. Repetto, and J. Tobacman (2003). "A Debt Puzzle", in eds. P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, *Knowledge, Information and Expectations in Modern Economics: In Honor of Edmund S. Phelps,* Princeton: Princeton University Press, 228-266.

McClure, S., D. Laibson, G. Loewensein, and J. Cohen. (2004). "Separate Neural Systems Value Immediate and Delayed Monetary Rewards", *Science*, 306, October 15.

Plelec D. (2004). ¨Decreasing Impatience: A Criterion for Non-stationary Time Preference and ¨Hyperbolic¨ Discounting¨, *Scandinavian Journal of Economics*, 106 (3), 511-532.

Rubinstein A. (2003). ¨Economics and Psychology? The Case of Hyperbolic Discounting¨, *International Economic Review*, 44(4), 1207-1216.

Weber B.J. and G.B. Chapman. (2005). ¨The Combined Effects of Risk and Time on Choice: Does Uncertainty Eliminate the Immediate Effect? Does Delay Eliminate the Certainty Effect?¨, *Organizational Behavior and Human Decision Process*, 96, 104-118.

Table 1: Comparing the Difference in Mean Responses Across Time Delay Treatments

Lottery Choice	T=8 vs. T=25	T=0 vs. T>0
1	0.194	3.383
	(0.423)	(0.001)
2	0.470	2.966
	(0.320)	(0.002)
3	0.230	3.189
	(0.409)	(0.001)
4	0.684	2.387
	(0.248)	(0.009)
5	1.042	2.355
	(0.150)	(0.010)

Notes: P-values in parentheses below
reported t-statistics.
T=8: n=63
T=25: n=45
T=0: n=43

Table 2: Nonlinear Least Squares Parameter Estimates for Delay Treatments

Model	Present Only (T=0)	Hyperbolic Model (T=0,8,25)	Stopping Rule Model (T=0,8,25)	Stopping Rule Model (T=0,8,25)
Risk r	0.397*	0.443*	0.500*	0.500*
	(0.026)	(0.031)	(0.040)	(0.040)
Present Bias α		0.674*		
		(0.049)		
Geometric β		0.996*[1]	1	0,95
		('0.004)		
Stopping δ			0.988*[2]	1.029*
			(0.003)	(0.002)
Convexity v			-0.264*	-0.264*
			(0.065)	(0.065)
$r \times$ male		-0.081*	-0.110*	-0.110*
		(0.039)	(0.041)	(0.041)
Adj. R-sq.	0.849	0.743	0.740	0.740
N	215	755	755	755

Notes: Standard errors in parentheses; individual controls not reported.
[1] 95% confidence interval contains 1 (0.986-1.005)
[2] 95% confidence interval contains 1 (0.989-1.003)
* Significant at 5%
** Significant at 10%

17

Table 3: Determinants of the Probability of Not Returning for the Lottery

Average WTP	-0.003**
	(0.034)
Lottery 2	-1.319*
	(0.648)
Lottery 3	-0.647
	(0.577)
Lottery 5	-1.776*
	(0.721)
Male	0.413
	(0.481)
Understanding the Experiment	-0.494
	(0.641)
Constant	0.380
	(0.767)
Psuedo R-sq.	0.172
N	54

Notes: Standard errors in parentheses.
* Significant at 5%
** Significant at 10%
Lottery 4 predicts failure perfectly

Table 4: Nonlinear Least Squares Parameter Estimates for Stopping Rule Treatments

Model	Present Only (T=0)	Hyperbolic Model (T=0,s8,s25)	Stopping Rule Model (T=0,s8,s25)
Risk r	0.397*	0.443*	0.500*
	(0.026)	(0.028)	(0.037)
Present Bias α		0.762*	
		(0.055)	
Geometric β		0.992*[1]	1
		(0.004)	
Stopping δ			0,99
Convexity v			-0.251*
			(0.047)
r x male		-0.041	-0.045
		(0.037)	(0.039)
Adj. R-sq.	0.849	0.793	0.795
N	215	665	665

Notes: Standard errors in parentheses; individual controls not reported.
[1]95% confidence interval contains 1 (0.986-1.005)
[2]95% confidence interval contains 1 (0.989-1.003)
* Significant at 5%
** Significant at 10%

Table 5: Estimates of Time Preference Using Risk Preference from No-Delay Treatment

		Parameter and Experimental Treatment			
	r	β (Delay)	β (Stopping)	β (Delay)	β (Stopping)
Lottery	T=0	T=8	T=8	T=25	T=25
1	0	0.940	0.956	0.979	0.981
2	0.249	0.951	0.962	0.981	0.981
3	0.393	0.953	0.973	0.983	0.985
4	0.383	0.952	0.958	0.978	0.988
5	0.438	0.963	1.010	0.973	0.997
Mean	0.397	0.948	0.962	0.981	0.982

Notes: Risk parameter estimated with T=0 data and used for all other estimations

Figure 1: The Experiment Decision Sheet

Figure 2: Predicted Willingness to Pay for Standard Discounting

Figure 3: Average Willingness to Pay in the Time Delay Treatments

|	Lottery Choice 1	Lottery Choice 2	Lottery Choice 3	Lottery Choice 4	Lottery Choice 5
t=0	24,4	21,1	17,6	11,2	5,7
t=8	18,9	16,6	12,7	8,1	3,4
t=25	18,5	15,7	12,3	7,0	2,3

Figure 4: Predicted Willingness to Pay for Stopping Rule Model

Figure 5: Average Willingness to Pay in the Stopping Rule Treatments

	Lottery Choice 1	Lottery Choice 2	Lottery Choice 3	Lottery Choice 4	Lottery Choice 5
t=0	24,4	21,1	17,6	11,2	5,7
t=stop8	21	17,9	14,7	8,5	4,9
t=stop25	19	15,5	12,7	8,7	4,2