Published in final edited form as:

J Exp Anal Behav. 2013 May ; 99(3): 335–345. doi:10.1002/jeab.24.

Sensitivity to Changing Contingencies in an Impulsivity Task

Michael E. Young,

Kansas State University

Tara L. Webb,

Southern Illinois University at Carbondale

Jillian M. Rung, and

Southern Illinois University at Carbondale

Eric A. Jacobs

Southern Illinois University at Carbondale

Abstract

Using a video-game-based escalating interest task, participants repeatedly encountered a reward that gradually increased in value over a 10-second interval. Responding early in the interval netted less immediate reward than responding later in the interval. Each participant experienced four different reward contingencies for waiting. These contingencies were changed three times as the experiment proceeded. Behavior tracked these changing contingencies, but wait times reflected long-term carryover from the previously assigned contingencies. Both the tendency to respond slowly and the optimality of behavior were affected by the order of contingencies experienced. Demographic variables only weakly predicted behavior, and delay discounting rate in a hypothetical money choice task predicted choice only when the contingencies in the game were weaker.

Choosing a smaller sooner reward over a larger later one can have benefits and costs depending on the conditions under which the choice is made. Research on delay discounting often focuses on the costs of choosing the smaller sooner outcome over the larger later one and defines this preference as impulsivity (Ainslie, 1974, 1975; Rachlin & Green, 1972). There are conditions, however, where one should take the smaller sooner outcome (cf. Ito & Nakamura, 1998; Logue, King, Chavarro, & Volpe, 1990). For example, if someone was offered the choice between \$100 now versus \$125 in five years, taking the smaller sooner reward is preferred if the money could be invested for those five years at an annual interest rate of 5% (this would create a net return of approximately \$128 at the end of five years). Similarly, the smaller sooner outcome may be preferred if it increases the number of opportunities to choose (i.e., if there is little or no imposed delay before a similar choice can be made again), or it is available at a moment when its utility is greatest (e.g., to pay this month's rent or to shoot an oncoming enemy).

Contact: Michael Young, 492 Bluemont Hall, Department of Psychology, Kansas State University, Manhattan, KS 66506-5302, 785/532-0602, michaelyoung@ksu.edu.

Michael E. Young, Department of Psychology, Kansas State University, 492 Bluemont Hall, Manhattan, KS, 66506-5302, Tara L. Webb, Jillian M. Rung, and Eric A. Jacobs, Department of Psychology, Mailcode 6502, Southern Illinois University, Carbondale, IL, 62901-6502.

Queries should be sent to the first author at michaelyoung@ksu.edu.

The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Experience with the utility of smaller sooner versus larger later outcomes is central to identifying which one is optimal under a given set of conditions. Furthermore, experience under a set of conditions where taking the sooner (or later) outcome is optimal may influence preference when the conditions dictate that each choice is of equal value or the later (or sooner) outcome is optimal. A behavioral preference that arises due to early contingencies is expected to carry over to a new situation when the contingency is changed. The goal of the present project was to identify the nature of this carryover when conditions changed multiple times within a single session.

Previous research in which the contingencies changed within a session always attempted to increase the relative preference for the larger later outcome (e.g., Logue & King, 1991; Mazur & Logue, 1978; Schweitzer & Sulzer-Azaroff, 1988). These studies sought to increase self-control in pigeons, adults or children through procedures in which (a) the tolerance for the longer delays was increased by gradually lengthening the delay to the larger later outcome or (b) exposure to the larger later outcome was produced by having both outcomes initially delayed but gradually decreasing the delay to the smaller outcome. The success of these procedures relative to control groups in which there was no gradual change in delay duration suggests that early experiences with different delays persists when the delay changes.

We were interested in the sequence of conditions that would both alter the overall likelihood of waiting as well as sensitivity to contingency changes. If particular contingency sequences produce more waiting or more optimal behavior, this result could be used to create better training schedules. Some situations necessitate short wait times for molar maximization, and these experiences likely carry over to situations where such behavior is detrimental. Furthermore, the impact of early experiences with different contingencies may not only persist into the subsequent conditions involving a new contingency but may produce cumulative effects that alter behavior in later conditions. It may seem self-evident that behavior *will* change, but how rapidly will it do so, and will early experiences linger and for how long? To study these effects, we used Young, Webb, and Jacobs' (2011) escalating interest (EI) paradigm because the contingencies for waiting can be easily modified within-session resulting in rapid shifts in behavior. Young et al. shifted contingencies very rapidly within a session with as many as 28 different contingencies experienced within an hour. In order to examine the unfolding of participants' behavior, longer term exposure to each contingency was necessary thus requiring an adaptation of their task.

The Escalating Interest Task

In Young et al.'s (2011) EI task, participants make a series of identical responses (in their experiment, firing a weapon in a video game) each of which leads to an outcome of varying magnitude (damage done to the player's target). Firing more rapidly corresponds to short inter response times (IRTs) and produces smaller sooner outcomes, whereas firing less rapidly corresponds to longer IRTs and produces larger later outcomes. Their experiment was designed such that a weapon charge bar visually indicates the amount of damage that will be done by pulling the weapon's trigger; the charge bar gradually fills as time passes and empties when a shot is made. There are no limits on the number of times a player can fire.

The key variable in their design is the mathematical parameter, *power*, of a super ellipsoid function that controls the way in which the weapon recharges:

$$damage = 100 \bullet \left(1 - \left(\frac{(10-t)}{10} \right)^{power} \right)^{1/power} \quad (1)$$

where t is the time elapsed since the last shot. As shown in the top of Figure 1, when power is 1.00, the weapon recharges linearly to its maximum charge (this occurred after 10 s in Young et al., 2011 but varied between-subjects from 5 s to 20 s in Young, Webb, Sutherland, & Jacobs, in press). When power is less than 1.00, the weapon recharges slowly early in the interval and more rapidly later (positive acceleration). When power is greater than 1.00, the weapon recharges rapidly early in the interval and more slowly later (negative acceleration).

In the EI task, behavior that maximizes the overall reinforcement density differs depending on the power parameter (see the bottom of Figure 1). When power is 1.00, any IRT between 0 s and 10 s produces the same rate of destruction of the player's targets, assuming that each shot impacts the target. However, when power is less than 1.00, damage density is maximal when the participant waits for the full charge and no longer. In contrast, when the power parameter is greater than 1.00, the damage density is maximal when the participant fires as rapidly as possible. The degree of departure from optimal density for suboptimal IRTs is greater as the power value is increasingly different from 1.00. For example, a participant who consistently waits 5 s between shots will destroy their targets at 52% of the optimal density for a power of 0.75, but the detrimental impact of these shorter-than-optimal IRTs is larger when the power is even lower. If the same participant waits 5 s between shots with a power of 0.50, the destruction of the targets is 17% of the optimal density.

Changing Contingencies

Because prior studies of human impulsive choice have either focused on a snapshot of current behavioral preference through the use of the delay discounting questionnaire (Bickel, Odum, & Madden, 1999; Madden, Begotka, Raiff, & Kastern, 2003; Rachlin, Raineri, & Cross, 1991) or on asymptotic or total performance (e.g., Flora & Pavlik, 1992; Marsh, Dougherty, Mathias, Moeller, & Hicks, 2002; Reynolds, Richards, & de Wit, 2006), very little is known about how behavior will track the changing reward contingencies in impulsivity tasks (cf. Mazur, Blake, & McManus, 2001; Robinson et al., 2009). Although Young et al. (2011; in press) documented that people's IRTs in the EI task will change as the power level changes, the exposure to each new contingency was so short that asymptotic behavior was unlikely to have been observed and there was little opportunity to examine behavior under transition.

To ensure adequate observation of behavior under different power contingencies, in our experiment the contingencies progressed through a series of 4 conditions each corresponding to a different value of the power parameter of Equation 1 rather than the 28 conditions experienced by most of Young et al.'s (2011) participants. We used a sequencing design for the four power values of 0.50, 0.75, 1.00, and 1.25 (see Figure 1) so that each value appeared in each ordinal position. These values were chosen because the 0.50 and 0.75 values tend to produce waiting times greater than 5 s whereas those of 1.00 and 1.25 tend to produce waiting times less than 5 s (Young et al., 2011). If particular sequences of exposure to these contingencies produce more waiting or greater optimality, there is an opportunity to leverage such a finding to produce desired behavioral change by sequencing experiences in specific ways.

Method

Participants

A total of 48 introductory psychology students at Southern Illinois University at Carbondale received course credit for their voluntary participation. There were 23 women and 17 men (8 did not report their sex).

Procedure

Participants entered the laboratory and either completed a set of questionnaires after playing our customized video game or played the game after completing the questionnaires. The video game is a first-person-shooter style game; an edited movie clip of the escalating interest task is available here: <http://www.k-state.edu/psych/research/young/suppmaterial.html>. Participants received verbal instructions covering navigation of the virtual environment and the nature of the task. Participants were directed to “destroy all the orcs in each level as effectively as possible.” Although participants were informed about the nature of the weapon charge and told to be mindful of it, they were given no advice on how long they should wait between shots.

The Torque Game Engine (obtained from www.garagegames.com) was adopted as the platform for game development. The game consisted of four levels, each containing seven separate regions with each region populated by three visually identical orcs. Once the 21 orcs in the gaming environment were destroyed (i.e., the first level was completed), the game reset to the same landscape repopulated with the same orcs but with the player starting at a random location within the region and the behavior of their weapon having changed (i.e., the power value had changed). This change represents the beginning of the second level of the game. Once these 21 orcs were destroyed, the game again reset thus starting the third level, etc.

For simplicity, the orcs were stationary and oriented toward a target region (e.g., a building) that the player was directed to protect. Every 4 s (on average), each orc fired its weapon. In each 3-orc region, the firing of one of the orcs produced contingent explosions in the target region; this design was used to make the player’s goal more realistic. Participants typically approached one trio of targets, fired until all were destroyed, then traveled to the next region containing three targets to destroy each of them, etc.

The player’s task was to destroy all of the orcs within each game level. Each orc required 2.7 fully-charged shots to be destroyed. Thus, a player who waited the full 10 s for every shot (which is not optimal behavior) would complete the task in 37.8 minutes (i.e., $2.7 \text{ shots} \times 10 \text{ sec} \times 3 \text{ orcs} \times 7 \text{ regions} \times 4 \text{ levels}$) plus travel time between regions and the time required to load each level. Four power values were used (0.50, 0.75, 1.00, and 1.25), one for each level, with condition sequences systematically varying across participants using a Latin square design. The condition sequences are depicted as part of Figures 2 and 3.

Students completed demographic sheets that asked their sex, age, cigarette smoking (using the Fagerström Test of Nicotine Dependence, Heatherton, Kozlowski, Frecker, & Fagerström, 1991), and previous video game experience (see Young et al., 2011, for details). Each student also completed a paper-and-pencil delay discounting questionnaire with hypothetical choices between smaller sooner and larger later monetary rewards (cf. Beck & Triplett, 2009; Smith & Hantula, 2008). Half of the students completed the questionnaires before the video game task and half completed the questionnaire after the task.

In our analyses, we excluded IRTs that exceeded 20 s in duration with the assumption that these were contaminated by inattention or excessive travel time. This resulted in a loss of

1.7% of the data with no more than 5.9% lost from any single participant. Because the distribution of IRTs in the EI task was strongly bimodal (with most IRTs less than 1 s or near 10 s), we dichotomized the IRTs into two groups: IRTs less than 5 s and those greater than or equal to 5 s. Although dichotomization can result in a significant loss of power, this loss is much smaller when the distribution is highly bimodal (Irwin & McClelland, 2003). The resulting binomial variable was analyzed using repeated measures logistic regression (specifically, linear mixed effects modeling in *R* with the specification of a binomial error distribution). The results of the logistic regression will characterize the individual participants' probability of waiting more than 5 s as a function of the *power* parameter (see Figure 1). Because the median IRT for IRTs that exceeded 10.0 s was 10.6 s, for simplicity in display we treated these longer IRTs as equivalent to 10.0 s.

Results

All participants completed the video game. The median time to complete the game was 38 minutes with 25th and 75th percentiles of 34 and 42 minutes and minimum and maximum times of 32 and 50 minutes.

Participants' firing rate varied in response to changes in the power value of the weapon in the predicted direction (Figures 2, 3, and 4). Figure 2 shows one representative participant from each of the four condition sequences along with a corresponding logistic curve fit. Behavior in each level was divided into deciles to illustrate behavioral change across the level. The proportion of non-dichotomized IRTs between 0 and 10.5 s for each individual is represented by the width of the contour plot for each decile. This method of plotting nicely illustrates the typically bimodal nature of the data, with participants usually responding very quickly (IRTs less than 3.0 s) or very slowly (IRTs greater than 7.0 s). Because this method of plotting is cumbersome for 48 participants, Figure 3 shows only the best logistic curve fits derived from a linear mixed effects model for individual participants (analysis details provided later).

In each of the four sequence conditions, participants initially tended to respond quickly. As the participant progressed through the first level, those who initially encountered the 0.50 and 0.75 power values increasingly shifted toward longer IRTs whereas those assigned to power values of 1.00 and 1.25 continued to show short IRTs. Once these patterns of responding quickly or slowly were established in the first level, they tended to persist through subsequent levels. Each participant in each row of Figure 3 was designated with a different line (see the on-line version of the manuscript to see these in color); participants who tended to wait longest between shots in the initial level also continued to wait longer in subsequent levels when the power value was changed. Although behavior tracked the changing contingencies in all four groups with more short IRTs in the 1.00 and 1.25 conditions and more long IRTs in the 0.50 and 0.75 conditions, it is clear that prior experience with different contingencies affected responding under subsequent contingencies.

To determine if participants who tended to produce the longest IRTs in one game level also tended to produce the longest IRTs in subsequent levels (i.e., to assess test-retest reliability), we calculated each participant's mean IRT in each game level and performed cross-level correlations. We used the mean IRT because the median IRT is very unstable for bimodal data. The resulting reliability measures are summarized in Table 1. Participants' relative mean IRTs in the first level were strongly related to their relative mean IRTs in the second and third levels but unrelated to that in the last level. Although mean IRTs in the third level were significantly related to those produced in the second level, they were more strongly related to those from the first. Mean IRTs in the fourth level of the game were largely unrelated to those in the previous levels except for the third.

To facilitate evaluation of sequence effects, Figure 4 plots spline fitted curves of grouped raw-data with different curves corresponding to different sequences of power values. IRTs for this fit were statistically sampled to balance the sample sizes – each participant contributed approximately 50 IRTs from each power level thus ensuring adequate sample size while achieving greater balance (the original range varied from 37 to 1278 responses in a condition, $Mdn = 91$; after sampling, the range varied from 35 to 67, $Mdn = 48$).

In the level in which power was 0.50 (the contingency that most strongly favors waiting), participants' IRTs approached 10 s by the end of the level. However, the initial wait times were quite short when this was the first contingency experienced or if the only prior experience was with a power of 1.00. In contrast, participants in the other two sequences initially had longer wait times.

A similar pattern was observed for the level in which the power was 0.75 and most participants showed wait times approaching 10 s by the end of the level. If the 0.75 contingency was the first experienced or followed levels in which the power values were 1.00 or 1.25, initial wait times were short but increased steadily over the level. When the previous (and only) level involved another contingency that discouraged short IRTs (power = 0.50), IRTs started long and remained so. More interestingly, the initial IRTs were longer when the previous experience included the 0.50 power value even when it had been experienced two levels before with a 1.25 power value intervening.

For the level involving the 1.00 power value that does not differentially reinforce IRTs between 0 and 10 s, IRTs were short if there was no prior experience with other values or if the only other experience was with a power value (1.25) for which short IRTs are differentially reinforced. If their prior levels had included the power values that reinforce longer IRTs (0.50 and 0.75), then the participants tended to wait longer between shots when subject to the 1.00 power value, especially if their prior experience did not include the power of 1.25.

Finally, for the level with the 1.25 power value, IRTs were generally short throughout the level especially if it was the first value experienced. Interestingly, in all of the experienced condition sequences there was relatively little behavioral change in the level for the average participant. Furthermore, Figure 3 reveals that this was not an averaging artifact; the low rate of change was very common at the individual participant level where fits were relatively flat.

Analysis Details

All effects were confirmed by a linear mixed effects analysis with condition sequence (1 through 4) and power value as categorical predictors and time-in-level as a continuous predictor (Pinheiro & Bates, 2004). Because the outcome variable, IRT, was strongly bimodal and thus violated the assumption of normality, we dichotomized the IRTs into those less than or equal to 5 s versus those greater than 5 s and specified the outcome variable as binomial (i.e., we performed a repeated measures logistic regression). The time-in-level predictor was centered so that the intercept estimates correspond to the mean IRT in the middle of the game level and in order to reduce multicollinearity between the interactions involving time-in-level and its main effect.

A mixed effects model (also known as a multilevel model) is the modern equivalent of a repeated measures analysis. Generalized mixed effects models, however, allow for the specification of non-Gaussian outcome distributions, use maximum likelihood estimation rather than error minimization due to its superiority in handling unbalanced data, and generate parameter estimates at the group and individual level simultaneously. If there is

sufficient evidence that a parameter estimate varies between individuals, this parameter is designated as both a fixed effect (thus producing an estimate at the group or condition level) and a random effect (producing a separate estimate for each individual participant).

The best fitting model (as determined by the Bayesian Information Criterion score) included all three variables and all of their interactions and time-in-level as a random slope effect that varied across participants. Table 2 shows the results in terms of the best fitting parameter estimates for the centered intercepts and slopes. Note that these estimates are in a transformed logistic space as necessitated by the use of a logistic regression. Thus, a centered intercept estimate of 0.0 corresponds to a predicted 50% likelihood of waiting more than 5 s at the midpoint of the game level. The full results of the regression are available from the first author.

It is evident from Table 2 that the mean IRTs (centered intercepts) systematically decreased as the power value increased thus conceptually replicating Young et al. (2011), but now with prolonged exposure to each power. It is also clear that behavior changed more dramatically from the beginning to the end of a game level (i.e., there were steeper slopes) for power values that differentially reinforced longer IRTs (0.50 and 0.75) than for the other power values. The outcome most germane to our interest in the influence of early experience was the much greater likelihood of generating shorter IRTs (Mean intercept estimates of -1.4 and -0.8) for participants whose first experience in the game was either the 1.25 or 1.00 power levels. However, due to the use of a Latin square, these participants also were the only ones to receive the 0.50 or 0.75 power levels as their last condition.

Deviation from Optimality

Participants' behavior clearly varied as a function of the sequence of conditions experienced, with those beginning with a contingency that encouraged short IRTs showing overall shorter IRTs throughout. However, an open question is whether a particular sequence of experiences tended to produce more or less optimal behavior. Thus, in our next analysis we examined whether people tended to wait when they should (power < 1.00) and not to wait when they should not (power > 1.00 , see bottom of Figure 1). Because of the bimodal distribution of the deviation scores, we judged a deviation from the optimal IRT that was less than 5 s as optimal whereas a deviation greater than 5 s was categorized as not optimal. Predictors were the time-in-level (both as a fixed and random effect), the condition sequence (labeled 1 through 4), and their interaction. Because the 1.00 power provides no differential incentive for IRTs between 0 and 10 s, levels involving a 1.00 value were not included in the optimality analysis.

There was no main effect of condition ($ts < 1.8$, $ps > .10$), but participants' responding became more optimal as a game level progressed, with the rate of learning varying across conditions. Participants assigned to the first sequence (0.50-0.75-1.00-1.25) and the fourth sequence (0.75-1.25-0.50-1.00) showed the fastest learning rates (slopes of 1.99 and 1.81, respectively) whereas those assigned to the second sequence (1.25-1.00-0.75-0.50) and third sequence (1.00-0.50-1.25-0.75) showed statistically significantly slower learning rates (slopes of 1.55 and 1.30, respectively), $ps < .01$.

Individual Differences Variables

When sex was included as a potential predictor of the overall likelihood of producing a short or long IRT, differences in this likelihood for each power level, and rate of change within level, the model was significantly improved ($\chi^2(5) = 3726$) with a statistically significant Sex $\times$ Time-in-Level $\times$ Power interaction ($ps < .05$). Overall, men showed slightly steeper IRT slopes than women, but the difference was quite small.

When discounting rate on the paper-and-pencil task ($\log(k)$) was included as a potential predictor (without sex), the model was also significantly improved ($\chi^2(5) = 2306$) with a statistically significant Discount Rate $\times$ Power interaction ($p < .05$). Specifically, participants with higher k values were more likely to produce shorter IRTs in the 0.75 and 1.00 conditions. The effect was quite small, however, generating about a 4–6% difference in the probability of waiting for those with the lowest versus highest discount rates. There was no relation between k and IRT for the 0.50 and 1.25 power conditions.

When task order was included as a potential predictor (without sex and $\log(k)$), the model was again significantly improved ($\chi^2(5) = 1982$) with a statistically significant Task Order $\times$ Power interaction ($p < .05$). Specifically, participants who played the game after completing the demographics form and the paper-and-pencil discounting task produced slightly more long IRTs in the 1.25 condition.

In general, the predictors of individual differences (sex, discount rate from a paper-and-pencil task, and task order) produced statistically significant, but relatively small, effects. Given that the effects were usually idiosyncratically related to specific power values, they should be treated with caution and replicated to ensure their robustness.

Discussion

A critical contribution of the present study was the ability to observe behavioral adaptation to the changing contingencies for short and long IRTs within 10 to 15 minutes of observation. In most studies of delay discounting in people, the tasks require a series of choices that collectively determine the indifference point thus preventing the research from identifying transitional behavior as contingencies change. In the EI task, each response produces an IRT that provides rich information about the participant's disposition toward waiting under the programmed conditions. The clarity with which we were able to observe behavioral change over small time scales revealed the behavioral dynamics in response to changing contingencies in an impulsivity task. Furthermore, people showed significant individual differences that persisted across conditions (see Table 1 and Figure 3). Although some other experiential impulsivity tasks are also structured to allow the study of people's behavioral change as the task proceeds, there are very few attempts to examine this finer level of analysis (for one example, see Moustafa, Cohen, Sherman, & Frank, 2009).

We also examined the relationship between behavior on the EI task and a common measure of impulsive choice, the discounting rate, k , derived from a paper-and-pencil task. When the contingencies strongly differentially reinforced short or long IRTs (the powers of 0.50 and 1.25), k did not predict game IRTs. But, when the contingencies were moderate or absent (powers of 0.75 and 1.00), k weakly predicted IRTs. The use of moderate contingencies may be more effective at documenting individual differences in discounting rates, and more extreme contingencies likely prompted greater sensitivity to the consequences of not behaving optimally regardless of individual differences.

The ability to examine behavioral change under rapidly changing contingencies revealed that the emerging sensitivity to a new set of conditions was clearly influenced by the previously experienced conditions (see Figure 4). In some cases, this effect of history was quite dramatic. For example, the 1.25 power condition produced significantly longer IRTs and little behavioral change within the condition when it was preceded only by the 0.75 power condition whereas in every other condition sequence the 1.25 produced a high proportion of short IRTs for all or nearly all participants by the end of the game level.

The persistent history effects suggest that there are a series of experiences that will produce the fastest behavioral adaptation to the consequences of choosing a smaller sooner versus a larger later reward. Alternatively, there may be experiences that tend to produce an enduring tendency toward choosing the sooner outcome (or the later one) regardless of the current contingencies. Our analyses revealed that participants who were assigned to the 1.25-1.00-0.75-0.50 and 1.00-0.50-1.25-0.75 condition sequences produced shorter IRTs overall (see Table 2) and to show slower behavioral adaptation toward optimal choice. We believe that these differences may be a function of the initially experienced contingency. Given that participants' natural behavioral propensity was toward rapid responding, early experience with a contingency that rewards slower responding may produce greater adaptation away from rapid responding when this type of contingency is reintroduced. This result suggests that early experiences with contingencies that produce waiting may not only produce proximate changes in behavior but may produce longer-term tendencies to adapt to those contingencies when they are reinstated. Whether these benefits are merely short term or longer lasting awaits further study (cf. Kellam et al., 2008).

Because of the consistently observed relationship between drug abuse and impulsivity (for a review, see MacKillop et al., 2011; Perry & Carroll, 2008), the study of individual differences in impulsivity using behavioral tasks has intensified over the past 20 years. Impulsivity tasks may serve as good assessment tools for the early identification of risk for later drug abuse. Ultimately, however, it is critical to identify the specific environmental variables that can give rise to more/less impulsive choice regardless of the individual's propensity because this knowledge provides a tool for behavioral change. If certain contingency sequences that produce self-controlled behavior continue to affect behavior under subsequent conditions, there is hope that the right training at the right time may have lasting effects that can benefit the organism.

Acknowledgments

This research was supported by NIDA 1R15-DA026290.

References

- Ainslie GW. Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior*. 1974; 21:485–489. [PubMed: 16811760]
- Ainslie GW. Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*. 1975; 82:463–396. [PubMed: 1099599]
- Beck RC, Triplett MF. Test–retest reliability of a group-administered paper–pencil measure of delay discounting. *Experimental and Clinical Psychopharmacology*. 2009; 17:345–355. [PubMed: 19803634]
- Bickel WK, Odum AL, Madden GJ. Impulsivity and cigarette smoking: delay discounting in current, never, and ex-smokers. *Psychopharmacology*. 1999; 146(4):447–454. [PubMed: 10550495]
- Flora SR, Pavlik WB. Human self-control and the density of reinforcement. *Journal of the Experimental Analysis of Behavior*. 1992; 57:201–208. [PubMed: 16812652]
- Heatherton TF, Kozlowski LT, Frecker RC, Fagerström KO. The Fagerström Test for Nicotine Dependence: A revision of the Fagerström Tolerance Questionnaire. *British Journal of Addiction*. 1991; 86:1119–1127. [PubMed: 1932883]
- Irwin JR, McClelland GH. Negative consequences of dichotomizing continuous predictor variables. *Journal of Marketing Research*. 2003; 40:366–371.
- Ito M, Nakamura K. Humans' choice in a self-control choice situation: Sensitivity to reinforcer amount, reinforcer delay, and overall reinforcer density. *Journal of the Experimental Analysis of Behavior*. 1998; 69:87–102. [PubMed: 9465415]

- Kellam SG, Brown CH, Poduska J, Ialongo N, Wang W, Toyinbo P, et al. Effects of a universal classroom behavior management program in first and second grades on young adult behavioral, psychiatric, and social outcomes. *Drug and Alcohol Dependence*. 2008; 95:S5–S28. [PubMed: 18343607]
- Logue AW, King GR. Self-control and impulsiveness in adult humans when food is the reinforcer. *Appetite*. 1991; 17:105–120. [PubMed: 1763904]
- Logue AW, King GR, Chavarro A, Volpe JS. Matching and maximizing in a self-control paradigm using human subjects. *Learning and Motivation*. 1990; 21:340–368.
- MacKillop J, Amlung M, Pryor LR, Ray LA, Sweet L, Munafo MR. Delay reward discounting and addiction: A meta-analysis. *Psychopharmacology*. 2011; 216:305–321. [PubMed: 21373791]
- Madden GJ, Begotka AM, Raiff BR, Kastern LL. Delay discounting of real and hypothetical rewards. *Experimental and Clinical Psychopharmacology*. 2003; 11:139–145. [PubMed: 12755458]
- Marsh DM, Dougherty DM, Mathias CW, Moeller FG, Hicks LR. Comparisons of women with high and low trait impulsivity using behavioral models of response-disinhibition and reward-choice. *Personality and Individual Differences*. 2002; 33:1291–1310.
- Mazur JE, Blake N, McManus C. Transitional choice behavior on concurrent-chain schedules. *Behavioural Processes*. 2001; 53:171–180. [PubMed: 11334705]
- Mazur JE, Logue AW. Choice in a "self-control" paradigm: Effects of a fading procedure. *Journal of the Experimental Analysis of Behavior*. 1978; 30:11–17. [PubMed: 16812082]
- Moustafa AA, Cohen MX, Sherman SJ, Frank MJ. A role for dopamine in temporal decision making and reward maximization in Parkinsonism. *Journal of Neuroscience*. 2009; 28:294–304.
- Perry JL, Carroll ME. The role of impulsive behavior in drug abuse. *Psychopharmacology*. 2008; 200:1–26. [PubMed: 18600315]
- Pinheiro, JC.; Bates, DM. *Mixed-effects models in S and S-PLUS*. New York: Springer; 2004.
- Rachlin HC, Green L. Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior*. 1972; 17:15–22. [PubMed: 16811561]
- Rachlin HC, Raineri A, Cross D. Subjective probability and delay. *Journal of the Experimental Analysis of Behavior*. 1991; 55:233–244. [PubMed: 2037827]
- Reynolds B, Richards JB, de Wit H. Acute-alcohol effects on the Experiential Discounting Task (EDT) and a question-based measure of delay discounting. *Pharmacology, Biochemistry and Behavior*. 2006; 83:194–202.
- Robinson ES, Eagle DM, Economidou D, Theobald DE, Mar AC, Murphy ER, et al. Behavioural characterisation of high impulsivity on the 5-choice serial reaction time task: Specific deficits in 'waiting' versus 'stopping'. *Behavioural Brain Research*. 2009; 196:310–316. [PubMed: 18940201]
- Schweitzer JB, Sulzer-Azaroff B. Self-control: Teaching tolerance for delay in impulsive choice. *Journal of the Experimental Analysis of Behavior*. 1988; 50:173–186. [PubMed: 3193054]
- Smith CL, Hantula DA. Methodological considerations in the study of delay discounting in intertemporal choice: A comparison of tasks and modes. *Behavior Research Methods*. 2008; 40:940–953. [PubMed: 19001385]
- Young ME, Webb T, Jacobs EA. Deciding when to "cash in" when outcomes are continuously improving: An escalating interest task. *Behavioural Processes*. 2011; 88:101–110. [PubMed: 21871951]
- Young ME, Webb TL, Sutherland SC, Jacobs EA. Magnitude effects for experienced rewards at short delays in the escalating interest task. *Psychonomic Bulletin and Review*. in press.

Figure 1.

Top: The four superellipsoid functions generated by Equation 1 using a power value of 1.25 (solid line, top curve), 1.00, 0.75, and 0.50 (dotted line, bottom curve). Bottom: The corresponding rate of damage produced by a consistent IRT duration for each power value.

Figure 2.

One representative participant for each of the condition sequences (indicated in each row of the figure). The probability of waiting longer than 5 s is plotted as a function of proportion of time in level across the four game levels (left to right). The curves are the best fitting curves derived from a mixed effects analysis that allowed slopes and intercepts to vary across participants.

Figure 3. Individual participants' probability of waiting longer than 5 s as a function of proportion of time in level across the four condition sequences. The curves are the best fitting curves derived from a mixed effects analysis that allowed slopes and intercepts to vary across participants.

Figure 4.

A smoothed spline fit of the average participants' IRT as a function of proportion of time in level for each power level for each of the four condition sequences (the latter are designated in the legend). Smoothing was based on a sampling of behavior to avoid bias due to differences in sample sizes across conditions and participants.

Table 1

Correlations of Mean IRT across Levels for Individual Participants.

Level	Level		
	1	2	3
2	0.62*		
3	0.58*	0.35*	
4	-0.18	-0.19	0.31*

Note: Statistically different from a 0.00 correlation is designated with an asterisk.

Best Fitting Parameter Estimates for the Centered Intercept Parameter (Top) and the Slope Parameter (Bottom) as Derived from a Repeated Measures Logistic Regression.

Table 2

Condition Sequence	0.50	0.75	1.00	1.25	Mean
0.50-0.75-1.00-1.25	1.1	2.8	0.4	-1.1	0.8
1.25-1.00-0.75-0.50	1.3	-0.6	-2.7	-3.4	-1.4
1.00-0.50-1.25-0.75	0.5	0.7	-2.4	-1.9	-0.8
0.75-1.25-0.50-1.00	1.8	-0.1	-0.3	-0.2	0.3
Mean Centered Intercept	1.2	0.7	-1.2	-1.6	

Condition Sequence	0.50	0.75	1.00	1.25	Mean
0.50-0.75-1.00-1.25	6.3	1.5	-1.2	0.9	1.9
1.25-1.00-0.75-0.50	-0.7	2.3	1.0	-0.6	0.5
1.00-0.50-1.25-0.75	3.4	2.1	-1.4	-1.5	0.6
0.75-1.25-0.50-1.00	1.4	5.1	0.9	-0.3	1.8
Mean Slope	2.6	2.7	-0.2	-0.4	

Note: Standard error is approximately 0.36 for the intercept estimates.

Note: Standard error is approximately 0.56 for the slope estimates.