Inconsistency Pays?: Time-inconsistent subjects and EU violators earn more

Berg, Nathan and Eckel, Catherine and Johnson, Cathleen

2010

Inconsistency Pays?: Time-Inconsistent Subjects and EU Violators Earn More

Nathan Berg
nberg@utdallas.edu

Catherine Eckel
eckelc@utdallas.edu

Cathleen Johnson
cjohnson@eller.arizona.edu

Abstract: Experimental choice data from 881 subjects based on 40 time-tradeoff items and 32 risky choice items reveal that most subjects are time-inconsistent and most violate the axioms of expected utility theory. These inconsistencies cannot be explained by well-known theories of behavioral inconsistency, such as hyperbolic discounting and cumulative prospect theory. Aggregating expected payoffs and the risk associated with each subjects' 72 choice items, the statistical links between inconsistency and total payoffs are reported. Time-inconsistent subjects and those who violate expected utility theory both earn substantially higher expected payoffs, and these positive associations survive largely undiminished when included together in total payoff regressions. Consistent subjects earn lower than average payoffs because most of them are consistently impatient or consistently risk averse. Positive payoffs from inconsistency cannot, however, be fully explained by greater risk taking. Controlling for the total risk of each subject's risk choices as well as for socio-economic differences among subjects, time-inconsistent subjects earn significantly more money, in statistical and economic terms. So do expected utility violators. Positive returns to inconsistency extend outside the domain in which inconsistencies occurs, with time-inconsistent subjects earning more on risky choice items, and expected utility violators earning more on time-tradeoff items. The results seem to call into question whether axioms of internal consistency—and violations of these axioms that behavioral economists frequently focus on—are economically relevant criteria for evaluating the quality of decision making in human populations.

Introduction

Given growing theoretical, empirical and policy literatures dealing with time-inconsistency, one could be forgiven for assuming—incorrectly—that there exists abundant evidence documenting significant economic costs that result from time-inconsistent choices. This lack of evidence is the subject of this paper, which presents experimental results showing the opposite: time-inconsistent subjects in laboratory experiments earn significantly higher total payoffs. This, of course, does not prove that time-inconsistent people are generally better off, or that myopic decision making is not a genuine problem in particular settings. Rather, this finding adds to a little known but important, and growing, collection of empirical results that pose a deep methodological challenge to what is perhaps the dominant research program in behavioral economics documenting deviations from axiomatic rationality (e.g., intransitivity, non-Nash play, violations of expected utility theory, and non-Bayesian behavior).[1]

[1] Although we are aware of no directly relevant evidence about time-inconsistency and total payoffs, there is some related evidence, none of which points conclusively in the direction of a positive correlation between consistency and payoffs. Burks, Carpenter, Gotte and Rustichini (2008) report that trainees learning to become truck drivers who have lower than average cognitive skills are more likely to lose money on their investment in the training program by leaving the trucking program before recouping their out-of-pocket costs of the training program. Similarly, Benjamin, Brown and Shapiro (2006) find a modest negative association between cognitive skills and the occurrence of preference anomalies, although even very-high-skill individuals frequently exhibit anomalies. Jacobson and Petrie (forthcoming) find zero pair-wise correlation between time-inconsistency on experimental time-preference elicitation instruments (exhibited by 55% of 181 subjects) and financial decision making in the field, and zero correlation between risk-aversion instruments and real-world financial decision making. Putting together time-inconsistency and risk-preference measures, they report an interaction effect suggesting a slight increase in time-inconsistent subjects' rate of using informal financial instruments, which these authors interpret as a mistake, although no real cost differentials are reported. Ashraf, Karlan and Yin (2006) combine experimental and field evidence suggesting that time-inconsistent bank customers are self-aware of the potential pitfalls of impulsivity and disproportionately choose to adopt new bank products that offer identical interest rates but added commitment devices which voluntarily shrink those savers' choice sets. This suggests that even significant numbers of people who are inconsistent in the lab are sophisticated in recognizing the vulnerabilities these inconsistencies may cause and preemptively deploying successful strategies aimed at accumulating greater wealth. Chu and Chu (1990) and Cherry, Crocker and Shogren (2003) report that subjects who are paid to avoid inconsistent choices quickly learn to be consistent. List and Millimet (2004) show that subjects in the field vary significantly in terms of consistency of choice patterns, and that market experience reduces the probability of inconsistent patterns of choice without showing, however, that inconsistency leads to reduced levels of economic performance.

Ask a behavioral economist what we learn from behavioral economics in applied work aimed at educating the public or designing institutions, and you will likely hear calls to help error-prone, biased, or irrational humans overcome the systematic pathologies built into their brains (e.g., Ariely's *Predictably Irrational*, 2008, or Sunstein and Thaler's *Nudge*, 2008). And yet, very little evidence exists linking violations of axiomatic rationality to high-stakes differences in real people's lives, such as earnings, physical health, lifespan, and happiness. In the normative behavioral economics literature (Camerer et al, 2003; O'Donoghue and Rabin, 2003; Berg, 2003; Loewenstein et al, 2007; Bernheim and Rangel, forthcoming) which sometimes articulates a need to educate and make policy aimed at "de-biasing" individuals (Jolls and Sunstein, 2006), there is ample grounds for concern that prescriptive advice based on behavioral economics takes as its goal to enforce internal consistency, including conformity with transitivity, the Savage axioms, Bayes Rule, and—yes—time consistency.[2]

This paper shows, however, that experimental subjects who violate time-consistency and make choices over risky gambles that cannot be rationalized with any expected-utility objective function wind up earning higher expected payoffs, after controlling for risk, demographic variables such as household income, and success in school. The results imply that decision-making processes which violate consistency may confer other, often unrecognized, benefits. The results also suggest that economists' normative analyses may be referencing the wrong

[2] Hubal et al (2007) suggest that behavioral biases in the tradition of Kahneman and Tversky can be used to understand recent intelligence failures in the lead-up to the Iraq War and propose to use the behavioral biases and anomalies literature to re-design US intelligence policy. Even the most sympathetic to the insights gleaned from psychology, however, may feel that simpler strategic explanations provide more precision and parsimony (e.g., following accounts in the popular press, which attribute false information used as a rationale for the invasion of Iraq to dishonesty in the executive branch and manipulation of classified information).

normative benchmarks if, as behavioral economists often claim, the goal is to add more empirical realism and real-world relevance to economic science.

Insofar as behavioral economics stands on the strength of empirical realism and real-world relevance, a clear methodological priority would seem to be collecting evidence that provides empirical tests of whether standard decision-theoretic axioms of self-consistency provide relevant normative benchmarks. After all, when one takes behavioral economics' stated priority of empirical realism seriously, it suggests a much needed follow-up question: If individuals do not conform to standard normative criteria found in axiomatic definitions of rationality, what then is the price?

Consider this real-world datum frequently observed along highways throughout the U.S.—billboards advertising "Vasectomy Reversal." According to the National Institutes of Health (NIH), approximately 500,000 men in the U.S. undergo a vasectomy each year, and nearly 10% of them later decide to undergo a vasectomy reversal. Vasectomy reversal costs between $5000 and $15,000, and it is an elective procedure not covered by most insurance plans, implying that the most men who choose to undergo the procedure pay full price, either in lump sum or over time with financing. From the perspective of the null hypothesis of perfect time consistency, one may ask how a sequential path of surgeries consisting of vasectomy at time t_1 and vasectomy reversal at t_2 ($t_1 < t_2$) might be rationalized as maximization of a stable preference ordering applying a single, time-consistent discount rate and exponential discounting.

Such a rationalization would seem to require two somewhat unlikely cost-benefit calculations. For a present-oriented and time-consistent man, the initial vasectomy would pass the cost-benefit calculus based mostly on anticipated short-run flows of increased sexual pleasure without worrying about birth control (which seems a likely motivation). The decision to undergo the reversal would have to also pass a cost-benefit test based primarily on the short-run benefits of fatherhood (which seems unlikely, or at least inconsistent with much that has been written on long-run motivations for procreation in terms of gene propagation).[3] On the other hand for a very patient future-oriented and time-consistent man, the initial vasectomy decision's benefits would be weighted toward flows of increased sexual pleasure arriving in the future, while realizing that this flow will stop or be reduced on the date of the vasectomy reversal. This does not strike us as a plausible account of the decision process leading to the observed sequence of choices.

There are doubtless other possible time-consistent rationalizations for choosing two surgeries, the second of which reverses the effect of the first. However, the prediction of any such theory is that, at the time of choosing the first surgery, the chooser considers it possible, or even likely, that he will elect to have the reversal at a later date. While we have no direct survey data on this point, doctors who offer vasectomies routinely warn patients that the procedure is permanent. According to the National Institute of Health (2008), reversing the vasectomy is "difficult, expensive, and often unsuccessful." Thus, a more parsimonious explanation for the choice of vasectomy reversal is simply a change of heart, that is, a change in preferences—where time

[3] Although some people may draw substantial short-run utility from parenting toddlers and teenagers, research at the intersection of economics and evolutionary biology would seem to suggest that far-off future benefit flows, perhaps extending beyond the parent's lifetime, are the most significant motivators for producing children.

preferences shift from placing more weight on the short-run to more weight on future flows of benefits.

If men choosing vasectomy reversals are time-inconsistent, then so what? Do these deviations from consistency cause systematic economic harm? We would caution anyone eager to label purchasers of vasectomy reversals as "irrational" that these men are, after all, financially well-off enough to afford the surgery. They have sufficient physical health and leisure to invest time and money in pursuing enhanced sexual pleasure (without worrying about birth control). Thus, one would be hard-pressed to label them as economic failures. Their apparent time-inconsistency allows a shift toward future-oriented priorities placing significant weight on procreation, which would likely receive positive evaluation according to many moral traditions as well as scientific norms of evolutionary fitness in evolutionary biology. We suspect, too, that many parents would strongly approve of the transformation of their adult child's preferences—from large to small discount rates.

Although further study of vasectomy patients and their time preferences in the field would be interesting to pursue, this paper turns to laboratory data to investigate an analogous tension between two distinct normative measures that arise in the example of men who choose vasectomy reversal—internal consistency on the one hand versus the level of payoffs (unconditionally and conditioning on risk-taking and demographic differences) on the other.

We present laboratory evidence using standard experimental instruments for eliciting time and risk preferences that document a surprising positive return on three distinct forms of inconsistency, the most dramatic of which is the positive return on time-inconsistency. Expected returns for 72 choices in a laboratory experiment (40 time-tradeoff and 32 risky choice items) were tabulated together with the standard deviation of all risky choices under the assumption that gambles across all risky choice items were uncorrelated. A variety of time-inconsistency measures were computed, the simplest of which is a binary indicator for any subject whose required compensation for waiting switched on pairs of time-tradeoff items with identical payoffs, identical duration between arrivals of payoffs switched, but different front-end delays. The second form of inconsistency is a kind frequently documented in experimental studies with multiple risky choice items, where subjects' choices cannot be rationalized as maximization of any expected utility function (e.g., a strictly risk-loving choice in one pair of choice items, and strictly risk-averse on others). The third form of inconsistency was measured by counting the number of non-matching responses for a subset of the sample who were invited back to repeat the same exact series of 72 choice items six months later, which we refer to as between-session inconsistency. Unconditionally and conditionally (controlling for risk-taking and a long vector of demographic information), maximally consistent subjects, on average, earned less money.

Positive returns on inconsistency appear to extend beyond the choice domain in which they are measured. For example, time-inconsistent subjects not only earn more on time-tradeoff choice items, they also earn more on risky choice problems with no time component. Similarly, subjects with inconsistent responses to risk-preference measures (i.e., those whose choices violate consistency axioms needed for an expected utility representation of risk preferences) earn more on time-trade-off choice items with no risky-choice component. Overall, subjects who violate time-consistency, violate expected utility theory, or make inconsistent choices on identical choice problems when repeating them on different days earn significantly higher payoffs than subjects who are perfectly consistent.

Calls to use behavioral economics as a prescriptive basis for institutional design, such as O'Donoghue and Rabin's (2003a, 2003b) suggestion to tax potato chips and subsidize carrots, or Thaler and colleague's (Thaler and Behartzi , 2004; Thaler and Sunstein, 2008) focus on changing defaults in savings plans, organ donation rules, and the positioning of dessert on the buffet line, naturally raise controversy. What seems clear, however, is the need for further investigation into normative behavioral economics in two directions—not only documenting real-world deviations from received normative benchmarks, but also investigating whether the normative measures we use are relevant to the economic problems we face.

If individuals fail the normative axiom of time-consistency while succeeding by other normative measures in a competitive American environment—at least surviving and, perhaps, thriving—then how is tension between decision-theoretic norms of internal consistency and other compelling economic norms—like how much money one has—to be resolved? (For example, men who can afford a $5,000 to $15,000 elective surgery may be internally inconsistent but economically well-off.) The dominant view in behavioral economics is that the norms of internal consistency are right and that human behavior is mistaken, even if those behaviors are manifestly well-performing by other metrics. An alternative view is that empirical normative science should seek to describe the characteristics of good decision making and discard theoretical constructs that do not describe what successful decision makers actually do. Why should we rely on commonly repeated folk reasoning in economics (with no empirical support) that there exist competitive forces in market economies that punish people whose choices violate internal consistency axioms (e.g., the money pump thought experiment as "proof" that intransitive preferences cannot survive, even though intransitivities are probably widespread in human populations)?

The biological literature provides positive theory as to why inconsistency can confer adaptive benefits. Buchstaber and Langsam (1985) argue that if nature tuned organisms optimally to any particular environment, it would lead to severe disadvantage in the face of changing environments, whereas less stringent satisficing rules promote fitness by allowing for much more flexible and faster adaptation to changing ecologies as they are buffeted by occasional random shocks (also see Schmidtz, 1995). There would seem to be obvious extensions of this intuition to human populations in contemporary environments, where being imaginative, entrepreneurial, and creative—or just flexible enough to survive life—sometimes requires one to inconsistently experiment, to change one's mind, or to use context-specific action rules rather than complete and self-consistent orderings of the elements in one's action space. It should not stretch one's imagination too far to envision humans who enjoy adaptive benefits by sometimes changing the weights they place on different factors out-competing those who dogmatically conform to the stricture of maximizing a time-separable utility function with exponential weighting.

While there is clear tension between economics' classical and neoclassical commitment to the Libertarian principle of consumer sovereignty and new rationales for interventions based on behavioral economics' re-discovery of human irrationality, there is another, more fundamental cause for concern about these paternalistic initiatives—even for writers like ourselves, who believe that neoclassical models of externalities and information asymmetries already provide very convincing rationales for government intervention to help market economies flourish.

The problem is this: the vast majority of studies claiming to have discovered irrationality take as their benchmark axioms of internal consistency. Psychologist Ken Hammond (1996) lays out an important distinction that is at the heart of this paper, contrasting two systems of evaluative criteria aimed at saying what it means to make good decisions and how to measure whose decision procedures are performing better. As everyone knows, the standard economic theory of time-consistent inter-temporal choice does not say how patient one should be, even though it is obvious that greater patience leads to larger accumulations of wealth, all else equal. The inter-temporal choice theory merely says that if I rank x dollars at time t_1 over y dollars at time t_2, with $x < y$ and $t_1 < t_2$, then I should always rank x dollars at time t_3 over y dollars at time t_4 as long as the wait between the arrivals of those payoffs, which we refer to as the *appreciation interval*, is at least as long: $t_3 - t_4 > t_1 - t_2$. According to this coherence norm, a person who spends his entire paycheck for a party on the day his paycheck arrives is rational (i.e., time consistent) as long as he remains equally impatient for the rest of his life–throwing parties that exhaust his entire paycheck each time a paycheck arrives. If, on the other hand, this highly impatient and perhaps impulsive person moderates his impatience and begins to save money, then the axiom of time-consistency deems his behavior as irrational.

The intuitive lack of appeal in this kind of normative analysis seems obvious. Yet much of our discipline proceeds as if on the basis of a firmly established law of social science that people should maximize an additively separable objective function with exponential discounting—or otherwise go to hell after living an unhappy and evolutionarily disadvantaged life. The tension here lies in the difference between coherence and correspondence—that is, between standard economic norms of internal consistency and alternative evaluative criteria, such as how much money is in one's bank account, physical health, happiness, or the accuracy of one's beliefs. In contrast to coherence norms based on internal consistency, correspondence norms evaluate decisions and inferences according to how well they are calibrated to the environments in which they are used (also referred to as ecological rationality by Gigerenzer et. al., 1999, and Smith, 2003)—normative evaluations that measure consequences of decision procedure in terms of how well-endowed decision makers are in terms of finance or fitness, how healthy, how happy, and how accurate. To perform well by these correspondence norms does not require internal consistency at all.

In Hammond's framework, decision makers can be internally consistent, but completely indebted, in miserable health, unhappy—and perhaps most surprising, one's beliefs can be entirely self-consistent in the sense of conforming to Bayes' Rule and the definition of conditional probability, and yet completely wrong about everything one believes!

Section 2 describes the data and different empirical measures of inconsistency used in this investigation. Section 3 reports unconditional and conditional statistical links between different forms of inconsistency and expected payoffs. Section 4 incorporates measures of risk, and Section 5 concludes with a discussion of how the empirical findings can be interpreted.

Section 2: Data

The data used in this study were collected from March, 2002, through the beginning of 2003 in the Canadian cities of Calgary, Ottawa, Vancouver, Halifax, Toronto, and several rural locations. The data analyzed here consist primarily of 40 time-tradeoff choices and 32 risky choices (31 of which were binary), together with demographic information from post-experiment surveys. The time-tradeoff data provide multiple measures of time preference from which individual-level measures of time-inconsistency can be derived. The risky choice data provide multiple measures of risk preferences from which individual-level measures of expected-utility violations (i.e., an indicator marking inconsistencies that cannot be rationalized as maximization of an expected utility objective function).

Of the 881 subjects for whom these data were collected in initial sessions, 156 were invited back, roughly six months after their initial sessions, to make the same exact time-tradeoff and risky choices a second time. Thus, in addition to within-session time inconsistency and within-session expected-utility violations that were measured for all 881 subjects, a third type of between-session inconsistency is available for the invited-back subsample of 156. For these invited-back subjects, between-session inconsistency can be measured simply by counting the number (ranging from 0 to 72) of time-tradeoff and risky choice items answered differently between sessions. Most of the analysis below focuses on relationships between the first two within-session inconsistencies taken from the initial experimental sessions and each subject's total risk-taking and expected earnings. In a later section, we analyze between-session inconsistencies and report its conditional effect on total payoffs controlling for the first two forms of inconsistency.

Time Inconsistency

Table 1 outlines how the 40 binary time-tradeoff items are grouped into eight choice sequences of five binary choices, where the binary choices in each choice sequence have fixed arrival dates of payoffs (i.e., identical front-end delay and appreciation interval), the same time-1 payoff, and contains binary choices that vary only by increasing sizes of the time-2 payoff (i.e., monotonic improvements in the net payoff for patience). First, Table 1 describes the general structure of all 40 binary time-tradeoff items: subjects choose between option 1, which consists of a smaller and sooner payoff x arriving at t_1, versus option 2, which consists of a larger and later payoff y arriving at t_2 ($x < y$, $t_1 < t_2$). Thus, the appreciation interval is $t_2 - t_1$ and the net payoff for patience is y – x. To avoid confounding levels and percentages, we parameterize y in terms of its implicit annualized rate of return (assuming no compounding), the solution in r to $y = x(1+ r)^{(t_2 - t_1)}$, where t_1 and t_2 are measured in years. The early payoff x is always stated to subjects as a level, and the later payoff y is also stated as a level, together with the implied annual percentage return r and the absolute return y - x.

A choice sequence can be represented as a binary string that directly reveals an individual's time preference, letting annualized percentage returns r_j range over $r_1 < r_2 < \ldots < r_J$ and recording the jth binary choice, C_j, as 0 if the sooner and smaller (i.e., impatient) payoff is chosen and 1 if the later and larger (i.e., patient) payoff is chosen. This produces binary choice sequences of the form $C_1C_2\ldots.C_J$ with the following property. The earlier in this sequence (in which payoffs for patience are increasing in j) that the subject chooses the patient option (i.e., the earlier in the sequence that the first 1 appears), the more patient that subject is. For subjects whose preferences are monotonically increasing in payoffs, a zero can never follow a one, because that would imply acceptance of $y_j - x_j$ as compensation for waiting the appreciation interval $t_2 - t_1$, but refusal of the larger net payoff $y_{j+1} - x_{j+1} > y_j - x_j$ as compensation for waiting an identical appreciation interval. Therefore, we say that sequence S is more *patient* than sequence S' iff, when evaluated as integers, S > S'.

The experimenter-controlled appreciation interval, $t_2 - t_1$, is fixed at one month in the first four sequences and one year in sequences five through eight. Within each of these sets of sequences with fixed $t_2 - t_1$, the four sequences differ only by increasing front end delays, as shown at the bottom of Table 1 and spelled out in full detail in the following list, where the five component choices of each of the eight sequences is formed by letting r_j range over the values 0.05, 0.20, 0.50, 1.00 and 2.00 :

S1: Do you prefer \$65 today or \$65(1+r_j/12) one month from today?

S2: Do you prefer \$65 tomorrow or \$65(1+r_j/12) one month and one day from today?

S3: Do you prefer \$65 one month from today or \$65(1+r_j/12) two months from today?

S4: Do you prefer \$65 one year from today or \$65(1+r_j/12) one month and one year from today.[4]

[4] The formula 65(1+r/12) is, of course, an approximation of $65(1+r)^{1/12}$. The approximation is quite good for small r and, in all cases, is slightly larger than the formula with the fractional exponent. This approximation is commonly used by finance industry professionals and provides a shortcut by which subjects could verify (if they wanted to) that the dollar payoffs did indeed correctly reflect the annual rates of return printed in the booklets that subjects used.

S5: Do you prefer \$65 today or \$65($1+r_j$) one year from today?

S6: Do you prefer \$65 tomorrow or \$65($1+r_j$) one year and one day from today?

S7: Do you prefer \$65 one month from today or \$65($1+r_j$) one year and one month from today?

S8: Do you prefer \$65 one year from today or \$65($1+r_j$) one month and one year from today.

Part c of Table 1 shows the sequence denoted as S5 in the list above as it was presented to subjects. Given the definition of "more patient sequence" (which is a complete order over any space of binary choice sequences elicited using increasing net payoffs for patience), the six monotonic (out of 2^5=32 possible) sequences are ordered from least to most patient as follows: 00000, 00001, 00011, 00111, 01111, 11111.[5]

As Table 1 indicates, we can now define time inconsistency empirically, in the coarsest or most inclusive sense, by an indicator variable that flags mismatches among any two sequences with the same appreciation interval $t_2 - t_1$. We say that a subject is *time-inconsistent* if there is one or more mismatching sequences among sequences S1, S2, S3 and S4, or one or more mismatching sequences among S5, S6, S7 and S8. This broad measure of time-inconsistency can be refined in

[5]Monotonicity refers to time preferences that always rank greater compensation for an equal or smaller between-arrival waiting duration over smaller compensation for an equal or larger between-arrival waiting duration. A few nonmonotonic responses were observed for each of the eight arrival-date treatments, as is typical in a variety of preference elicitation techniques that do not impose monotonicity *a priori* (Jacobson and Petiei, forthcoming).

several ways. One way is to measure inconsistency in one-month condition (i.e., one or more inconsistencies among sequences S1, S2, S3 and S4) separately from inconsistency in one-year condition (i.e., among S5, S6, S7 and S8). We also record a count variable tallying instances of time-inconsistency, which ranges from 0 to 6, since the maximum number of mismatches among four sequences is 3 (the number of unique pairs from the combinatoric formula, 4!/(2!2!), for 4 choose 2). In addition, we report data on the direction of time-inconsistency, which connects the evidence reported here to a large theoretical and empirical literature on hyperbolic discounting and temptation (Laibson, 1997; O'Donoghue and Rabin, 1999; Coller, Harrison and Rutström, 2005).

In our framework, an individual is a *hyperbolic discounter* iff his or her sequences satisfy the inequalities $S1 \leq S2 \leq S3 \leq S4$ and $S5 \leq S6 \leq S7 \leq S8$, with at least one weak inequality holding strictly—a definition that requires all instances of time-inconsistency to be shifts from less patient to more patient as the front end delays move further into the future. Similarly, an individual is a *hypo-bolic discounter* if he or she shifts in the opposite direction, from patient in the short-run to impatient in the long-run: $S1 \geq S2 \geq S3 \geq S4$ and $S5 \geq S6 \geq S7 \geq S8$, with at least one inequality holding strictly. Although hyperbolic and hypo-bolic discounters are nonempty subsets among the time-inconsistent, most time-inconsistent subjects are shift in both directions at least once, implying that they are neither (strictly) hyperbolic or hypo-bolic discounters. We construct a count variable on the *net* number of hyperbolic minus hypo-bolic shifts to look for evidence of a systematic direction of time-inconsistency.

Table 2 shows empirical distributions for the time-tradeoff sequences S1 and S4 to examine one among the six pairs of sequences in which instances of time-inconsistency can be observed. Both S1 and S4 have a one month appreciation intervals ($t_2 - t_1 = 1/12$). In S1, the first arrival date t_1 is "today" (i.e., zero front-end delay). In S4, the first arrival date is one year from today (i.e., front-end delay of one year). Of 881 subjects, 579 are time-inconsistent based on these two sequences only. Among time-inconsistent subjects, shifts toward increasing patience are three to four times more likely than shifts toward impatience, which is the strongest evidence for hyperbolic discounting in our data but does not survive when directions of shifts are pooled over all six pairs. The next-to-last column shows the empirical distribution among subjects who are consistent, revealing that the modal sequence chosen by consistent subject is maximally impatient: 00000. Nearly 60 percent of consistent subjects are in the impatient half of the empirical distribution (00000, 00001 or 00011).

In contrast, the final column shows the empirical distribution from S4 among those who switched, showing a startling difference in the distribution of destinations to where these subjects switched compared to the choice sequences of consistent subjects. The modal destination choice among time-inconsistent subjects is the maximally patient choice sequence 11111, chosen by 42.5 percent of the time-inconsistent subjects.

These very different distributions in the final two columns of Table 2 illustrate an interesting challenge to normative interpretations of time-inconsistency. If most time-consistent subjects are consistently impatient, while time-inconsistent subjects produce a time-preference distribution that is on average more patient, then would anyone really suggest intervening pedagogically (e.g., teaching MBAs to be time-consistent) or with institutional change aimed at achieving consistency? Table 2 shows one at least one choice environment in which it would seem reasonable to conjecture that time-inconsistent subjects earn higher payoffs on average, after transforming all cash flows to present value using reasonable market discount rates. Before we pursue this empirical link between payoffs and consistency, we report detailed statistics for time-inconsistency in one-month and one-year appreciation interval conditions.

Table 3 tabulates aggregate numbers of time-inconsistent choices between pairs of sequences in the order in which they were implemented in the laboratory. The number of subjects providing inconsistent time-tradeoff sequences varies substantially, ranging from 152 subjects whose S1 and S2 choice sequences are mismatched to 469 subjects whose S7 and S8 choice sequences are mismatched. Recall from Table 2 that even greater numbers of mismatched choice sequences are observed among other pairs of sequences (e.g., S1 and S4). Taking the union of all individuals with one or more time-inconsistencies in the one-month treatments leads to 636 subjects recorded as ones using the one-month-condition-specific indicator TI_month. Taking the union of all individuals with one or more mismatched sequences on one-year condition leads to 614 subjects indicated by TI_year = 1.

Several observations from Table 3 are important for seeing that time-inconsistency should not be interpreted as random error (i.e., a tremble). First, the number of inconsistent responses grows substantially from S1 to S4, whereas they would be uniformly distributed (or decreasing with experience in models where learning is hypothesized to reduce inconsistency). In S5 through S8, inconsistent responses once again grow substantially. And the fact that the number of inconsistencies in the later one-year treatments (S5-S8) is not significantly less than in earlier one-month treatments (S1-S4) is yet another pattern that speaks against the random error or learning-reduces-inconsistency models. If learning reduced time inconsistency, then one would expect to see fewer instances of it in the last 20 binary time-tradeoff choices than the first 20, yet we observe roughly equal numbers in the first and last halves of the time-tradeoff data. Another piece of evidence against the random tremble explanation is that we would expect many more nonmonotonic sequences among the sequences that subjects switched to. If inconsistent choice sequences were the result of random trembles, we would expect the fraction of all switched-to sequences that are monotonic of only 5/31, whereas the observed fraction is approaching one: more than 99 percent of all switches are switches from one monotonic sequence to another.

Table 4 reveals another surprising pattern regarding time-inconsistency, which is its surprising lack of overlap in one-month and one-year conditions. Counting the number of subjects in the union of TI_month=1 or TI_year=1 (i.e., those who were time-inconsistent in the one-month or one-year treatments) reveals a total of 758 out of 881 time-inconsistent subjects. Looking at the cross-tabulation of TI_month and TI_year, which shows counts on consistency and inconsistency in one-month and one-year conditions, one finds that although 492 subjects were time-inconsistent in both conditions and 123 were time-consistent in both, a surprisingly large number—266 subjects—were time-inconsistent in one or the other sets of conditions, but not both.

Table 5 presents empirical distributions for count variables that tally instances of different kinds of time-inconsistent shifts. The first count variable in Table 5 tallies all forms of time-inconsistency, ranging from a minimum of zero (which occurs when S1=S2=S3=S4 and S5=S6=S7=S8) to a maximum of six (which occurs when S1, S2, S3 and S4 are all unequal, and S5, S6, S7 and S8 are all unequal). With a modal time-inconsistency count of 2 and more than 140 individuals tallying four to six time-inconsistencies, it is clear that the typical time-inconsistent individual is time-inconsistent more than once. The columns of Table 5 labeled "1-month" and "1-year" present empirical distributions for analogously defined count variables restricted to pairs in one-month (S1-S4) and one-year conditions (S5-S8), respectively, revealing more than 50 individuals in each case (more than 100 unique individuals: 53 + 59 – 6 = 106) who were maximally time-inconsistent in one of the two conditions. Although the link between payoffs and inconsistency is documented below, we note here (since these groups are not analyzed separately below) that subjects who were maximally time-inconsistent in one-month condition earned significantly higher payoffs (details to follow); subjects who were maximally time-inconsistent in one-year conditions earned significantly higher payoffs; and pooled together, of course, the 106 maximally time-inconsistent subjects earned higher-than-average payoffs. These bivariate correlations indicate that the scope of beneficial time-inconsistency goes well beyond switching once or twice in the direction of increased patience.

As defined earlier, hyperbolic discounting refers to a particular form of time-inconsistency in which the compensation required for waiting (a fixed appreciation interval in order to receive a larger payoff) decreases, the further forward into the future the waiting begins. Hyperbolic discounting occurs, for example, when a return of 50% or more is required to induce waiting one year versus receiving the first payoff today, but only 5 or 20% is required to induce waiting one year from tomorrow versus receiving the sooner payoff tomorrow. In both cases, the wait is precisely one year, but it matters to many people whether that wait begins sooner rather than later.

Counts on hyperbolic and hypo-bolic shifts show an interesting pattern. Although slightly more hyperbolic shifts were observed in the one-month treatment, slightly more hypo-bolic shifts were observed in the one-year treatment. Thus, no clear directional pattern of time-inconsistent shifts emerges from these data.

The three columns under the heading "NetHYP_count" count the number of hyperbolic shifts minus hypo-bolic shifts. The high frequency of subjects who shift in opposite directions at least once can be seen by comparing corresponding columns under NetHYP_count and HYP_count and HYPO_count, respectively. For example, the column under the heading "HYPO_count" shows that 5 subjects have 4 hypo-bolic shifts, whereas under the heading NetHYP_count one finds only 2 subjects that have a net count of -4 (i.e., 4 hypo-bolic shifts). The apparent discrepancy reflects the fact that 3 among the 5 subjects (who made 4 hypo-bolic shifts) also made one or two hyperbolic shifts as well, which changed their NetHYPcount values from -4 toward zero. The final two columns of Table 5 are counts on two indicator variables that mark subjects who make one or more shifts in the same direction and no shifts in the opposite direction. Of the 778 time-inconsistent subjects, only 211 are pure hyperbolic shifters and 144 are pure hypo-bolic shifters, which leaves 403 subjects who shift in opposite directions at least once.

Despite intense interest in hyperbolic discounting in some literatures, these experimental data provide no empirical support for hyperbolic discounting as a general explanation for observed time-inconsistency. Figure 1 shows three empirical distributions over subjects' counts of hyperbolic shifts minus the number of hypo-bolic shifts— NetHYPcount. These distributions appear to be centered at zero and rather symmetrical in all three cases. If the hyperbolic discounting model were the mechanism behind time-inconsistency, then we would expect these distributions to be substantially shifted to the right of zero.

One can examine whether the choice items with options for same-day payoffs drive most instances of time-inconsistency, and whether the hyperbolic discounting theory enjoys more empirical support after eliminating these pairs from consideration when counting instances of time-inconsistency. Table 5B show that the data do not support this interpretation—that time-inconsistencies are mostly generated by same-day choice items or that, after removing same-day-payoff choice items, most time-inconsistency can be explained as hyperbolic shifts. Eliminating same-day-payoff choice items from consideration increases the number of time-consistent subjects by only 23, from 123 to 146. And removing pairs of choice sequences with same-day-payoffs increases the number of pure hyperbolic discounters by only 28, from 211 to 239.

Risky Choice Data

This subsection describes how inconsistency with respect to expected-utility theory was measured. Subjects were asked to make 31 choices between pairs of gambles (one of which was often a sure thing), and one choice among six gambles. Two widely used experimental instruments for measuring risk-aversion (Holt-Laury (xxx) and Eckel-Grossman (xxx)) were among these, as well as an ambiguity aversion instrument consisting of choices over gambles with unknown probabilities.

Two of the risky choice items were mean-preserving spreads often used to categorize subjects as strictly risk-averse versus weakly risk-loving. A cross-tabulation of these choices is shown in Table 6. In choice 1, subjects choose between \$60 for sure versus \$120 with probability 5/10 and \$0 otherwise. In choice 2, subjects choose between \$60 for sure versus \$80 with probability 5/10 and \$40 otherwise. The 409 + 53 = 462 subjects in the off-diagonal cells appear to switch from risk-averse to risk-loving (or the reverse) and cannot be rationalized as having maximized an expected utility objective unless it is perfectly risk neutral. Other risky choice items gave subjects an opportunity to more sharply reveal risk neutrality (by providing risk-neutral responses on the Holt-Laury and Eckel-Grossman instruments). But only 37 of the 462 off-diagonal subjects in Table 6 are risk-neutral on the Holt-Laury instrument, and only 25 are risk-neutral on both Holt-Laury and Eckel-Grossman. Informal data collection on this pair of gambles from hundreds of our students reveal that many people, often a majority, choose inconsistently on these gambles (from the point of view of expected utility theory), and justify these choices, not in terms of risk neutrality, but by expressing a strong preference in favor of risk-taking on choice 2 and against risk-taking in choice 1. We feel that this is not necessarily crazy or irrational behavior, although it surely is a violation of expected utility theory (see Rabin (2000) for more examples of reasonable preferences in risky choice that admit no expected-utility representation).

The risky choice data provide additional opportunities for subjects to step outside the expected-utility preference model, as documented in Table 7. The two choice items from Table 6 indicate (conservatively) 422 EU violators. Combining these with the Eckel-Grossman measure raises the tally of EU violators to 496.

Another 79 subjects violated a basic monotonicity property implied by expected utility theory. These subjects preferred the gamble G (payoff H with probability p and payoff L with probability 1-p) over the sure thing S, but ranked the higher-probability-of-winning G' (payoff H with probability $p + \varepsilon$ and payoff L with probability $1 - p - \varepsilon$, $1 - p > \varepsilon > 0$) as inferior to S. A similar nonmonotonicity can be seen in the choice data from 16 ambiguity aversion items, which take the form: Gamble A (60 to 90% chance at winning $50) versus Gamble B ($s cash today), where s ranges from $18 to $48 in increasing $2 increments. If a subject were maximizing an expected-utility objective function with any prior distribution on the probability of winning, then monotonicity would require that, if B with sure-thing payoff s is preferred over gamble A, then B with sure-thing payoff s+2 should also be preferred over A. There were 57 subjects in our sample who violated monotonicity in this way. Taking the union of all EU-violation indicators leads to 553 EU violators in the sample.

Section 3: Results on Inconsistency and Total Payoffs

This section reports unconditional and conditional differences in expected payoffs as a function of different forms of inconsistency.

Payoff and Risk Measures

Each subject's total expected payoff was computed as follows. The 40 time-tradeoff items were assigned payoffs equal to the present value of subjects' choices (from the perspective of a Canadian subject in 2002-2003) using a discount factor of 0.05.[6] Risky choice items were mapped into the expected value of each subject's choice. Ambiguous gambles of the form "60 to 90 percent chance at winning $50" were translated to expected value using a uniform prior on the chance at winning (expected chance of winning equal to 75 percent). The variable we refer to as "total expected payoff" is simply the sum of present values and expected values across these 72 choice items.

Total risk was computed under the assumption of zero correlation among gambles in separate choice items. The variance of each gamble was computed and summed before taking the square root to produce a total standard deviation, which is the total risk measure referred to in the remainder of the paper. Time-tradeoff items contributed zero to total risk.

Payoffs and Inconsistency

Table 8 presents unconditional differences in total expected payoffs among time-inconsistent and EU-violator subsamples, respectively. The left-hand side of Table 8 presents means and ranges for total payoffs, total risk, and the sum of payoffs on different subsets of the choice items: 1-month appreciation interval, 1-year appreciation interval, all risky choices, the Eckel-Grossman choice item alone, the 10 Holt-Laury items, five other risky-choice items, and 16 ambiguity-

[6] According to the Bank of Canada's Department of Monetary and Financial Analysis, rates on 91-day treasuries varied between 2.25 and 3 percent from March, 2002, through January, 2003. Commercial borrowing rates were probably above 5 percent during this time, whereas interest rates on time deposits were probably less than 2 percent. Thus, the appropriate discount rate for future cash flows is not precisely known. Fortunately, none of the cash flows in our experimental items have a time horizon longer than two years, and the reported differences in payoffs are not even remotely sensitive to changes in the discount rate between 1 and 10 percent.

aversion choice items. The empirical ranges and sizes of those ranges (reported in the left-hand block of Table 8) are important for judging the economic significance of the differences in expected payoffs reported in the far right-hand column.

The right-hand block of Table 8 reports differences in (not levels of) mean payoffs for inconsistent subjects relative to consistent subjects. Notice that all these changes are positive, indicating that inconsistent subjects in all subsets of choice items achieved higher expected payoffs although these increases were not always statistically significant. Below each change in expected payoffs is the p-value associated with a two-sided unconditional t-test of equality of means between the inconsistent and consistent populations.

The average time-inconsistent subject earns $213 more than the average time-consistent subject. And the average EU-violator earns $112 more than the average non-EU-violator. Both differences are statistically significant and, as conditional effects from regressions reported below show, these differences are independently significant even after adding controls for risk-taking and demographic information including household income. The size of these effects should be judged as economically significant because, together, they cover more than one fourth of the entire total payoff range of $1,159 (maximum payoff minus minimum payoff =). One observation from Table 8 that might explain why EU violators earn higher payoffs is that they take somewhat more risk, 21.5 additional standard deviations on a range of size 201. However, this is far from the entire story, as Table 8 also shows that the average EU violator earns $93 more on time-tradeoff items and the average time-inconsistent subject earns an extra $14 or $15 on risky choice items.

Most of the gains from time-inconsistency show up as the $199 positive differential in payoffs on time-tradeoff items, and most this difference comes from one-year appreciation interval items where the gains from patience are largest. Beyond the raw earnings advantages among inconsistent subjects, we want to underscore that inconsistency in the risky-choice domain appears to have positive predictive power for earnings in the time domain, just as inconsistency in the time domain has predictive power for earnings in the risk-choice domain. This cross-domain predictive power of inconsistency survives in regression analysis, suggesting an interesting normative puzzle as to why inconsistency might yield benefits both inside the domain in which inconsistent choices are committed and beyond.

Before turning to the payoff regressions, we want to first consider the role of risk-taking as it is conventionally measured. Figure 2 presents a scatterplot of total risk and total payoffs. Plotted together, the northern-most points (i.e., convex closure of the points) in the figure represent an empirical risk-reward envelope across all subjects' choices in the experiment. Points along this empirical risk-reward envelope show the maximum payoff achieved among subjects for each level of risk or, equivalently, the minimum risk achieved at each payoff level.

In Figure 2, perfectly consistent subjects (time-consistent non-EU-violators) are plotted as squares, and everyone else as dots. "Everyone else" consists of subjects who are time-inconsistent or EU-violators (or both). While there are some squares located along or near the risk-reward envelope, most of them are deep inside the interior of the feasible risk-payoff choice set. The reason is that most of the consistent subjects are consistently impatient and consistently risk-averse, implying lower present-value payoffs on both time-tradeoff and risky-choice items. In contrast, while we observe some dots also located deep in the interior (i.e., far from the efficient envelope), most are clustered near the efficient frontier. Thus, the scatterplot implies that inconsistent subjects achieve more efficiency according to this very standard risk-return benchmark than consistent subjects do. A similar scatterplot with the y-axis replaced by expected payoffs on risky-choice items alone produces a similar result. Yet another statistical test of this difference in efficiency can be computed by fitting regressions of total payoffs (or payoffs from risky-choice items alone) on a quadratic in variance. Based on the data in the figure, this quadratic regression line among consistent subjects lies strictly below that of consistent subjects on the relevant range of variances, and this difference is statistically significant.

Table 9 reports the main results regarding the conditional effects of inconsistency on expected payoffs, using three regression models of total expected payoffs as a function of time-inconsistency and EU-violator status. Model 1 demonstrates that these different forms of inconsistency have independent predictive power. Furthermore, the effects of inconsistency on payoffs retain nearly the full magnitude of their respective unconditional effects, even after allowing for correlation between the two inconsistencies to absorb a portion of the other's effect.

Model 2 in Table 9 adds total risk-taking to the regression, which hardly reduces the effect size of time-inconsistency at all. After controlling for risk-taking, there remains a large effect size for EU-violator status as well—of $56, which is more than half the unconditional effect size reported in Table 8. This means that EU-violators achieve higher earnings, partly because they take more risk, but also because of something else that correlates with EU-inconsistency, which is not correlated with time-inconsistency.

In Model 3 of Table 9, another 20 demographic and survey-item variables are included as controls (for age, gender, marital status, geography, personal debt, household income, attitudes toward school, and success in school). Model 3 shows that these positive, large-magnitude and independently significant effects of inconsistency on payoffs are robust to a variety of other sources of interpersonal variation. Adding demographic controls makes the time-inconsistency and risk-taking coefficients increase slightly, while the EU-violator coefficient declines slightly, but remains large. Almost none of the time-inconsistency effect can be made to go away with the inclusion of risk-taking and demographic information. A little more than half of the EU-violator effect goes away with the inclusion of controls. In no case does one form of inconsistency appear to absorb much of the effect size of the other, possibly suggesting distinct mechanisms generating these two statistical links.

Between-Session Inconsistency

As mentioned earlier, 156 of 881 subjects returned to the lab after six months, plus or minus a few weeks. These 156 subjects repeated the same exact 72 choice items from the initial session. Whereas all results reported until now are based on the full sample of 881 using only information collected from subjects' initial sessions, we now analyze within-person differences in responses between sessions for the 156 repeat subjects.

The primary measure of between-session inconsistency considered here is a count of the number of switches among 72 choice items. Although this count of switches in each subject's responses ranges, in theory, from 0 to 72, its empirical range is 31 to 72, as shown in the histogram in Figure 3. We attempt to use the information provided by this extensive variation among subjects in the number of switches to estimate the conditional effect of between-session inconsistency on payoffs, revealing two more surprising results. The total payoff measure is computed just as before, but this time using second-session choice data.

Table 10 presents three regression models of total expected payoffs as a quadratic function of the number of switches. In Models 1 through 4, the conditional mean payoff is first increasing and then decreasing in the number of switches, indicated by the concave quadratic. We want to calculate the ranges of the switch count variable over which expected payoffs are increasing and decreasing. After fitting a quadratic conditional mean, it is straightforward to compute the "optimal" number of switches as the number of switches that maximizes the conditional mean expected payoff. Interestingly, the optimal number of switches according to the quadratic regression lines in Models 1, 2 and 3 are 56.8, 59.5 and 57.9, respectively--very near the sample median of the switch count variable, which is 57.5. Thus, in the first three models, the median subject is optimally inconsistent between sessions from the point of view of expected payoffs.

According to Table 10, between-session inconsistency is a statistically significant predictor of payoffs in Models 1, 2 and 3. In Model 2 with risk but not with other forms of inconsistency, the number of switches remains correlated with payoffs. In Model 3, between-session inconsistency retains its predictive power after controlling for time-inconsistency and inconsistency with respect to expected-utility theory. In Model 4, however, with risk and other inconsistency controls included, the coefficients on between-session inconsistency become statistically insignificant. However, the magnitude of the effect is sometimes large, with the conditional mean covering a range of more than $190 as the number of switches varies from 31 to 72 while holding all other regressors constant. Perhaps the most important result from Table 10 is that, after controlling for between-session inconsistency (in Models 3 and 4), time-inconsistency and EU-violator status continue to function as robust predictors of total payoffs.

Section 4: Conclusion

Inconsistency isn't always good, isn't always bad

We have some evidence that impatience often leads to economic outcomes that seem not to be desirable by objective evaluative criteria. For example, students who choose partying now over studying tend to get worse grades, drop out of school, earn less income, get sick with a higher likelihood, and die sooner. However, the normative economics of behavioral economics remains stuck on internal consistency, rather than context-specific measures of absolute performance (e.g., test scores, graduation, earnings, incidence of disease, and lifespan). The standard normative economics (also the normative perspective of perhaps most behavioral economists) is fixated on consumer sovereignty. Who's to criticize a student who values parties now over good grades at the end of the semester? Who can question trading off more income later for more fun in the present—that's the essence of economic man. And people who eat bad food now and get sick and die sooner than average are simply revealing a preference for immediate food pleasure over extra years of life.

References

Ariely, D. (2008) Predictably Irrational, New York: HarperCollins.

Ashraf, Nava, Dean Karlan, and Wesley Yin. (2006). "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines," Quarterly Journal of Economics, 121(2), 673-697.

Benjamin, Daniel, Sebastian A. Brown, Jesse M. Shapiro. (2006). "Who is 'Behavioral'? Cognitive Ability and Anomalous Preferences," Working Paper, University of Chicago.

Berg, N. (2003) Normative behavioral economics, Journal of Socio-Economics 32, 411-427.

Bernheim, B.D. and Rangel, A. (forthcoming), Behavioral Public Economics: Welfare and Policy Analysis with Fallible Decision-Makers, in Peter Diamond and Hannu Vartianen (eds.), Institutions and Behavioral Economics.

Bookstaber, R. and Langsam, J., 1985, On the optimality of coarse behavior rules, Journal of Theoretical Biology 116, 161-193.

Camerer, C., Issacharoff, S., Loewenstein, G., O'Donoghue and Rabin, M. (2003) Regulation for conservatives: Behavioral economics and the case for asymmetric paternalism, University of Pennsylvania Law Review 151, 1211-1254.

Caplan, B. (2001) Rational ignorance versus rational irrationality, Kyklos 54(1): 3-26.

Castillo, Marco, Paul Ferraro, Jeffrey Jordan, and Ragan Petrie. (2008). "The Today and Tomorrow of Kids," Working Paper, Georgia Tech.

Coller, Maribeth, Harrison, Glenn W., and Rutström, Elisabet E. (2005), "Does Everyone Have Quasi-Hyperbolic Preferences?" Working Paper, University of Central Florida.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde. (2007). "Are Risk Aversion and Impatience Related to Cognitive Ability?" IZA Discussion Paper No. 2735.

Gigerenzer, G., Todd, P.M., ABC Research Group, (1999), Simple Heuristics that Make Us Smart. New York: Oxford University Press.

Holt, Charles A. (1986). "Preference Reversals and the Independence Axiom," American Economic Review, 76, 508-15.

Hubal, R., Staszewski, J. and Marrin, S. (2007) Overcoming Decision Making Bias: Training Implications for Intelligence and Leadership, The Interservice/Industry Training, Simulation & Education Conference (I/ITSEC)Volume: 2007 (Conference Theme: Maintaining the Edge...Transforming the Force.)

Jacobson, S and Petrie, R. (forthcoming), Inconsistent Choices in Lottery Experiments: Evidence from Rwanda. Journal of Risk and Uncertainty.

Jolls, Christine and Cass R. Sunstein. 2006. Debiasing Through Law, Journal of Legal Studies, 35:199-241.

Laibson, David. "Golden Eggs and Hyperbolic Discounting." Quarterly Journal of Economics, 1997, 112(2), pp. 443–477.

Loewenstein G, Brennan T, Volpp KG. 2007. Asymmetric paternalism to improve health behaviors. Journal of the American Medical Association. Nov 28;298(20):2415-7.

Meier, Stephan and Charles Sprenger. (2007). "Impatience and Credit Behavior: Evidence from a Field Experiment," Federal Reserve Bank of Boston Working Paper Series, No. 07-3.

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

O'Donoghue, T. and Rabin, M. (2003) Studying optimal paternalism, illustrated by a model of sin taxes, American Economic Review 93(2), 186-191.

O'Donoghue, Ted, and Matthew Rabin. 2003. A note on optimal paternalism and health capital subsidies. Economics Letters. Volume 101, Issue 3, December 2008, Pages 241-242.

National Institute of Health (2008) http://www.nichd.nih.gov/publications/pubs/vasectomy_safety.cfm#xdisad Retrieved December 9, 2008.

Rabin, Matthew. 2000. Risk Aversion and Expected-Utility Theory: A Calibration Theorem. Econometrica. 68(5): 1281–292.

Schmidtz, D. (1995). Rational Choice and Moral Agency. Princeton: Princeton University Press.

Smith, Vernon L. 2003. Constructivist and Ecological Rationality in Economics. American Economic Review 93(3): 465-508.

Thaler, Richard H., and Shlomo Behartzi. Save More Tomorrow™: Using Behavioral Economics to Increase Employee Saving. Journal of Political Economy, 2004, vol. 112, 164-187.

Thaler, R.H. and Sunstein, C.R. (2008) Nudge: Improving Decisions About Health, Wealth, and Happiness, New Haven, CT: Yale University Press.

Table 1: Construction of Time-Tradeoff Data

A single choice

a) Let C denote a binary choice between a sooner payoff versus a larger payoff (i.e., x at t_1 versus y at t_2, with x < y and $t_1 < t_2$):

		code choice C as:
C	→ sooner and smaller payoff: x arriving at date	0
	→ later and larger payoff: y arriving at t_2	1

Definition of a time-tradeoff sequence

b) Fix the sooner and smaller payoff at x dollars and the two arrival dates measured in years, t_1 and t_2. Then parameterize the later and larger payoff y in terms of annualized rates of return r_j, $y_j = x(1 + r_j)^{(t2 - t1)}$, where r_j ranges over $r_1, r_2, \ldots, r_J$, to produce the following sequence of binary choices:

Choice 1 x arriving at t_1 v. $x(1 + r_1)^{(t2 - t1)}$ arriving at t_2, (code choice as C_1 in {0, 1})

Choice 2 x arriving at t_1 v. $x(1 + r_2)^{(t2 - t1)}$ arriving at t_2, (code choice as C_2 in {0, 1})

⋮

Choice J x arriving at t_1 v. $x(1 + r_J)^{(t2 - t1)}$ arriving at t_2, (code choice as C_J in {0, 1})

A realization of a single choice sequence can be represented as a 1xJ string of 0s and 1s: $C_1C_2\ldots C_J$.

Definition: The time-choice sequence S is more *patient* than S' if S > S', evaluated as integers.

One of the eight choice sequences that subjects faced (denoted S5 below)

c) 5 binary choices (J=5), where the first option's payoff is x = 65, the first option's arrival date is t_1 = today, the second option's arrival date is t_2 = one year from today, and the between-arrival duration t_2-t_1 = 1 year:

Choice 1: → 65 arriving today
→ 68.25 [5% more] arriving in one year

Choice 2: → 65 arriving today
→ 78 [20% more] arriving in one year

Choice 3: → 65 arriving today
→ 98.5 [50% more] arriving in one year

Choice 4: → 65 arriving today
→ 130 [100% more] arriving in one year

Choice 5: → 65 arriving today
→ 195 [200% more] arriving in one year

The sequence of five choices is coded as a 1x5 strings of 0s and 1s: $C_1C_2C_3C_4C_5$.

The following sequences are ordered from least to most patient: 00000, 00001, 00011, 00111, 01111, 11111.

Eight arrival-date conditions for each subject, first, holding between-arrival waiting duration constant and shifting arrival dates into the future for four conditions, and then repeating with between-arrival waiting duration changed from 1 month to 1 year

d) Time-choice data in this study consist of eight choice sequences of length five (40 binary choices in total), each with fixed x, and annualized rates of return: 5, 20, 50, 100 and 200 percent. Between-arrival waiting duration, t2 - t1, is 1 month for the first four sequence and 1 month for the last four sequences:

arrival of first payoff	*between-arrival waiting durations*	
t_1	t_2-t_1 = 1 month	t_2-t_1 = 1 year
today	S1	S5
tomorrow	S2	S6
in 1 month	S3	S7
in 1 year	S4	S8

Definition of time inconsistency: one or more mismatches among S1, S2, S3 and S4, or one or more mismatches among S5, S6, S7 and S8.

Table 2: Empirical Distributions* for a Pair of Choice Sequences

	Choice sequences: now versus one month from now			Choice sequences: 12 months versus 13 months from now		*empirical distributions of consistent v. inconsistent*			
	choice sequences	# subjects		choice sequences	# subjects	# time-consistent subjects	# time-inconsistent subjects**	percentage among the consistent	percentage among the inconsistent
	00000	272		00000	171	109	62	36.1	10.7
	00001	160	*579 time inconsistent*	00001	90	35	55	11.6	9.5
	00011	213	<-------≠------->	00011	173	63	110	20.9	19.0
	00101	1							
	00110	1		00110	1		1		0.2
|	00111	74		00111	66	10	56	3.3	9.7
|	01011	1							
|	01100	1							
V	01111	58		01111	60	16	44	5.3	7.6
more	10011	1							
patient				10101	1		1		0.2
	10111	1		10111	1		1		0.2
	11000	1							
				11011	1		1		0.2
				11100	1		1		0.2
				11110	1		1		0.2
	11111	97		11111	315	69	246	22.8	42.5
Total:		881			881	302	579	100.0	100.0

*Empirical distributions tabulate the number of individuals who choose different sequences over 5 binary decisions. The left-most distribution is parameterized by arrival times of today and one month from today. The second distribution has arrival times of one year from today and 13 months from today. In both cases, the five binary choices are defined by increasing rewards for waiting an extra month for the later payoff, ranging from annual rates of return of 5, 20, 50, 100 and 200 percent. The 302 subjects who choose two identical sequences are referred to as consistent (at least for this pair of choice sequences). The 579 whose choice sequences do not match are referred to as inconsistent. Empirical frequency and percentage distributions for both subsamples are presented to the right.

**Among consistent subjects, the empirical distributions are by definition identical in both choice sequences. Among the inconsistent, the distribution presented in the table shows choices for the second sequence (t_1 = 12 months from today versus t_2 = 13 months from today)--in other words, where subjects shifted to.

Table 3: Construction of TI, an Indicator of Time Inconsistency (Unequal Choice Sequences Over Evenly Spaced Cashflows with Different Starting Dates)

TI = 1 if TI_month==1 or TI_year==1, 0 otherwise.

In other words, TI = 1 *if any of the following six pairs of choice sequences involving time trade-offs (indicated by curving arrows) fails to match. By this definition, 758 of 881 subjects are time inconsistent.*

Duration between arrivals of payoffs is one month

sequence	sequence 2	sequence 3	sequence 4
Now v. 1 month	Tomorrow v. 1day+1month	1month v. 2months	1year v. 1year+1month

mismatching choice sequence?	mismatching choice sequence?	mismatching choice sequence?
152	385	449

union of subjects with one or more mismatches = TI_month, which includes 636 individuals

Duration between arrivals of payoffs is one year

sequence	sequence 6	sequence 7	sequence 8
Now v. 1 year	Tomorrow v. 1day+1year	1month v. 13 months	1year v. 2 years

mismatching choice sequence?	mismatching choice sequence?	mismatching choice sequence?
187	254	469

union of subjects with one or more mismatches = TI_year, which includes 614 individuals

Table 4: Cross-Tabulation of Time Inconsistency in 1-Month versus 1-Year Between-Duration Treatments

	1-year treatment		
1-month treatment	consistent	inconsistent	Total
consistent	123	122	255
inconsistent	144	492	636
Total	267	614	881

Table 5A: Frequency Distributions for Time Inconsistency, Hyper- and Hypo-bolic Discounting

sample	TI_count*			HYP_count*			HYPO_count*			NetHYP_count**			Pure***	
frequency	all	1-month	1-year	all	1-month	1-year	all	1-month	1-year	all	1-month	1-year	HYPER	HYPO
-4							5			2				
-3							57	1	3	22	1	3		
-2							159	37	67	63	19	36		
-1							326	243	372	156	106	262		144
0	123	245	267	267	353	537	334	600	439	285	373	384	670	737
1	164	339	377	290	397	294				188	278	170	211	
2	245	244	178	221	124	49				112	97	25		
3	206	53	59	84	7	1				43	7	1		
4	97			18						9				
5	40			1						1				
6	6													

*The variable TI_count tallies the number of time-inconsistencies (i.e., mismatching choice sequences illustrated in Table 2) observed out of a possible six comparable pairs. HYP_count (HYPO_count) begins from zero and adds "1" ("-1") every time a subject shifts in the direction of patience (impatience) as payoff arrival dates move further into the future holding between-arrival-waiting duration fixed.

**Net_HYP_count is the sum of HYP_count and HYPO_count, tallying a net count on shifts toward patience after allowing shifts in opposite directions to cancel.

***Pure_HYP (Pure_HYPO) is an indicator = 1 if at least one shift toward patience (impatience) and zero shifts toward impatience (patience) were observed. These variables follow the definitions of hyperbolic and hypo-bolic discounting (allowing for no shifts in the opposite direction) given in the text.

Table 5B: Frequency Distributions for Time Inconsistency, Hyper- and Hypo-bolic Discounting with Same-Day-Payoff Choice Items Removed*

sample	TI_count*			HYP_count*			HYPO_count*			NetHYP_count**			Pure***	
frequency	all	1-month	1-year	all	1-month	1-year	all	1-month	1-year	all	1-month	1-year	HYPER	HYPO
-4							1			1				
-3							24			17				
-2							136	18	28	57	18	28		
-1							335	238	353	169	104	250		168
0	146	270	314	314	392	592	385	625	500	295	404	417	642	713
1	211	388	411	319	418	264				204	284	161	239	
2	266	223	156	194	71	25				94	71	25		
3	218			49						39				
4	40			5						5				

*Table 5B is a modified version of Table 5A after removing counts of inconsistency from pairs of time-tradeoff sequences with any options to receive same-day payoffs. The data do not support the interpretation that time-inconsistencies are mostly generated by same-day choice items or that, after removing same-day-payoff choice items, most time-inconsistency can be explained as hyperbolic shifts. Eliminating same-day-payoff choice items from consideration increases the number of time-consistent subjects by only 23, from 123 to 146. And removing those items increases the number of pure hyperbolic discounters by only 28, from 211 to 239.

**Net_HYP_count is the sum of HYP_count and HYPO_count, tallying a net count on shifts toward patience after allowing shifts in opposite directions to cancel.

***Pure_HYP (Pure_HYPO) is an indicator = 1 if at least one shift toward patience (impatience) and zero shifts toward impatience (patience) were observed. These variables follow the definitions of hyperbolic and hypo-bolic discounting (allowing for no shifts in the opposite direction) given in the text.

Figure 1: No Evidence for Hyperbolic Discounting in Net Number of Shifts Toward Patience (as arrival dates move further into the future)

1-month treatments

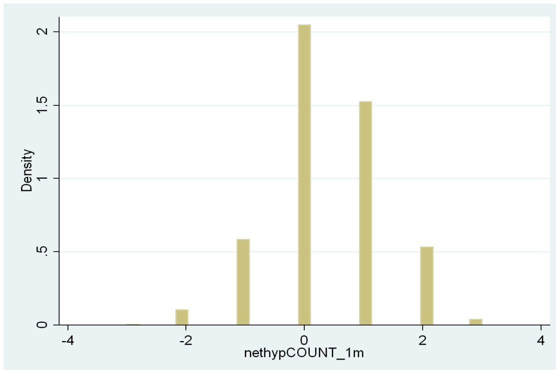

1-year treatments

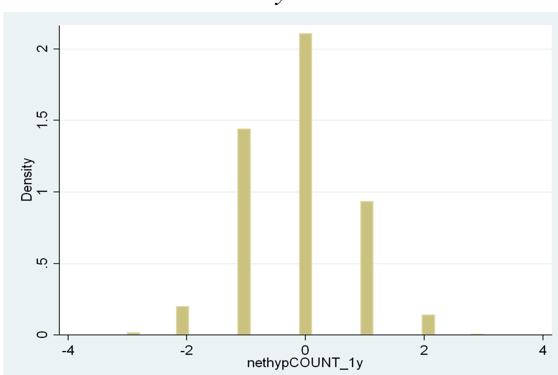

all treatments

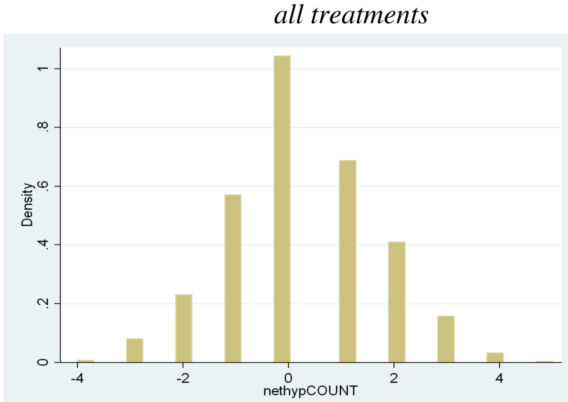

Table 6: Conflicting* Choices among Two Mean-Preserving Spreads

choice 2	*choice 1* A: $60 for sure	B: $120 with prob 5/10, $0 otherwise	Total
C: $60 for sure	359	53	412
D: $80 with prob 5/10, $40 otherwise	409	60	469
Total	768	113	881

*Because a risk-neutral agent is indifferent between A and B, and indifferent between C and D, the 409 + 53 = 462 subjects on the off-diagonal cells above are not necessarily violating expected utility theory. However, of these 462 subjects, only 37 provide risk-neutral responses on the Holt-Laury instrument measuring risk preferences, and only 25 provide responses on both Holt-Laury and Eckel-Grossman that are consistent with risk neutrality.

Table 7: Violations of Expected Utility Theory in 38 Risky Choices

expected-utility violations	# violators out of 881	
RiskLoving&Averse in Two Mean-Preserving Spreads	462	(422 if risk-neutral HL excluded)
Inconsistent responses on two binary choice items involving mean-preserving spreads: sure-thing $60 v. $80 with probability 0.5 and $40 otherwise; and sure-thing $60 v. $120 with probability 0.5 and 0 otherwise. Inconsistency on these items do not violate EU theory in case of risk-neutral preferences. In the broadest, all-inclusive indicator of non-EU behavior (the variable "All EU Violations" at the bottom of this list), inconsistent responses on these mean-preserving spreads are *not* counted if the same subject's responses to the Holt-Laury items are perfectly risk neutral		
RiskLoving&Averse+EG	521	(496 if risk-neutral HL excluded)
Inconsistent responses among three items, including previous two, and 6-gamble choice problem (Eckel-Grossman) which includes one gamble that is a mean-preserving spread of another gamble. An EU maximizer (who is not risk neutral) must answer these three items consistently.		
Holt-Laury Nonmonotonic	79	
10 binary choices between gambles with fixed payoffs and probabilities of winning ranging from 0.10 through 1.00. Nonmonotonic responders have risk preferences with no EU representation. The last binary choice in the sequence is between two sure-thing payoffs, $40 v $77, revealing 6 subjects who apparently prefer $40 over $77.		
Ambiguity Nonmonotonic	57	
16 binary choices between an ambiguous gamble (probability of winning between 0.60 and 0.90) and increasing sure-thing payoffs. EU maximizers with any subjective probability of winning cannot have nonmonotonic choice sequences.		
Yes-to-Big/No-to-Low Risk	66	
Choose gamble with EV=82.5 and sd=126 over 75 for sure, and choose 120 for sure over EV=140 and sd=70. These subjects take 82.5-75 = 7.5 as compensation for bearing risk of sd=126 in the first choice, but refuse an EV premium of 140-120 = 20 to bear risk of sd=70 in the second choice.		
All EU Violations	553	(576 if we don't exclude HL risk-neutral)
Indicator = 1 for any of the above EU violations, except that the violators from the mean-preserving spread items are *not* counted for the 84 subjects with perfectly risk-neutral responses to Holt Laury, since inconsistency on choices over mean-preserving spreads is consistent with risk neutrality.		

Table 8: Increases in Expected Dollar Payoffs among Time- and EU-Inconsistent Subjects

summary statistics in ***levels*** *among the entire sample*					*unconditional* ***difference*** *in mean payoffs: inconsistent versus consistent subsamples*		
payoff measure	min	mean	max	size of range		time-inconsistent v. time-consistent	EU-violators v. non-violators
total payoff	3764.3	4573.8	4923.4	1159.1	Δ E[total payoff]	213.5	112.3
					p-value*	0.0000	0.0000
individual σ	10.3	89.7	211.5	201.3	Δσ	4.7	22.5
					p-value	0.2807	0.0000
time payoffs	2566.4	3246.9	3496.5	930.2	Δ time payoffs	199.0	93.1
					p-value	0.0000	0.0000
1-month time payoff	1283.2	1329.2	1357.9	74.7	Δ 1-month payoffs	11.4	6.7
					p-value	0.0000	0.0001
1-year time payoff	1283.2	1917.7	2138.6	855.5	Δ 1-year payoffs	187.5	86.4
					p-value	0.0000	0.0000
risky payoffs	1165.4	1326.9	1427.9	262.5	Δ risky payoffs	14.5	19.3
					p-value	0.0162	0.0000
Eckel-Grossman payoff	28.0	32.6	36.0	8.0	Δ Eckel Grossman	0.5	0.8
					p-value	0.0325	0.0000
Holt-Laury payoffs	370.4	459.1	491.9	121.5	Δ Holt Laury	2.2	4.2
					p-value	0.4373	0.0391
ambiguity payoffs	498.0	592.7	633.0	135.0	Δ Ambiguity	10.8	12.2
					p-value	0.0036	0.0000
other gambles	355.0	362.6	387.5	32.5	Δ Other Gambles	1.1	2.1
					p-value	0.2956	0.0038

*Unconditional t tests of the equality of means among inconsistent versus consistent subpopulations produced the p-values in this table.

Figure 2: The Risk-Reward Envelope (consistent subjects represented by squares, and time-inconsistent and/or EU-violators represented by dots)

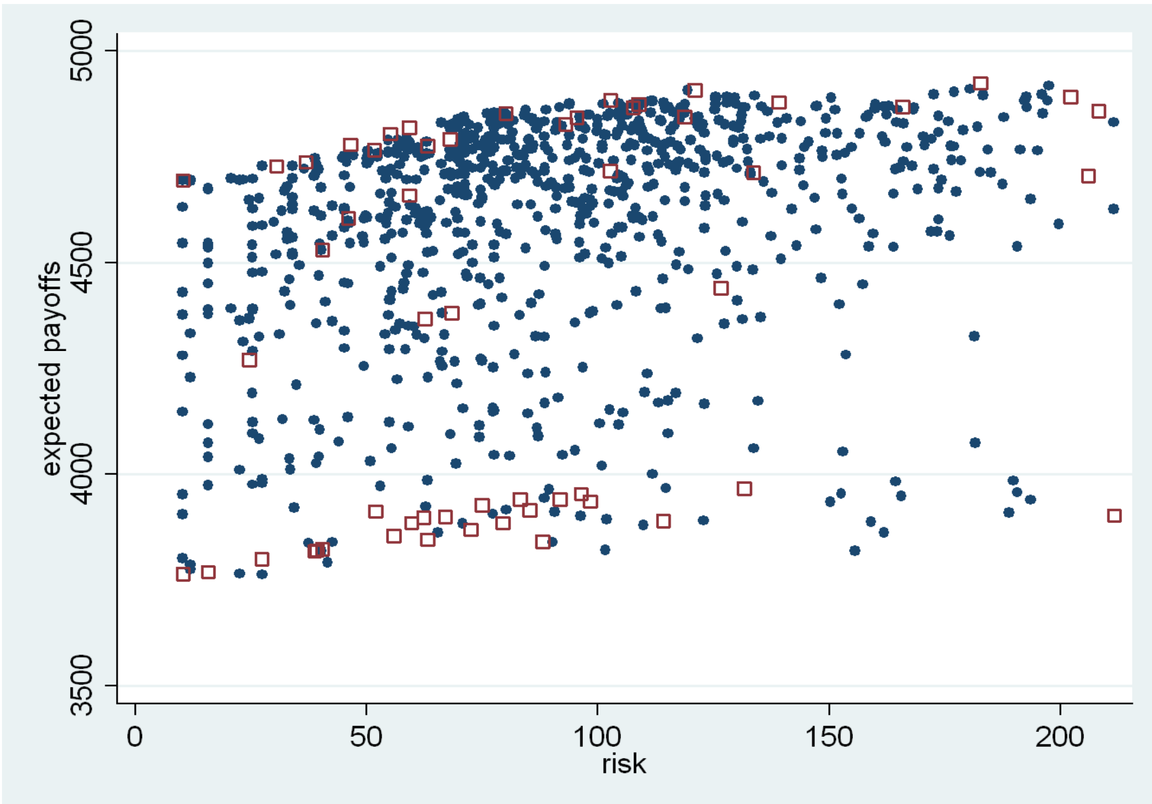

Table 9: Regressions of expected payoffs on inconsistency, risk taking, and other controls (N = 881)

	Model 1		Model 2		Model 3	
variables	coeff	t stat	coeff	t stat	coeff	t stat
TI	198.70	7.2	196.50	7.5	197.35	7.7
EU-violator	96.68	4.9	56.59	2.9	47.23	2.5
Individual σ			1.79	8.5	1.93	9.2
Under 25					28.89	1.2
Female					15.52	0.8
Immigrant					0.00	0.1
Married					-16.46	-0.8
Ontario					-138.58	-3.3
British Columbia					-69.67	-1.5
Nova Scotia					-79.84	-1.8
Alberta					-124.09	-2.7
Native Person					-22.88	-0.4
Disabled					8.03	0.3
French Speaker					-22.95	-0.4
Burdened by Debt					-20.62	-1.1
Sell Asset to Pay Debt					-68.38	-2.4
Medium Household Income					48.80	2.2
High Household Income					90.45	3.7
Not Working					-41.58	-1.5
Completed High School					-132.13	-2.6
Liked School					-11.67	-0.5
Peers Liked School					46.07	2.2
Performed Well in School					57.07	3.0
Constant	4342.17	160.2	4208.89	138.2	4208.08	61.5
R squared	0.0636		0.1575		0.2285	

Figure 3: Empirical distribution of between-session switches in 72 time-tradeoff and risky choices

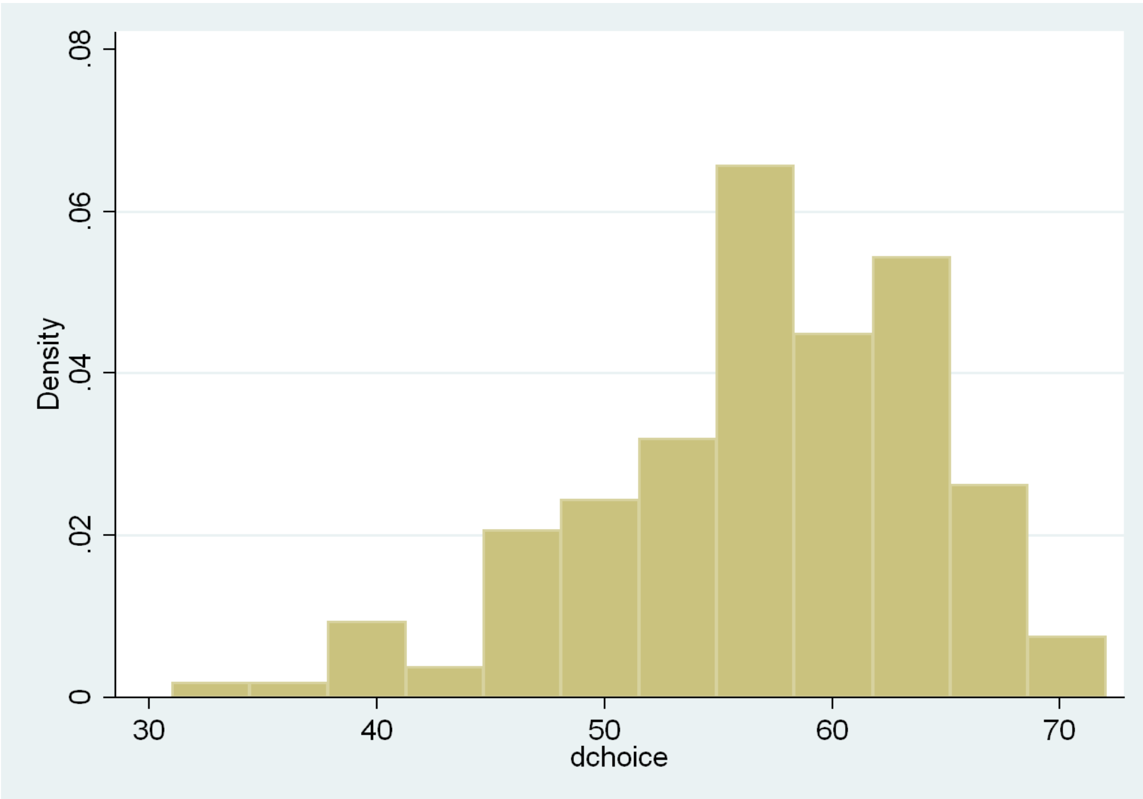

Table 10: Regression of Second-Session* Expected Payoffs as a Quadratic Function of the Number of Between-Session Switches

	Model 1		Model 2		Model 3		Model 4	
variables	coeff	t stat	coeff	t stat	coeff	t stat	coeff	t stat
#switches	119.25	3.2	81.73	2.3	96.21	2.5	42.81	1.3
#switches^2	-1.05	-3.1	-0.69	-2.1	-0.83	-2.4	-0.42	-1.3
Individual σ			2.84	4.5			1.29	2.2
TI					148.47	1.9	191.32	2.8
EU-violator					132.03	2.5	121.43	2.6
constant	1112.55	1.1	1846.73	1.9	1500.46	1.5	3154.22	3.5
#obs	156							

*A subsample of 156 subjects was invited back six months after their initial sessions to exactly repeat the 72 decisions they had made earlier. The variable #switches counts the number of switches, which in theory ranges from 0 to 72. The empirical range is 31 to 72. The maximizers of the quadratic regression lines in Models 1, 2 and 3 are 56.8, 59.5 and 57.9, respectively, which is very near the sample median, 57.5. The maximizer in Model 4 is 50.8.