

Running head: INCREASED TOLERANCE FOR DELAY WITH BUNDLED REWARDS

**Building blocks of self-control:
Increased tolerance for delay with bundled rewards**

George Ainslie

Coatesville Veterans Affairs Medical Center and Temple University

John R. Monterosso

University of Pennsylvania

For the JOURNAL OF THE EXPERIMENTAL
ANALYSIS OF BEHAVIOR: 79, 37-48, 2003

Final draft

Corresponding author:

George Ainslie

Department of Psychiatry

Coatesville VA Medical Center

Coatesville, PA 19320, USA

Phone: (610) 384-7711 X4260

Fax: (610) 380-4377

Email: George.Ainslie@med.va.gov

Building blocks of self-control: Increased tolerance for delay with bundled rewards

George Ainslie and John R. Monterosso

Abstract

Impulsive choice can be defined as temporary preference for a smaller-sooner reward (SS) over a larger-later reward (LL). Hyperbolic discounting implies that impulsive choices will occur less when organisms choose between a series of SSs versus LLs all at once than when they choose between single SS versus LL pairs. Eight rats were exposed to two conditions of an intertemporal choice paradigm using sucrose solution as reward. In both conditions, the LL was 150 μ l delayed by 3 s, while the SS was an immediate reward that ranged from 25 – 150 μ l across sessions. Preference for the LL was greater when the chosen reward was automatically delivered three times in succession (*bundled*) than when it was delivered after each choice. For each of the 8 rats, the estimated SS amount that produced indifference was higher in the bundled condition than in the single condition. Because bundling in humans may be based on the perception that one's current choice is predictive of future choices, the data presented here may demonstrate an important building block of self-control.

Keywords: hyperbolic discounting, intertemporal conflict, self-control, reward bundling, willpower, lever press, rats

There is now considerable evidence that reward-seeking organisms devalue delayed rewards at a steeper rate when the rewards are close than when they are distant. The longer the delay already associated with a reward, the less the reward is further devalued with additional delay. Mathematically, this is represented by temporal discount functions in which value is inversely proportional to delay (Baum & Rachlin 1969; Grace, 1996; Mazur 1984; Myerson and Green 1995). Mazur's formula is simplest:

$$V = \frac{A}{1 + kD} \quad (1)$$

where V is value, D is delay, and k is a constant describing the individual subject's degree of impatience. This hyperbolic discount function contrasts with exponential discount functions, in which reward is devalued at a fixed rate over time. Laboratory data from both humans and non-humans are better fitted by hyperbolic models of discounting than exponential models (Kirby, 1997; Mazur, 2001).

The apparent fact that rewards are discounted hyperbolically rather than by the constant rate described in exponential models has radical implications. Unlike the exponential discounter, the hyperbolic discounter's preference between a smaller-sooner reward (SS) and a larger-later reward (LL) will vary depending on the organism's distance to the choice pair. When the choice pair is relatively distant, the LL will be preferred; but when the choice pair is more immediate, preference will often reverse in favor of the SS. This has been demonstrated empirically in several species. For example, in one study, pigeons chose between a SS that varied in proximity across conditions (from 0.01 to 12 s) and a LL that was twice as large but that was always delayed by 4 s more than the SS. All subjects demonstrated reversal of preference, choosing the SS when the delay was short but the LL when it was long, despite the fact that the difference between the alternatives was constant (Ainslie and Herrnstein, 1981). Similar results have been obtained in rats (Deluty, 1978) and in humans (Kirby and Herrnstein, 1995).

Although hyperbolic discounting induces temporary preference for inferior rewards, it also provides mechanisms for consistency of preference. The tails of hyperbolic curves are higher than those of exponential ones, suggesting that the more a choice can be made under the influence of relatively delayed rewards, the less impulsive it will be. The organism that is vulnerable to temporary preferences will have an incentive to commit itself to choosing an LL alternative when the commitment can be made far enough in advance. For instance, Rachlin and Green (1972) reported that pigeons preferred a key that led to a terminal link containing only the LL of an SS—LL choice pair, provided that the onset of the terminal link was sufficiently delayed. Similarly, Ainslie (1974) showed that some pigeons would consistently peck a key that had the effect of suppressing the future availability of a SS alternative, despite the fact that the SS was virtually always chosen over the LL when it was available.

Another way that LLs might come to have more influence would be if choices were made in whole series at once. Because the rewarding effects of multiple rewards has been shown to be the simple sum of each reward discounted proportionally to its delay (Mazur 1986), consistent

preference might be achieved if choices between temporally extended series of SS—LL alternatives were *bundled* together rather than made case by case (Ainslie, 1975; Ainslie 2001). This is because such a choice must by definition be made before the series begins, and only in the first pair of alternatives is the SS strongly overvalued; subsequent pairs, if sufficiently distant, add weight on the side of the LL alternative. For instance, consider a woman who discounted hyperbolically according to Equation 1. And for simplicity, assume she discounted at a steepness such that her discount parameter k was equal to 1 when delay was computed in days, thereby making *Value* equal to *Amount* (*Delay* + 1). Given the choice on Monday between an immediate \$60 and \$108 delayed by a day, the woman would choose the immediate \$60, because the 1 day delay would reduce the value of the \$108 by half. If the same alternatives were presented again on Tuesday, and again on Wednesday, the woman would always be expected to choose the \$60, gaining a total of \$180. However, suppose that on Monday she had the same alternatives, except this time with the additional consideration that Tuesday's and Wednesday's choices were bundled together with the current choice such that whatever choice was made on Monday bound her to the same alternative on Tuesday and Wednesday. In this case, the value of the \$60 bundle would be $60/1 + 60/2 + 60/3$, or 110. The value of the \$108 bundle would be $108/2 + 108/3 + 108/4$, or 117. The same individual given these alternatives would prefer the more delayed \$108 bundle, this time gaining a total of \$324.

By contrast, bundling pairs of exponentially discounted alternative rewards together in series will not change their relative values. Exponentially discounted rewards are devalued at a fixed rate per unit of time, and so the ratio between an immediate alternative and a discounted delayed alternative does not change if a quantity of delay is added to both options. Figures 1a and 1b illustrate the predicted effect of bundling series of choices for a hyperbolic discounter, and the corresponding noneffect for an exponential discounter.




Figure 1a

Figure 1b

Fig. 1. a. Summed hyperbolic curves from a series of hypothetical larger-later rewards and a series of hypothetical smaller-sooner rewards. The vertical rectangles represent the value of the reward when immediate, and each discount curve represents the discounted value of that alternative when summed with all other like rewards occurring later in time (to the right). At the beginning of the series, preference for the series of larger rewards is consistent. By contrast, the curves from just the final pair of rewards indicate a period of temporary preference for the

smaller-sooner reward when it is imminent. b. Summed exponential curves from the two series of rewards shown in Figure 1a. Again, the vertical rectangles represent the value of the reward when immediate, and each discount curve represents the discounted value of that alternative when summed with all other like rewards occurring later in time (to the right). Summing does not change the relative heights of the curves.

Testing the bundling hypothesis, Kirby and Guastello (2001) reported that college students showed greater preference for LLs when amounts of money were bundled into series rather than chosen individually. These authors also observed this phenomenon with amounts of primary reward (pizza). In any experiment with human subjects, however, there is always the possibility that the results are a product of demand characteristics or acculturation rather than a universal property of choice. An experimenter could rule out this kind of explanation by bundling choices for nonhumans and measuring the effect that this bundling had on preference. The present research compares single *standalone* choices to *bundled* choices by rats to determine whether the bundling effect also occurs in this nonprimate species.

Method

Subjects

Subjects were 8 experimentally naive Sprague-Dawley rats purchased from Charles River Laboratories. Rats were housed one per cage. Lights were off in the colony room from 7:00 a.m. to 7:00 p.m. Testing took place Monday through Friday between 10:00 a.m. and 2:00 p.m. Water was available continuously, whereas caloric intake was limited to the sucrose delivered during testing and four additional 5-g pellets given after each test session.

Apparatus

Two experimental chambers (30 cm wide by 25 cm deep by 30 cm high) were used throughout the experiment. The chambers had stainless steel grid floors, aluminum front and back walls, and Plexiglas side walls. A front wall through which two levers could be extended served as the test panel. The two levers were located 8 cm on either side of the center of the test panel wall, 2 cm above the cage floor. Each lever was 3 cm in width. When retracted, both levers were flush with the wall and inaccessible. When extended, levers protruded 2 cm into the cage and required approximately 0.25 N of force to register a press. Above each retractable lever was a light that was on only when that lever was accessible. A spout for dispensing liquid was located at the midpoint of the test panel wall, 7 cm from the cage floor. Liquid was dispensed through metal-coated plastic tubing into a metal reservoir which extended approximately 2 cm into the cage. The reservoir had a capacity of approximately 600 μ l. The concentration of the sucrose solution was 12.5%.

Experimental protocols were conducted using a 120 MHz microcomputer that operated a syringe pump (Harvard Apparatus, PHD2000), activated two 5-volt solenoids used to gate the passage of sucrose solution through tygon tubing, and controlled the levers and lights in each experimental

chamber. The computer interfaced with all equipment via the Coulbourn Universal LabLinc®, and all experimental protocols were written in Quick BASIC®.

Procedure

The procedure was designed to establish the amount of immediate reward that was equally preferred to a standard delayed reward under two conditions—a *standalone* fixed-ratio 1 (FR1) schedule and a *bundled* FR1 schedule. In the bundled schedule, the chosen reward was delivered 3 consecutive times following each press (described below).

On each choice trial of the experiment, both levers were available. The left lever produced immediate delivery of sucrose solution, whereas the right lever produced the sucrose solution delayed by 3 s. The reward amount of the delayed alternative was set to 150 μ l throughout the experiment (the LL), while the reward amount of the immediate alternative varied across sessions between 25 and 150 μ l (the SS). At the beginning of each trial the houselight was turned on and the levers were extended. When either lever was pressed, both levers were retracted, the houselight was turned off, and sucrose solution was delivered after the specified delay (0 s or 3 s) through the central spout. The houselights remained off and levers remained retracted during the interval (ITI), which was set to make the trial duration 6 s. That is, the ITI was 6 s when the delay was 0 s, and 3 s when the delay was 3 s. To ensure sampling of alternatives, four forced-choice trials were scheduled between sets of four choice trials. Forced-choice trials were identical to choice trials except that only one lever was available. Each block of four-forced choice trials contained two trials with each lever, with order based on computerized random selection without replacement. Sessions were continued until satiation, as indicated by a subject's failure to press an available lever within 60 s. Thus the number of trials varied across sessions.

Trials in the bundled condition were identical to the standalone condition in terms of the presentation of alternatives, the initial delay period, the initial reward delivery, and the period of darkness to complete the 6-s cycle. In the bundled condition, however, without further presentation of the levers, a second 6-s cycle of the chosen delay + reward + darkness followed directly, followed by a third cycle. Only after the end of the third cycle did the houselight turn on and the lever or levers extend to allow the beginning of the next trial. The general layout of trials in both the standalone and the bundled condition are shown in Figure 2. Because bundled trials delivered three times as much sucrose solution as standalone trials, only a third as many could be conducted before satiation. The mean number of choice trials before satiation in the bundled condition was 18.6 (55.8 rewards), as compared with 53.7 for the standalone condition.

Standalone Condition




Fig. 2. Schedule of reinforcement delivery in the standalone and bundled conditions. The amount (X) of immediate sucrose solution (Su) varied across sessions between $25 \mu\text{L}$ and $150 \mu\text{L}$.

The delayed (3 s) reward was $150 \mu\text{L}$ throughout the experiment. The amount of the immediate reward varied across sessions among the following six alternatives: $150 \mu\text{L}$, $125 \mu\text{L}$, $100 \mu\text{L}$, $75 \mu\text{L}$, $50 \mu\text{L}$, or $25 \mu\text{L}$. Standalone and bundled choices were presented in an ABBA design over 72 sessions on separate days as follows: three sessions of standalone choice with each of the six immediate reward amounts in descending order (total = 18); then three sessions of bundled choice with each immediate reward amount in ascending order (total = 18); then 3 sessions of bundled choice with each immediate reward amount in descending order (total = 18); then three sessions of standalone choices with each immediate reward amount in ascending order (total = 18).

Analysis of Individual Choices

To examine overall preference across the group of 8 rats, the mean percentage choice was computed for each session and the data were analyzed as a repeated measure analysis of covariance (ANCOVA) with both amount and experimental condition as within subject variables. This was done (a) over the entire experiment, (b) separately for the first 36 (AB) and second 36 (BA) sessions, (c) using just the first 15 trials of all sessions (to equalize the number of trials across conditions), and (d) using only the first 30 trials from each standalone session and first 10 trials from each bundled session (to minimize the effects of satiation while holding mean consumption constant across conditions).

To examine the generality of the results, the effect of condition was considered individually for each rat. This was done by computing the percent of trials that the delayed reward was chosen in

each session and then performing a repeated measure ANCOVA for each rat with percentage choice as the dependent variable, condition as the fixed factor, and delay as a covariate.

Analysis by Session

Given sufficient exposure to concurrent fixed-ratio schedules, near exclusivity of preference would be expected to emerge. As such, one might be cautious in interpreting quantitative differences in preference such as the difference between 90% LL responses in one condition and 80% LL responses in another. An alternative approach is to consider preference for each session only categorically. A binomial test was performed comparing SS versus LL responses within each session to a test proportion of 0.50. If the 0.50 test proportion was rejected with a certainty of $p < .05$, a preference was concluded to have been expressed in the session. The results of session analysis are presented, indicating for each session the direction of preference if the above criterion was met, or the absence of preference if it was not.

Results

Figure 3 presents the overall percentage choice of the LL reward for all sessions (collapsing across subjects) as a function of the size of the SS and the experimental condition, separately comparing the AB (top) and the BA (bottom) blocks of the experiment. At 50 μl , 75 μl , and 100 μl — the swing values of SS reward, where preference was not overwhelmingly in one direction — the lines signifying preference in the bundled condition are substantially above those signifying preference in the standalone condition. This indicates greater preference for the LL in the bundled conditions. The downward slope of the lines indicates the effect that reward amount had no preference. In all conditions the larger the immediate reward, the less frequently subjects chose the delayed 150 μl alternative. Repeated measure ANCOVAs were conducted separately for the AB and BA blocks, modeling percentage preference for the LL reward based on experimental condition (standalone vs. bundled) and amount of the SS. There was significantly greater preference for the LL in the bundled condition in both the AB ($F [1, 7] = 31.9, p < .001$) and BA analyses ($F [1, 7] = 7.0, p = .03$). As would be expected, amount was a highly significant predictor of preference in both the AB and BA analyses ($F [5, 35] = 284.0, p < .001$, and $F [5, 35] = 194.0, p < .001$ respectively).

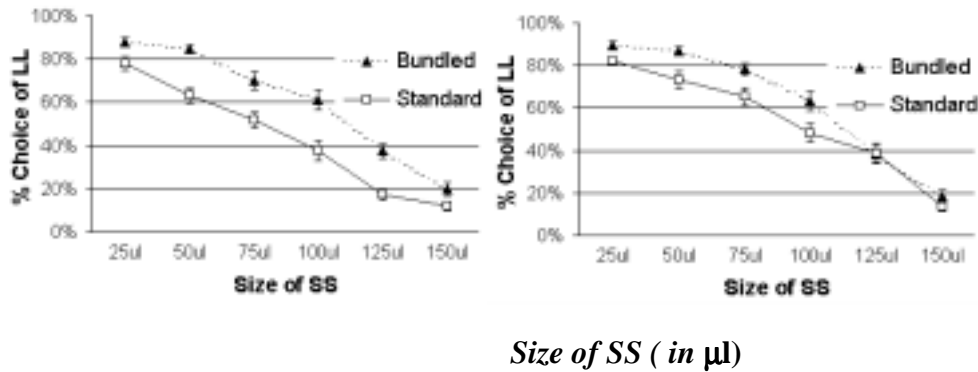


Fig. 3. Overall percentage of trials that the larger-later alternative was chosen (150µl delayed by 3 s) as a function of the size of the immediate alternative, and separated by condition. Data are presented for the AB portion of the experiment (top) and the BA portion of the experiment (bottom). Bars indicate ± 1 standard error of the mean.

Although significant in both comparisons, inspection of Figure 3 suggests that the effect of bundling was more robust in the AB condition than in the BA condition. One possible cause for this could have been a general increased tolerance for delay over the course of the experiment. To test this possibility, we examined the residuals from a repeated measure ANOVA in which amount and condition were used as independent variables to predict the percentage of preference for the LL over the 72 sessions. A significant trend was observed, indicating an increasing tolerance for delay during the course of the experiment $F(71, 497) = 2.3, p < .01$. Importantly, the effect of session in the group data was highly linear with no significant quadratic or cubic component ($p > .77$ and $p > .16$, respectively). Thus, the true effect of the experimental manipulation can be reasonably estimated as halfway between that observed in the AB condition (where the observed effect of time augmented the greater preference for LL of the bundled condition) and that observed in the BA condition (where the observed effect of time attenuated the greater preference for LL of the bundled condition). Furthermore, because the effect of time was linear, we were able to reasonably collapse the two A conditions and the two B conditions in some subsequent analyses.

Because the number of trials per session differed considerably across the two conditions, and because sessions were run to satiation, the overall analysis was repeated in two alternative ways: (a) using only the first 15 trials of each session so that the number of trials was equal in the two conditions, and (b) using only the first 30 trials from each standalone session and the first 10 trials from each bundled session. The number of trials retained in this second analysis was generally sufficiently small to exclude trials in which rats were approaching satiation, while the 3:1 proportion of retained trials kept the average level of deprivation constant across conditions by equating reinforcers delivered per trial. Both these reanalyses yielded results similar to those of the overall analyses. Preference for LL alternatives was greater in the bundled condition when using only the first 15 trials ($F [1, 7] = 20.2, p = .003$), and when including only the first 30 trials from the standalone condition and first 10 from the bundled condition ($F [1, 7] = 22.8, p =$

.002).

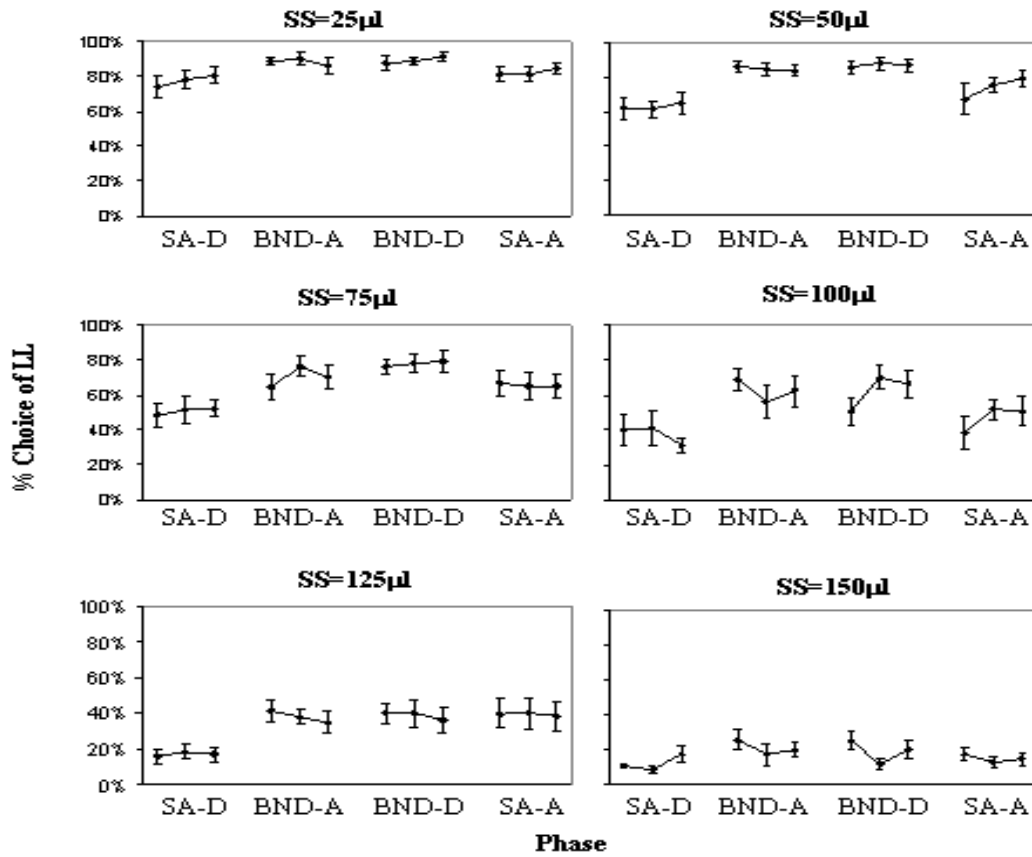



Fig. 4. Percentages of trials that subjects chose the larger-later alternative (150μl delayed by 3s). Each panel provides the summarized data for a particular SS. The four lines on each panel indicate mean preference for the 8 subjects (± 1 standard error) during the four phases of the experiment: (a) standalone descending (SA Dsn), (b) bundled ascending (BND Asn), (c) bundled descending (BND Dsn), and (d) ascending SA Asn). The three points on each line represent the three consecutive sessions at each SS value. For these data, only the first 15 trials from each

session were included. Bars indicate ± 1 standard error of the mean.

Because all choice schedules were conducted for three consecutive sessions, order effects and hysteresis could be examined. Figure 4 presents the preference data, collapsing over subjects, for each of the 72 experimental sessions. In that preference for the LL in general declined over ascending SS conditions and increased over descending SS conditions, hysteresis would be expected in the form of greater preference for LL in ascending phases than descending phases. Also, over the course of the three consecutive sessions at each SS amount, preference for the LL would be expected to decrease in the ascending phase and increase in the descending phase. Although preference for the LL does appear consistently higher in the ascending standalone condition than in the descending standalone condition, the same is not true for the bundling condition. As such, the difference between the descending standalone trials (first 18 sessions) and ascending standalone trials (final 18 sessions) is more likely to be the product of the observed general increase in preference for the LL over the course of the experiment than of hysteresis. Also, hysteresis was not discernable in the examination of the three within-phase sessions at each schedule, and there was little if any tendency for LL preference to be higher in the first than third session at each SS amount during ascending phases, or lower in the first than third session during descending phases.

Figure 5 presents mean choice data for individual subjects as a function of both the amount of the immediate reward and of experimental condition. These data are combined over phases (ascending and descending) and thus each data point is the mean preference of six sessions. The general trends in Figure 3 are present here at the individual level. Lines generally slope downwards, indicating diminishing preference for the LL as the size of the SS increased. And again, at most SS amounts, the lines representing the bundled condition are higher than the lines representing preference in the standalone condition. Using a repeated measure ANCOVA with size of the immediate reward included as a covariate, this differences was individually significant at $p < .05$ in 6 of the 8 subjects (all but 1 and 4).



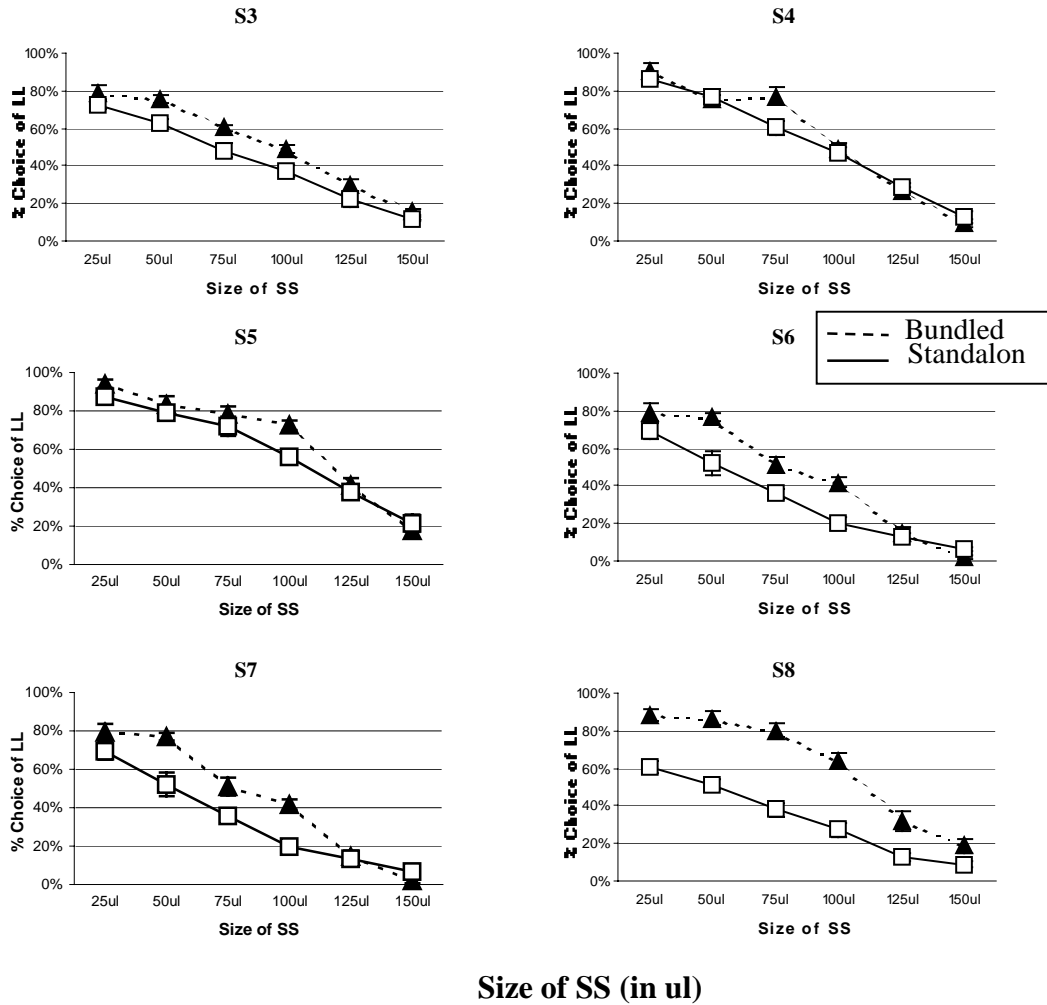


Fig. 5. Individual subject data for percentages of trials that the larger-later alternative was chosen (150ul delayed by 3 s). Data are presented as a function of the size of the immediate alternative and separated by condition. Bars indicate ± 1 standard error of the mean.

Because percentage choice varied with amount in a nearly monotonic fashion across all subjects, indifference points were estimated as the amount at which percentage choice intersected 50%. This was estimated by linear interpolation using the two data points that straddled the 50% level. These amounts are presented for the standalone conditions in the first column of Table 1. The second column of Table 1 displays the discount parameter k implied by the single indifference point when the data from the standalone condition are fitted using Equation 1.

Table 1

Amounts of reward (ul) at the indifference point for individual subjects, discount parameter fit based on Equation 1, predicted indifference amounts in the bundled condition assuming

additivity of multiple rewards, and actual indifference amounts in the bundled condition.

	Indifference amount (in μl) standalone condition	Best-fit k - value	Predicted indifference amount (in μl) for bundled condition	Observed indifference amount (in μl) for bundled condition
S1	122	.08	128	140
S2	97	.18	109	130
S3	75	.33	90	100
S4	97	.18	109	100
S5	110	.12	119	120
S6	82	.28	96	109
S7	57	.54	73	79
S8	53	.61	68	111

The third column of Table 1 presents the amounts at the indifference point predicted by Equation 1 for the bundled condition (assuming additivity of rewards), and the final column presents the actual amounts at the indifference point observed in the bundled condition. Consistent with our hypothesis, the indifference amount for the bundled condition was greater than that of the standalone condition in all 8 rats. Based on a paired t test, the difference in indifference points between the bundled and standalone conditions was highly significant ($t[7] = 4.2, p = .004$). There was some suggestion that the effect of bundling was greater than Equation 1 predicted (nearly significant by paired t test; $t[7] = 2.2, p = .06$).

Based on binomial tests of choices made within each individual test session, the null hypothesis of equal preference was rejected (with $\alpha = .05$) in 66.7% of sessions in the standalone condition, and 50.3% of sessions in the bundled condition, the lower percentage probably being attributable to the lower power of the test in the condition where only a third as many choices were made. Table 2 presents the results of all by-session analyses. For the swing values of 50 μl , 75 μl , and 100 μl , the total number of sessions in which the LL was significantly preferred was greater in the bundled condition. This pattern appeared despite the lower probability of obtaining significant preference in the bundled condition due to the smaller number of trials in each session.

Table 2

Preferences during each experimental session. “S” indicates preference for the smaller-sooner alternative, “L” indicates preference for the later-larger alternative, and “-” indicates that indifference could not be rejected (p > .05). For each cell, the top row indicates preference during the standalone condition and the bottom row during the bundled condition.

		25 ml	50 ml	75 ml	100 ml	125 ml	150 ml
S1	S	LLLLLL	LLLLLS	LLLLLL	LLSLLS	SSSLLL	SS-SSS
	B	LLLLLL	LLLLLL	LL-LLL	LSLSLL	-----	----S-
S2	S	LLLLLL	LLLLLL	---LLL	-----	SSS---	SSSSSS
	B	LLLLLL	LLLLLL	-LLL-	LLLLLL	----L-	-----
S3	S	L-LLLL	---L-	-----	-SS---	SSSSSS	SSSSSS
	B	-LLLLL	LLLLLL	-----	-----	-----	-SSSSS
S4	S	LLLLLL	--LLLL	---LL-	-----	SSSS-	SSSSSS
	B	LLL-LL	--LL-	L---LL	-----	-----	SSSSSS
S5	S	LLLLLL	LLLLLL	---LLL	---L-	SSS---	SSSSS-
	B	LLLLLL	----LL	-L---L	---L-	-----	-S-SS
S6	S	LLLLLL	---LLL	-----	S-----	SSS---	SSSSSS
	B	-LLLLL	-LLLLL	-----	-----	----S	SSSSSS
S7	S	-LLLLL	---LL-	SS-S-	SSSSSS	SSSSSS	SSSSSS
	B	--L-LL	L---LL	-----	-----	S-SSSS	SSSSSS
S8	S	---L-	-----	SS----	SSSS-S	SSSSSS	SSSSSS
	B	LLL-LL	LLLL-L	-L-LLL	-----	-----	SS-SS
Tot(%)	S	L85, S0	L56, S2	L29, S10	L10, S33	L6, S73	L0, S96
	B	L85, S0	L73, S0	L37, S0	L23, S4	L2, S13	L0, S65

Discussion

The results support the hypothesis that bundling pairs of SS-LL choices results in a greater preference for LLs than when choices are made singly. The amount of immediate reward that was equally preferred to a fixed reward quantity delayed by 3 s was significantly greater in the bundled condition than in the standalone condition (Table 1). This primary finding is consistent with prior research demonstrating (a) that temporal discounting occurs according to a hyperbolic function, and (b) that the value of multiple rewards occurring at different delays is roughly the sum of the discounted values of each of those rewards individually.

As noted in the results (see Table 1), the increase in indifference amounts in the bundled condition was generally greater than predicted by Equation 1. One way of reducing this discrepancy would be to raise the $(1 + kD)$ term to a power, fitted to individual subjects. An exponent of less than 1.0 would result in a flattening of the discount curve at longer delays, and thus predict larger shifts towards LLs given bundling. Such a solution would be in accord with parametric analyses of individual human data by Myerson and Green (1995), who found that when an exponent was included in the denominator, the best fit value for the parameter was indeed typically less than 1.0.

The increase in preference for LL's during the course of the experiment is also noteworthy. Mazur and Logue (1978) found increased preference for LL's during a fading procedure in which the delay of the SS was initially equal to the LL, but was gradually reduced. This procedure is analogous to the ascending phases of this experiment in which the SS started small and was increased in 25 μ l increments every 3 sessions. An analysis of the residuals of session preference after controlling for amount and condition did not, however, support this explanation of the increased tolerance for delay over time observed here. When, in addition to session day, a categorical variable was included indicating whether the session occurred in an ascending or descending phase of the experiment, no interaction was observed between the day and phase predictors ($p > .3$). Hence the observed effect of time cannot be attributed to fading.

Several possibilities remain as to the cause of the observed shift towards greater preference for LLs over the course of the experiment. A decrease in discounting may have occurred as a function of (a) aging (the experiment lasted 15 weeks), (b) experience with SS and LL trials generally, or (c) experience with the bundled conditions in particular, which may have caused a decrease in discounting that persisted when the standalone condition was reinstated. The design of the experiment was such that these three factors were each highly or perfectly confounded with one another. Additional data thus would be required to untangle these factors. It is emphasized that this observed shift towards preference for LLs, whatever its cause, cannot account for the primary finding, because greater preference for the LL was observed not only in the AB phases of the experiment, but in the BA phases as well.

The observed increase in preference for LLs when choices are bundled into series is additional evidence that discount curves from delayed rewards are more deeply bowed than exponential curves. The observation of this phenomenon presently in rats and previously in humans (Kirby

and Guastello 2001) suggests that it is not confined to a particular species, nor is it the product of human culture. Rather it seems to reflect a basic property of how organisms respond to reward.

This finding also suggests an answer to the question of why humans often, but not always, achieve consistency of choice over time. Unlike the exponential curve, the hyperbolic curve predicts an incentive for an individual to commit future choices to the course that currently promises the most reward, discounted for delay. Such tactics are sometimes observed—people put their money in investments that are less liquid than would seem to be optimal (Laibson 1997), avoid information about the availability of temptations (Carillo 1999; Metcalfe and Mischel 1999), and avoid arousing their appetites (Mischel and Mischel 1983). However, tactics like these are distinguished by their scarcity—by how little people seem to use them in the everyday exercise of self-control. In ordinary speech people *resolve on* or *intend* courses of action, and, if they are aware of a mechanism at all, refer to it as *willpower*. There has been a large piece missing in the puzzle of how people achieve consistency of choice.

Since antiquity, authors have advocated that impulses could be controlled by choosing according to principle; that is, choosing in categories containing a number of expectable choices rather than just the choice at hand. Aristotle argued that incontinence (*akrasia*) was the result of choosing according to "particulars" instead of "universals" (in Barnes, 1984). Kant argued that the highest kind of decision-making involved making all choices as if they defined universal rules (the "categorical imperative," [(Kant, 1793/1960)]). The Victorian psychologist Sully wrote that will consisted of uniting "particular actions... under a common rule" so that "they are viewed as members of a class of actions subserving one comprehensive end" (Sully, 1884). In recent years behavioral psychologists have followed this approach to decrease pigeons' preference for SSs—Heyman & Tanz (1995) by giving them extra reward for choosing according to "overall" rather than "local" maxima; Siegel & Rachlin (1995) by making choice depend only on every 31st peck—thus arguably creating a molar rather than molecular choice pattern. This latter theory of choice is explored in depth in Rachlin's treatment of the topic (Rachlin, 2000).

Ainslie (1975, 2001) suggested that sufficiently intelligent organisms would come to make choices between LLs and SSs in whole bundles insofar as they interpreted their current choices as cues predicting what they were most likely to choose in the future. Lacking an innate organ for consistent choice, a person will get her best information about her own prospective choices from behavioral observation of herself in similar situations, with her current choice the most germane. The incentives bearing on her current choice will then include not only its direct consequences but also the expected consequences of the bundle of LLs versus SSs that this choice predicts. A current choice of LL will come to predict a whole bundle of LLs and thus be valued more than it would be by itself. More importantly, the discounted value of the whole bundle of LLs may come to exceed the discounted value of the whole bundle of SSs, even though the discounted value of the most imminent SS exceeds the discounted value of its alternative LL (see Figure 1a).

The above hypothesis depends on two component processes: (a) that an increase in the value of LL over SS alternatives can result from bundling expected rewards, and (b) that subjects form bundles through interpretation of current choices as predictive of future ones. Although clearly predicted from prior research, direct experimental evidence that choosing in series reduces

impulsiveness has been lacking until recently. The findings we have reported here, in combination with those of Kirby and Guastello (2001), support the effectiveness of bundling in preventing temporary preferences for SSs.

Evidence for the second proposed component process—the spontaneous formation of bundles of choices based on perceived precedent—remains primarily indirect, based on thought experiments (see Ainslie, 2001, pp. 126-139), phenomenological accounts of the will (Ainslie, 2001, pp. 117-125), and experimental study of interpersonal bargaining games that can serve as models of intertemporal bargaining (Monterosso, Ainslie, Toppi-Mullen, & Gault, 2002). Although the details of these approaches are beyond the scope of this discussion, a simple thought experiment provides suggestive evidence that precedent-based bundling is common in humans: Consider a smoker who is trying to quit, but who craves a cigarette (Monterosso & Ainslie, 1999). Suppose that an angel whispers to her that, regardless of whether or not she smokes the desired cigarette, she is destined to smoke a pack a day from tomorrow on. Given this certainty, she would have no incentive to turn down the cigarette—the effort would seem pointless. What if the angel whispers instead that she is destined never to smoke again after today, regardless of her current choice? Here, too, there seems to be little incentive to turn down the cigarette—it would be harmless. Fixing future smoking choices in either direction (or anywhere in between) evidently makes smoking the dominant current choice. Only if future smoking is in doubt does a current abstention seem worth the effort. But the importance of her current choice cannot come from any physical consequences for future choices; hence the conclusion that it matters as a precedent. Accordingly, when Kirby and Guastello (2001) merely suggested to student subjects that the subjects' current choices might serve as predictions of their future choices, preference for LLs increased, although not as much as when the experimenters directly bundled the choices.

The data of the present study demonstrate that the bundling of even a small temporally extended series of SS and LL pairs can significantly shift preference towards LL choices. With larger series, even greater effects would be expected, although with diminishing returns for each more distant choice (e.g. pairs added to the right in Figure 1a). In conjunction with a mechanism for the spontaneous bundling of choices, the phenomenon demonstrated here offers an account of *willpower* within a deterministic framework.

References

- Ainslie, G. (1974). Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior*, 21, 485-489.
- Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82, 463-496.
- Ainslie, G. (2001). *Breakdown of Will*. New York: Cambridge University Press.
- Ainslie, G., & Herrnstein, R.J. (1981). Preference reversal and delayed reinforcement. *Animal Learning & Behavior*, 9, 476-482.
- Barnes, J. (Ed.). (1984). *The Complete Works of Aristotle: Vol. 1. The revised Oxford translation*. Princeton, NJ: Princeton University Press.
- Baum, W., & Rachlin, H. (1969). Choice as time allocation. *Journal of the Experimental Analysis of Behavior*, 12, 861-874.
- Carillo, J. J. (1999). *Self-control, moderate consumption, and craving*. Unpublished manuscript. Universite Libre de Bruxelles.
- Deluty, M. (1978). Self-control and impulsiveness involving aversive events. *Journal of Experimental Psychology: Animal Behavior Processes*, 4, 250-266.
- Grace, R. (1996). Choice between fixed and variable delays to reinforcement in the adjusting-delay procedure and concurrent chains. *Journal of Experimental Psychology: Animal Processes*, 22, 362-383.
- Heyman, G. M., & Tanz, L. (1995). How to teach a pigeon to maximize overall reinforcement rate. *Journal of the Experimental Analysis of Behavior*, 64, 277-297.
- Kant, I. (1960). *Religion Within the Limits of Reason Alone*. (T. Green & H. Hucken, Trans.). New York: Harper & Row. (Original Work published 1793)
- Kirby, K. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General*, 126, 54-70.
- Kirby, K. N., & Guastello, B. (2001). Making choices in anticipation of similar future choices can increase self-control. *Journal of Experimental Psychology: Applied*, 7, 154-164.
- Kirby, K. N., & Herrnstein, R.J. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science*, 6, 83-89.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 62, 443-479.

- Mazur, J. E. (1984). Tests of an equivalence rule for fixed and variable reinforcer delays. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 426-436.
- Mazur, J. E. (1986). Choice between single and multiple delayed reinforcers. *Journal of the Experimental Analysis of Behavior*, *46*, 67-77.
- Mazur, J.E. (2001). Hyperbolic value addition and general models of animal choice. *Psychological Review*, *108*, 96-112.
- Mazur, J., & Logue, A. (1978). Choice in a "self-control" paradigm: Effects of a fading procedure. *Journal of the Experimental Analysis of Behavior*, *30*, 11-17.
- Metcalf, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review*, *106*, 3-19.
- Mischel, H., & Mischel, W. (1983). The development of children's knowledge of self-control strategies. *Child Development*, *54*, 603-619.
- Monterosso, J., & Ainslie, G. (1999). Beyond discounting: Possible experimental models of impulse control. *Psychopharmacology*, *146*, 339-347.
- Monterosso, J., Ainslie, G., Toppi-Mullen, P., & Gault, B. (2002). The fragility of cooperation: An empirical study employing false-feedback in a sequential iterated prisoner's dilemma. *Journal of Economic Psychology*, *23*, 437-448.
- Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior*, *64*, 263-276.
- Rachlin, H. (2000). *The science of self-control*. Cambridge, MA: Harvard University Press.
- Rachlin, H., & Green, L. (1972). Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior*, *17*, 15-22.
- Siegel, E., & Rachlin, (1995). Soft commitment: Self-control achieved by response persistence. *Journal of the Experimental Analysis of Behavior*, *64*, 117-128.
- Sully, J. (1884). *Outlines of Psychology*. New York: Appleton-Century-Crofts.

Author Note

George Ainslie, Department of Psychiatry, Veterans Affairs Medical Center, Coatesville, PA, and Department of Psychiatry, Temple Medical College.

John Monterosso, Department of Psychiatry, University of Pennsylvania.

We thank Chae Kim, Kathy Meeker, Dawn Lovejoy, Stephanie Kirylyck for technical support and data collection. This research was supported by a Network 4 Competitive Pilot Project Fund

Award from the Department of Veterans Affairs.

Correspondence should be directed to the first author.