Estimating Child Time Preferences: Aiding Rural Schools in Improving Capital Formation

Jeff Jordan
Department of Agricultural and Applied Economics, University of Georgia, jjordan@uga.edu
Marco Castillo
School of Public Policy, Georgia Institute of Technology, marco.castillo@pubpolicy.gatech.edu
Paul J. Ferraro
Regan Petrie
Department of Economics, Andrew Young School of Policy Studies, Georgia State University, pferraro@gsu.edu, rpetrie@gsu.edu

Selected Paper prepared for presentation at the American Agricultural Economics Association Annual Meeting , Orlando, FL, July 27-29, 2008

1. Introduction

Most important economic decisions involve inter-temporal tradeoffs. Thus an understanding of time preferences, typically modeled through discount rates, is central to economic theory and empirical analyses. In fact, the rate and form of discounting are often among the most important parameters that affect outcomes in models of individual or social decision-making. Although studies have estimated discount rates from revealed preferences among adults (see next section), little work has been done to infer such rates among children (a notable exception is Bettinger and Slonim, 2007). The absence of such work would cause few problems if children's discount rates were similar to adults' discount rates, or if important inter-temporal decisions are not made until the age of eighteen.

Clearly, however, individuals make important inter-temporal decisions before they turn eighteen. Should they drop out or stay in school? Should they keep the baby or give it up for adoption? Understanding the time preferences of children, particularly adolescents, is thus critical for understanding human capital formation and economic outcomes, as well as for formulating policies that affect children. For example, students who discount the future more heavily should be less likely to invest in school-related efforts relative to other uses of their time that provide more immediate benefits (e.g., leisure).[1] Thus policies targeted towards increasing immediate returns to children with high discount rates may improve educational outcomes more than alternative policies.

Whether children's discount rates are similar to adults' discount rates is an empirical question. Although psychologists, and a few economists, have studied children's patience (see Eigsti et al., 2007, for references of studies in psychology), our paper is the first to formally

[1] Becker and Mulligan (1997) argue that discount rates may be endogenously determined by factors that include education. More education makes one more capable of envisioning the future, which lowers the discount rate.

estimate subjective discount rates among adolescents from a broad range of backgrounds. Moreover, we examine the factors that affect the heterogeneity of discount rates and explore the degree to which discount rates are stable or well-formed among adolescents (the latter was also done in Bettinger and Slonim, 2007). We also extend the literature on eliciting discount rates experimentally by introducing an alternative way of drawing inferences about discount rates from observations of "inconsistent/irrational" decisions, which are often observed in time preference experiments.

We find that: 1) the mean annual discount rate from our sample of two-hundred 8th graders is approximately double the estimates of mean adult discount rates; 2) there is substantial heterogeneity of discount rates, with the first quartile having annual discount rates lower than 35%, and the top quartile having rates greater than 100%; 3) adolescents from poor and non-poor households do not differ in their discount rates, but there is evidence of heterogeneity across sex, race and cognitive skills; and 4) the proportion of adolescents who exhibit inconsistent preferences is much higher in our sample than in comparable adult samples.

2. Experimental Design

In the economics literature, four methods have been used to estimate discount rates among adults. Three are revealed preference methods: 1) econometric estimation from observations of the use of financial instruments (e.g., Ausubel 1991) or of the purchase of durable consumer goods (e.g., Gately, 1980; Hartman and Doane, 1986; Hausman 1979; Ruderman et al., 1986); 2) natural experiments in which individuals are forced to choose among alternative payoffs with differential time dimensions (e.g., Warner and Pleeter (2001), which took advantage of data generated from an early retirement program in the U.S. military to estimate discount rates for

enlisted men and officers); and 3) controlled experiments in which subjects are offered real monetary payoffs that vary in their timing (Holcomb and Nelson, 1992; Pender, 1996; Coller and Williams; 1999; Harrison et al., 2002; Eckel et al., 2003; Meier and Sprenger, 2006; Bettinger and Slonim, 2007). Stated preference methods, in which discount rates are elicited by asking individuals to make hypothetical choices in the revealed preference settings described above, are also used (Thaler, 1981; Loewenstein, 1988; Benzion et al., 1989; Shelley, 1993; Curtis 2002; Bradford et al. 2004).

Given the potential sources of bias inherent in stated preference methods, and the difficulty in observing the consumption and investment decisions of children, we use a controlled experiment. Psychologists, and more recently, economists, have used experiments to study time preferences among children. However, these studies look at the factors that affect "patience," which is defined as a binary choice to forgo short-term benefits for larger and longer-term rewards. None of the studies explicitly define and characterize discount rates. To do this, we adopt the front-end delay method used by Harrison et al. (2002).

In our experiment, subjects are asked, orally and in writing, to make twenty decisions in total. For each decision, subjects are asked if they would prefer $49 one month from now or $49+$X seven months from now. The amount of money, $X, is strictly positive and increases over the twenty decisions. The decision sheet that the subject sees is shown in Table 1 without the last column indicating the implied annual discount rate. For example, in the first decision, a subject is asked if she would prefer $49 one month from now or $50.83 seven months from now. And, in the ninth decision, a subject is asked if she would prefer $49 one month from now or $67.61 seven months from now. Subjects are asked to make one choice for each of the twenty decisions on the decision sheet. Based on discussions with teachers and students at other schools,

we determined that the range of $50 to $99 would be considered by adolescents to be "large" payoffs, but not so large as to potentially cause problems with their parents.

Table 1: Subject Decision Sheet
(The last column of implied annual discount rates was not shown to subjects)

Decision	Paid one month from now	Paid seven months from now	Implied Annual Discount Rate
1	$49.00	$50.83	7.35%
2	$49.00	$52.71	14.7%
3	$49.00	$54.66	22.05%
4	$49.00	$56.66	29.40%
5	$49.00	$58.72	36.75%
6	$49.00	$60.85	44.10%
7	$49.00	$63.04	51.45%
8	$49.00	$65.29	58.80%
9	$49.00	$67.61	66.15%
10	$49.00	$70.00	73.50%
11	$49.00	$72.46	80.25%
12	$49.00	$74.99	88.20%
13	$49.00	$77.59	95.55%
14	$49.00	$80.27	102.90%
15	$49.00	$83.03	110.25%
16	$49.00	$85.86	117.60%
17	$49.00	$88.78	124.95%
18	$49.00	$91.77	132.30%
19	$49.00	$94.85	139.65%
20	$49.00	$98.02	147.00%

Harrison et al.'s (2002) decision sheet includes the implied annual interest rate and annual effective interest rate associated with each delayed payment option. However, our discussions with teachers at the study site and with similar aged students at other schools led us to believe that students do not price field investments in terms of interest rates. Thus information on rates would simply confuse students. Coller and Williams (1999) and Harrison et al. (2002) also argue that one should elicit the market rates of interest that subjects face so that one can control for arbitrage opportunities (field censoring) in the econometric analysis. Because our subjects are children, we feel comfortable assuming that subjects do not incorporate credit

market options into their experimental decision task. If subjects were to have access to credit markets, and these interest rates were binding in the experiment, our estimates would be lower bounds on the true discount rates.

Economic theories of discounting predict that an individual faced with the decision sheet in Table 1 would either choose (a) $49 for all decisions, (b) the higher payment for all decisions, or (c) $49 for a certain number of decisions starting with Decision 1 and then switch to the higher payment for the remaining decisions. In other words, if an individual chose to receive $Y in seven months rather than $49 in one month, then the individual will prefer any amount $Z > $Y in seven months rather than $49 in one month. Following Harrison et al. (2002), we call these individuals "consistent" decision-makers (Bettinger and Slonim called them "rational").

However, in experiments using decision sheets like the one in Table 1, some individuals are "inconsistent" decision-makers: they choose $Y in seven months rather than $49 in one month, but then choose $49 in one month rather than $Z>$Y in seven months. Harrison et al. (2002) and Meier and Sprenger (2006) found that 4% and 11%, respectively, of their adult subjects were inconsistent in their choices. Bettinger and Slonim (2007), whose subjects were between 5 and 16 years old, found that 34% of their sample were inconsistent decision-makers.[2] We return to the problem of inconsistent decision-makers in Section 4.

In each room, subjects are assigned a unique identification code. This code is private, and subjects do not know the identification codes of other subjects. Subjects make their decisions by circling one amount, either $49 or $49+$X, on their decision sheet. After subjects make their decisions, each subject puts her decision sheet in an envelope and the envelopes are collected.

[2] The actual rate may be a little higher because the first decision offers a subject $10 in two months or $9 in four months, and 4% of the sample (classified as "always patient") chose to receive the lower value in four months.

One decision out of the twenty decisions is randomly chosen for payment. This is done by taking 20 index cards with the numbers 1-20 written on them, shuffling them in front of the subjects, and asking a subject to choose one card. The number on the card is the decision number to be paid for each of the six subjects who are chosen to receive payment. So, for example, if decision 15 is chosen for payment and one of the winning subjects circled $83.03, the subject would receive $83.03 in seven months. If another subject circled $49, that subject would receive $49 in one month.

After determining the decision to be paid, all the envelopes are shuffled in front of the subjects, and six envelopes per room are chosen for payment. The identification codes of those chosen to receive payment are written on the blackboard. Because identification codes are kept private by each subject, no other subject knows which subject have been chosen to receive payment. Subjects who are chosen to receive payment are paid with a Wal-Mart gift card by the school principal on the specific date for the decision chosen. The school principal keeps in her office the Wal-Mart gift cards and the names of the subjects who are chosen for payment. Within a week of the experiment, the subjects stop by the principal's office to verify winning the gift card. On or within a week of the payment date, the subjects go privately to the principal's office to pick up their gift cards. For the subjects chosen to be paid, their names and the amount of payment are kept private. Subjects know all of these procedures before making their decisions.

In our experiment, 202 8^{th} grade students participated (ages 13 to 15). The subjects were divided into four different rooms, so there were roughly 50 subjects in each room. The average payment was $69.94 (SD = $19.46), with a total pay-out of $1,657.01. Nine of the twenty-four chosen students received gift cards of $49 one month after the experiment. Seven months after the experiment, five students received cards of $67.61, four received $72.46, and six received the

highest pay-out of $98.02. All experiments were conducted on the same day (September 16, 2006). Subject and study site characteristics are detailed in the next section.

Although Bettinger and Slonim (2007) use a similar experimental approach with children, their analysis is different. Rather than estimate discount rates for each subject, they treat each decision separately and determine the factors that make the choice of Option B more likely in each decision (e.g., boys are 10% less likely than girls to choose Option B). Their analysis uses a probit model in which the dependent variable equals one if Option B is chosen for that decision. Thus discount rates are not explicitly defined, and the problems of inconsistent decision-makers and censoring (34% and 5%, respectively, of the sample) are not relevant. Moreover, their sample is restricted to children on free or reduced lunch, whereas our sample represents a broader socio-economic range of children.

3. Subject Pool

The setting for our study is Spalding County, Georgia, located on the southern end of the Atlanta MSA. Although part of the vibrant metropolitan area of Atlanta, demographic data on Spalding County resembles less the exponential growth of the Atlanta area and more the persistently poor counties of southern Georgia. In Spalding County, the child poverty rate is 21.7% (17.1% in Georgia) and per capita income is $16,791 ($21,154 in Georgia).

Thirty-two percent of persons over 25 in Spalding County have not completed high school in 2000 — over 50% higher than for Georgia. Only 8% of adults completed a Bachelor's degree or higher (24% in Georgia). The high school dropout rate in school year 2000/01 was 15.6 per 100 enrolled, more than double the state average of 6.4. Less than half (46.8%) of the

class of 2001 that entered in ninth grade, graduated (71.1% rate in Georgia). By 2004 the official non-completion rate was 46%,

Our experiment was conducted at Cowan Road Middle School, one of four middle schools in the Spalding County School District. Whereas Bettinger and Slonim (2007) focused on free and reduced-price lunch students ranging from five to sixteen years of age (n=191), we focus on a broader range of socio-economic backgrounds but a narrower age range (n=202). At the time of the experiment, 95% of our subjects were 13 or 14 years old (mean=13.90, SD=0.56), while the remaining 5% were 15 years old. We chose this age group because, nationally, 35% of students lost in the high school pipeline are lost at the end of 9th grade and in Georgia, students can make the decision to drop out of school at the age of 16. Thus, we wanted to elicit discount rates in the period prior to when this important decision would be made.

Table 2. Summary Statistics

	FULL SAMPLE (N=202)		**CONSISTENT SUBJECTS** (N=117)		**INCONSISTENT SUBJECTS** (N=85)	
Variables	***N***	***Percent of Sample***	***N***	***Percent of Group***	***N***	***Percent of Group***
Males	87	43.07%	54	46.15%	33	38.82%
Black Males	34	16.83%	14	11.97%	20	23.53%
White Males	48	23.76%	36	30.77%	12	14.12%
White Females	44	21.78%	27	23.08%	17	20.00%
Black Females	60	29.70%	27	23.08 %	33	38.82%
Gifted	16	7.92%	13	11.11%	3	3.53%
Special Education	41	20.30%	21	17.95%	20	23.53%
*Poor Reader**	44	21.78%	20	17.09%	24	28.24%
*Poor Math**	26	12.87%	10	8.55%	16	18.82%
Free Lunch	112	55.45%	57	48.71%	55	64.71%
Reduced Lunch	25	12.38%	15	12.82%	10	11.77%

* Scoring level 1 on CRCT exam "does not meet standards."

Table 2 presents summary statistics of the students' characteristics.[3] The table presents statistics for the full sample, as well as broken down by those who made consistent and inconsistent choices in the experiment. Slightly less than half of the students are male, and 93% are either Black or White.[4] Because we have few observations on racial/cultural groups other than Whites and Blacks, we focus our analysis on Black-White differences. Sixty-eight percent of students are eligible for free or reduced-price lunches, which is a proxy for family income. Twenty percent are enrolled in special education courses designed for students with learning disabilities, and 8% are enrolled in a gifted learning program. Twenty-two percent of students failed to meet the state standards in reading at the end of the previous academic year, and 13% failed to meet the standards for mathematics. Nine percent failed to meet both standards. We examine the differences between consistent and inconsistent decision-makers in Section 5.

4. Consistent Adolescents

4.1 Summary Statistics: Consistent Adolescents

We begin the analysis by focusing on the decisions of consistent decision-makers. These results are comparable to studies that drop inconsistent decision-makers from the analysis (e.g., Harrison et al., 2002; Meier and Sprenger, 2006). In the next section, we take a closer look at the inconsistent decision-makers, introduce two ways in which inferences can be drawn about their discount rates, and analyze the entire sample.

Figure 1 displays the distribution of elicited annual discount rates. The X axis lists the lower bound of the interval. Thus "0" implies the discount rate is between 0% and 7.5%, while

[3] We do not break down our results by age in any of the subsequent analyses because, given the lack of variability of this variable in our subject pool, it has no effect on any of the outcomes of interest.

[4] The sample also includes three Asian females (two are consistent), one Asian male (inconsistent), one Native-American female (consistent), four Hispanic females (all consistent), four multi-racial males (all consistent), and three multi-racial females (two are consistent).

"147" implies the discount rate is something greater than or equal to 147%. There is substantial heterogeneity of discount rates. The bottom quartile has discount rates less than or equal to 29.4%, while the top quartile has discount rates greater than or equal to 102.9%. Half the subjects have discount rates above 44.1%.

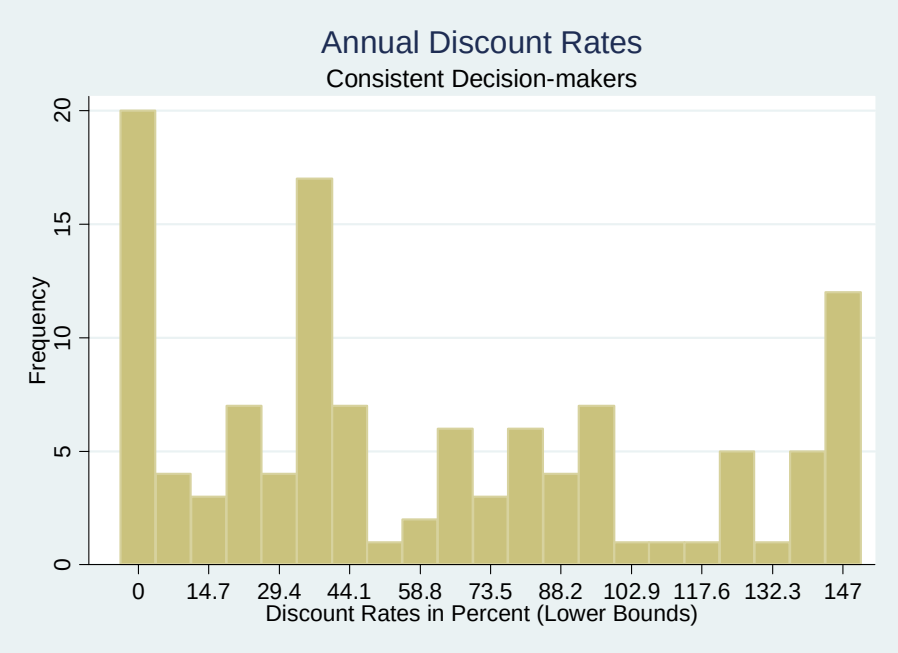

Figure 1. Annual Discount Rates: Consistent Decision-makers

Table 3 presents estimates of unconditional, mean discount rates from our sub-sample of consistent decision-makers. To estimate these means for the overall sample and sub-groups of consistent decision-makers, we use an interval regression model that accounts for the interval nature of the data and censoring at the 20th decision. The mean discount rate is 66.6%, which is substantially higher than the mean discount rates estimated from adult populations, whether they were estimated through experimental or non-experimental methods. For example, Harrison et al., using similar experimental and econometric methods, estimate the mean discount rate of Danish adults to be about 30% (Coller and Williams, 1999, report lower rates). Hausman (1979), using

observations of durable good purchases, infers mean discount rates in the range of 15%-26%, depending on the assumptions and econometric model used.

Table 3. Mean Discount Rates for Consistent Decision-makers
Interval Regressions

GROUP	MEAN DISCOUNT RATE
Overall	66.60%
Males	77.31%
Black Males	95.50%
White Males	70.06%
Females	57.38%
Black Females	54.98%
White Females	58.29%
Gifted	40.99%
Special Education	50.71%
Poor Reader	38.59%
Poor Math	53.26%
Free or Reduced Lunch[1]	65.53%

[1] For ease of exposition, we combine free and reduced lunch into a single variable because neither played a significant role separately in any of our subsequent analyses.

We also find interesting differences, and lack of differences, among different sub-groups. Males have higher discount rates than females. Black males have the highest discount rates, which are substantially higher than White males. There is little difference, however, between Black and White females. Likewise, there is little difference between the overall mean discount rate and mean discount rate of adolescents receiving free or reduced fare lunches. Gifted students have lower than average discount rates, but so do special education students, albeit not as low as gifted students.

In their study using a similar experimental procedure, Bettinger and Slonim (2007) also find that boys are less patient than girls (i.e., boys are less likely to wait for future rewards), but note that the evidence on gender effects on discount rates among adults is mixed (most studies find no effect). Duckworth and Seligman (2005), using a delay choice task for 8th graders, also

find girls to be more patient. Bettinger and Slonim (2007) find that family income has no effect on their measure of children's patience. However, in contrast to our results, they find no relationship between taking gifted classes and their measure of patience.

The lowest estimated discount rate is among students who failed to meet the state standards in reading. In fact, the mean discount rate for this sub-group is about half the mean for students who meet or exceed the standards (72.59 percent). Why would this be the case? Human capital theory suggests that students with high discount rates would invest less in building up skills such as reading (Lang and Ruud, 1986). Considering the other causal direction, Becker and Mulligan (1997) argue that better educated people should have lower discount rates because they can better imagine the future. Although the low estimated discount rates for gifted children affirm these hypotheses, how might one interpret our finding that low discount rates are associated with poor reading skills and that math skills are uncorrelated with elicited discount rates? It may imply a problem with experimentally eliciting discount rates. If subjects have poor verbal comprehension skills (oral and written), they may be less likely to understand the experimental decision-making environment. For example, a subject with poor verbal comprehension skills may simply be attracted to the higher numbers in option B and thus more likely to switch earlier regardless of their time preferences. More work is needed to understand the role of verbal comprehension skills in experimental elicitation of discount rates, but in the regression analysis below we control for such skills by using a dummy variable for subjects who failed to meet the state standards in reading.

4.2 Regression: Consistent Adolescents

To further explore the factors that affect discount rates among the adolescents in our sample, we run two regression models: an interval regression model and a probit model.[5] The interval regression model uses all twenty discount rate intervals in the decision sheet and takes into account the censoring in the final interval (>147%).[6] The probit model defines the variable "Impatient," which is equal to one if the student's discount rate is above 80% and equal to zero otherwise (i.e. switched at Decision 10 or later). We chose 80% as a cutoff because it is a high rate and reflects the second-half of the decision table. The results are not sensitive to changing the threshold by 7.35% in either direction. We also defined a variable with four categories – very patient (0-36.75%), patient (36.75-73.5%), impatient (73.5-110.25%), very impatient (110.25% - ∞) – and ran an ordered probit analysis. The results were not qualitatively different from the binary probit analysis and thus are not reported here.

In all models, we use the following (binary) covariates: male, black, gifted, poor reader (i.e. failed to meet state standards), free or reduced lunch. A dummy variable for special education students added no further explanatory power to the model when poor reader was in the model and thus is not included in the results reported below. Similarly, a dummy for failing to meet the state standards in math is excluded (70% of math poor students are also reading poor). We control for fixed session (room) effects in all regression models (not controlling for these effects has little effect on coefficient estimates). Although the inclusion of interaction terms

[5] We were concerned that the overdispersion of the discount rate distribution in Figure 1 would be a problem in the interval regression model. We thus also estimated a generalized linear model that uses the decision switch points as the dependent variable (0 to 20, with 0 indicating a switch at the first decision and 20 indicating only option B was chosen), with a Poisson family distribution, and a scaling of the standard errors using the square root of the Pearson chi-square dispersion (to compensate for the overdispersion). The results are close to those from the interval regression and change none of our conclusions.

[6] There is no censoring at the lowest interval (0-7.35 percent). Discount rates, in principal, cannot assume negative values. In other words, the first interval is not a limit interval and thus there is no qualitative difference between the first interval and the next eighteen intervals.

generated coefficients that reflect the differences observed in Table 3, no interaction term was significantly different from zero at conventional significance levels. Thus we do not include interaction terms in our presentation of the results. We group the Asians with the Whites because their behavior in the experiments is similar and because their academic performances are known to be similar. We create a single dummy variable, "Other race/culture," for the other cultural groups. Our results are unaffected if we instead separate them, or include Asians in "Other race/culture."

Table 4. Factors Affecting Discount Rates: Consistent Decision-makers

	PROBIT REGRESSION[1] (IMPATIENT= 1 IF DISCOUNT RATE>80%)			**INTERVAL REGRESSION[2]**	
Covariate	***Marginal Effect (SE)***	***p-value***		***Marginal Effect (SE)***	***p-value***
Male	0.139 (0.095)	0.143		21.48 (9.37)	0.022
Black	0.302 (0.126)	0.017		14.29 (12.09)	0.237
Other race/culture	0.134 (0.188)	0.472		17.15 (17.37)	0.324
Gifted	-0.309 (0.088)	<0.001		-35.32 (15.06)	0.019
Poor Reader	-0.263 (0.092)	0.004		-40.17 (12.49)	0.001
Free/Reduced Lunch	-0.229 (0.116)	0.048		-12.17 (11.08)	0.272
	Controls for Fixed Session Effects			Controls for Fixed Session Effects	

[1] Marginal effects estimated for discrete change of variables from 0 to 1 evaluated at the means. The Likelihood Ratio chi-squared statistic is 18.72, and McKelvey & Zavoina's R2 is 0.256.
[2] The constant in this model is 80.98 (SE: 12.65; p<0.001). The Likelihood Ratio chi-squared statistic is 24.20 and McKelvey & Zavoina's R2 is 0.171. Twelve observations are right-censored.

The estimated marginal effects and their corresponding standard errors and p-values are reported in Table 4. Looking at the probit regression, Black students are 30% more likely to be impatient (discount rate greater than 80%), holding other characteristics constant. Males are 14% more likely to be impatient, but this effect is not precisely estimated and is not significantly

different from zero in our model. Gifted students, in contrast, are 31% *less* likely to be impatient. Students who failed to meet state standards in reading and students on free and reduced lunch are also less likely to be impatient (26% and 23%, respectively).

Looking at the interval regression, the precision with which we estimate the marginal effects are different in some cases from the probit model, but the qualitative interpretation of the marginal effects are similar. Holding other factors constant, males have significantly higher discount rates (by 22%), while gifted students and poor readers have significantly lower discount rates (35% and 40%, respectively). Black students have higher discount rates and free and reduced lunch students have lower discount rates, but the effects are not significantly different from zero for both.

5. Inconsistent Adolescents

5.1 Who are the Inconsistent Decision-makers and What are They Doing?

As noted in Section 2, Bettinger and Slonim's (2007) results suggest that children are more likely to be inconsistent in their inter-temporal decisions than adults. Our results support this conclusion. In our sample, 46 percent of the students make inconsistent choices. Looking at the simple statistics in Table 2, one can see that fewer than 20% of the gifted children were inconsistent, whereas 44% of the other students were. Less than one-third of Whites were inconsistent, whereas just over half of Blacks were. There are no male-female differences. Figure 2 shows the distribution of inconsistent choices by sex and race. About 45% of black boys and black girls made no inconsistent choices and 75% of white boys and 61% of white girls made no inconsistent choices. Blacks, and particularly black boys, make more inconsistent choices than non-blacks.

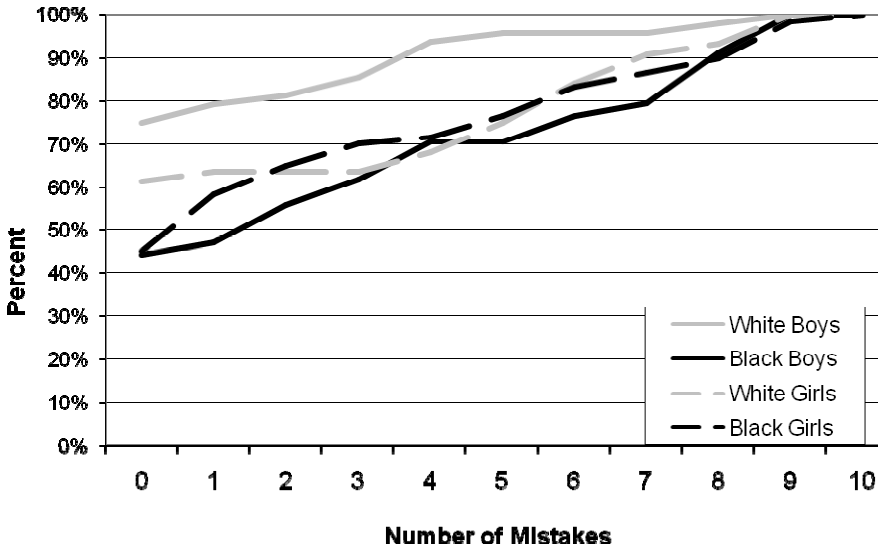

Figure 2: Distribution of Inconsistent Choices

We estimated several probit model specifications that regress a binary variable, which equals 1 if the subject was an inconsistent decision-maker, on demographic characteristics from Table 2. The only variable that consistently showed a significant effect was being Black. Black adolescents were about 20% more likely to be inconsistent. This result contrasts with that of Bettinger and Slonim (2007), who find in some of their specifications that Blacks are less likely to be inconsistent. In some of our specifications, but not all, being gifted was marginally significant ($p<0.15$) and negative (19% less likely to be inconsistent), while a specification that included a dummy variable for being poor in both reading and math had a marginally significant ($p<0.10$) and positive effect (22% more likely to be inconsistent). The separate effects of poor reading or math skills were not significant.

5.2 What Can We Say About the Discount Rates of Inconsistent Decision-makers?

As noted in Section 4, many studies simply drop inconsistent subjects from the analysis. An alternative is to keep them in the sample, and use their first switch point as indicating their true preferences. [7] This is akin to the outcome that would arise in experiments that ask a subject to simply pick one switch point (thus assuming monotonicity of the underlying preferences and automatically filling out the remaining choices; e.g., reference). [8] However, as can be seen in Figure 2, a substantial number of the inconsistent subjects made more than two switches.

An alternative approach, and a superior one in our view, is to minimize the costs of making an inconsistent subject's choices consistent. This is done by calculating the cost to switch decisions to all possible consistent choice profiles over the twenty decisions. The cost to switch a decision is the absolute value of the difference between $49 and the higher value. The rational-choice profile that costs the least is the profile that is chosen as the minimum cost. Note that the cost to make a student rational will vary depending on the number of inconsistent choices the student made and also at what point the student made inconsistent choices. Mistakes are more costly when the payoff in seven months is larger.

Figure 3 shows the distribution of costs to make students consistent. The costs are zero for consistent decision-makers. Costs differ by sex and race. The distribution of costs for black boy's first-order stochastically dominates that for white boys. The same holds for black girls compared to white girls. The cost distribution for black boys is significantly higher than for white boys, but there is no significant difference across girls (Kolmogorov-Smirnov test).

[7] In all of the subsequent analyses, we drop one (male) subject, labeled as inconsistent in our analysis, because he did not answer most of the choices, and for the choices he did answer, he often circled both options.

[8] An alternative approach is to take use the first switch point in a regression with a dummy variable for inconsistent decision-makers (Meier and Sprenger).

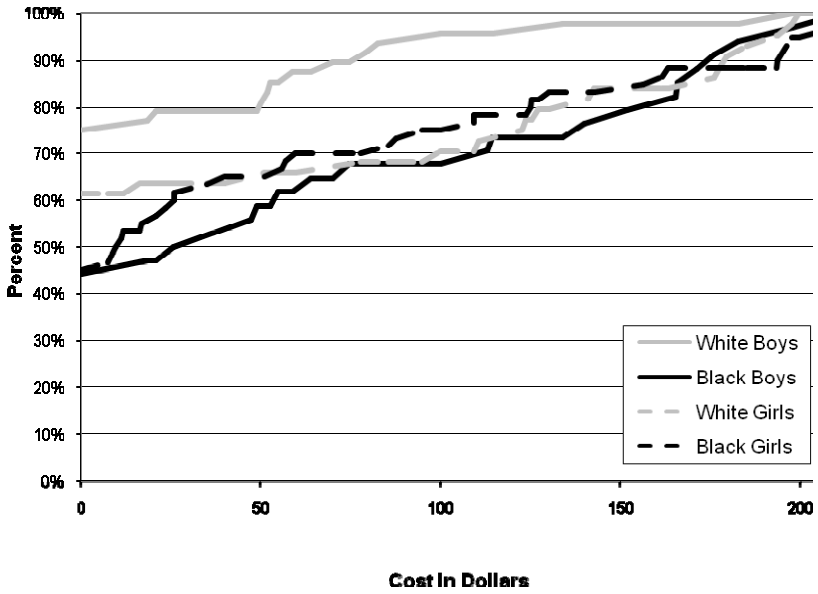

Figure 3: Distribution of Costs to Rationality

5.3 Summary Statistics: Full Sample

We now analyze the entire sample, which combines the consistent decision-makers and the inconsistent decision-makers. We re-do the analysis of Section 3 using discount rate intervals that are based on the subject's first switch point or based on the interval that minimizes the cost of making a subject consistent. We expect very different results from these two methods of handling inconsistent decision-makers. The mean estimated discount rate of the inconsistent decision-makers is 23.0% using the first switch point, and 97.4% using the cost-minimizing adjustment.

Figures 4 and 5 display the distribution of elicited discount rates. As in Figure 1, there is substantial heterogeneity of discount rates. As expected, the distribution in Figure 4 is substantially shifted to the left, with less than one-third of subjects with discount rates above 50%, and the distribution in Figure 3 is shifted to the right, with more than half of the subjects (58%) with discount rates above 50%.

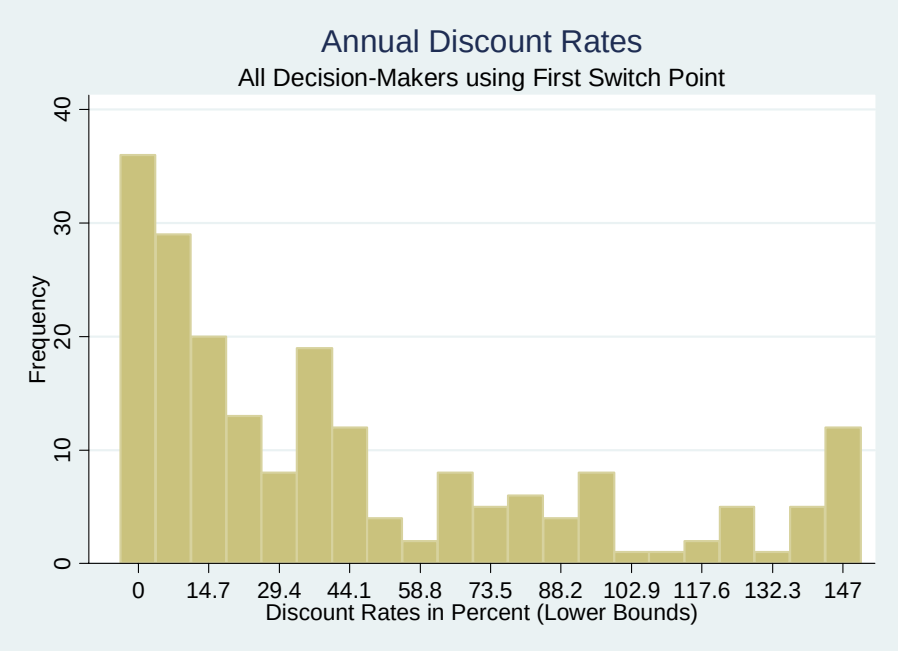

Figure 4. Annual Discount Rates: All Decision-makers (using first switch point)

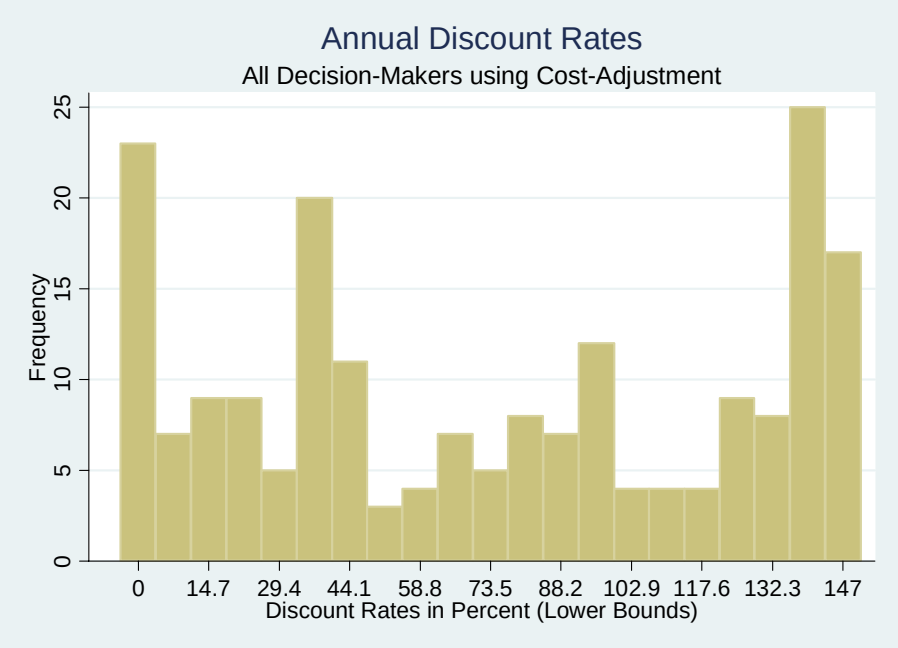

Figure 5. Annual Discount Rates: All Decision-makers (using cost-adjustment)

Table 5 presents estimates of unconditional, mean discount rates for the entire sample using the different methods of handling inconsistent decision-makers. The two methods come to quite different conclusions about the mean discount rates as well as the differences among sub-

groups. Using the first switch point, the mean estimated discount rate is 48%, which, although on the high end, is not dramatically different from mean estimates of adult discount rates. Using the cost-minimizing adjusted discount rates, however, the mean estimated discount rate is 79.6%.

In contrast to the results in Table 2, the results using the first switch point suggest that black males have lower discount rates than white males, and black females have lower discount rates than white females. Gifted children and poor readers, however, continue to have lower than average discount rates, but the gifted children are no longer substantially below average.

Table 5. Mean Discount Rates for Pooled Sample (Consistent and Inconsistent)

	MEAN DISCOUNT RATE	
Group	***Using First Switch Point***	***Using Cost-Minimizing Adjustment***
Overall	48.01%	79.64%
Males	57.67%	85.90%
Black Males	51.51%	100.66%
White Males	60.80%	76.14%
Females	40.78%	74.94%
Black Females	37.69%	78.33%
White Females	41.48%	71.96%
Gifted	44.10%	53.75%
Special Education	34.85%	67.90%
Poor Reader	28.63%	61.79%
Poor Math	34.45%	70.69%
Free or Reduced Lunch	44.89%	82.04%

The cost-minimizing adjusted discount rate estimates are much closer to those in Table 2, albeit higher. Black males still have the highest discount rates, and the rates of gifted children and poor readers are still substantially below average. There is, however, little difference in the rates among white males, white females and black females.[9]

[9] Estimating the mean discount rates using the first switch point, but controlling for inconsistency with a dummy variable, yields estimates that are lower than in Table 2, but much closer than those that only use the first switch point in Table 5.

5.4 Regression: Full Sample

Table 6 presents estimates of the marginal effects and their corresponding standard errors and p-values for the full sample using the two methods for dealing with inconsistent decision-makers. Similar to the results for consistent students, the probit regressions shows that those who failed to meet state standards in reading are less likely to be impatient (17%). All other factors are insignificant. Using the cost-minimizing adjustment method, Blacks are 20% more likely to be impatient while gifted students and poor readers are less likely to be impatient (27% and 24%, respectively).

Table 6. Factors Affecting Discount Rates: Pooled Sample

	PROBIT REGRESSION[1] (= 1 IF DISCOUNT RATE>80%)			**INTERVAL REGRESSION[2]**	
Covariate	***Marginal Effect (SE) p-value***			***Marginal Effect (SE) p-value***	
	Using First Switch Point[1]	***Using Cost-Minimizing Adjustment***		***Using First Switch Point***	***Using Cost-Minimizing Adjustment***
Male	0.095 (0.061) p=0.122	0.091 (0.075) p=0.225		17.13 (6.42) p=0.008	13.80 (7.67) p=0.072
Black	0.050 (0.072) p=0.490	0.201 (0.083) p=0.015		-1.68 (7.32) p=0.819	16.93 (8.76) p=0.053
Other race/culture	0.169 (0.160) p=0.291	0.038 (0.165) p=0.820		28.38 (14.38) p=0.048	13.76 (17.11) p=0.421
Gifted	-0.104 (0.08) p=0.192	-0.269 (0.119) p=0.023		-14.60 (12.00) p=0.223	-34.68 (14.27) p=0.015
Poor Reader	-0.169 (0.055) p=0.002	-0.240 (0.084) p=0.004		-23.89 (7.87) p=0.002	-28.39 (9.37) p=0.002
Free/Reduced Lunch	-0.109 (0.077) p=0.157	-0.055 (0.088) p=0.532		-9.01 (7.57) p=0.234	-1.78 (9.03) p=0.844
	Controls for Fixed Session Effects			Controls for Fixed Session Effects	

[1] Marginal effects estimated for discrete change of variables from 0 to 1 evaluated at the means. The Likelihood Ratio chi-squared statistic is 16.39, and McKelvey-Zavoina R2 is 0.158.

Looking at the interval regression, holding other factors constant, males and other races/cultures have significantly higher discount rates (by 17% and 28%, respectively) while poor readers have significantly lower discount rates by 24%. All other effects are not significantly different from zero. Using the cost-minimizing adjustment method, males and blacks have significantly higher discount rates (12% and 17%, respectively) while gifted and poor reading students have lower discount rates (35% and 28%, respectively).

5.5 School Outcomes

The purpose of this study was not to examine causal relationships between discount rates and school outcomes. However, our results are consistent with some of the general patterns of drop-out rates observed in the United States, with males more likely than females to drop out, and black males being the most likely to drop out. We can also look at observations of disciplinary actions taken against the students in our sample during the period September 16, 2006 – March 15, 2007. Disciplinary actions do not include simple referrals to the principal's office, but comprise actions taken against the students and entered into the school accounting system, such as suspensions from school or in-school suspensions.

Disciplinary actions are over-dispersed count data (75% of the observations are zero or one), and thus we use a negative binomial model to regress disciplinary actions on student characteristics. We include all of the characteristics in Table 6, as well as whether the child is in a special education program because it is well known that such children are more likely to be disciplined in school. The variables Male (0.55; p=0.045), Black (0.66; p=0.022), and Special Education (0.55; p=0.087) have a significantly positive effect on disciplinary actions, while Gifted (-2.04; p=0.002) and Poor Reader (-0.62; p=0.06) have a significantly negative effect (the

others are not significantly different from zero at $p<0.10$).[10] Recall that, on average, males and Blacks have higher discount rates in our sample, while gifted children and poor readers have lower discount rates. This suggests a relationship between impatience and discipline rates and the likelihood of remaining in school.

These relationships in the data lead us to hypothesize that differences in educational outcomes may be partially explained by heterogeneity in time preferences across students. If true, our hypothesis also implies that policies targeted towards using monetary incentives to encourage students to get better grades or stay in school may have greater impacts on certain groups. Such a prediction is consistent with the findings of Angrist and Levy (2002), for example, who found that an incentive program in Israel that paid students conditional on their performance on university entrance exams had a greater effect on girls than boys. However, further research, particularly the development of panel data sets starting at an early age, will be necessary to establish causal relationships between time preferences and observed outcomes among children.

6. Conclusion

We experimentally elicit discount rates among more than two-hundred 8th grade students (13-15 years old). We estimate that the mean annual discount rate among these adolescents falls between 60 and 80%, which is approximately double the estimates of mean adult discount rates. We also find that there is substantial heterogeneity of discount rates, with the first quartile having annual discount rates lower than 35%, and the top quartile having rates greater than 100%. Although we do not find any evidence that adolescents from poor and non-poor households differ in their discount rates, we do find evidence that males have higher discount rates, Blacks

[10] A model that includes a dummy for an inconsistent decision-maker offers no further explanatory power.

have higher discount rates (specifically black males), gifted program children have lower discount rates, and poor readers have lower discount rates. Although the negative relationship between poor reading skills and discount rates is puzzling and requires more research to understand, we believe it may reflect the difficulty in using experimental methods to elicit the time preferences of students with poor verbal comprehension skills (written or oral).

As observed in Bettinger and Slonim (2007), our analysis also indicates that the proportion of adolescents who exhibit inconsistent time preferences is much higher than in comparable adult samples using similar methodologies. We extend the literature on eliciting discount rates experimentally by introducing an alternative way of drawing inferences about discount rates from inconsistent decision-makers. We choose the discount rate interval that minimizes the costs of making an inconsistent subject's choices consistent. Using this method, the results using the entire sample are similar to the results of using consistent decision-makers with only a few exceptions. Specifically, the estimated discount rates are higher and the male-female difference for Whites becomes small.

There are several avenues for future research. An obvious one is to expand the subject pool to include other ages and other geographic locales. Second, as done by Bettinger and Slonim (2007), one can further characterize time preferences among adolescents by varying the length of the front-end delay, including removing it entirely. This will allow one to examine whether adolescent preferences are better described by hyperbolic or quasi-hyperbolic rather than by exponential discounting. Third, future research should examine the effect of correcting the elicited discount rate for the concavity of the utility function. This could be done by using a risk-specific experimental instrument (Holt and Laury, 2002, as discussed in Harrison et al. 2005 and Bettinger and Slonim, 2007) or use a method that jointly elicits risk and time preferences.

Fourth, in studies that experimentally manipulate incentives for educational performance, the elicitation of discount rates would help understand the measured treatment effects.

The observed similarity between the factors that explain disciplinary actions and time preferences is promising in that it invites further research on the relationship between investment in human capital and time preferences. While our results are only suggestive of a relationship, these observations, combined with large observed heterogeneity of discount rates across sex, racial and cognitive groups, suggest that incentive programs targeted toward adolescents may need to be designed with heterogeneity of time preferences in mind.

References

Ausubel, L. M. 1991. The Failure of Competition in the Credit Card Market. *American Economic Review* 81(1): 50–81.

Becker, G.S. and C.B. Mulligan. 1997. The Endogenous Determinants of Time Preferences. *The Quarterly Journal of Economics* 112(3): 729-758.

Benzion, U., A. Rapoport, and J. Yagil. 1989. Discount Rates Inferred from Decisions: an experimental study. *Management Science* **35**: 270-284.

Bettinger, E. and R. Slonim. 2007. Patience among Children. *Journal of Public Economics* 91(1-2): 343-363.

Coller, M. and M.B. Williams. 1999. Eliciting Individual Discount Rates. *Experimental Economics* 2(2): 107-127.

Eckel, C., C. Johnson and C. Montmarquette. 2005. Human Capital Investment by the Poor: Calibrating Policy with Laboratory Experiments. Working Paper.

Eigsti, I., V. Zayas, W. Mischel, Y. Shoda, O. Ayduk, M.B. Dadlani, M.C. Davidson, J.L. Aber, and B.J. Casey. 2006. Predictive Cognitive Control from Preschool to Late Adolescence and Young Adulthood. *Psychological Science* 17(6): 478-484.

Gately, D. 1980. Individual discount rates and the purchase and utilization of energy-using durables: comment. *Bell Journal of Economics* 11: (373-374).

Harrison, G.W., M. Lau and M. Williams. 2002. Estimating Individual Discount Rates for Denmark: a field experiment. *American Economic Review* 92(5): 1606-1617.

Harrison, G.W., M. Lau, E.E. Rutstrom, and M.B. Sullivan. 2005. Eliciting Risk and Time Preferences Using Field Experiments: Some Methodological Issues. In *Field Experiments in Economics,* J. Carpenter, G. W. Harrison and J. A. List, eds. Greenwich, CT: JAI Press, Research in Experimental Economics, Volume 10, pp. 125-218.

Hartman, R. S., and M. J. Doane. 1986. Household Discount Rates Revisited. *The Energy Journal* 7: 139-48.

Hausman, J.A, 1979. Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables. *Bell Journal of Economics* 10: 33-54.

Holcomb, J.H. and P.S. Nelson. 1992. Another Experimental Look at Individual Time Preference." *Rationality and Society* 4: 199–220.

Holt, C.A. and S.K. Laury. 2002. Risk Aversion and Incentive Effects. *American Economic Review* 92: 1644–1655.

Lang, K. and P.A. Ruud. 1986. Returns to Schooling, Implicit Discount Rates and Black-White Wage Differentials. *The Review of Economics and Statistics* 68(1): 41-47.

Lowenstein, G.F. 1988. Frames of Mind in Intertemporal Choice. *Management Science* 34: 200–214.

Meier, S. and C. Sprenger. 2006. Impatience and Credit Behavior: evidence from a field experiment. Working Paper, October.

Pender, J.L. 1996. Discount Rates and Credit Markets: Theory and Evidence from Rural India. *Journal of Development Economics* 50: 257–296.

Ruderman, H., M. Levine, and J. McMahon. 1986. Energy Efficiency Choice in the Purchase of Residential Appliances. In *Energy Efficiency: Perspectives on Individual Behavior*, W. Kempton and M. Neiman, eds. American Council for an Energy Efficient Economy, pp. 41-50.

Shelley, M.K. 1993. Outcome Signs, Question Frames and Discount Rates. *Management Science* 39: 806–815.

Thaler, R. H. 1981. Some Empirical Evidence on Dynamic Inconsistency. *Economics Letters* 8: 201-207.

Warner, J. T. and S. Pleeter. 2001. The Personal Discount Rate: evidence from military downsizing programmes. *American Economic Review* 91: 33-53.