Magnitude effects for experienced rewards at short delays in the escalating interest task

Michael E. Young, Tara L. Webb, Steven C. Sutherland, Eric A. Jacobs

Magnitude Effects for Experienced Rewards at Short Delays in the Escalating Interest Task

Michael E. Young

Kansas State University

Tara L. Webb, Steven C. Sutherland, Eric A. Jacobs

Southern Illinois University at Carbondale

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

Abstract

A first-person shooter video game was adapted for the study of choice between smaller sooner and larger later rewards. Participants chose when to fire a weapon that increased in damage potential over a short interval. When the delay to maximum damage was shorter (5 – 8 s), people showed greater sensitivity to the consequences of their choices than when the delay was longer (17 – 20 s). Participants also evidenced a magnitude effect by waiting proportionally longer when the damage magnitudes were doubled for all rewards. The experiment replicated the standard magnitude effect with this new video game preparation over time scales similar to those typically used in nonhuman animal studies and without complications due to satiation or cost.

No one likes to wait, but the willingness to do so certainly depends on the length of the wait and what will be received once the wait is over. In empirical studies of the willingness to wait, participants are given a series of choices involving the tradeoff between a smaller immediate reward versus a larger delayed reward (for a collection of articles on the topic, see Madden & Bickel, 2010). As the delay to the larger reward increases, people are more likely to choose the smaller sooner reward over the larger later one. In empirical studies of the effect of reward magnitude, research has revealed that people demonstrate a greater willingness to wait when both the sooner and later outcomes have greater value (Green, Myerson, & McFadden, 1997; Jimura, Myerson, Hilgard, Braver, & Green, 2009). For example, most people would rather have $10 now than $20 in 6 months (the $20 is discounted to less than half its immediate value due to the delay), but people are much more likely to wait if the tradeoff is between $1 million now and $2 million in 6 months. This magnitude effect has been robust in studies of humans but difficult to replicate in nonhuman species (e.g., Calvert, Green, & Myerson, 2010; Richards, Mitchell, de Wit, & Seiden, 1997).

It would not be surprising if human and nonhuman species show different behaviors and tactics in response to a particular environmental challenge like choosing between smaller sooner and larger later outcomes. Their ecological niches are quite different, their neural machinery varies, and their basic biological differences produce dissimilar demands on the importance of immediacy (e.g., metabolism, locomotive speed, and prey vulnerability). Interpreting this species difference, however, has been complicated by the fact that human studies typically involve hypothetical rewards and imagined delays on the order of days to years (e.g., Estle, Green, Myerson, & Holt, 2007; Hinvest & Anderson, 2010; Madden, Begotka, Raiff, & Kastern, 2003; Rachlin, Raineri, & Cross, 1991), whereas nonhuman studies always involve consumable

rewards with very short delays on the order of seconds to minutes (e.g., Green, Myerson, Holt, Slevin, & Estle, 2004; Mazur, 1987; Richards et al., 1997). In two studies involving humans and consumable rewards, however, people showed a magnitude effect even when the delays were experienced, not hypothetical, and on the order of seconds (Jimura et al., 2009; Logue & King, 1991). The use of consumables, however, provides a natural limit on the number of choices before participants are satiated. Jimura et al. and Logue and King's designs involved fewer than 20 trials in a session and maximum delays of 60 s.

The need to use consumable rewards in order to observe discounting across short delays seemed odd in light of emerging technologies that may condition people to expect immediate outcomes. Small delays in a computer's response to a command, the sending of a text message, or the loading of a web page prompt common complaints among technology users. This apparent intolerance for small delays in interactions with technology was leveraged by Young, Webb, and Jacobs (2011) who recently pioneered a novel video game procedure for studying the tradeoff between smaller sooner and larger later rewards. In the game, people may either fire sooner and do little damage or later and do more damage. The amount of damage increased as time passed and reached a maximum after 10 s using a paradigm that they called an *escalating interest* task. The key independent variable in their method is a parameter that dictates how the weapon damage potential (its charge) approaches this maximum. The left side of Figure 1 shows the precise variations in recharge acceleration across three values of a parameter, *power*, that alter the recharge behavior, and the right side of Figure 1 shows the resulting rate of damage for a consistent waiting time between shots for each of these values. For power values less than 1.0, the maximal damage rate is achieved by waiting the full 10 s; for power values greater than 1.0, the maximal damage rate is achieved by firing as rapidly as possible.

The advantage to a video game preparation is that it allows the study of delay discounting in humans under short delays without a trial limit created by satiation. This ability to examine behavior across many choices allows better study of how choice evolves with time. The drawback to the preparation is that there is only one published study using it (Young et al., 2011) which raises questions as to whether behavior in the video game responds to manipulations, like a change in magnitude, in a way that is similar to that used in published studies involving hypothetical choice or consumables.

For the purposes of the present study, we wished to determine whether there would be a magnitude effect similar to that found in other studies involving humans or if the use of short delays and immediate rewards would result in the absence of a magnitude effect like that observed in nonhuman species. Furthermore, we wished to determine if the length of the delay to the maximum amount would systematically affect performance. In Young et al.'s (2011) previous study, the maximal reward was achieved after 10 s. Although published studies vary quite dramatically in the length of the delay to the larger later outcome, we are unaware of any research that has systematically varied this variable between-subject to determine if discount behavior changes as a function of the range of delays used (a form of context effect, cf. Percoco & Nijkamp, 2009).

The examination of both the effects of magnitude and the length of the delay would help determine whether waiting in the escalating interest task shows the same sensitivity to these variables as do hypothetical money choice tasks. Given that recent research has revealed that experience-based decisions can sometimes produce results opposite from those studied using described hypothetical choices (e.g., Hertwig & Erev, 2009; Rakow & Newell, 2010; Shafir,

Reich, Tsur, Erev, & Lotem, 2008), it is imperative that magnitude and delay effects be revisited using this paradigm for studying impulsivity.

In order to assess the magnitude effect, we adapted the Young et al. (2011) video game task to assign participants to either the standard magnitude used in their earlier study or a condition in which the maximum weapon damage was doubled. If participants show a magnitude effect consistent with prior reports involving hypothetical choice, their behavior should reveal longer interresponse times (IRTs) in the double magnitude condition regardless of the other variables being manipulated. In order to examine the effect of the length of the delay until maximum value is achieved, participants were assigned a randomly chosen delay between 5 s and 20 s in order to map out the functional relationship (see Young, Cole, & Sutherland, 2012). The manipulation of delay could be accomplished by having the same maximum achieved in less time, but doing so produces a situation in which the density of reinforcement necessarily increases. Thus, the maximum damage amount decreased as the delay decreased to ensure that the density of reinforcement would be identical across different delays to maximum reward if the participant waited the full delay between shots. This relationship is depicted in Figure 2 in which we show the change in damage across four 5 s periods, two 10 s periods, or one 20 s period when 5, 10, and 20 s is the assigned delay to maximum damage. By changing the damage potential as a function of the delay to maximum, we ensure that the same total damage is done for each of these delay conditions over a 20 s period for participants who wait until the maximum is available.

Method

Participants

A total of 80 undergraduate students enrolled in an introduction to psychology course at Southern Illinois University at Carbondale received course credit for their voluntary participation. There were 51 women and 27 men; 2 did not report their sex.

Procedure

The game world included four levels each containing seven separate regions with each region populated by two orcs as described by Young et al. (2011); see http://www.k-state.edu/psych/research/young/suppmaterial.html for a video clip of game play. The player's task was to destroy all of the orcs within each game level. Thus, a game level is isomorphic to a block of training in a conventional study of choice.

The player's weapon reached its maximum charge after an assigned delay. Each participant experienced the same assigned delay for the entire game. We used a random sampling design in which the delay to maximum charge for each participant was randomly chosen from the 5 to 20 s range (uniformly). The amount of damage a fully charged weapon could inflict was varied as a function of the programmed delay in order to hold constant the reinforcement density (i.e., average delay) that could be achieved for players waiting for the maximum damage. For the standard magnitude condition (43.33 points per 10 s, Young et al., 2011), equality was achieved by multiplying this density by the assigned delay to maximum charge (5 to 20 s) divided by 10 s. Thus, for a programmed delay of 5 s, a fully charged shot inflicted 21.66 points on the enemy orc (43.33 * 5/10), and for a programmed delay of 20 s, a fully charged shot inflicted 86.66 damage points on the enemy orc (43.33 * 20/10). In the double magnitude condition, all damage was doubled. Half of the participants were assigned to the

standard magnitude condition and the other half were assigned to the double magnitude condition. Note that it takes 100 points of damage to destroy a target (2.308 fully-charged shots in the standard magnitude condition).

The amount of damage a weapon inflicted during the delay to maximum charge was determined by the power value of that weapon (see Equation 1 of Young et al., 2011). The power value of the weapon changed each time the participant destroyed two orcs. We used a random sampling design in which the power value was randomly chosen from the 0.5 to 1.5 range (uniformly). The change in a weapon's power was accompanied by a 1250 ms three-tone sequence with a pitch that was correlated with the new power level with participants told that the pitch of the tone indicated the way that their weapon would recharge (higher pitch for higher power value). Because there were 14 orcs in each level and the power value changed when 2 orcs were destroyed, each player experienced seven power values in each game level.

After completing the game, participants completed a demographics program which asked their sex and self-reported amount of video game play for each of nine types (see Cole, 2011, Appendix D). Amounts of play were solicited on scale of 0 (indicating none) to 6 (indicating daily) with the middle of the scale indicating monthly play.

Results

All but one of the participants completed all four levels in the allotted time of one hour; one participant's session was terminated in the middle of the second level. Of the 80 participants, 78 completed all of the demographics questions.

Most participants produced a bimodal IRT distribution either firing as quickly as possible or waiting until the maximum damage was available (which differed between participants). To avoid statistical issues with the bimodal nature of the data, we dichotomized participant behavior

into short responses (IRTs less than half of the programmed delay to maximum damage) and long responses (IRTs greater than or equal to half of the programmed delay to maximum damage) and used a logit link function in a generalized linear mixed effects analysis, thus assuming a logistic relationship between our predictors and the likelihood of a long response. All responses greater than twice the delay to maximum damage, however, were dropped due to their high likelihood of being contaminated by inattention; this criterion resulted in dropping 1.4% of the responses. Because behavior in the first level was distinctly different from that in the subsequent levels, we treated Level 1 as categorically different from Levels 2 through 4 in our analyses.

Figure 3 shows the mean likelihood of waiting more than half the programmed delay for each participant at each experienced power value for the two magnitude conditions; Level 1 performance is excluded. The height of the smooth spline fit corresponds to a participant's overall likelihood of waiting (higher lines indicate more waiting) whereas the slope of the line signifies whether a participant's waiting is sensitive to changes in the power value (as it should, with higher power values making short IRTs more efficient thus producing a negative slope). There were clear individual differences in the overall likelihood of waiting and in sensitivity to the manipulation of power. Furthermore, there is some visual evidence that there was a higher proportion of participants who waited for the maximum reward in the double magnitude condition than in the standard magnitude condition.

One of the challenges of analyzing data from free operant procedures with continuous predictors is the visual presentation of the results. Given that some situations (e.g., high values of the power parameter) produce many more responses, it is essential that this be considered in describing the results to ensure that the graphs are not biased by this oversampling of behavior in

some conditions. This issue is complicated by the fact that some participants produce much shorter IRTs than others that can bias graphs by producing means that over-represent the behavior of these individuals. As a result, we statistically sampled from our participants' behavior at each power value and in each level to approximate a balanced design for our figures of raw data (Figures 3, 4, and 5). Our analytical approach involved linear mixed effects modeling of the complete data set (Pinheiro & Bates, 2004) which appropriately incorporates the sample size imbalances that produce significant problems for traditional general linear modeling approaches like ANOVA and regression.

In our examination of the distribution of a sample of scaled IRTs as a function of magnitude, Figure 4 reveals that participants produced bimodal IRTs for both the standard magnitude and the double magnitude conditions, but no significant difference between the two in Level 1. In Levels 2 through 4, the lower peak was shifted toward longer IRTs (about 30% of the delay to maximum charge) for the double magnitude condition, thus producing longer waiting times.

In the distribution of a sample of scaled IRTs as a function of programmed delay to maximum charge (here grouped into four categories only for illustration), Figure 5 shows that the shortest delays (5 – 8 s) produced a greater likelihood of waiting the full duration in both Level 1 and Levels 2 through 4. In contrast, the longest delays (17 – 20 s) produced a greater likelihood of rapid firing for both Level 1 and Levels 2 through 4. In Level 1, the intermediate delays (9 – 12 s and 13 – 16 s) showed IRT distributions very similar to that for the longest delays. In Levels 2 through 4, the distribution for the 9 – 12 s delays was generally shifted toward longer IRTs for the lower mode and the distribution for the 13 – 16 s delays was similar to that for the longest delays except with a less pronounced bimodal distribution.

In the linear mixed effects analysis, we specified a binomial distribution (thus producing *z*s, not *t*s, in the later analyses) and centered the continuous predictors in the interactions to avoid multicollinearity. The mixed effects analysis allowed intercepts, level slope effects, and power slope effects to vary across participants. Only those interactions that produced a better model (lower AIC) were retained. A logarithmic relation for delay was used because doing so produced a better fit. The results are shown in Table 1.

First, participants showed a strong sensitivity to the power manipulation thus replicating Young et al. (2011). Sensitivity was roughly twice as strong in Levels 2 through 4 ($b = -2.04$) as in Level 1 ($b = -1.13$). In Level 1, the modeled likelihood of a "long" response (i.e., waiting more than half the delay to maximum damage before firing) was 55% for the lowest power value of 0.5 and 31% for the highest power value of 1.5 ($z = 5.01, p < .01$). This difference was much larger in Levels 2 through 4 (73% vs. 28%, $z = 10.45, p < .01$).

In the standard magnitude condition (that used by Young et al., 2011), the best fitting model shown in Table 1 predicted 40% "long" responses for Levels 2 through 4 whereas in the double magnitude condition the model predicted 61% "long" responses ($z = 2.71, p < .01$, see Table 1). The model revealed no statistically significant difference as a function of overall magnitude in Level 1 (41% vs. 44%, $z = -.06, p > .25$). These modeled estimates are consistent with those shown in Figure 4.

The overall effects of the length of the delay to maximal damage appeared large (Figure 5) but were undermined by inconsistency across participants and game level. In Level 1, the model predicted 52% "long" responses for the shortest delay and 32% for the longest delay ($z = -1.98, p < .05$, see Table 1). The effect was not sustained in Levels 2 through 4 where the

predicted preferences for "long" responses were 49% and 53% for the shortest and longest delay, respectively ($z = .66, p > .05$).

Importantly, there was a statistically significant effect of the length of the delay to maximal damage on sensitivity to the power parameter (i.e., a Power × log(Delay) interaction, $z = 2.02, p < .05$, see Table 1). For example, in Levels 2 through 4 participants were more sensitive to the current power value when the delay to maximum was only 5 s (predicted "long" response Ms = 77% for a power of 0.5 vs. 21% for a power of 1.5) than when it was 20 s (predicted "long" response Ms = 67% for a power of 0.5 vs. 38% for a power of 1.5).

The individual differences variables (sex and self-reported video game play) had little effect on performance. Only the reported amount of third-person video game play was retained as a predictor of the likelihood of producing a long IRT (greater than half the delay to maximum) with those who reported the most (a score of 6) having a 22% chance of producing a "long" response and those who reported the least (a score of 0) having a 54% chance of producing a "long" response ($z = -3.11, p < .01$).

Discussion

The observation of a magnitude effect in this very different preparation for the study of impulsivity both broadens the generality of the original reports and further validates the escalating interest task as a measure of impulsive responding. The effect of the length of the delay on sensitivity to the consequences of shorter versus longer IRTs (via the power parameter) was small but somewhat surprising. Our participants' behavior showed greater differentiation as a function of the power value for shorter delays than for longer delays.

The discounting of reward magnitude is often well-described by Mazur's (1987) hyperbolic model:

$$V = \frac{Damage_{IRT}}{1 + k \; \times \; IRT} \tag{1}$$

where V is the discounted value of the programmed amount of damage delivered ($Damage_{IRT}$; see Figure 1) for a given IRT and k is a free parameter that varies directly with sensitivity to delay; k is often estimated independently for each individual. Using Equation 1, the left side of Figure 6 depicts hyperbolic discounting across a 20 s span with a k of 0.25. An outcome that reaches its maximum value at 5 s will have been significantly discounted relative to its value if it had been available immediately. However, an outcome that reaches its maximum value at 20 s would have been discounted even further. The relative amount of discounting for these two delays is more clearly evident on the right side of Figure 6 where each curve is scaled by dividing by its maximum (5 s vs. 20 s). This latter figure suggests that in general, participants should wait a greater proportion of the delay for short delays than for long delays. As Table 1 and Figure 5 show, this result was observed only in Level 1 suggesting that participants assigned to conditions with longer delays adapted to them. Indeed, Figure 5 reveals that the principal change in behavior was a shift toward longer IRTs in later game levels for the participants assigned to the longer delay conditions consistent with their learning to exhibit less impulsive responding in the presence of longer delays.

With regard to the weaker sensitivity to the power contingency for longer delays to maximum, initially this behavior may have been driven by the generally low likelihood of waiting in the presence of long delays (see Figure 5). However, when the delay to maximum was short, the dynamics of the reward magnitude (i.e., how quickly the charge bar changed) may be more evident than when the delay to maximum was long. As Figure 2 shows, the difference that a few seconds of waiting makes to the outcome magnitude is much greater for short delays than for long delays when the power is less than 1.00. The opposite is true for power greater

than 1.00, but this difference is less dramatic due to the asymmetric nature of the superellipsoid function that was governing the recharge rate.

Our findings suggest that both magnitude effects and an increased sensitivity to discounting contingencies for short delays occur within the context of an escalating interest video game task. The first result suggests that the difficulty with observing magnitude effects in nonhuman species is not merely a byproduct of the shorter experienced delays in the tasks used in nonhuman studies nor the use of hypothetical monetary rewards in human studies. Furthermore, the magnitude effect did not interact with our delay contingencies (via the power parameter) nor the delay to maximum damage, thus establishing its independent effect on behavior. Finally, our manipulation of the delay to maximum reward served to generalize the original report of Young et al. (2011) by establishing control by the power parameter at delays to maximum reward that were both shorter and longer than those originally reported.

Although we have placed the present results within the context of research on delay discounting, the escalating interest task more closely resembles a delay of gratification (DG) task (Mischel, 1966; Mischel & Ebbesen, 1970) in which the smaller sooner reward is available throughout the delay thus allowing for the possibility of defection while waiting. Given that we have not uncovered any research on magnitude effects using a DG task, our results extend demonstrations of participants' waiting longer for larger magnitudes into related paradigms involving a delay to reward. Furthermore, the DG paradigm usually involves delays of many minutes (Forzano, Michels, Carapella, Conway, & Chelonis, 2011; Mischel & Ebbesen, 1970) although studies using other species show experimental control with much shorter delays (e.g., Reynolds, de Wit, & Richards, 2002).

Differences among species are readily apparent and not surprising, but it is too easy to jump to conclusions that imply a qualitative difference. It is important to choose tasks that conform to the abilities and ecological niches to which a species has adapted. Because most adult humans are not currently living in environments where a few seconds of delay of food, money, or an opportunity to mate has survival significance, it should not be surprising that there is little discounting of these rewards compared to that of other species where competition for basic needs is much more fierce. Thus, using an environment where seconds do matter for humans provides a better test case for identifying whether certain variables are differentially affecting humans and nonhumans.

Author Note

References

Calvert, A. L., Green, L., & Myerson, J. (2010). Delay discounting of qualitatively different reinforcers in rats. *Journal of the Experimental Analysis of Behavior, 93*, 171-184.

Cole, J. J. (2011). *Temporal discounting in an uncertain world: The absence of probability and magnitude effects in an experience-based task* (Unpublished Master's thesis). Southern Illinois University, Carbondale.

Estle, S. J., Green, L., Myerson, J., & Holt, D. D. (2007). Discounting of monetary and directly consumable rewards. *Psychological Science, 18*, 58-63.

Forzano, L. B., Michels, J. L., Carapella, R. K., Conway, P., & Chelonis, J. J. (2011). Self-control and impulsivity in children: Multiple behavioral measures. *Psychological Record, 61*, 425-448.

Green, L., Myerson, J., Holt, D. D., Slevin, J. R., & Estle, S. J. (2004). Discounting of delayed food rewards in pigeons and rats: Is there a magnitude effect? *Journal of the Experimental Analysis of Behavior, 81*, 39-50.

Green, L., Myerson, J., & McFadden, E. (1997). Rate of temporal discounting decreases with amount of reward. *Memory & Cognition, 25*, 715-723.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences, 13*, 517-523.

Hinvest, N. S., & Anderson, I. M. (2010). The effects of real versus hypothetical reward on delay and probability discounting. *Quarterly Journal of Experimental Psychology, 63*, 1072-1084.

Jimura, K., Myerson, J., Hilgard, J., Braver, T. S., & Green, L. (2009). Are people really more patient than other animals? Evidence from human discounting of real liquid rewards. *Psychonomic Bulletin & Review, 16*, 1071-1075.

Logue, A. W., & King, G. R. (1991). Self-control and impulsiveness in adult humans when food is the reinforcer. *Appetite, 17*, 105-120.

Madden, G. J., Begotka, A. M., Raiff, B. R., & Kastern, L. L. (2003). Delay discounting of real and hypothetical rewards. *Experimental and Clinical Psychopharmacology, 11*, 139-145.

Madden, G. J., & Bickel, W. K. (2010). *Impulsivity: The behavioral and neurological science of discounting*. Washington, DC: American Psychological Assocation.

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

Mischel, W. (1966). Theory and research on antecedents of self-imposed delay of reward. In B. A. Maher (Ed.), *Progress in Experimental Personality Research*. New York: Academic Press.

Mischel, W., & Ebbesen, E. B. (1970). Attention in delay of gratification. *Journal of Personality and Social Psychology*(16), 329-337.

Percoco, M., & Nijkamp, P. (2009). Estimating individual rates of discount: A meta-analysis. *Applied Economics Letters, 16*, 1235-1239.

Pinheiro, J. C., & Bates, D. M. (2004). *Mixed-effects models in S and S-PLUS*. New York: Springer.

Rachlin, H. C., Raineri, A., & Cross, D. (1991). Subjective probability and delay. *Journal of the Experimental Analysis of Behavior, 55*, 233-244.

Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for

    future research on experience-based choice. *Journal of Behavioral Decision Making, 23*,

    1-14.

Reynolds, B., de Wit, H., & Richards, J. B. (2002). Delay of gratification and delay discounting

    in rats. *Behavioural Processes, 59*, 157-168.

Richards, J. B., Mitchell, S. H., de Wit, H., & Seiden, L. S. (1997). Determination of discount

    functions in rats with an adjusting-amount procedure. *Journal of the Experimental*

    Analysis of Behavior, 67, 353-366.

Shafir, S., Reich, T., Tsur, E., Erev, I., & Lotem, A. (2008). Perceptual accuracy and conflicting

    effects of certainty on risk-taking behaviour. *Nature, 453*, 917-920.

Young, M. E., Cole, J. J., & Sutherland, S. C. (2012). Rich stimulus sampling for between-

    subjects designs improves model selection. *Behavior Research Methods, 44*, 176-188.

Young, M. E., Webb, T., & Jacobs, E. A. (2011). Deciding when to "cash in" when outcomes are

    continuously improving: An escalating interest task. *Behavioural Processes, 88*, 101-110.

Table 1. *Regression Weights for the Best Fitting Generalized Linear Mixed Effects Model of Dichotomized Interresponse Times.*

	Level 1	SE	Levels 2-4	SE
(Intercept)	2.23*	0.81	2.26*	0.18
Power	-1.13*	0.22	-2.04*	0.17
Log(Delay)	-0.62*	0.31	0.13	0.29
Magnitude	-0.02	0.29	0.71*	0.27
Power × Log(Delay)	0.91*	0.38	0.89*	0.34

Note: The link function was the logit. *$p < .05$.
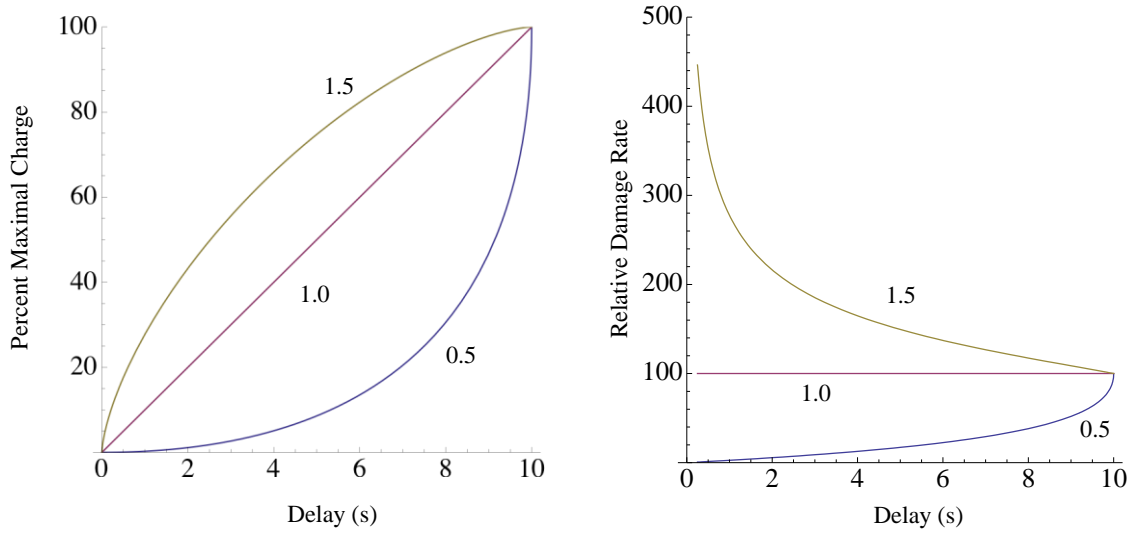
Figure 1. Left: Three examples of the superellipsoid functions generated by Equation 1 of Young et al. (2011) using a power value of 1.5 (top curve), 1.0, and 0.5 (bottom curve). Right: The damage rates produced by these power values for a consistent delay between responses.
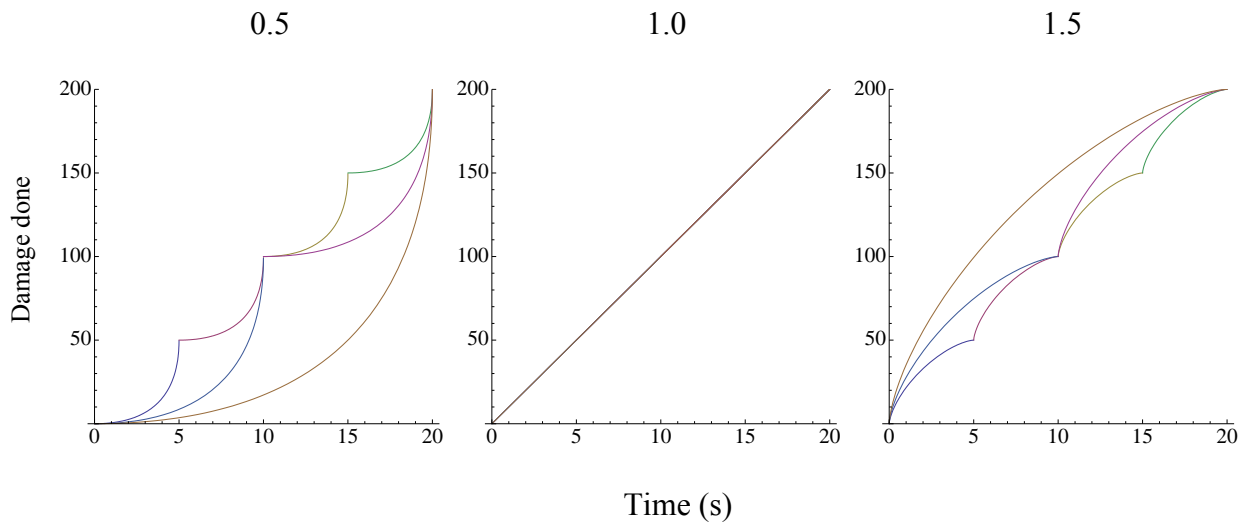
Figure 2. Examples of the growth in damage for short, intermediate, and long delays (5 s, 10s,

and 20 s) across 20 s if the participant waited until full charge for each choice. The left graph is

for positive acceleration (power = 0.5), the middle graph for linear acceleration (power = 1.0),

and right graph for negative acceleration (power = 1.5).   These curves are those produced by the

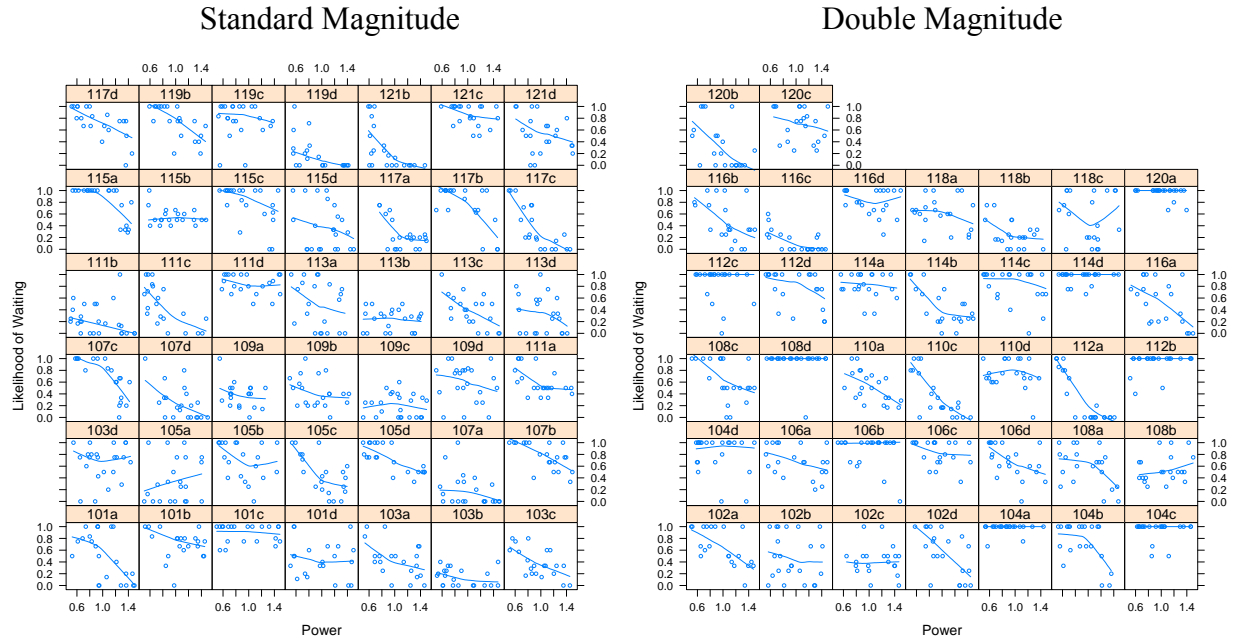standard magnitude condition with damage scaled to 100% of the standard (43.33 pts).

Standard Magnitude   Double Magnitude

Figure 3. Mean likelihood of waiting more than half the programmed delay to maximum reward for each participant in the standard and double magnitude conditions. Smoothed spline fits are included to ease readability; higher lines indicate a higher probability of waiting for that participant and a more negative slope indicates a greater sensitivity to the changing contingencies of waiting.
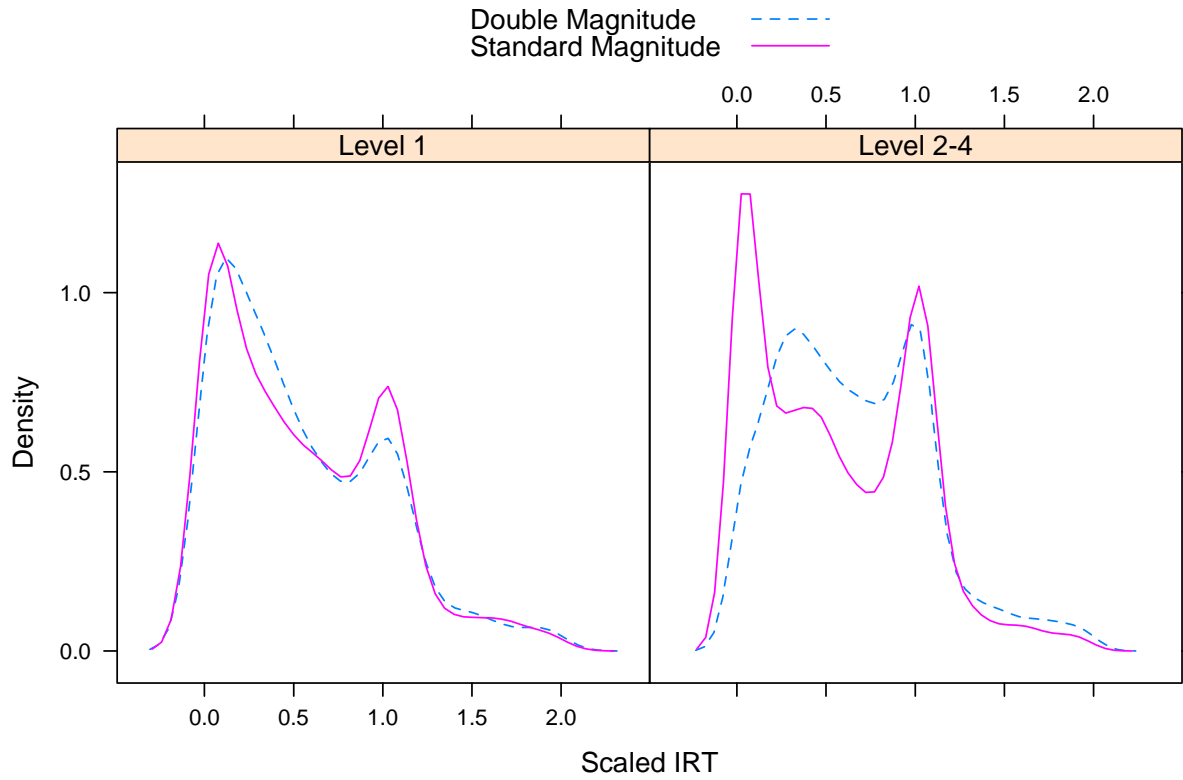
Figure 4. Probability density curves of sample of scaled IRTs (1.0 = delay to maximum charge) for the standard magnitude condition and the double magnitude condition in Level 1 and across Levels 2 through 4. Note that these plots are functionally equivalent to smoothed frequency histograms, thus the plots go below zero due to smoothing.
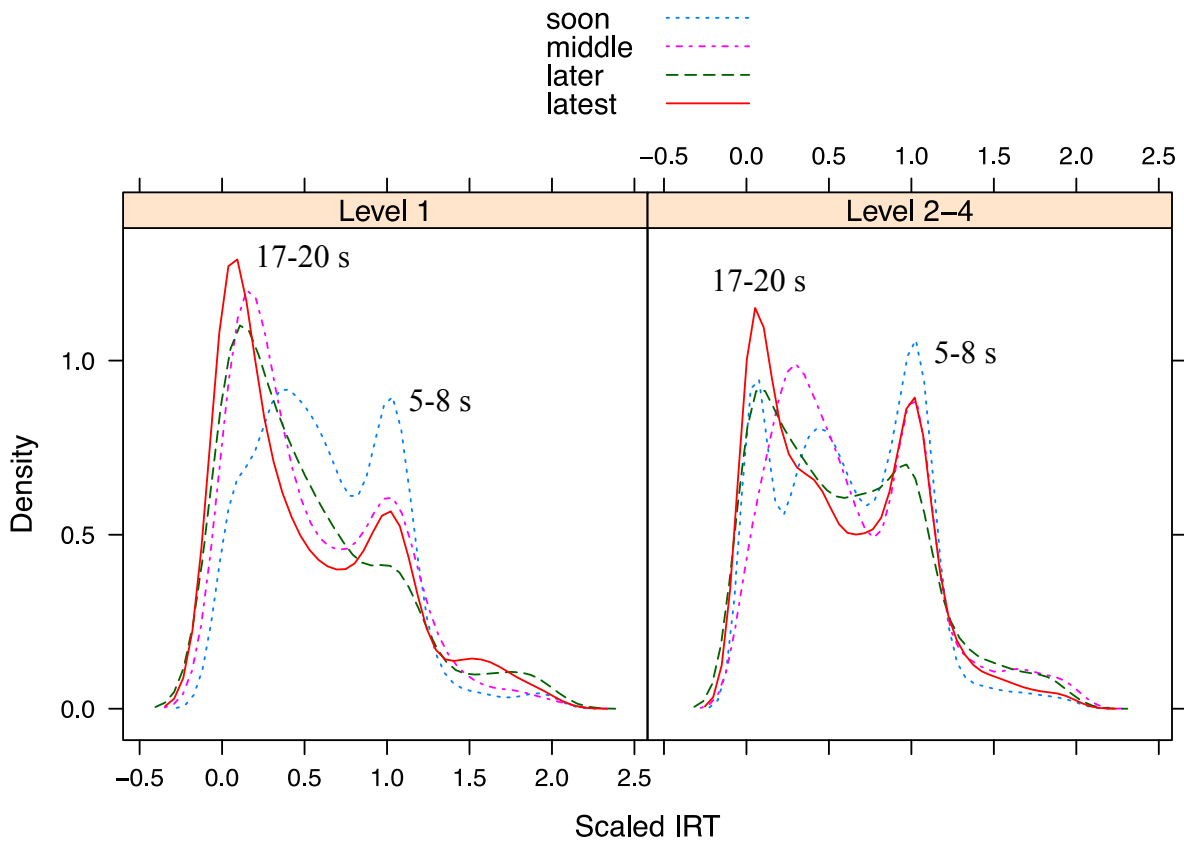
Figure 5. Probability density curves of a sample of scaled IRTs (1.0 = delay to maximum charge) for the 5 – 8 s delays (short dashes), 9 – 12 s delays (short-long dashes), the 13 – 16 s delays (long dashes), and 17 – 20 s delays (solid line) in Level 1 and across Levels 2 through 4. Note that these plots are functionally equivalent to smoothed frequency histograms, thus the plots go below zero due to smoothing.
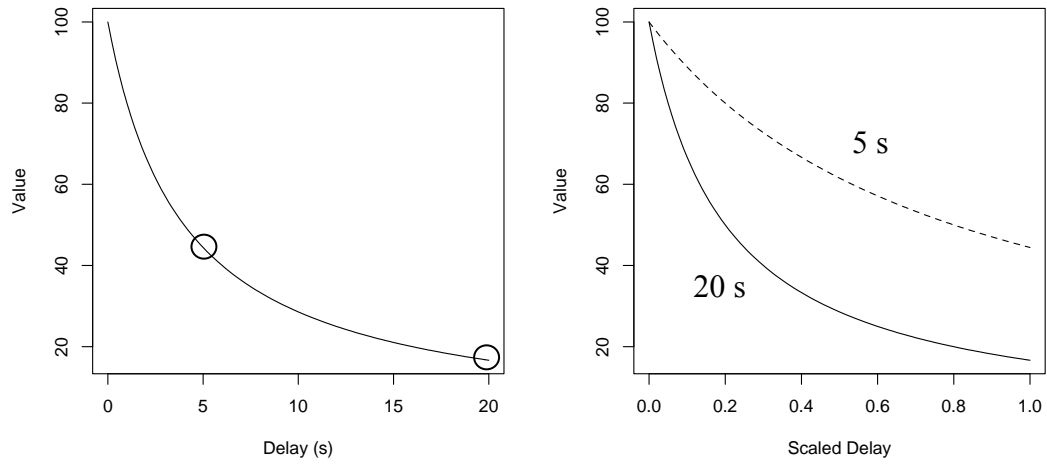
Figure 6. Plot of hyperbolic discounting (left) and a replot of these discounting effects when the

delay to maximum value is plotted on a relative scale (right). See text for a more complete

description.