The Psychological Record, 2009, 59, 247–258

TEMPORAL DISCOUNTING OF HYPOTHETICAL MONETARY REWARDS BY ADOLESCENTS, ADULTS, AND OLDER ADULTS

Robert Whelan
Trinity College Dublin

Louise A. McHugh
Swansea University

The present experiment examined temporal discounting across 3 different age bands: adolescents, adults, and older adults (mean ages 14, 46, and 73 years, respectively). A computerized task was employed in which participants were asked to choose between larger rewards available at a specified time in the future—either £100 or £1,000 (approximately $200 and $2,000, respectively)—or a smaller reward available immediately. The subjective value of the reward decreased with increasing delay for each of the 3 age groups. A hyperbola-like function adequately described the group discounting data. The adolescent group discounted significantly more than the adult group when the larger later reward (LLR) was £100 but not when the LLR was £1,000. The adolescents discounted significantly more than the older adult group when the LLR was either £100 or £1,000. There were no significant differences in discounting between the adult and older adult groups. The results of the present study suggest that the rate of temporal discounting is higher in adolescents than in adults but is stable from middle adulthood to older adulthood. Furthermore, the process of temporal discounting appears to be quantitatively similar across the life span.

The term *temporal discounting* refers to the tendency of individuals to prefer smaller sooner rewards (SSRs) over larger later rewards (LLRs). For example, if presented with the choice between $85 now and $100 in 1 year, some people would choose $85 now. Those who choose the SSR can be considered impulsive; those who prefer the LLR can be said to have exercised self-control (e.g., Rachlin & Green, 1972).

One widely used procedure to measure the rate of temporal discounting involves a choice option between an amount of money, hypothetical or real, that is available immediately and a larger amount that is available at a later date. The SSR is manipulated over successive trials, while the LLR is kept constant. The aim of this procedure is to identify the current subjective value of the temporal rewards, which is defined as the magnitude of SSR that

Correspondence concerning this article should be addressed to Louise McHugh, Department of Psychology, Swansea University, SA2 8PP, UK. E-mail: l.mchugh@swansea.ac.uk

generates indifference in a choice against the LLR (Critchfield & Kollins, 2001). We refer to this value as the *indifference point.* Findings to date indicate that behavior patterns are not substantially affected by real versus hypothetical rewards (e.g., Johnson & Bickel, 2002).

Several mathematical equations have been proposed to analyze the underlying mechanisms of discounting. An account of the benefits of the different equations is beyond the remit of the current article (for a more detailed discussion see Green, Myerson, & Ostaszewski, 1999). In the discounting literature, Equation 1 (a hyperbolic function) is suggested to provide the best account of choice patterns of probabilistic and delayed rewards (Kirby & Marakovic, 1995; Mazur, 1987; Myerson & Green, 1995).

$$V = \frac{A}{(1 + kD)^s} \qquad (1)$$

In this equation, V is the discounted value of the delayed reward, A is the amount of the delayed reward, k describes the rate of discounting, and s represents the sensitivity to delay. Using this equation, a high value of k means a larger decrease in the value of reward as a function of increasing delay. The addition of the s parameter has been employed by several researchers including Green, Fry, and Myerson, 1994, and Rachlin, 1989. A smaller value of s indicates greater sensitivity to short delays, whereas a larger value of s indicates greater sensitivity to longer delays.

The ability to produce behavior that results in large, delayed reinforcers rather than behavior that results in small, more immediate reinforcers is a hallmark of development. To date, studies indicate that the rate of temporal discounting does indeed change over the life span, and quantifying these differences across the life span could yield important insights into human development. One such study was conducted by Green et al. (1994; see also Green et al., 1999), who examined temporal discounting across three different age groups of participants (12-, 20-, and 70-year-olds). A comparison of the median indifference points for each group demonstrated a developmental trend in terms of discounting. That is, levels of temporal discounting decreased as a function of age. Specifically, the 12-year-olds were more impulsive than both the 20-year-olds and the 70-year-olds, and the 20-year-olds were more impulsive than the 70-year-olds. Furthermore, the sensitivity-to-delay parameter allowed Green et al. to demonstrate that the children were more sensitive to short delays than were the adults. That is, the children perceived longer delays to be approximately equivalent.

One aspect of the Green et al. (1994) study was that the 12- and 20-year-old groups were closer in age to each other than to the 70-year-old group. Some longitudinal studies have indicated that impulsivity declines throughout young adulthood (e.g., Roberts, Caspi, & Moffitt, 2001), and therefore it is possible that young adults (i.e., 20-year-olds) discount at a different rate than middle-age adults (i.e., 40- to 60-year-olds). It is interesting to examine if impulsivity declines linearly across the life span or rather stabilizes in middle adulthood. The maximum delay employed by Green et al. was 25 years. This may have affected some of the findings for both the adolescent group, who have no experience of such long delays, and the older adult group, as it is statistically unlikely that participants would survive to 95 years of age.

The present experiment aimed to extend the findings of Green et al. (1994) by examining temporal discounting across three groups of participants who were approximately equidistant in age. The maximum delay in the current study was 1 year, and we employed an automated procedure, rather than the conventional tabletop procedure, in order to remove some of the methodological issues associated with tabletop experiments (Dymond, Rehfeldt, & Schenk, 2005). This automated procedure employed the choice algorithm described by Johnson and Bickel (2002). We predicted that adolescents would discount more steeply than their adult and older adult counterparts but that there would not be a difference between the adult and older adult groups.

Method

Participants and Design

The final number of participants in each group was as follows (means and standard deviations of the group age are in parentheses): 12 adolescents (8 female; $M = 14$; $SD = 2.8$), 16 adults (5 female; $M = 46$; $SD = 3.4$), and 10 older adults (6 female; $M = 73$; $SD = 4.12$). The results of some participants are not reported here because their responses were not consistent with temporal discounting (as explained in the Results section). The adolescents who participated were selected from volunteers following classroom announcements made in various schools within the Swansea area. They were chosen on the basis that neither their mainstream teachers nor their parents had identified them as presenting a learning difficulty. The consent of parents and teachers was obtained prior to participation. All of the adult participants were recruited through faculty board announcements within the Department of Psychology at Swansea University. The older adult participants were recruited through personal contacts of the experimenters and night classes in Swansea University. All participants in the current study were from a middle-class socioeconomic background (classification based on the participants' postcodes). Participants were categorized in terms of income, either their own (adults and older adults) or their parents' (adolescents).

Setting and Materials

The study was conducted in a quiet room containing a desk, chair, and personal computer with a 36-cm color monitor and a standard computer mouse. Each participant was exposed to the protocol individually.

All trial presentations were controlled by the computer, using a choice algorithm described by Johnson and Bickel (2002; see also Richards, Zhang, Mitchell, & de Wit, 1999). The participant was presented briefly with an SSR and an LLR on the computer screen and asked to choose between the two. The LLR remained constant, whereas the SSR varied across responses. For each LLR, the algorithm gradually converged on the indifference point by using a random adjusting-amount procedure. This procedure used the answers to the previous questions to narrow the range of values from which the value for the next comparison was selected. An important feature of this method is that, in order to minimize the effects of subject error (e.g., due to inattention), the computer varied the magnitude of the smaller, more immediate reward according to a double limit procedure, which precluded any single answer from controlling the convergence toward an indifference point (see the Appendix for more details).

Procedure

Each participant was taken to the experiment room, and when he or she indicated readiness to begin, the experimenter instructed the participant simply to follow the instructions that appeared on the screen. Participants were assured that the computer would present all of the information they required to conduct the study but if additional assistance was needed, they could seek the experimenter's attention. At the beginning of the experiment, participants were provided the following instructions:

> For this study, you will be given a series of choices. There are no right or wrong answers, and there is no time limit to make choices. Please just answer honestly.
>
> For each choice, you will see one button on the left side of the screen, and another button on the right side of the screen. Each button will contain a money amount and also whether that money would be available now or after a delay. For example, you may be asked to choose between £10 now, and £20 in one month. Using the mouse, simply click on the relevant choice to select it. It will not be possible to make a choice until the text on the button turns from grey to black.

After answering any questions, the experimenter instructed the participant to attend to the screen.

A box at the top of the screen contained the caption "Which would you prefer" for all trials. Below this statement two other labels were presented, one displaying the SSR and the word "now" (e.g., "£50 now") and the other displaying the LLR and a specified delay (e.g., "£100 in one month"). The text in these boxes was grey for 0.5 s, and any response during this time was ineffective. After this delay, the text color changed to black and clicking on either box cleared the screen. Due to a computer error, 3 indifference points (of the total 608 indifference points that were recorded across all participants) were not saved. This did not affect the subsequent analyses to any substantial degree.

Indifference points were calculated for two LLRs (£100 and £1,000; approximately $200 and $2,000, respectively, at the time of the study) across eight delays (1 day, 2 days, 1 week, 1 month, 3 months, 6 months, 9 months, and 1 year). The participant was presented with a particular LLR/delay combination until an indifference point was reached for that combination. All LLR/delay combinations were presented in a randomized order. All trials were presented in a single session, and the time to complete the study was approximately 45 min. After the last test screen had been completed, each participant was thanked for participating and was debriefed.

Results

The objectives of our data analyses were to (a) determine if Equation 1 could adequately describe the discounting of all groups, (b) produce the parameters for rate of discounting and sensitivity to delay, and (c) determine if the amount of discounting differed significantly among the groups.

There were some departures from the theoretically ideal pattern of discounting. The data recorded from some participants are not discussed here. Some participants always chose the LLR (7 adolescents and 3 adults); thus their data were excluded. The remaining data were further examined for departures from normality by checking that the mean of the indifference points from the three shortest delay conditions did not exceed the mean of the indifference points from the three longest delay conditions (see also Dixon, Jacobs, & Sanders, 2006). The results of 1 adolescent participant were excluded on this basis alone (i.e., the previously excluded participants also met this criterion). The final number of participants included in the following analyses was as follows: 12 adolescents, 16 adults, and 10 older adults.

The median and interquartile range of the number of trials until computation of the indifference point for each time delay are presented in Table 1. The data were first analyzed by applying Equation 1 to the median indifference point for each time point in each group. The parameters were fit using the Matlab Curve Fitting Toolbox running under Matlab 7 (Mathworks, Sherborn, MA). Figure 1 displays the resulting fit of Equation 1 to these data. Table 2 presents the parameter values of these best fits and the variance explained by Equation 1. Equation 1 accounted for a high proportion of the variance (> 90%), with the exception of the older adult group when the LLR was equal to £100 (86.5%). Table 2 displays the *k* (rate of discounting) and *s* (sensitivity to delay) values.

Table 1
Median (and Interquartile Range) for Number of Trials Until Computation of the Indifference Point Across Time Delays

LLR (£)	Delay (days)							
	1	2	7	30	90	180	270	365
100	14.5 (7.7)	17 (10)	17.5 (4)	16 (11.2)	18 (10.7)	18 (14.5)	20.5 (15.7)	20.5 (13.2)
1,000	11 (11)	12 (7)	14.5 (9)	18.5 (12.7)	19.5 (7.5)	19 (11)	20 (13.5)	20 (8.2)

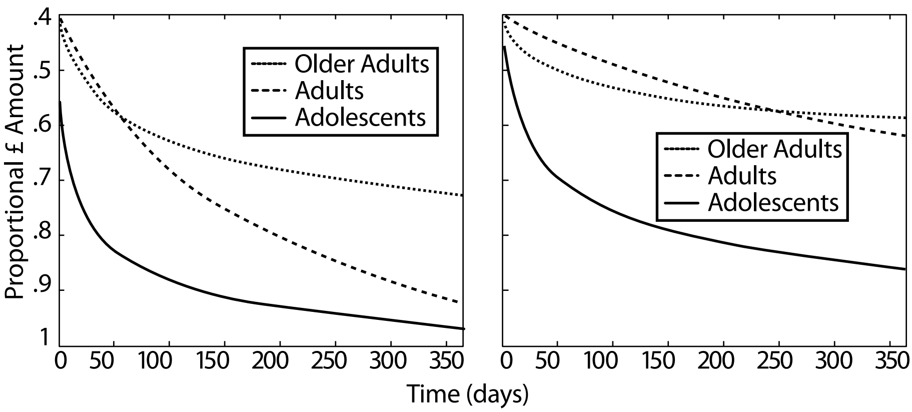

Figure 1. Equation 1 fit to the data from the adolescent, adult, and older adult group.

Table 2
Summary of the k *and* s *Values and the Proportions of Variance for the Fits to Equation 1*

Group	LLR (£)	Parameters		
		k	s	R^2
Adolescent	100	0.91	0.146	0.976
	1,000	0.212	0.140	0.946
Adult	100	0.009	0.524	0.929
	1,000	0.004	0.276	0.920
Older adult	100	0.087	0.114	0.865
	1,000	0.118	0.055	0.907

In order to test for differences among the groups, we employed the theoretically neutral method of calculating discounting proposed by Myerson, Green, and Warusawitharana (2001), the area under the empirical discounting function (also termed the *area under the curve* [AUC]). To calculate this area, the delay and subjective value for each data point were first normalized. The delay was expressed as a proportion of the maximum delay, and the subjective value was expressed as a proportion of the nominal amount. These normalized values were used to construct a series of trapezoids. The area of each trapezoid was equal to $(x_2 - x_1)[(y_1 + y_2)/2]$, where x_1 and x_2 were successive delays and y_1 and y_2 were the subjective values associated with these delays. (For the first trapezoid, the value of x_1 and y_1 were set at zero and one.) The area under the empirical discounting function is equal to the sum of the areas of these trapezoids. The amount of discounting ranges between one (no discounting) and zero (total discounting). Complete details of this method are described in Myerson et al. (2001).

Table 3 presents the AUC values for each participant. Figure 2 presents the AUC value of the immediate rewards that the adolescents, adults, and older adults selected. When the LLR was equal to £100, the most discounting was observed in the adolescent group (M = .54), followed by the adult group (M = .62), with the least discounting occurring in the older adult group (M = .73). When the LLR was equal to £1,000, the adolescents discounted most (M = .60), whereas the older adults discounted slightly less than the adults (M = .76 vs. M = .77, respectively).

A series of planned comparisons were conducted to examine if the participants' rate of discounting differed across age groups. There was no significant difference between discounting levels of the adult group and the older adults for either the £100 or the £1,000 LLR. There was no significant difference between the adolescents and the adults on the £1,000 LLR. However, the adolescents discounted significantly more than the adults when the LLR was £100 ($t(26) = 1.91$, $p < .05$, one-tailed). The adolescents discounted significantly more than the older adults both when the LLR was equal to £100 ($t(20) = 3.53$, $p < .001$, one-tailed) and when the LLR was equal to £1,000 ($t(20) = 2.67$, $p < .01$, one-tailed).

To summarize, all three age groups (adolescents, adults, and older adults) discounted the LLR more as the time delay increased. The adolescents discounted significantly more than the older adults across both LLRs and significantly more than the adults when the LLR was equal to £100. There were no significant differences between the adults and older adults across either LLR.

Table 3
Individual Data for the Area Under the Curve Measure

Group					
Adolescent		Adult		Older Adult	
LLR (£)					
100	1,000	100	1,000	100	1,000
0.219	0.283	0.914	0.900	0.955	0.965
0.466	0.589	0.281	0.857	0.946	0.623
0.394	0.381	0.640	0.926	0.738	0.849
0.442	0.605	0.721	0.872	0.530	0.646
0.646	0.686	0.754	0.848	0.970	0.982
0.572	0.844	0.550	0.637	0.856	0.848
0.509	0.591	0.910	0.945	0.508	0.539
0.814	0.926	0.764	0.915	0.557	0.636
0.508	0.597	0.625	0.409	0.531	0.755
0.307	0.597	0.518	0.753	0.713	0.745
0.631	0.650	0.229	0.469		
0.509	0.676	0.294	0.658		
		0.688	0.905		
		0.845	0.960		
		0.238	0.349		
		0.927	0.973		

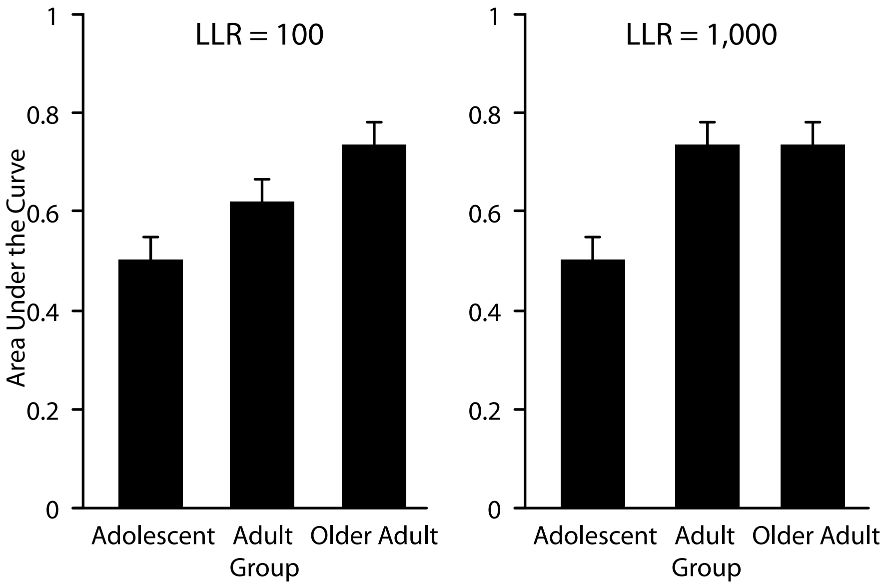

Figure 2. A bar chart, including standard error bars, depicting the area under the curve across the three LLRs for the adolescent, adult, and older adult groups.

Discussion

The present study has shown once again that the subjective value of a future reward is less than an immediate reward of equal value, in this case across three age groups (adolescents, adults, and older adults). The results of the present experiment broadly support previous research demonstrating that the rate of temporal discounting is greater in adolescents than in adults (Green et al., 1994; Green et al., 1999). There were a number of features that distinguished the present experiment from previous, similar research (Green et al., 1994; Green et al., 1999). The delays presented in this study were within the experience of all participants and were within the likely life span of all participants (the maximum delay was 1 year). The age range of the participants was approximately evenly spread (cf. Green et al., 1994). Finally, an automated, as opposed to tabletop, procedure was employed.

Equation 1 was an adequate fit to the data, with the variance explained generally exceeding 90%, the exception being when the LLR was equal to £100 in the adult group. It should be noted that other studies have reported higher values (e.g., Green et al., 1994, reported a range of 94.5%–99.9%). Group data rather than individual data were analyzed in order to calculate the discounting function; this is consistent with previous temporal discounting studies (e.g., Dixon et al., 2006) and often is necessary due to the between-subject variability. One possibility for subject variability is that the present study employed shorter delays; consequently there was greater variability between successive time delays (i.e., there is a smaller difference between 6 months and 9 months than between 5 years and 10 years). It is possible that longer delays would reduce this variability.

The *k* value was highest for adolescents, across both LLRs, indicating that this group had the highest discounting rate. The AUC also was lowest for the adolescent group across both LLRs, again demonstrating that this group tended to discount more. Perhaps surprisingly, the older adults had higher *k* values than the adults across both LLRs. However, when the LLR was equal to £100, the AUC result suggested that the adult group discounted more than the older adult group (the AUC was similar for both the adult and older adult groups when the LLR was equal to £1,000). These seemingly different results perhaps can be explained by an examination of the sensitivity-to-delay parameter (*s*). An examination of the *s* value indicates that the adolescents and older adults both were more insensitive to differences between long delays than were the adults. That is, longer delays were perceived as being approximately equivalent for adolescents and older adults. In contrast, the adults perceived longer delays as different and continued to discount as delay increased when the LLR was equal to £100.

Green et al. (1994) reported a developmental trend across their three groups (12-, 20-, and 68-year-olds). That is, the younger participants discounted the most, followed by the 20-year-old group, who in turn discounted more than the older adult group. In contrast, we did not find any difference in the rate of discounting between the adult and older adult groups. The most obvious reason for these differences between the studies is that the participants in the adult group in the present study were, on average, 46 years old, rather than 20 years old as reported in Green et al. Indeed, one of the motivations for this research was to examine if impulsivity would stabilize in middle adulthood, given that impulsivity declines throughout young adulthood (e.g., Roberts et al., 2001). The findings of the present study suggest that impulsivity may decrease between young and middle adulthood. An interesting future study would include even more adult groups, in order to more precisely describe when impulsivity begins to decline.

Although temporal discounting procedures often omit contact with real consequences, most participants will have prior experience with money. In addition, some studies employing real, rather than hypothetical, consequences have produced results much like those described. For example, Richards et al. (1999) allowed healthy volunteers to choose between various SSRs and an LLR of $10, available at delays up to 365 days. For each delay, the indifference points indicative of current subjective value were estimated in an adjusting procedure. Similar to participants in studies that employed real money consequences, participants in the Richards et al. study discounted the LLR as delay increased. Furthermore, a hyperbolic model of discounting proved a good fit to the data. The findings suggest direct parallels in human discounting of hypothetical and real consequences.

There are some limitations to the present study. The groups were not balanced in terms of gender; there were proportionally more males in the adult group than in either the adolescent group or the older adult group. This disparity may have had some influence on the outcome. A large meta-analysis, however, has indicated that there is little difference in impulsivity between males and females, with a median effect size of 0.03 (Feingold, 1994). Consistent with previous studies of temporal discounting, it was necessary to disregard some data due to inappropriate responding. That is, some participants always chose the LLR over the SSR, thus rendering it impossible to calculate an indifference point. According to Critchfield and Kollins (2001), the aim of temporal discounting procedures is to calculate the indifference point, and thus when no subjective value can be derived, the data should be disregarded. In general, however, the automated procedure was successful and may be a useful replacement for tabletop temporal discounting procedures. It is possible that some genuine differences between groups were not observed because of insufficient power in the present study ($n = 38$). Future studies could add to the current findings by employing more participants (reducing between-subject variability) or perhaps by obtaining more than one indifference point for each LLR for each participant and averaging across those indifference points (reducing within-subject variability).

The present study replicates and extends the findings of previous work in this area. Indeed, in their 1994 paper, Green et al. stated, “[I]nterpretation of the present results must be tempered until age differences in delay discounting have been examined in additional populations and with different procedures” (p. 36). The present study has gone some way in addressing both these issues by the inclusion of a middle-age adult group and two LLR amounts common to all groups, and by the use of an automated procedure. Potential future studies in this area could incorporate more age ranges (e.g., young adolescents), longer delays, or different LLRs. The present data add to the self-control/impulsivity literature (e.g., Binder, Dixon, & Ghezzi, 2000) by indicating that the amount of temporal discounting varies systematically as a function of development. Furthermore, a hyperbola-like function described the behavior of age groups, indicating that developmental differences in temporal discounting are primarily quantitative, rather than qualitative, in nature.

References

BINDER, L. M., DIXON, M. R., & GHEZZI, P. M. (2000). A procedure to teach self-control to children with attention deficit hyperactivity disorder. *Journal of Applied Behavior Analysis, 33,* 233–237.

CRITCHFIELD, T. S., & KOLLINS, S. H. (2001). Temporal discounting: Basic research and the analysis of socially important behavior. *Journal of Applied Behavior Analysis, 34,* 101–122.

DIXON, M. R., JACOBS, E. A., & SANDERS, S. (2006). Contextual control of delay discounting by pathological gamblers. *Journal of Applied Behavior Analysis, 39,* 413–422.

DYMOND, S., REHFELDT, R. A., & SCHENK, J. J. (2005). Nonautomated procedures in derived stimulus relations research: A methodological note. *The Psychological Record, 55,* 461–482.

FEINGOLD, A. (1994). Gender differences in personality: A meta-analysis. *Psychological Bulletin, 116,* 429–456.

GREEN, L., FRY, A. F., & MYERSON, J. (1994). Discounting of delayed rewards: A life-span comparison. *Psychological Science, 5,* 33–36.

GREEN, L., MYERSON, J., & OSTASZEWSKI, P. (1999). Discounting of delayed rewards across the life span: Age differences in individual discounting functions. *Behavioural Processes, 46,* 89–96.

JOHNSON, W., & BICKEL, W. (2002). Within subjects comparison of real and hypothetical money rewards in delay discounting. *Journal of the Experimental Analysis of Behavior, 77,* 129–146.

KIRBY, K. N., & MARAKOVIC, N. N. (1995). Modelling myopic decisions: Evidence for hyperbolic delay-discounting within subjects and amounts. *Organisational Behavior and Human Decision Processes, 64,* 22–30.

MAZUR, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior*: *Vol. 5. The effect of delay and of intervening events on reinforcement value* (pp. 55–73). Hillsdale, NJ: Erlbaum.

MYERSON, J., & GREEN, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of Experimental Analysis of Behavior, 64,* 263–276.

MYERSON, J., GREEN, L., & WARUSAWITHARANA, M. (2001). Area under the curve as a measure of discounting. *Journal of the Experimental Analysis of Behavior, 76,* 235–243.

RACHLIN, H. (1989). *Judgment, decision, and choice: A cognitive/behavioral synthesis.* New York: W. H. Freeman.

RACHLIN, H., & GREEN, L. (1972). Commitment, choice and self-control. *Journal of the Experimental Analysis of Behaviour, 17,* 15–22.

RICHARDS, J. B., ZHANG, L., MITCHELL, S. H., & DE WIT, H. (1999). Delay or probability discounting in a model of impulsive behaviour: Effect of alcohol. *Journal of the Experimental Analysis of Behaviour, 71,* 121–143.

ROBERTS, B. W., CASPI, A., & MOFFITT, T. (2001). The kids are alright: Growth and stability in personality development from adolescence to adulthood. *Journal of Personality and Social Psychology, 81,* 670–683.

Appendix

All trials consisted of a choice between the smaller sooner reward (SSR) and the larger later reward (LLR). The software program set the LLR equal to the monetary amount for which the discounting rate was being determined. The program varied the magnitude of the SSR according to the double limit procedure described by Johnson and Bickel (2002). The program contained values for four limits, two outer limits and two inner. The outer upper limit (OUL) and the outer lower limit (OLL) bound an area representing the program's "best guess" for the location of the indifference point. The OUL and the OLL adjusted and eventually converged on an indifference point. In order to protect against any one response controlling the movement of these outer limits, an inner upper limit (IUL) and an inner lower limit (ILL) also were employed. The values of these inner limits remained within the outer-limit-bound region and functioned to buffer the effect of each response on the outer limits. That is, a single response immediately changed the inner-limit-bound range but only changed the outer limits if the choice confirmed or disconfirmed the best guess represented by the inner-limit-bound range.

On the first trial of each indifference-point determination the two upper limits were set to the value of the LLR, and the two lower limits were set to a value of zero. Throughout all trials, the OUL was either greater than or equal to the IUL, and the OLL was either less than or equal to the ILL. The magnitude of the SSR was a multiple of 2% of the LLR and was randomly picked from within the inclusive range of the upper and lower outer limits. After each response by a participant, the computer adjusted one or more of the four limits, and the next SSR was picked randomly from within the inclusive range of the outer limits. Figure A displays a schematic of the program that was used in the present study.

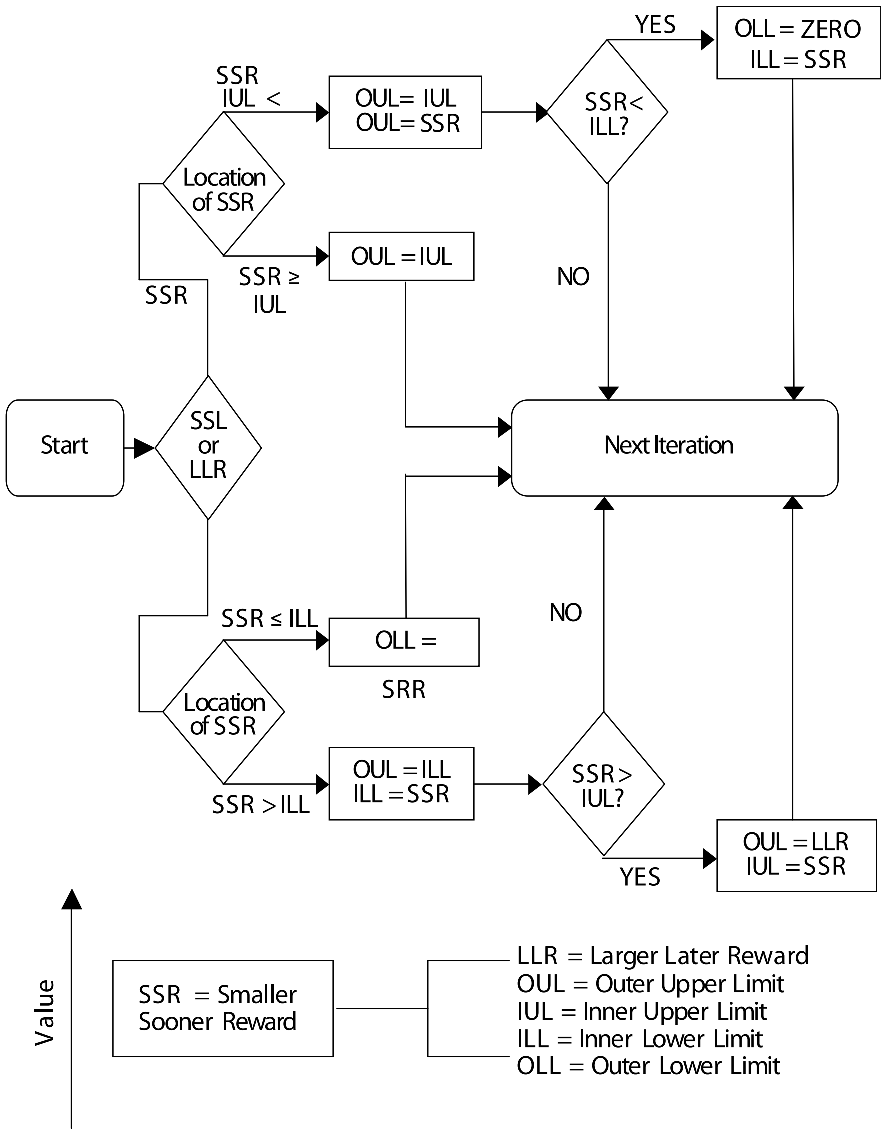

Figure A. A schematic of the double limit procedure.