Inconsistency Pays?: Time-Inconsistent EU Violators Earn More

Nathan Berg[*], Catherine Eckel[**] and Cathleen Johnson[***]

Abstract: Experimental choice data from 881 participants based on 40 time-tradeoff and 32 risky choice items indicate that most participants are time-inconsistent and also violate the consistency requirements of expected utility theory. These inconsistencies cannot be explained by theories such as hyperbolic discounting and cumulative prospect theory, nor are they inversely correlated with math ability. Aggregating expected payoffs and risk associated with the 72 choice items, statistical associations between inconsistency and expected payoffs are reported. Time-inconsistent participants and those who violate expected utility theory both earn substantially higher expected payoffs, and these positive associations survive the presence of various controls in total payoff regressions. Consistent participants earn lower than average payoffs because most are consistently impatient or consistently risk averse. Positive payoff effects from inconsistency are not fully explained by greater risk taking, however. Controlling for the total risk, math ability and socio-economic differences, both time-inconsistent and expected-utility-violating participants earn significantly more money. Positive returns to inconsistency extend beyond the domain in which the inconsistencies occurs, with time-inconsistent participants earning more on risky choice items, and expected utility violators earning more on time-tradeoff items. The results raise questions regarding the normative status of rationality axioms based on internal consistency and common interpretations that draw on insights from behavioral economics to form a mistaken premise: that inconsistency is evidence of pathological decision making which motivates paternalistic intervention.

JEL Codes: D83 (Information and Knowledge and Communication) [please add and fix]

Key Words: Bounded Rationality, Behavioral Economics, Paternalism, De-bias, Dynamic Inconsistency, Savage Axioms, Within-Person, Between-Session, Axiomatic, Rationality

[*] Nathan Berg, Associate Prof. of Economics, University of Otago, PO Box 56, Dunedin 9054, New Zealand,

phone: (+64) 3 479 8648, nathan.berg@otago.ac.nz

[**] Catherine Eckel, xxx Professor, Department of Economics, xxx

[***] Cathleen Johnson, xxx

Inconsistency Pays?: Time-Inconsistent EU Violators Earn More

Introduction

Given large theoretical, empirical and policy literatures dealing with time-inconsistency and violations of expected utility theory, one could be forgiven for assuming—incorrectly—that there exists well-founded evidence linking internally inconsistent choice patterns to significant economic costs. This paper reports evidence that, at least within the confines of common decision tasks used to elicit time and risk preferences, inconsistency may be positively associated with payoffs (i.e., changes in wealth). Our data reveal that participants who violate time consistency and expected utility theory regularly exit lab studies with significantly more money than participants whose choice data are internally consistent. Expected utility violations in our choice data are positively correlated with math ability, whereas time inconsistency's correlation with math ability (although positive) is indistinguishable from zero. Positive correlation between inconsistency and payoffs easily survives inclusion of controls for math ability, socio-economic status, demographics and geography.[1]

---

[1] Some previously reported evidence conflicts with the findings we report. None (as far as we know), however, points conclusively to positive correlations between consistency and payoffs. Burks, Carpenter, Gotte and Rustichini (2008) report that trainees learning to become truck drivers who had lower-than-average cognitive skills were more likely to lose money on their investment in the training program (by leaving the training program before recouping out-of-pocket costs). Understandably, there are no data on changes in wealth among those truck drivers who left the program measuring how well they performed (financially) in the alternatives they pursued (other work, training, or receipt of government transfers). Benjamin, Brown and Shapiro (2006) report a modest negative association between cognitive skills and the incidence of preference anomalies. Their data show, however, that those with very high cognitive skills nevertheless frequently exhibit anomalies. Jacobson and Petrie (2009) find zero correlation between time-inconsistency on experimental time-preference elicitation instruments (exhibited by 55% of 181 participants) and real-world financial decisions; similarly for risk aversion. Jacobson and Petrie do find an

Of course, this finding of positive associations between inconsistency and payoffs does not prove that inconsistent people are universally better off, or that myopic decision making is not a genuine problem in particular decision settings. Instead, normative characterization of choice data may essentially require vector- rather than scalar-valued measures to produce meaningful inter-personal comparisons of performance (or wellbeing) calibrated to context. The positive correlation between inconsistency and payoffs that we report addresses an infrequently discussed evidential gap concerning how frequently studied violations of normative decision theory based on consistency link to classical normative comparisons based on wealth (e.g., as in the writings of Smith, Ricardo and Marshall). If different normative metrics give conflicting answers to the question "Which decision procedure should I use?," then the descriptive empirics of those multiple performance metrics become relevant. We provide a characterization of correlational structure among five frequently applied normative metrics that sometimes generate

---

interaction effect (time-inconsistency and risk-preference measures) which suggests a slight increase in time-inconsistent participants' rate of using informal financial instruments; Jacobson and Petrie interpret their finding to mean that preference anomalies are correlated with pathological financial decision making, although no dollar costs of time inconsistency are reported. Ashraf, Karlan and Yin (2006) combine experimental and field evidence to show that time-inconsistent bank customers are self-aware about the potential pitfalls of succumbing to their impulsivity, which motivates them to choose bank products with less flexibility (offering identical interest rates) as a sometimes costly commitment device. Ashraf, Karlan and Yin's finding suggests that people who are inconsistent in the lab may be sophisticated in recognizing their vulnerability to making errors, leading some (perhaps many?) to preemptively deploy successful strategies aimed at accumulating greater wealth. Chu and Chu (1990) and Cherry, Crocker and Shogren (2003) report that participants who were paid to avoid inconsistent choices quickly learned to be consistent. List and Millimet (2004) show that participants in the field vary significantly in terms of consistency of choice, and that market experience reduces the probability of inconsistent choice patterns—without showing, however, that inconsistency leads to reduced economic performance.

conflicting prescriptive advice: (i) expected utility (EU) violations are observed in risky choice patterns that cannot be rationalized as expected utility maximization; (ii) time inconsistency (TI) is observed in time trade-offs that cannot be rationalized as maximization of a time-separable objective function with exponential discounting; (iii) between-session inconsistency is observed (and quantified as an increasing function of the number of switched responses) when identical decision tasks elicit inconsistent responses after returning to the experimental lab three to six months later; (iv) math ability; and (v) cumulative change in wealth during a lab session.

The evidential gap linking violations of consistency axioms and substantial economic losses is noted (to varying degrees) by those cautioning that there are missing links in the logical chain that starts from the vast empirical literature documenting many instances of choice data violating consistency-based norms and ends in prescriptive recommendations for interventions that encourage closer conformity with consistency axioms (Sugden (xxx)[2], Güth and Kliemt, xxx), Gilboa, xxx and postlewait, xxx). The prescriptive goal of bringing inconsistent people into closer conformity with rationality, defined in terms of consistency norms, rests on the (largely untested) assumption that inconsistency leads to costs that harm people's performance or wellbeing. If individuals do not conform to standard rationality axioms, what then is the economic cost?[3] Our goal is to begin filling that evidential gap by estimating the conditional predictive power of different inconsistencies on payoffs in lab experiments.

The flexibility of consistency norms (i.e., consumer sovereignty, formalized as a methodological commitment to not make interpersonal comparisons based on dollar- or utility-

[2] "[I]ndividuals who are known to be fully rational in the conventional sense may be less successful in reaching their objectives than they would have been, had they been (and been known to be) less rational" (Sugden, 1991, p. 782).

[3] Caplan's (2001) model argues that axiomatic irrationality should be concentrated in decision domains where its costs are least.

based metrics) is attractive for rationalizing different people's choices. Consistency norms can, however, be both too strong and too weak. Consistency rules out choice procedures as irrational that are, by other performance metrics, well performing, reasonable, adaptive, or likely to achieve good chances of survival. The adaptive benefits of inconsistency may be particularly likely to show up, for example, in unstable environments where action sets (and mappings from action profiles into payoffs) change unpredictably, one reason why random experimentation may be beneficial.[4] Consistency also "rules in" as rationalizable many behaviors that many observers and other disciplines consider pathological. If one accepts that violations of consistency are in fact widespread in human populations (as we believe the empirical behavioral economics literature persuasively demonstrates), then how much adaptive pressure penalizing inconsistent choice behavior—in both present and past environments—can there be? Perhaps, one worthwhile refinement of our empirical research program would ask: Which environments penalize, and which reward inconsistency?

The paper proceeds as follows. Section 2 describes the data and summarizes the empirical measures of inconsistency. Section 3 reports both unconditional and conditional correlations among different measures of inconsistency and expected payoffs. Section 4 presents the main

[4] Rubinstein (1998) describes an intransitive choice rule when selecting an item from a long catalog: scan items only on the first and last pages, and then choose the best item from within this small subset of the catalog. This choice rule violates *order invariance* and, of course, gives rise to intransitivities when the pages of the catalog are re-ordered. Rubinstein argues that such an inconsistent choice rule is in fact reasonable and should not be automatically dismissed as uninteresting or pathological simply because it violates consistency. Sugden (2008, 2009) argues in favor of induction in environments with profound uncertainty rather than narrow consistency norms in environments where actions and payoffs are exhaustively known. Gilboa, Postlewaite and Schmeidler (2009, p. 285) advocate "a view of rationality that requires a compromise between internal coherence and justification…."

results that express mean cumulative payoffs conditional on inconsistency while controling for risk and math ability. Section 5 concludes with a discussion.

Section 2: Data

The data record 40 choices involving time tradeoffs and 32 choices over risky lotteries (as well as demographic information from post-experiment surveys) for 881 Canadian participants. Time-tradeoff choice items consist of eight sequences of binary choices between an earlier and smaller payment or a later and larger payment: two appreciation intervals (one month, or one year) × four front-end delays (no delay, one day, one month, or one year) × five annual rates of return for the later payment (5%, 10%, 20%, 50%, 100% or 200%). Risky choice items included Holt-Laury and Eckel-Grossman instruments for measuring risk preferences, as well as several mean-preserving spreads.

Because the eight sets of time-trade-off choice sequences consist of two subsets of four with identical appreciation intervals (one month, or one year), a maximum of six time inconsistencies were observable for each participant. The risky choice items generated several opportunities to observe inconsistency with respect to expected utility theory (described in detail below) in the form of alternating risk-averse and risk-loving choices. Of the 881 participants, 156 were invited back roughly six months after their initial sessions to face exactly the same 72 choice items a second time. Our time inconsistency and expected-utility violations were all observed within a single experimental session and are used to present our main empirical results regarding their effects on expected cumulative earnings. Based on the subsample of 156 who were invited back for a second session six months after the initial session, a third type of between-session inconsistency was observable, which we report as the number of switches in response to identical choice items (ranging theoretically from 0 to 72, but with an empirical range of 3xxx to xxx). Most of our analysis focuses on relationships between the first two within-session inconsistency measures—time inconsistency and EU violations—and the expected present value of each participant's cumulative earnings in these 72 items. Participants were told before the sessions began that they would be paid only for one randomly drawn round, implying that they were incentivized to respond to each choice item as if it were a one-off. In Section xxx, we return to between-session inconsistency, estimating its effect on cumulative earnings while controlling for time inconsistency and EU violations.

Time Inconsistency

A choice sequence can be represented as a binary string revealing an individual's time preference as follows. Let annualized percentage returns r_j range over $r_1 < r_2 < \ldots < r_J$ and record 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 coding scheme produces binary choice sequences of the form $S \equiv C_1C_2\ldots C_J$ with convenient properties. Noting that payoffs for patience are increasing in j, if the later and larger option is chosen earlier in the sequence (i.e., the first 1 appears earlier in the sequence), then the participant reveals herself to be more patient.[5] Therefore, we can say that sequence S is a *more patient* than sequence S' iff S > S' when both sequences are evaluated as integers. Furthermore, if preferences over deterministic cash

[5] Given this definition of *more patient* sequence (which is a complete order over any space of binary choice sequences elicited using increasing net payoffs for patience) of length J = 5, there are only six monotonic (out of 2^5 = 32 possible) sequences, ordered from least to most patient: 00000, 00001, 00011, 00111, 01111, 11111.

flows are monotonically increasing in present value for any discount rate, then a zero can never follow a one (although we do observe a handful of non-monotonic sequences as show below).[6]

The appreciation interval is fixed at one month in the first four sequences and at one year in the second four. The front-end delay takes on values of zero, one day, one month and one year. And for each of the eight pairs of appreciation intervals and front-end delays, five binary choices force choices between a fixed earlier payment of $65 and later payments (at a fixed date) with increasing rewards for patience corresponding to ascending annualized[7] rates of return (r_1 , r_2, r_3, r_4, r_5) = (0.05, 0.20, 0.50, 1.00, 2.00). The following eight sequences are five-digit binary numbers in which each digit represents binary choices between an earlier payment of $65 versus a later and larger payment with annualized rates of return indexed by j = 1, 2, ..., 5:

S_1: $65 today or $65(1+$r_j$/12) one month from today?

S_2: $65 tomorrow or $65(1+$r_j$/12) one month and one day from today?

S_3: $65 one month from today or $65(1+$r_j$/12) two months from today?

S_4: $65 one year from today or $65(1+$r_j$/12) one month and one year from today.

---

[6] In Appendix 1, x denotes the earlier and smaller payoff arriving at t_1; y denotes the later and larger payoff arriving at t_2 ($x < y$, $t_1 < t_2$); the appreciation interval is $t_2 - t_1$; and the net payoff for patience (in dollars) is $y - x$. To avoid confounding levels and percentages, y is parameterized by implicit annualized rates of return (assuming no compounding), taking on the values r = 0.05, 0.10, 0.20, 0.50, 1.00 and 2.00 in the future value formula $y = x(1+r)^{(t2 - t1)}$, where t_1 and t_2 are measured in years. A non-monotonic sequence implies acceptance of $y_j - x$ as compensation for waiting the appreciation interval $t_2 - t_1$, but refusal of the larger net payoff $y_{j+1} - x > y_j - x$ as compensation for waiting an identical appreciation interval.

[7] The formula 65(1+r/12) is an approximation of $65(1+r)^{1/12}$, which is good for small r and, in all cases, slightly larger than the correct future value formula with fractional exponent.

S_5: \$65 today or \$65(1+r_j) one year from today?

S_6: \$65 tomorrow or \$65(1+r_j) one year and one day from today?

S_7: \$65 one month from today or \$65(1+r_j) one year and one month from today?

S_8: \$65 one year from today or \$65(1+r_j) one month and one year from today.

Table 1 shows how time-inconsistencies among the eight 5-digit sequences are coded. In the coarsest or most inclusive sense, time inconsistency is coded as an indicator for one or more mismatches among any two of the four sequences with the same appreciation interval. We say that an individual's choice data are *time-inconsistent* if there is one or more mismatching one-month-appreciation sequence among {S_1, S_2, S_3, S_4} (TI_month=1, 0 otherwise) or one or more mismatching one-year-appreciation sequences among {S_5, S_6, S_7, S_8} (TI_year=1, 0 otherwise). The most inclusive binary indicator TI = 1, 0 otherwise, codes the union one-month-appreciation and one-year-appreciation time inconstancies. We also record a count variable tallying instances of time-inconsistency, which ranges from 0 to 6, because the maximum number of mismatches among four sequences with the same appreciation interval is $3 = \binom{4}{2}$, and there are two appreciation intervals.[8]

In addition, we report data on the direction of time-inconsistency, linking the evidence reported here to the large theoretical and empirical literatures on hyperbolic discounting and

[8] Appendix 2 shows a surprising pattern regarding observed time-inconsistencies, which is its surprising lack of overlap in one-month and one-year conditions. Table 1 showed counts on the number of participants 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), which revealed a total time-inconsistency frequency of 758 out of 881 individuals. The cross-tabulation of TI_month and TI_year in Appendix 2 indicates that although 492 were time-inconsistent in both conditions and 123 were time-consistent in both, a surprisingly large number—266 participants—were time-inconsistent in only one or the other accumulation intervals but not both.

temptation (Laibson, 1997; O'Donoghue and Rabin, 1999; Coller, Harrison and Rutström, 2005). An instance of *hyperbolic discounting* is observed whenever one or more of the following six inequalities hold: $S_1 < S_2 < S_3 < S_4$ or $S_5 < S_6 < S_7 < S_8$, which indicates shifting toward greater patience as the front-end delay (indexed by the subscripts denoting the choice sequences) increases, while holding the accumulation interval constant (among S_1 through S_4, or S_5 through S_8). We count the number of time inconsistencies that satisfy this definition and net out counts of inconsistencies in the opposite direction (with all six inequalities reversed, which defines an instance of *hypobolic discounting*, shifting from patient in the short-run to impatient in the long-run). Although the events of hyperbolic and hypobolic discounting are nonempty subsets among all observed time inconsistencies, most time-inconsistent individuals in our sample shift in both directions at least once, implying that they are neither (strictly) hyperbolic or hypo-bolic discounters.

Figure 1 shows the empirical distributions of the *net* number of hyperbolic minus hypobolic shifts, broken out by the two accumulation intervals and pooled, to look for evidence of any systematic directionality among the observed time inconsistencies. If the hyperbolic discounting model were the mechanism underlying observed time inconsistencies, then we would expect these distributions to be substantially shifted to the right of zero. Little, if any, evidence linking front-end delay to the directionality of time-inconsistent shifts can be found in Figure 1 or in Appendices 3A and 3B, which tabulate net counts with and without inclusion of inconsistencies that involve zero front-end delay.[9]

Table 2 shows empirical distributions for sequences S_1 and S_4 as an example of just one among the six pairs of sequences in which time-inconsistencies were observed. This particular

[9] Of the 778 time-inconsistent individuals, only 211 are pure hyperbolic shifters and 144 are pure hypobolic shifters, which leaves 403 who shift in opposite directions at least once.

pair was chosen because it exhibits the strongest tendency toward hyperbolic discounting (i.e., shifting in the direction of greater patience when the front-end delay increases: $S_1 < S_4$). Both S_1 and S_4 have one-month appreciation intervals ($t_2 - t_1 = 1/12$ years). In S_1, the earlier payoff's arrival date t_1 is "today" (i.e., zero front-end delay). In S_4, the earlier payoff's arrival date is "one year from today" (i.e., front-end delay of one year). Among the 881 observations of S_1 and S_4, 579 are time-inconsistent based on these two sequences only.[10] The next-to-last column shows the empirical distribution among individuals who are time-consistent in S_1 and S_4 (i.e., $S_1 \neq S_4$), which shows that the modal time-consistent sequence is maximally impatient: 00000. Nearly 60 percent of consistent observations are in the impatient half of the empirical distribution (00000, 00001 or 00011). In contrast, the final column shows the empirical distribution from S_4 among those who were inconsistent, revealing a very different distribution. The modal observation of S_4 is maximally patient, 11111, chosen in 42.5 percent of time-inconsistent cases.

If, as in Table 2, most time-consistent individuals are consistently impatient while time-inconsistent individuals are more patient, then is there a compelling case to intervene pedagogically (e.g., teaching MBAs to be time-consistent) or with institutional changes aimed at achieving consistency? By the performance metric of net present value (at any discount rate), it would seem that many who advocate intervening to encourage greater time consistency might instead want to directly focus on encouraging patience (whether time consistent or otherwise). Table 2 shows just one choice environment in which it would seem reasonable to conjecture that

[10] Among time-inconsistent participants, 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. The tendency toward greater patience as front-end delay increases (or changes from zero to strictly positive) does not survive, however, when data from all eight sequences are considered, as in Figure 1 and Appendices 3A and 3B.

time-inconsistent participants earn higher payoffs on average, after transforming all cash flows to present value using market discount rates.

Risky Choice Data

This subsection describes how inconsistency with respect to expected-utility theory was measured. Participants were asked to make 32 risky choices: 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 choices over mean-preserving spreads.

Table 3 presents a cross-tabulation of two binary choices over mean-preserving spreads with a common mean. In choice 1, participants chose between \$60 for sure versus \$120 with probability 5/10 and \$0 otherwise. In choice 2, participants chose between \$60 for sure versus \$80 with probability 5/10 and \$40 otherwise. The 409 + 53 = 462 observations in the two 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 perfectly risk neutral. Nearly all of those counted in the off-diagonal cells revealed themselves to be strictly risk-averse on at least on other choice item, ruling out the perfect risk neutrality explanation.

For example, only 37 of the 462 off-diagonal individuals participants were perfectly risk neutral as measured by the Holt-Laury instrument. And only 25 of the 462 were risk-neutral on both Holt-Laury and Eckel-Grossman. We collected hundreds more observations of these two choice items subsequently including open-ended survey items asking participants to explain their reasoning. Very few alternators on choices 1 and 2 in Table 3 expressed indifference. Instead, the most frequent explanation focuses on worst-case payoffs. In gamble D in choice 1, the worst-case payoff is \$40, which many participants pointed to as a good enough guarantee to justify gambling and having at a chance at \$80. In choice 2, however, the threat of leaving \$60 on the table and receiving \$0 made gamble B unattractive relative to the sure-thing A. Other participants alternated in the opposite direction on the basis of the best-case payoff. Reasoning on the basis of the best-case payoff suggests to some that \$120 is a sufficiently large upside potential that warrants risking exiting the lab with \$0, whereas \$80 relative to \$60 in choice 1 is not sufficiently thrilling to shift from a default of risk-aversion. As others have theorized, we contend that these reasoned preferences which are situationally strictly risk-averse in some settings depending on the particular attributes and risk-loving when other conditions are met cannot be dismissed as irrational. In the world of commerce and in many other settings, it seems possible that we benefit from having these heterogeneous ways of reasoning about risk in our population and perhaps within-person. Such reasoning violates expected utility theory under the common assumption of global 2nd-order risk preferences. (Rabin, 2000, provides further examples of reasonable preferences in risky choice that admit no expected-utility representation).

Combining the choice data in Table 3 with the Eckel-Grossman measure raises the tally of EU violators to 496. Another 79 participants violated a basic monotonicity property implied by expected utility theory by preferring a gamble G (payoff H with probability p and payoff L with probability 1-p, H > L) over the sure thing S while ranking the higher-probability-of-winning gamble 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 dollars for sure), where s ranges from \$18 to \$48 in increasing \$2 increments. If participants were maximizing an expected-utility objective function with any prior distribution on the probability of winning, then monotonicity would require that, if Gable 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 participants in our sample who violated monotonicity in this way. Taking the union of all EU-violation indicators produces our primary measure of EU violation, with 553 EU violators in the sample.[11]

Section 3: Results on Primary Measures of 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 participant's total expected payoff was computed as follows. The 40 time-tradeoff items were assigned payoffs computed as present value using a discount factor of 0.05.[12] Risky choice items were mapped into expected values. 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, implying an expected chance of winning equal to 75 percent (midpoint of 60 and 90). The variable we refer to as "total expected payoff" is the sum of present values and expected values across the 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

---

[11] Appendix 4 provides more detail and summarizes counts for each component EU violation whose union is coded by our variable EU violator.

[12] 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 during the 10-month period in which our data were collected. The relevant commercial borrowing rates for individuals in our sample were above 5 percent during this time, whereas interest rates on time deposits were less than 2 percent. Thus, the appropriate discount rate is not precise but is range-bound. Fortunately, none of the cash flows in our experimental items had a time horizon longer than two years. Given that the annualized returns for patient choices ranged from 5% to 200%, the differences in the time-discounted payoffs we report are not sensitive to changes in the discount rate between 1 and 10 percent.

the square root to produce a total standard deviation, which is the total risk measure σ as referred to in the remainder of the paper. Time-tradeoff items contributed zero to total risk measure σ.

Payoffs and Inconsistency

Table 5 presents unconditional differences in total expected payoffs among time-inconsistent and EU-violator subsamples, respectively. The left-hand side of Table 5 presents means and empirical 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-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. For example, Table 5 shows that more payoff variation was generated by the 40 time choice items than the 32 risky choice items.

The right-hand block of Table 5 reports differences (not levels) in mean payoffs for inconsistent and consistent individuals. Notice that all these changes are positive, indicating that inconsistent participants in all subsets of choice items achieved higher expected payoffs (although 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.

The average time-inconsistent participant earns $213 more than the average time-consistent participant. 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 will show, these differences are independently significant even after including controls for risk-taking, demographic information including household income, total risk, and math scores. 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 5 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 difference is at best an incomplete explanation, as Table 5 also shows that the average EU violator earns $93 more on time-tradeoff items and the average time-inconsistent participant 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 of this difference comes from the one-year appreciation interval items where gains from being patient are largest. Beyond these raw earnings advantages among inconsistent participants where the inconsistency is matched with the choice domain, Table 5 reveals that inconsistency in the risky-choice domain is associated with an earnings premium in the time domain, just as inconsistency in the time domain is associated with an earnings premium in the risk-choice domain. These cross-domain or out-of-domain correlations between different inconsistencies and earnings will be seen to survive in multiple regression analysis using the full sample.

Before turning to the payoff regressions, we want to first look directly at risk-taking and payoffs. Figure 2 presents a scatterplot of total risk and total payoffs: an empirical version of the risk and expected return "Markowitz Bullet" in finance textbook. The northern-most points (i.e., convex closure) in Figure 2 represent a risk-reward envelope across the observed range of participants' risk taking. Points along this empirical risk-reward envelope show maximum payoffs achieved among participants (with variation along each vertical) for each level of risk; or, equivalently, the minimum risk achieved at each payoff level.

Perfectly consistent individuals (i.e., time-consistent non-EU-violators) are plotted as squares in Figure 2, and everyone else as dots. "Everyone else" consists of individuals 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 participants are consistently impatient and consistently risk-averse, resulting in lower present-value payoffs on both time-tradeoff and risky-choice items. While some dots also located deep in the interior (i.e., far from the efficient envelope), Figure 2 shows they are clustered near the efficient frontier. The scatterplot implies that inconsistent participants achieve more efficiency according to the standard risk-return benchmark than consistent participants do. A similar scatterplot with the y-axis replaced by expected payoffs on risky-choice items alone produces a similar result.

Table 5 reports the main results regarding the conditional effects of inconsistency on expected payoffs, using six regression models of total expected payoffs as a function of time-inconsistency and EU-violator status. Model 1 indicates that these two forms of inconsistency have independent predictive power. Furthermore, the effect of inconsistency on payoffs retains nearly its full magnitude even after allowing for correlation between the two inconsistencies to absorb a portion of the other's effect.

Model 2 in Table 5 adds total risk-taking to the regression, which reduces the effect size of time-inconsistency hardly at all. After controlling for risk-taking, there remains a large effect size of $56 for EU-violator status, which is more than half the unconditional effect size reported in Table 4. This result means that EU-violators achieve higher earnings partly because they take more risk, but also because of something else that correlates with EU violation but is uncorrelated with time-inconsistency.

In Model 3 of Table 5, another 20 demographic and survey items are included as controls: for age, gender, marital status, geography, personal debt, household income, attitudes toward school, and self-reported success in school. Model 3 shows that 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 modestly but remains large. Almost none of the time-inconsistency effect disappears with the inclusion of risk-taking, demographics or other information in Model 3. A little more than half of the EU-violator effect goes away with the inclusion of these same controls. In no case, however, does one form of inconsistency appear to absorb much of the effect size of the other. This observation suggests that distinct mechanisms are responsible for generating these two positive effects of inconsistencies on payoffs.

Math Test Scores

Models 4 through 6 in Table 5 include nationally standardized math scores, further rural/urban information (that is not captured by Provincial indicator variables) and, finally, an interaction term between math scores and both indicators of inconsistency.[13] Including math ability further attenuates the EU-violator premium, although it remains significant at more than a third of the effect size in Model 1. The main effect of time inconsistency on payoffs remains virtually unchanged. In Model 6, the indicator term implies that time inconsistency is most beneficial for those with low math scores and less so for those who perhaps apply a different procedure using math skills to achieve high payoffs. Nothing in Table 5 suggests that the

[13] Models 4 through 6 also include self-described math ability on a Likert scale; the effects of this variable (included for completeness to show how the data respond to filtering using all available controls) are minimal and removing it from the model has no visible effects on other coefficients.

inconsistency premiums we report are fragile or badly confounded by other observable characteristics.

Table 6 shows pairwise correlations among inconsistencies and math scores. Time inconsistency and EU violation have a significant correlation coefficient of around 0.11. This rather small magnitude reinforces evidence seen in Table 5 of their independent effects on payoffs. Time inconsistency is uncorrelated with math scores. And EU violation is positively correlated with math scores, once again running counter to standard normative interpretations in which conformity with expected utility theory, math ability, and payoffs provide a harmonious set of overlapping metrics for classifying people as less or more rational.

Between-Session Inconsistency

As mentioned earlier, 156 of 881 participants returned to the lab after six months (plus or minus a few weeks). These 156 participants repeated the same 72 choice items from the initial session. All results reported subsequently were based on the full sample of 881 using only information collected from participants' initial sessions. We now analyze within-person differences in responses between sessions for the 156 repeat-participant subsample.

The primary measure of between-session inconsistency is a count of the number of switches among 72 choice items, #Switches.[14] Although this count ranges, in theory, from 0 to 72, its empirical range is 31 to 72, as shown in the histogram in Figure 3. #Switches captures a potentially distinct form of inconsistency, and we try using this information (with substantial variation among participants) to estimate the conditional effect of between-session inconsistency on payoffs. The total payoff measure is computed just as before, but this time using second-session choice data. The results reveal two more surprising results.

[14] We experimented with various increasing transformations and discretized versions of #Switches without overturning anything reported in this section.

Table 7 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 #Switches over its empirical range (as indicated by the concave quadratic). After fitting a quadratic conditional mean, it is of course straightforward to compute the expected-payoff-maximizing number of switches. 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, which is near the sample median of #Switches of 57.5. Thus, in the first three models, the median participant is optimally between-session inconsistent (by the metric of expected payoffs).

According to Table 7, between-session inconsistency is a statistically significant predictor of payoffs in Models 1, 2 and 3. In Model 2 (adding control for risk but not for other forms of inconsistency), the number of switches remains correlated with payoffs. In Model 3, between-session inconsistency, once again, retains its predictive power after controlling for time-inconsistency and EU-violator status (but excluding risk). In Model 4, however, with risk and other inconsistency controls included, the coefficients on between-session inconsistency become statistically insignificant.[15] We note from Table 7 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

We present evidence from lab data collected from 881 individuals showing that time inconsistencies, EU violations and between-session switches on identical decision tasks are all

[15] The magnitude of the effect of changes in #Switches is sometimes nevertheless large: the conditional mean covers a range of more than $190 as the number of switches varies from 31 to 72 (holding all other regressors fixed).

positively associated with the expected present value of cumulative payoffs. The positive effects of time inconsistency and EU violation on payoffs are large and robust to the inclusion of risk, math scores and a variety of other controls. Different forms of inconsistency exert surprisingly independent effects: it is not easy to make these effects go away by including another one as a regressor.

Different time-inconsistency measures were computed, the simplest of which was a binary indicator for any participant 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. Most participants were time inconsistent (xxx out of 881). Most time-inconsistent participants exhibited both hyperbolic (more patient) and hypobolic (less patient) shifts in their required rewards for enduring a fixed appreciation interval as front-end delay was increased. Therefore, the data provide little support for quasi-hyperbolic or hyperbolic discounting theories as explanations for the mechanism generating observed time inconsistencies.

The second form of inconsistency is frequently documented in experimental studies with multiple risky choice items: participants' choices over risky gambles cannot be rationalized (e.g., a strictly risk-loving choice in one pair of choice items, and strictly risk-averse on others) as maximization of any globally concave or convex expected utility function. The third form of inconsistency was measured by counting the number of switches in responses for a subsample of individuals who were invited back to repeat the same 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 participants, on average, earned less money.

Positive payoff premiums for inconsistency extended beyond the choice domain in which they were measured. For example, time-inconsistent participants earned more not only on time-tradeoff choice items, but also on risky choice items (with no time variation). Similarly, EU violators earned more on time-trade-off choice items (with no risky-choice component). Participants who violated time-consistency, violated expected utility theory, or switched responses to identical decision tasks between sessions earned significantly higher payoffs than participants who were perfectly (or more) consistent.

[we could close the paper right here...but let's consider continuing as follows:)]

Normative Behavioral Economics and Ecological Rationality

Ask a behavioral economist what behavioral economics teaches the social sciences for purposes of applied policy work or institutional design, and one likely will hear prescriptive calls to help error-prone, biased, and irrational humans overcome systematic pathologies built into their brains (e.g., Ariely's *Predictably Irrational*, 2008, or Sunstein and Thaler's *Nudge*, 2008). And yet, little evidence exists linking frequently observed violations of axiomatic rationality to differences in other performance norms that economists commonly interpret as proxies for wellbeing, such as earnings, physical health, lifespan, and happiness. The normative behavioral economics literature (Camerer et al, 2003; O'Donoghue and Rabin, 2003; Berg, 2003; Loewenstein et al, 2007; Bernheim and Rangel, forthcoming) increasingly includes discussions of policy measures designed to "de-bias" those who violate consistency norms (Jolls and Sunstein, 2006).[16] The prescriptive goal of bringing inconsistent people into closer conformity

[16] For example, 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 re-design US intelligence policy based on this literature and the de-biasing program it suggests. This point of view rests on the

with rationality, defined in terms of consistency norms, rests on the (largely untested) assumption that inconsistency leads to real costs that harm people's performance or wellbeing.

This paper shows that experimental participants 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, with and without controls for risk-taking, demographic differences (including household income), and success in school. The data we present challenge the frequent normative interpretation of consistency violations as generally pathological, whether referred to as "irrational," "a-rational," or simply as "mistakes." Our results also suggest that, at least in some important decision environments, behavioral economists' normative analyses may be referencing normative benchmarks that are not particularly relevant to real-world decision makers' own assessments of success (i.e., what it means to make a good decision).

Psychologists Hastie and Rasinski (1987) introduced the coherence-correspondence distinction (later expanded upon by Hammond, 1996): a dichotomous taxonomy partitioning the universe of normative criteria. Coherence is what this paper refers to as a *consistency norm* (e.g., transitivity, the Savage axioms, dynamic consistency, Bayes' Rule). Consistency norms impose restrictions only on two or more decisions (i.e., can rationalize any single decision or inference considered in isolation) and do not produce a metric of performance in dollars that naturally supports interpersonal comparisons. In contrast, correspondence is what this paper refers to as a *level norm*, which uses a scale whose units can be used to evaluate performance based on a single observation and make interpersonal comparisons (e.g., wealth, health, lifespan, self-

---

premise that failures in US foreign policy could have been avoided if only they were made to be consistent with maximizing a scalar-valued objective function.

reported happiness). Hammond observes that decision makers can be internally consistent yet perform poorly (e.g., indebted, unhealthy, unhappy, etc.), just as one's beliefs may be entirely self-consistent in the sense of conforming to Bayes' Rule (i.e., probabilistic logic based on the definition of conditional probability) and yet completely wrong (i.e., badly mismatching objective frequency distributions).

Biologists provide some ideas as to why organisms can be inconsistent and successful. Bookstaber and Langsam (1985), for example, argue that if nature "tuned" organisms optimally to any particular environment, it would lead to maladaptive (in the long-run) evolutionary disadvantage in the face of changing environments; whereas less stringent satisficing rules promote fitness by allowing for much greater flexibility and faster adaptation to changing ecologies as they are buffeted by 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—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 our imagination too far to envision humans who enjoy adaptive benefits by flexibly changing the weights they place on different factors and competing favorably with those who dogmatically conform to the stricture of maximizing a time-separable utility function with exponential discounting.

Consider the money-pump argument used in many textbooks to justify transitivity as a reasonable consistency requirement: people with intransitive choice patterns can be induced to make a sequence of trades that leaves them with no money and therefore without influence on important economic variables that economists routinely analyze.[17] By making reference to how much money people have, however, the money pump argument draws on a *level norm* (i.e., wealth levels) to justify a *consistency norm* (i.e., transitivity). These are different normative criteria in need of empirical evidence to show that conforming to one (transitivity) correlates with performing well by the other (wealth). Money pump arguments ask readers to infer (based on logical deduction, not on empirical evidence) that these distinct normative criteria are in harmony with each other—that rational decision making should perform well, by both consistency norms such as transitivity and level norms such as wealth. A minute's reflection on transitive decisions over labor versus leisure, however, should quickly dispel the hypothesis that those who conform to transitivity will necessarily accumulate more wealth than those who violate it. Thus, the question of convergence versus divergence among multiple normative criteria becomes an empirical question, one which this paper attempts to address.

[we could close the paper here, too. i was tempted, however, to continue:]

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, referred to as the *appreciation interval*, is at least as long: $t_3 – t_4 > t_1 – t_2$. According to this *consistency norm*, a person who spends his entire paycheck for a party on the day his paycheck arrives is rational

[17] Rubinstein and Spiegler (2008) critique money pump arguments on the grounds that actually carrying out exploitative transactions requires face-to-face contact that very likely triggers an attitude of caution or suspicion among the potentially exploitable. As Rubinstein and Spiegler (2008, p. 237) put it, “We tend to think strategically about the situation and suspect that there is a ‘catch,’ even if we cannot pinpoint it.”

(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 that behavior irrational.

The intuitive lack of appeal in this kind of normative analysis seems obvious. Yet many analyses drawing on insights from behavioral economics proceed 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 live an unhappy and evolutionarily disadvantaged life. The tension here lies in the difference between consistency norms and level norms, 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, level norms evaluate the performance achieved by competing decision making procedures according to how well calibrated they are to the environments in which they are used (also referred to as ecological rationality by Gigerenzer et. al., 1999, and Smith, 2003). Evaluating decision procedures by asking How wealthy?, How healthy?, How happy?, and How accurate? should complement many important and widely shared goals of economic analysis. Perhaps normative analysis requires a multivariate characterization of multiple performance metrics.

Acknowledging that performing well (as measured by one of the level norms just mentioned) does not require internal consistency need not undermine the rigor of economic analysis. Rather, it should add to its relevance and range of applicability when characterizing not only potential pitfalls of decision making, but best practices, too: when decision procedures are well matched to the particular environments in which they are used.

References

Ariely, D. (2008), Predictably Irrational, New York: HarperCollins.
Ashraf, N., Karlan, D. and Yin, W. (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, Journal of Public Economics, 95(11-12), 1377-1385

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.

Güth, W., & Kliemt, H. (2001). From full to bounded rationality: The limits of unlimited rationality. Discussion Papers, Interdisciplinary Research Project 373: Quantification and Simulation of Economic Processes.

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. (2009), Inconsistent Choices in Lottery Experiments: Evidence from Rwanda. Journal of Risk and Uncertainty, 38, 143-158.

Jolls, C. and Sunstein, C. R. (2006), Debiasing Through Law, Journal of Legal Studies, 35, 199-241.

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

Loewenstein G, Brennan T, Volpp K. G. (2007). Asymmetric paternalism to improve health behaviors. Journal of the American Medical Association, 298(20), 2415-2417.

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.

Sugden, R. (1991), Rational choice: A survey of contributions from economics and philosophy. Economic Journal, 101, 751-785.

Sugden, R. (1998). The metric of opportunity. Economics and Philosophy, 14, 307-337.

Sugden, R. (2003). Opportunity as a space for individuality: Its value and the impossibility of measuring it. Ethics, 113(4), 783-809.

xxxSugden, R. (2004). The opportunity criterion: Consumer sovereignty without the assumption of coherent preferences. The American Economic Review, 94(4), 1014-1033.

Sugden, R. (2006a). Taking unconsidered preferences seriously. Royal Institute of Philosophy Supplements, 59, 209-232.

Sugden, R. (2006b). What we desire, what we have reason to desire, whatever we might desire: Mill and Sen on the value of opportunity. Utilitas, 18(01), 33-51.

Sugden, R. (2008), Why incoherent preferences do not justify paternalism, Constitutional Political Economy, 19(3), 226-248.

Sugden, R. (2009). On nudging: A review of nudge: Improving decisions about health, wealth and happiness by Richard H. Thaler and Cass R. Sunstein, International Journal of the Economics of Business, 16(3), 365-373.

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 TI, an Indicator of Time Inconsistency Defined as Unequal Choice Sequences Over Cash Flows with Identical Appreciation Intervals but Different Front-End Delays (i.e., Evenly Spaced Cashflows with Different Starting Dates and Identical Present Values on Their Respective Starting Dates)

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

TI = 1 *if one or more mismatching pairs of sequences among* $\{S_1,\ldots,S_4\}$ *or among* $\{S_5,\ldots,S_8\}$ *(indicated by curving arrows) fails to match. By this definition, 758 of 881 participants are time inconsistent, most of whom generate more than one mismatching pair of sequences.*

1-month appreciation intervals

sequence S_1	sequence S_2	sequence S_3	sequence S_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

1-year appreciation intervals

sequence S_5	sequence S_6	sequence S_7	sequence S_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 2: Empirical Distributions* for Choice Sequences S_1 (today v. one month from today) and S_4 (12 v. 13 months from now)

	Choice sequence S_1: today versus one month from now			Choice sequence S_4: 12 months versus 13 months from now		*empirical distributions among individuals who are consistent v. inconsistent*			
	choice sequences	# individuals		choice sequences	# individuals	# time-consistent	# time-inconsistent**	percentage among the consistent	percentage among the inconsistent
	00000	272		00000	171	109	62	36.1	10.7
	00001	160	*579 time inconsistencies*	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 sequences over 5 binary decisions. The left-most distribution for sequence S_1 is parameterized by arrival times of today and one month from today. The second distribution S_4 has arrival times of one year from today and 13 months from today. In both cases, the five choices represented by digits in the choice sequences indicate more patient choices as 1 in the digit position corresponding to annual rates of return of 5, 20, 50, 100 and 200 percent. The 302 individuals who chose two identical sequences are referred to as consistent (at least for this pair of choice sequences, S_1 and S_4). The 579 whose choice sequences did not match are referred to as time inconsistent. Frequency and percentage distributions for both subsamples are presented.

**Among consistent individuals, the empirical distributions are by definition identical in both choice sequences. Among the inconsistent, the distributions presented in the table show only the distribution of S_4 or destination choices to which participants shifted.

Table 3: Conflicting* Choices among Two Mean-Preserving Spreads

	choice 1		
choice 2	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 individuals on the off-diagonal cells above are not necessarily violating expected utility theory. However, of these 462 individuals, 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 4: Increases in Expected Dollar Payoffs among Time- and EU-Inconsistent Individuals

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.

Table 5: Regressions of expected payoffs on inconsistency, risk taking, and other controls (N = 881)

	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6	
	time and EU inconsistencies only		add risk control		add demographics		add math (self-described & tested)		add urban/rural		inconsistency-math interactions	
variables	coeff	t stat	coeff	t stat	coeff	t stat	coeff	t stat	coeff	t stat	coeff	t stat
TI	198.70	7.2	196.50	7.5	197.35	7.7	194.61	8.0	194.55	7.9	194.27	8.1
EU-violator	96.68	4.9	56.59	2.9	47.23	2.5	33.17	1.8	33.59	1.8	33.53	1.9
Individual σ			1.79	8.5	1.93	9.2	2.03	10.2	2.03	10.2	1.96	10.0
Under 25					28.89	1.2	26.55	1.1	26.36	1.1	37.13	1.6
Female					15.52	0.8	34.95	2.0	34.47	1.9	27.35	1.6
Immigrant					0.00	0.1	-0.02	-0.4	-0.02	-0.4	0.00	-0.1
Married					-16.46	-0.8	-25.80	-1.3	-27.72	-1.4	-15.60	-0.8
Ontario					-138.58	-3.3	-120.28	-3.0	-123.01	-3.0	-117.07	-3.0
British Columbia					-69.67	-1.5	-88.73	-2.0	-91.46	-2.1	-82.57	-1.9
Nova Scotia					-79.84	-1.8	-69.02	-1.6	-73.52	-1.7	-67.03	-1.6
Alberta					-124.09	-2.7	-104.07	-2.4	-104.95	-2.4	-105.24	-2.5
Native Person					-22.88	-0.4	-23.77	-0.5	-22.88	-0.5	-37.23	-0.8
Disabled					8.03	0.3	-5.73	-0.3	-4.34	-0.2	-9.16	-0.4
French Speaker					-22.95	-0.4	23.62	0.4	21.82	0.4	30.76	0.5
Burdened by Debt					-20.62	-1.1	-18.22	-1.0	-18.67	-1.1	-13.14	-0.8
Sell Asset to Pay Debt					-68.38	-2.4	-61.18	-2.2	-61.84	-2.2	-67.42	-2.5
Medium Household Income					48.80	2.2	24.18	1.1	25.48	1.2	23.37	1.1
High Household Income					90.45	3.7	60.43	2.6	62.45	2.7	54.03	2.3
Not Working					-41.58	-1.5	-21.59	-0.8	-22.26	-0.8	-12.05	-0.5
Completed High School					-132.13	-2.6	-89.79	-1.8	-90.52	-1.9	-86.78	-1.8
Liked School					-11.67	-0.5	-0.54	0.0	-0.56	0.0	2.38	0.1
Peers Liked School					46.07	2.2	38.52	2.0	39.92	2.0	34.81	1.8
Performed Well in School					57.07	3.0	13.78	0.8	13.84	0.8	17.14	1.0
Self-Described Math (7-point)							2.04	0.2	2.05	0.2	3.95	0.4
Math Testscore (std = 1)							88.11	9.7	87.94	9.6	225.13	9.9
Rural									13.18	0.5	6.30	0.3
TI_x_Math											-143.42	-6.5
EU-violator_x_Math											-28.64	-1.6
Constant	4342.17	160.2	4208.89	138.2	4208.08	61.5	4253.29	65.0	4253.34	64.9	4,247.10	66.3
R squared	0.0636		0.1575		0.2285		0.3048		0.3050		0.3406	

Table 6: Correlations among time-inconsistency, expected utility violator status, and math test scores

	TI	EU-violator	Math Testscore
TI	1		
EU-violator	0.1098*	1	
Math Testscore	0.0027	0.0783*	1

Table 7: 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 individuals 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.

Figure 1: Histograms of net number of time-inconsistent shifts toward patience as front-end delay increases (for fixed 1-month, fixed 1-year, and pooled appreciation intervals)

1-month treatments

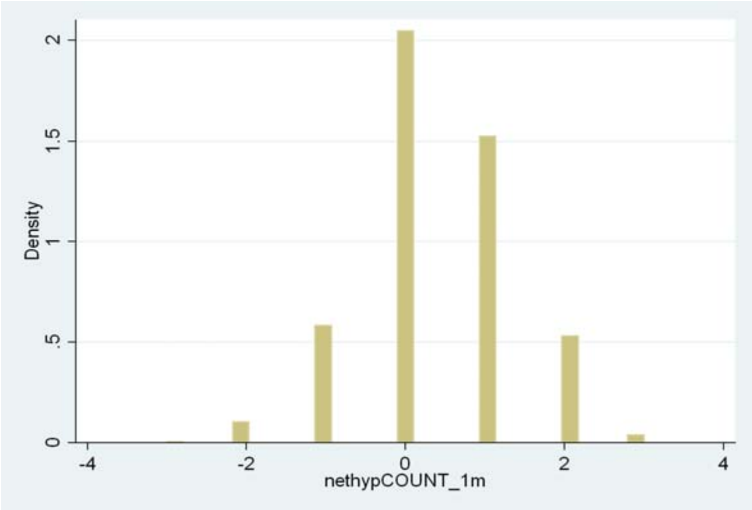

1-year treatments

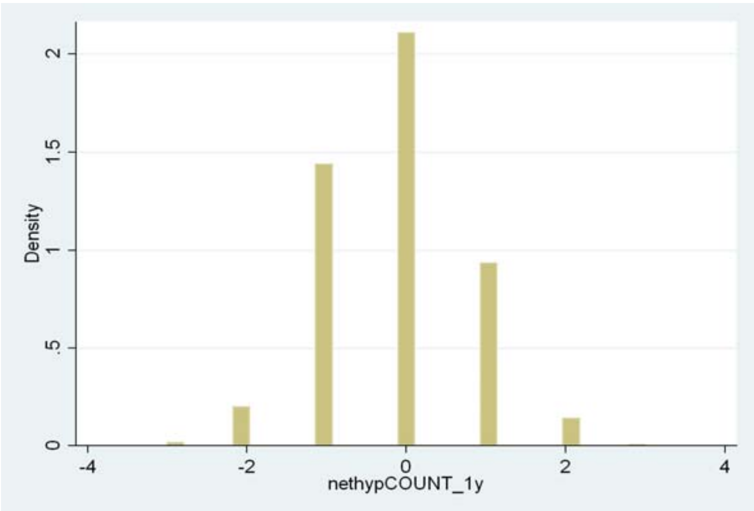

all treatments

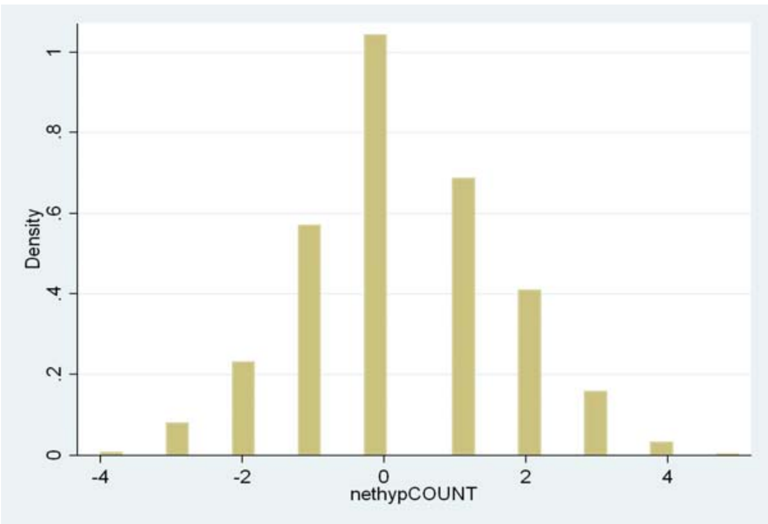

Figure 2: The Risk-Reward Envelope (consistent subjects represented by squares, and time-inconsistent and/or EU-violators represented by dots)

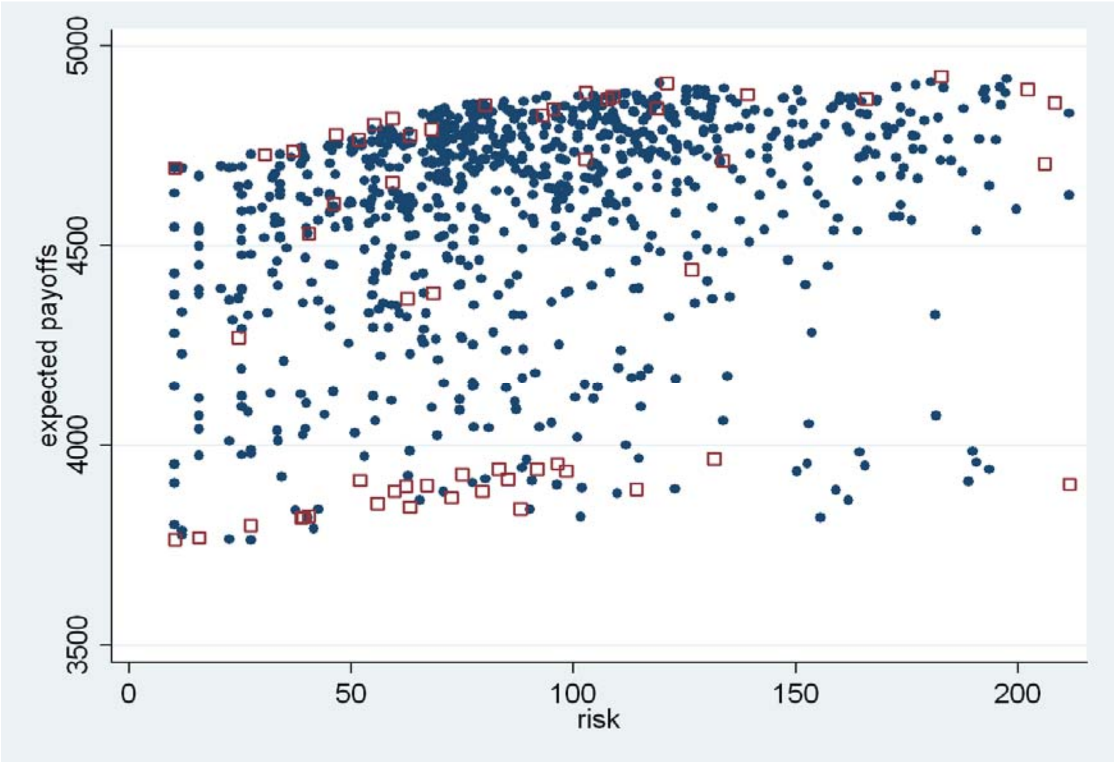

Figure 3: Empirical distribution of between-session switches in 72 time-tradeoff and risky choices

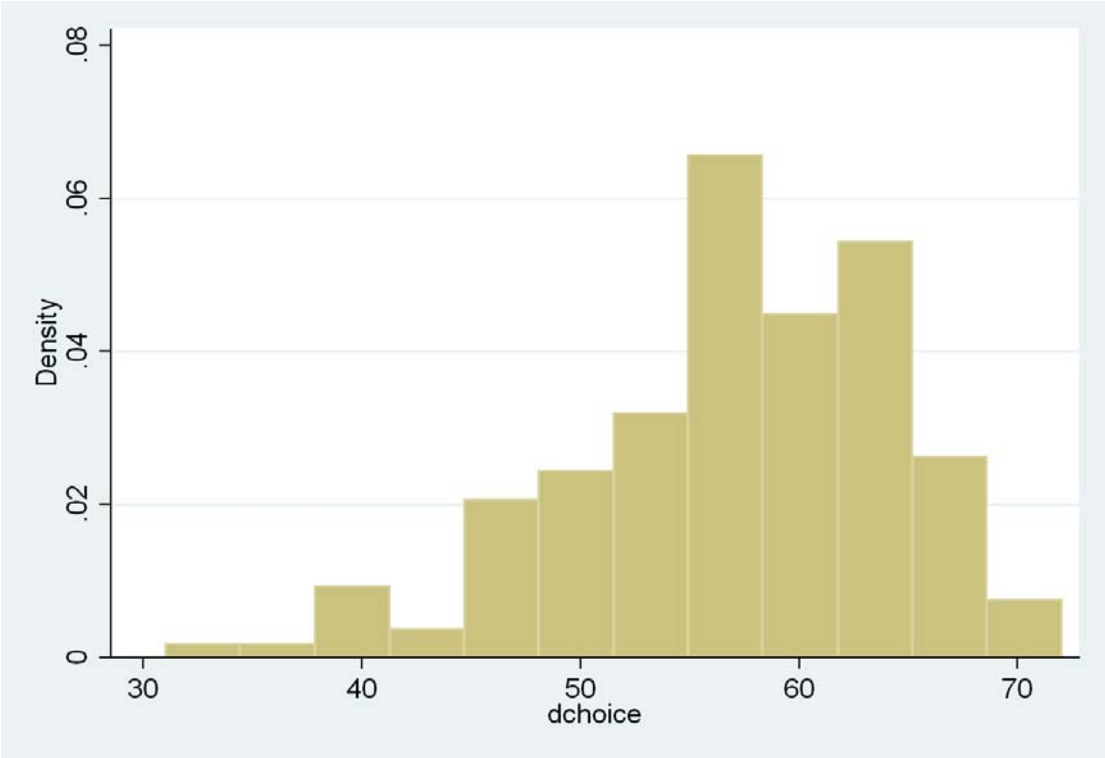

A single choice

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 t_1 0
C → later and larger payoff y arriving at t_2 1

Definition of a time-tradeoff sequence

b) Fix the sooner and smaller payoff x and the two arrival dates 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 said to be more patient than S' if, evaluating strings as integers, S > S'.

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 front-end delay is zero (t_1 = today), the second option's arrival date t_2 is one year from today, and the appreciation interval, t_2-t_1, is 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 different 5-choice sequences (40 binary choices in total), each with fixed x, and annualized rates of return: 5, 20, 50, 100 and 200 percent. The appreciation interval t2 - t1 = 1 year for four of the 5-choice sequences and 1 month for the other four sequences:

front-end delay	*appreciation intervals*	
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.

Appendix 2: Cross-Tabulation of Time Inconsistency in 1-Month versus 1-Year Appreciation Intervals

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

Appendx 3A: 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.

Appendix 3B: Frequency Distributions for Time Inconsistency, Hyper- and Hypo-bolic Discounting with Same-Day-Payoff Choice Items Removed*, not predicted by Laibson's β-δ quasi-hyperbolic discounting model

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-trade-off 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 disappear. Eliminating same-day-payoff choice items from consideration increases the number of time-consistent subjects by only 23, from 123 to 146.

**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 hypobolic discounting (allowing for no shifts in the opposite direction) given in the text.

Appendix 3C

Table 3A presents empirical distributions for count variables that tally instances of different directions for time-inconsistent shifts. Hyperbolic shifts are those where more patience is revealed, i.e., a smaller premium over the earlier and smaller payoff is accepted when choosing the later and larger payoff, when the front-end delay is longer. Hypobolic shifts are the opposite: from more patient at short front-end delays to less patient for longer front-end delays.

The first count variable in Appendix 3A tallies all forms of time-inconsistency, ranging from a minimum of zero (which occurs when $S_1 = S_2 = S_3 = S_4$ and $S_5 = S_6 = S_7 = S_8$) to a maximum of six (which occurs when S_1, S_2, S_3 and S_4 are all unequal, and S_5, S_6, S_7 and S_8 are all unequal). With a modal time-inconsistency count of 2 and more than 140 individuals with four to six time-inconsistencies, it is clear that the typical time-inconsistent individual is time-inconsistent more than once. The columns of Appendix 3A labeled "1-month" and "1-year" present empirical distributions for analogously defined count variables restricted to pairs in one-month (S_1-S_4) and one-year (S_5-S_8) accumulation intervals, respectively. These columns reveal that more than 50 individuals in each case (more than 100 unique individuals: 53 + 59 – 6 = 106) were maximally time-inconsistent in one of the two conditions.

Although the statistical association between payoffs and inconsistency is analyzed in detail in the main results sections of this paper, we note here (because these subsamples are not analyzed separately below) that individuals who were maximally time-inconsistent in the 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 associations between status as being maximally time-inconsistent and payoffs 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 an annualized rate of return of 50% or more is required to induce waiting one year as opposed to receiving the earlier payoff today, but only 5 or 20% is required to induce waiting one year from tomorrow as opposed to receiving the sooner payoff tomorrow. In both cases, the wait is precisely one year. In theoretical models of hyperbolic and quasi-hyperbolic discounting, it matters whether that wait begins sooner rather than later.

Counts on hyperbolic and hypobolic shifts show several interesting patterns. Although slightly more hyperbolic shifts were observed in the one-month accumulation interval, slightly more hypobolic shifts were observed in the one-year accumulation interval, as seen in Appendix 3A and in Figure 1. Thus, no clear directional pattern of time-inconsistent shifts emerges from our data.

The three columns in Appendix 3A under the heading "NetHYP_count" count the number of hyperbolic shifts minus hypobolic shifts. The high frequency of individuals 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 individuals have 4 hypobolic shifts, whereas under the heading NetHYP_count one finds only 2 individuals that have a net count of -4 (i.e., 4 hypobolic shifts). The apparent discrepancy reflects the fact that 3 among the 5 individuals (who made 4 hypobolic shifts) also made one or two hyperbolic shifts, which changed their NetHYPcount values from -4 toward zero.

The final two columns of Appendix 3A 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 individuals, only 211 are pure hyperbolic shifters and

144 are pure hypobolic shifters, which leaves 403 who shift in opposite directions at least once.

Our data provide no empirical support for hyperbolic discounting as a general explanation for observed time-inconsistency. Figure 1 shows three empirical distributions over individuals' counts of hyperbolic shifts minus the number of hypobolic 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 then examine whether the choice items with options for same-day payoffs drive most instances of time-inconsistency as predicted by quasi-hyperbolic discounting; and whether the more general hyperbolic discounting theory enjoys any empirical support after focusing exclusively on time choices involving same-day payoffs or after eliminating these pairs from consideration when counting instances of time inconsistency. Comparing Appendices 3A and 3B we find that the data do not provide evidence in support either of these theories. Time-inconsistencies are *not* mostly generated by same-day choice items. After removing same-day-payoff choice items, most time-inconsistency *cannot* be explained as individuals who follow a pattern of making hyperbolic shifts. Eliminating same-day-payoff choice items from consideration increases the number of perfectly time-consistent subjects by only 23, from 123 to 146. These 23 are those whose only time inconsistency involved same-day payoffs. Removing pairs of choice sequences with same-day-payoffs increases the number of pure hyperbolic discounters by only 28, from 211 to 239. The majority of time-inconsistent individuals in both Appendices 3A and 3B shift at least once in both directions.

Appendix 4: 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.

Appendix 5 Cross-Tabulation of Time-Inconsistency and Expected Utility Violations

	expected utility		
time	consistent	inconsistent	Total
consistent	62	61	123
inconsistent	266	492	758
Total	328	553	881

Appendix 6: Empirical Distribution of Self-Described Math Skills on 7-point Scale

Value	Freq.
-3	18
-2	31
-1	205
0	478
1	116
2	25
3	8