JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 2011, **95,** 175–187 NUMBER 2 (MARCH)

MODELING THE EFFECT OF REWARD AMOUNT ON PROBABILITY DISCOUNTING

JOEL MYERSON, LEONARD GREEN, AND JOSHUA MORRIS

WASHINGTON UNIVERSITY

The present study with college students examined the effect of amount on the discounting of probabilistic monetary rewards. A hyperboloid function accurately described the discounting of hypothetical rewards ranging in amount from \$20 to \$10,000,000. The degree of discounting increased continuously with amount of probabilistic reward. This effect of amount was not due to changes in the rate parameter of the discounting function, but rather was due to increases in the exponent. These results stand in contrast to those observed with the discounting of delayed monetary rewards, in which the degree of discounting decreases with reward amount due to amount-dependent decreases in the rate parameter. Taken together, this pattern of results suggests that delay and probability discounting reflect different underlying mechanisms. That is, the fact that the exponent in the delay discounting function is independent of amount is consistent with a psychophysical scaling interpretation, whereas the finding that the exponent of the probability-discounting function is amount-dependent is inconsistent with such an interpretation. Instead, the present results are consistent with the idea that the probability-discounting function is itself the product of a value function and a weighting function. This idea was first suggested by Kahneman and Tversky (1979), although their prospect theory does not predict amount effects like those observed. The effect of amount on probability discounting was parsimoniously incorporated into our hyperboloid discounting function by assuming that the exponent was proportional to the amount raised to a power. The amount-dependent exponent of the probability-discounting function may be viewed as reflecting the effect of amount on the weighting of the probability with which the reward will be received.

Key words: discounting, probability, amount, scaling, decision weight, prospect theory, humans

---

Individuals often prefer a smaller immediate reward over a larger delayed reward. It is frequently assumed that this is because the subjective value of the delayed reward is discounted whereas the value of the immediate reward is not. A similar phenomenon is observed with respect to probabilistic rewards: Individuals often prefer a smaller certain reward over a larger probabilistic reward, and it is frequently assumed that this occurs because the value of the probabilistic reward is discounted whereas the value of the certain reward is not. Indeed, the same hyperboloid function describes both delay and probability discounting:

$$V = A/(1 + bX)^s, \qquad (1)$$

where V is the subjective value of the delayed or probabilistic reward, A is the amount of reward, and X is either the delay until, or the odds against, receipt of the reward (for a review, see Green & Myerson, 2004). Equation 1 has two parameters: b, which governs the rate of discounting, and s, which has been interpreted as reflecting the psychophysical scaling of amount, time, and likelihood. The present study concerns the validity of this interpretation of the exponent, s, particularly as it applies to the discounting of different amounts of probabilistic reward.

To facilitate comparisons involving the discounting of different amounts of reward, Equation 1 is often written as

$$V = 1/(1 + bX)^s, \qquad (2)$$

where V is the relative subjective value (i.e., subjective value of a reward expressed as a proportion of its nominal value). The psychophysical scaling interpretation of the exponent in the hyperboloid discounting function, Equations 1 and 2, assumes that amount, time, and likelihood are all nonlinearly scaled (Green & Myerson, 2004; see Myerson & Green, 1995, and the first section of the

The research was supported by NIH Grant R01 MH055308. Josh Morris was supported by an award from the Office of Undergraduate Research at Washington University. Portions of the work were presented at the Nano-Conference on Discounting at Jacksonville State University organized by Bill Palya (Jacksonville, AL, November 2008) and at the Society for the Quantitative Analysis of Behavior (San Antonio, TX, May 2010).

Address correspondence to Joel Myerson, Department of Psychology, Campus Box 1125, Washington University, St. Louis, MO 63130 (e-mail: jmyerson@wustl.edu).

doi: 10/1901/jeab.2011.95-175

Appendix for derivations that explicate how nonlinear scaling of amount affects the exponent in delay and probability discounting, respectively). For example, the relation between the subjective value of a monetary reward and the actual amount of that reward is frequently assumed to be a negatively accelerated power function (Gonzalez & Wu, 1999; Myerson & Green, 1995; Tversky & Kahneman, 1992). Such a function captures the fact that as the amount of reward increases, its subjective value increases at an ever-decreasing rate. For similar reasons, both the relation between subjective and actual duration and the relation between subjective and actual likelihood may be assumed to be negatively accelerated power functions.

In both cases (i.e., delay and probability discounting), it is assumed that the determinants of the exponent, *s*, in the hyperboloid discounting function remain constant when either the amount of money and the duration of the delay or the amount of money and the likelihood of reward are varied. This is because a psychophysical scale is a rule that describes the way that changes in some objective quantity affect the corresponding perceived quantity (Stevens, 1957), and although the perceived quantity will change with the objective quantity, the rule itself (and thus the exponent that characterizes the rule) remains constant. As a consequence, the psychophysical scaling interpretation predicts that the exponent in the discounting function will remain constant for a given individual and type of reward. With respect to delay discounting, the results of a number of studies support this prediction (Estle, Green, Myerson, & Holt, 2006; Myerson & Green, 1995; Myerson, Green, Hanson, Holt, & Estle, 2003). With respect to probability discounting, however, the prediction that the exponent will remain constant as the amount of probabilistic reward is varied has been called into question (Estle et al., 2006; Myerson et al., 2003).

It should be noted that although both delay and probability discounting can be well described by the same hyperboloid function form, there is ample evidence that they involve at least some different processes (Green & Myerson, 2010). Most prominently, the magnitude effects for delay and probability discounting are opposite in direction. Whereas smaller delayed amounts are discounted more steeply than larger delayed amounts, smaller probabilistic amounts are discounted *less* steeply than larger probabilistic amounts (e.g., Green, Myerson, & Ostaszewski, 1999; for a review, see Green & Myerson, 2004). Given the differential effects of amount on the degree to which delayed and probabilistic rewards are discounted, it may well be the case that amount affects the parameters of the corresponding discounting functions in qualitatively different ways. In particular, if the exponent in the probability-discounting function were to change with reward amount, such a finding would indicate that the psychophysical scaling interpretation of the exponent does not apply to probability discounting, in which case the probability-discounting function itself may need to be reinterpreted.

Accordingly, the present study examines the effect of amount on probability discounting by varying the amount of the probabilistic reward over an extremely wide range ($20 to $10,000,000). Of interest is how increases in amount affect the parameters of the probability-discounting function, and in particular, whether the exponent, *s*, will remain constant as predicted by the psychophysical scaling interpretation or whether it will change systematically with reward amount. The answer to this empirical question will have important implications for several related theoretical questions: Is it possible to go beyond describing changes in parameters of the discounting function as increases or decreases and actually specify mathematically how amount affects these parameters? If so, then how should the effect of amount on probability discounting be modeled? And finally, what does such modeling reveal about the decision-making processes involved? Prospect theory, the major account of probabilistic decision making in economics, does not deal with how such decision making is affected by the absolute magnitude of the amounts involved (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). Thus, the results of the present effort, particularly the attempt to model the effect of amount on probability discounting, may have implications for prospect theory as well as for our understanding of discounting in general.

METHOD

Participants

Forty undergraduate students (16 males and 24 females, mean age = 20.6 years) at

Washington University were recruited through the Department of Psychology's participant pool and received 10 dollars each for their participation.

Procedure

Participants were tested individually in a quiet room using a computer-administered, adjusting-amount discounting procedure. Participants read the instructions on the monitor, informing them that they would be choosing between hypothetical amounts of money, one of which could be received for sure whereas the other could be received with a given probability. They also were informed that there were no correct or incorrect choices, and that they should make their choices as if real money were involved.

For each choice that they had to make, participants saw the amount of the certain reward presented on the screen next to both the amount of the probabilistic reward and the probability of receiving that reward. Participants were informed that if they changed their mind, they would have an opportunity to change their response. They were then given several practice trials and offered the opportunity to ask questions before beginning the actual experiment. The experimenter then left the testing room and sat in another room from which he could observe the participant.

The experiment consisted of 45 conditions: nine amounts of probabilistic reward ($20, $250, $3000, $20,000, $50,000, $100,000, $500,000, $2,000,000, and $10,000,000) crossed with five probabilities (.80, .50, .25, .10, and .05). A computer program randomly selected a probabilistic reward amount (without replacement) and then administered all five probability conditions for that amount in a random order. The side of the computer screen on which the probabilistic reward was presented alternated randomly across the 45 conditions. Within each condition, an adjusting-amount procedure was used to obtain an estimate of the amount of the smaller, certain reward that the participant judged to be equivalent in value to the larger, probabilistic amount.

In each amount/probability condition, participants were presented with a series of six choice trials. On the first trial, the participant chose between receiving either the probabilistic amount or half of that amount ''for sure.'' On each subsequent trial, the amount of the certain reward was adjusted based on the participant's choice on the preceding trial. Specifically, if the participant had chosen the certain reward on the previous trial, the amount of the next certain reward was decreased; if the participant had chosen the probabilistic reward, the amount of the next certain reward was increased. The size of the adjustment, either increase or decrease, itself decreased with successive choices. The first adjustment was half of the difference between the amounts of the certain and probabilistic rewards presented on the first trial, and the size of each subsequent adjustment was half that of the preceding adjustment, rounded off to the nearest dollar.

For example, in the condition with $250 at p = .25, the choice on the first trial would be between ''$250 with a 25% chance'' and ''$125 for sure.'' If the participant chose the ''$250 with a 25% chance,'' the choice on the second trial would be between ''$250 with a 25% chance'' and ''$188 for sure.'' If the participant then chose the ''$188 for sure,'' the choice on the third trial would be between ''$250 with a 25% chance'' and ''$156 for sure.'' Three more trials with the same probabilistic reward were presented before the program switched to a new amount and probability. The adjusting-amount procedure that was used rapidly converges on an estimate of the subjective value of the probabilistic reward, which was calculated by taking the amount of certain reward presented on the sixth trial and either adding or subtracting 1.56% (i.e., $100/2^6$ percent) of the probabilistic reward amount, depending on whether the participant chose the probabilistic or certain reward, respectively.

RESULTS

Figure 1 plots the group median relative subjective value as a function of the odds against receiving the probabilistic reward for each of the nine amounts studied. Relative subjective value (i.e., subjective value divided by the actual amount of the probabilistic reward) is used as the dependent variable in order to facilitate comparisons across the different reward amounts. In each panel, the solid curve represents the best fitting hyperboloid function (Eq. 2) for that amount

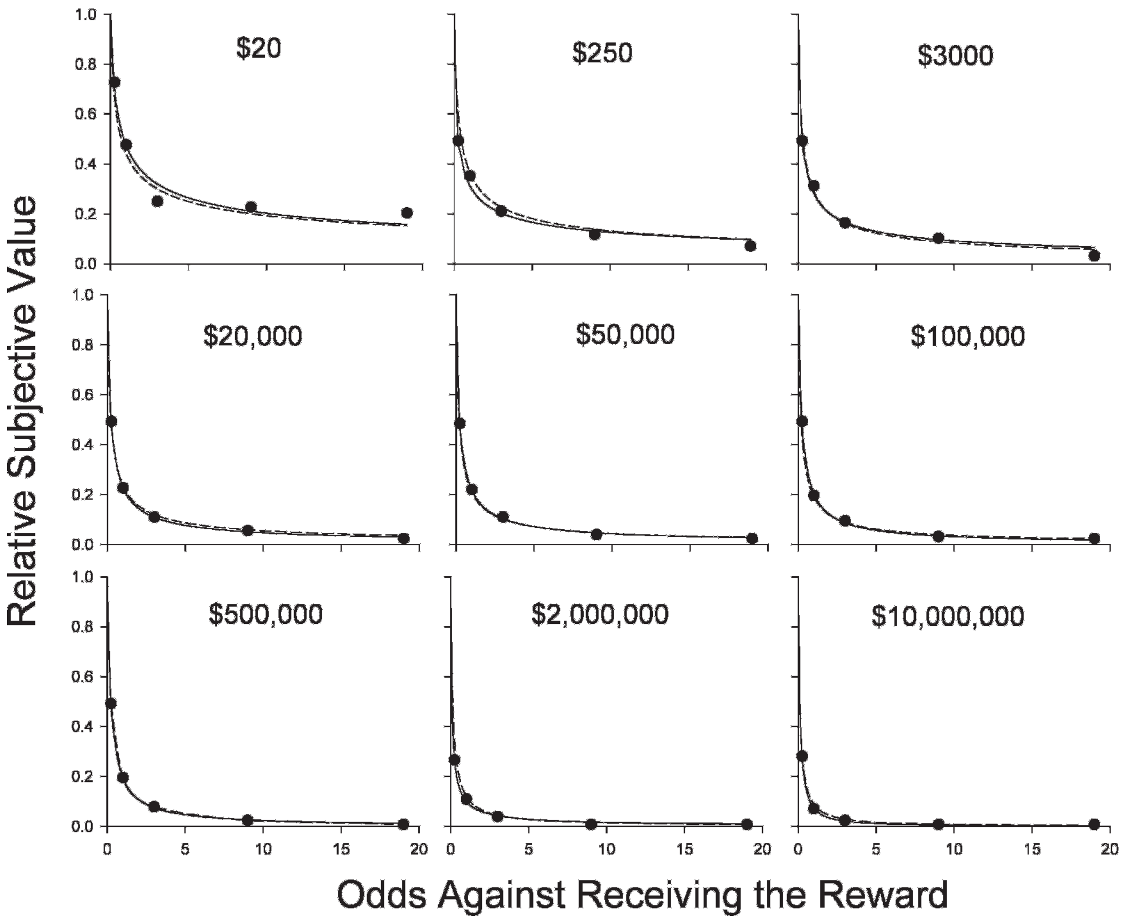

Fig. 1. Relative subjective value as a function of the odds against receiving a reward. Each panel depicts the group median subjective values for a different amount of probabilistic reward. Solid curves represent Equation 2 fitted to the data; dashed curves represent Equation 3 rewritten in terms of relative subjective value.

condition. (The dashed curve represents the fit of a power model to be described later.) The hyperboloid function provided excellent fits to the data from all nine amount conditions; all of the R^2s were greater than .96.

For the most part, Equation 2 also provided good fits to the data at the individual level. For each participant, the hyperboloid function was fitted simultaneously to the relative subjective values from all nine amount conditions, resulting in separate estimates of the b and s parameters for each condition and providing a single overall fit measure (R^2). Across the 40 participants, the mean and median R^2s were .906 and .918, respectively. In order to select representative individuals to depict, we first ranked all of the participants based on their R^2s and then chose 3 whose R^2s were at the 25^{th}, 50^{th}, and 75^{th} percentiles. The data for these 3 representative individuals are shown in Figure 2. For ease of presentation, only the data for the $20, $20,000, and $500,000 reward amounts are shown. (The complete set of individual data, consisting of the indifference points for each probabilistic amount at each probability, are presented in an online supplement, available in the supplemental section of this article at PubMed Central.)

The Effects of Amount of Probabilistic Reward

As may be seen in Figures 1 and 2, larger probabilistic rewards were consistently discounted more steeply than smaller probabilistic rewards. The relation between steepness of discounting and amount of probabilistic reward may be seen more clearly in Figure 3,

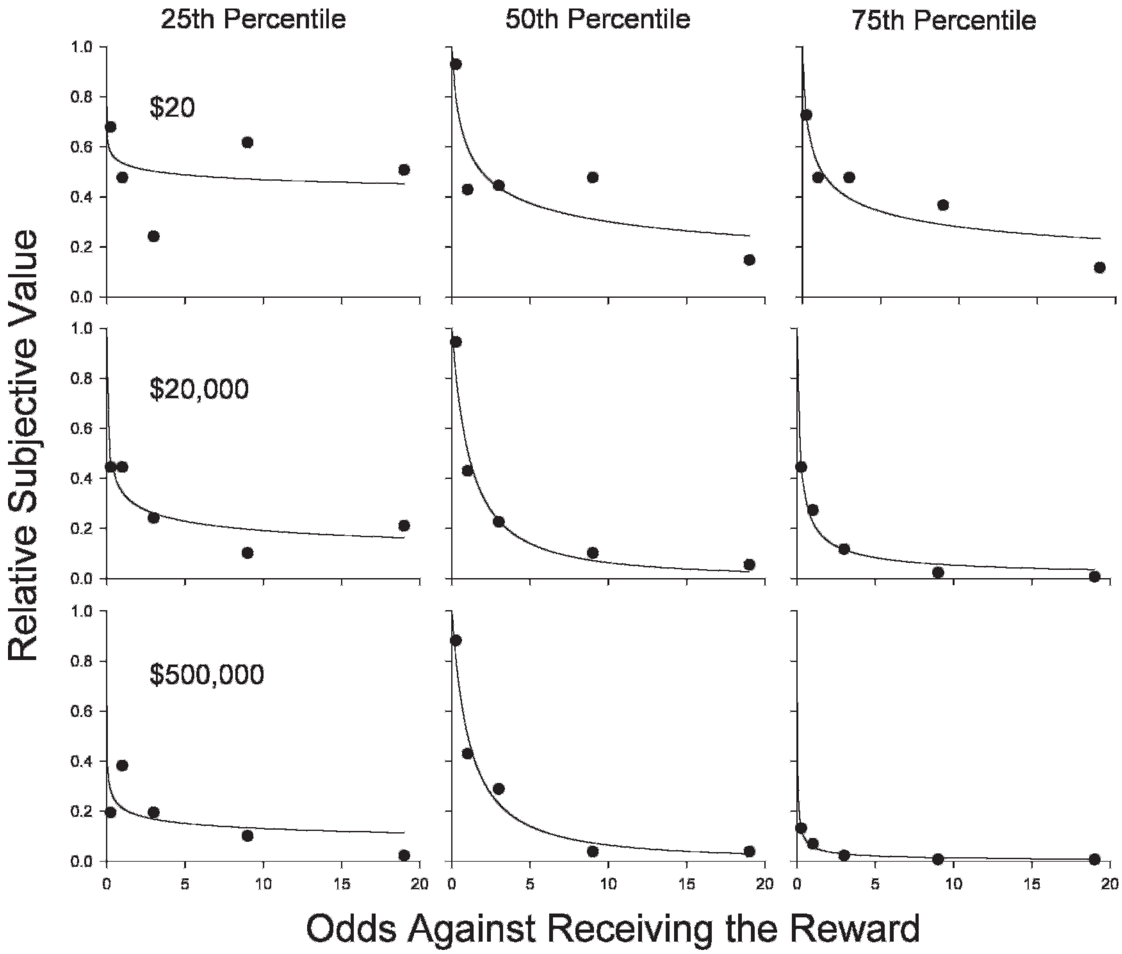

Fig. 2. Relative subjective value as a function of the odds against receiving a reward for 3 representative participants. The left, middle, and right columns depict the data for the participants whose R^2s were at the 25th, 50th, and 75th percentiles. Solid curves represent Equation 2 fitted to the data.

which shows both the area under the group median data (top panel) and the median of the areas calculated from the data for each individual (bottom panel) plotted as a function of amount on log–log coordinates. The area-under-the-curve (AuC) measure (Myerson, Green, & Warusawitharana, 2001) represents the area under the observed subjective values and provides a single, theoretically neutral measure of the degree of discounting. Because the independent and dependent variables are normalized for purposes of calculating the AuC, values may range between 0.0 (maximally steep discounting) and 1.0 (no discounting). The degree of discounting, as measured by the AuC, was strongly correlated with the amount of the probabilistic reward: *r*s = −.991 and −.983 for the group medians and the medians of the individuals, respectively.

Given the observed strong effect of amount, we next sought to determine what changes in the parameters of the hyperboloid discounting function underlie this magnitude effect. Figure 4 shows the parameter estimates from the fits of the hyperboloid discounting function depicted in Figure 1. As may be seen, the logarithm of the value of the *s* parameter increased linearly as a function of the logarithm of the amount of probabilistic reward (top panel of Figure 4; $r = .948$, $p < .001$), suggesting that *s* is a power function of amount. In contrast, log *b* and log amount were not significantly correlated (bottom panel; $r = -.117$, $p = .76$).

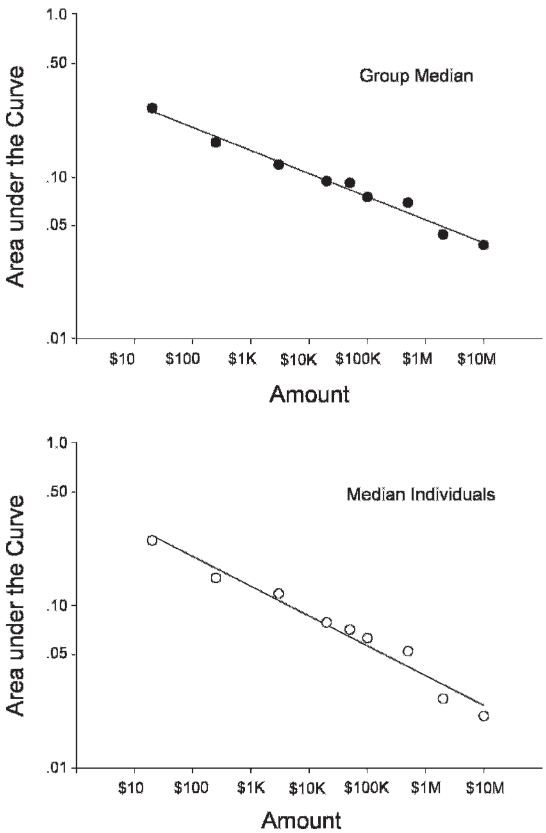

Fig. 3. Area-under-the-curve (AuC) as a function of amount of probabilistic reward. The AuCs in the upper panel are based on the group median data at each amount, and the lower panel depicts the median individual AuC at each amount. (Note the logarithmic scales on the axes.)

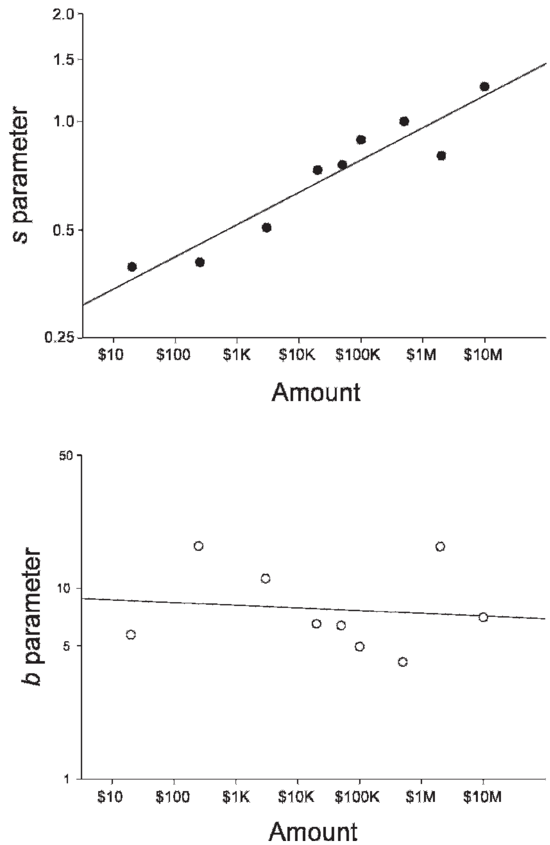

Fig. 4. Values of the *s* and *b* parameters of Equation 2 as a function of amount of probabilistic reward. Values are based on fits of Equation 2 to the group median data. (Note the logarithmic scales on the axes.)

Similar patterns of results were observed at the individual level, although there was considerable variation among individual participants. This variation may be seen in Figure 5, in which the top panel depicts the frequency distribution of the correlations between log *s* and log amount, and the bottom panel depicts the frequency distribution of the correlations between log *b* and log amount. The correlations between log *s* and log amount at the individual level tended to be negatively skewed with a peak at about .50, and an average *r* (back-transformed from the mean Fisher *z*) of .38. Importantly, the 95% confidence interval about the mean Fisher *z* (0.173 ± 0.030) did not include zero, indicating that the correlation between log *s* and log amount was significant. In contrast, the correlations between log *b* and log amount showed little evidence of a clear central tendency and had a mean (back-transformed from the mean Fisher *z*) of .15. In this case, the confidence interval did include zero (0.064 ± 0.082), indicating that the correlation between log *b* and log amount was not significant.

It should be noted that when $s = 1.0$, Equation 2 is a simple hyperbola, which provides a useful benchmark against which to assess the form of the probability-discounting function. The effect of amount on *s* at the individual level is reflected in the fact that the percentage of cases in which *s* was significantly less than 1.0 decreased from 63% at the smallest ($20) amount to 30% at the largest ($10 million) amount. Overall, the *s* parameter was less than 1.0 in 256 out of 360 possible cases at the individual level (40 participants × 9 amounts), and significantly so in more than

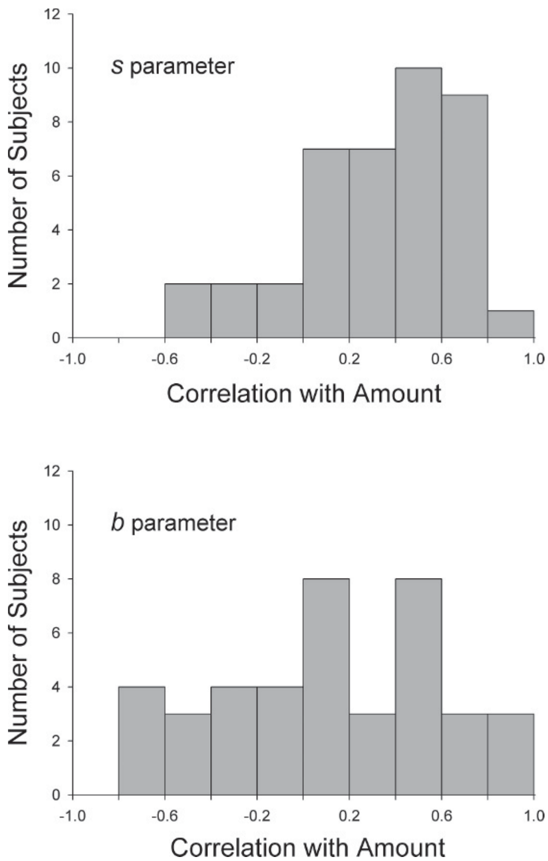

Fig. 5. Frequency distributions for the correlations of the s and b parameters (upper and lower panels, respectively) with the corresponding probabilistic reward amounts. Each correlation was calculated based on the logarithm of an individual's parameter estimates for the nine different amount conditions and the logarithm of the corresponding reward amounts.

half (138) of those cases. In spite of the fact that s increased as a function of the amount of the probabilistic reward, s was significantly greater than 1.0 in only 1 out of the 360 possible cases at the individual level, and this at the largest amount.

Modeling the Effects of Amount

The preceding analyses were based on fits of the hyperboloid discounting function (Eq. 2) in which parameters were estimated for each amount condition separately. It is possible, however, that a more parsimonious description of these data would be sufficient. Given that only the s parameter (and not the b parameter) was systematically affected by amount, one possible model assumes that although a different value of s may be needed for different amount conditions, a single value of b will suffice for all amounts. A model in which s is assumed to be amount-independent whereas b varies with amount is conceivable (indeed, such a model does describe delay discounting), but given our finding that only the s parameter changed systematically with amount of probabilistic reward, we considered only regression models in which s varied with amount but b did not.

One can fit the present data estimating a separate s for each amount condition, but based on the observed linear relation between log s and log amount (see Fig. 4), it seems reasonable to consider a power model in which s is assumed to increase as a power function of the amount of the probabilistic reward, that is, $V = A/(1 + bX)^s$, where $s = aA^c$, or equivalently,

$$V = A/(1 + bX)^{aA^c}. \tag{3}$$

As was the case with Equation 1, the power model can be rewritten in terms of relative subjective value (i.e., the subjective value of the probabilistic reward expressed as a proportion of its nominal value) for purposes of comparing the discounting of different probabilistic amounts.

When a regression model (based on Eq. 2) in which separate s and b parameters were estimated for each amount was fitted to the group median relative subjective values from all nine amount conditions, it accounted for 99.1% of the total variance. This 18-parameter model (two parameters for each of the nine amount conditions) may be compared with a 10-parameter model that incorporates a single amount-independent b parameter and nine separate s parameters (one for each amount condition), which accounted for 98.3% of the total variance. Notably, the 3-parameter power model (Eq. 3), in which the exponent of the hyperboloid is a power function of amount, also provided a very good fit, accounting for 96.4% of the total variance with $b = 7.75$, $a = 0.295$, and $c = 0.082$. Indeed, inspection of the fit of the power model to the group median data, represented by the dashed curves in Figure 1, reveals how little there may be to gain from making the model more complicated, at least in this application.

DISCUSSION

The present results demonstrate that the amount of a probabilistic reward has a profound effect on the degree to which its value is discounted. More specifically, the degree of probability discounting in the current study increased continuously as reward amount increased over more than five orders of magnitude (i.e., from $20 to $10M). These results are in contrast to those observed with delayed monetary rewards, for which the degree of discounting decreases with amount. Moreover, the observed changes in degree of probability discounting appear to have been driven by changes in the value of the exponent, the *s* parameter, which increased as a function of the amount of the probabilistic reward, whereas the value of the rate parameter, *b*, did not change systematically. Again, these results stand in contrast to those obtained with delayed rewards, for which the *b* parameter decreases with amount while the *s* parameter remains constant. The theoretical significance of the present findings stems in part from the fact that according to the psychophysical scaling interpretation of the exponent in the discounting function, the exponent should be constant across variations in the amount of a particular type of reward. Although previous findings with delayed rewards are consistent with this interpretation, the present results clearly are not, and argue that the value of the exponent in the probability-discounting function is not a simple scaling parameter.

Reward Amount and Probability Discounting

The current study is not the first to show that amount of reward affects probability discounting. Previous studies have reported that larger probabilistic rewards are discounted more steeply than smaller ones (e.g., Estle et al., 2006; Green et al., 1999; Myerson et al., 2003; Rachlin, Brown, & Cross, 2000), but this is the first study to show that this effect of amount on individuals' choices between certain and probabilistic rewards holds over a range extending up to millions of dollars.

For the most part, the possibility of an effect of amount on probability discounting has received little attention (for example, see the review by Wu, Zhang, & Gonzalez, 2004). This has been true even in studies that manipulated reward amount. For example, Gonzalez and Wu (1999) gave participants choices between certain and probabilistic monetary rewards, and determined the certain equivalents (subjective values) of the probabilistic rewards. In the conditions relevant to the current study, the amount of probabilistic reward ranged between $25 and $800 and the probabilities ranged from .99 to .01, corresponding to odds against receiving the reward of 0.01 to 99.0. The data for the smallest ($25) and largest ($800) amounts used in their study, as well as for an intermediate ($100) amount, are presented in Figure 6. As may be seen, amount effects are apparent even over this relatively narrow range. Perhaps because the study examined choice from the perspective of prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), which implicitly assumes that amount does not affect the degree of discounting, Gonzalez and Wu did not analyze their data in a way that could reveal such magnitude effects. Nevertheless, their data provide further support for the generality of the finding that amount affects how steeply probabilistic rewards are discounted and does so in a way opposite to that observed with delayed rewards.

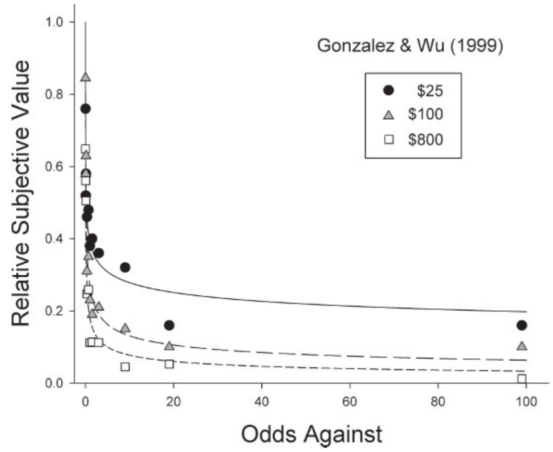

Fig. 6. Relative subjective value as a function of the odds against receiving a reward. Data are from Gonzalez and Wu (1999) and represent the means for the smallest and largest probabilistic reward amounts, as well as an intermediate amount. Curves represent the fit of Equation 2 to the data (all R^2s > .91).

Given that amount of reward affects the degree of probability discounting, the question arises as to what mechanisms are responsible. One approach to answering this question is to examine the way in which the parameters of the discounting function behave when the amount of probabilistic reward is varied. In the current study, the value of the s parameter increased systematically as a function of amount whereas the value of the b parameter did not, a finding consistent with previous studies that also examined the effects of amount but did so over a much narrower range (e.g., Estle et al., 2006; Myerson et al., 2003). Indeed, McKerchar, Green, and Myerson (2010) showed that amount appears to affect the exponent of the probability-discounting function regardless of which form of hyperboloid discounting function is used, be it the form proposed by Rachlin (2006), in which the odds against receiving a probabilistic reward are raised to a power, or the form examined here (Eq. 2), in which the entire denominator is raised to a power. As already noted, although the value of the s parameter of individuals' discounting functions increased systematically as the amount of probabilistic reward was increased from $20 to $10 million, it was significantly greater than 1.0 in only one out of a possible 360 cases. Given how large some of the reward amounts were, these results suggest that there may be a ceiling on the value of the exponent in the hyperboloid discounting function.

As noted previously, the present results raise problems for an interpretation of the exponent in the probability-discounting function that is based on the psychophysical scaling of reward likelihood. This is because if the exponent did reflect the scaling of reward likelihood, then (contrary to the present results) it would have remained constant across amounts of probabilistic reward. The present results, however, are not inconsistent with the idea that the exponent reflects the scaling of amount, although this would not explain the increase in the exponent as the amount of the probabilistic reward increases (as is shown in the derivation provided in the Appendix).

Weighting Reward Likelihood

If the standard interpretation of the probability-discounting function, in which the exponent reflects psychophysical scaling, is incorrect, then how should the function be interpreted? One way to think about the form of the probability-discounting function is to note its similarity to the equation for the expected value (EV) of a probabilistic outcome. The EV is the amount (A) of the reward (or other outcome) multiplied by its probability (P). That is, $EV = A * P$. To better see this similarity, we may rewrite the equation for expected value in terms of the odds against receiving the reward (X), $EV = A * [1/(1 + X)]$, and then introduce a coefficient, b, that makes the expression in brackets in effect a biased statement of the likelihood of reward:

$$EV = A * [1/(1 + bX)]. \tag{4}$$

It may be seen that Equation 4 is a simple hyperbola, where b governs the rate at which EV decreases as the odds against increase, and is similar to Equation 1, the hyperboloid discounting function, but without the exponent, s. Thus, the probability-discounting function may be seen to involve something similar to expected value in that the amount of a reward is multiplied by a function of the likelihood of the reward. But what are we to make of the exponent from this perspective?

To answer this question, we would note that prospect theory (Kahneman & Tversky, 1979) takes a similar approach to establishing subjective value. That is, it too assumes that the value of a prospect or gamble is the product of a function of amount, termed the *value function*, multiplied by a function of probability, termed the *weighting function*. The value function is typically approximated using a power function of amount; a number of mathematical forms have been suggested for the probability-weighting function (e.g., Gonzalez & Wu, 1999; Prelec, 1998; Tversky & Fox, 1995), but no consensus has yet emerged. (It may be noted that Kahneman and Tversky also assume that determination of a prospect's value follows an ''editing phase,'' but this need not be considered here.)

For our immediate purposes, what is important is that Kahneman and Tversky (1979) clearly distinguished between a subjective probability and a decision weight. That is, an individual may know that the probability of a reward is very low, yet give little weight to that probability when deciding whether or not to gamble. Thus, according to prospect theory, which emphasizes the role of decision weights rather than subjective probabilities, the fact that people buy lottery tickets does not necessarily mean that they think the odds are better than those printed on the tickets; it may just mean that they give more weight to the amount of the possible win and less weight to the odds when making their decisions to buy tickets.

To some extent, this weighting of the odds is captured in the *b* parameter of the discounting function. This effect of *b* may be seen by considering the case where the exponent s = 1.0. Under this condition, choice will be what economists term *risk-averse* if *b* is greater than 1.0 and *risk-taking* if *b* is less than 1.0. The *b* parameter captures an aspect of weighting that, as the present results show, does not vary systematically with reward amount. In contrast, the exponent does vary with amount, implying that the weight one places on the probability, relative to the amount, varies with what is at stake. More specifically, as the exponent increases with the amount of reward, there is a corresponding increase in the ''weighted probability,'' as given by the expression $[1/(1 + bX)]^{aA^c}$ in the power model (Eq. 3). This expression corresponds to what Kahneman and Tversky (1979) termed a *probability-weighting function*, except that our probability-weighting function incorporates an effect of amount whereas prospect theory does not. The implications of this amount-dependent weighting function may be made apparent by considering specific examples.

We begin by rewriting the power model (Eq. 3), substituting the estimates of the *b*, *a*, and *c* parameters obtained from fitting the power model to the group medians (see Fig. 1), as

$$V = A * [1/(1+7.75\,X)]^{0.295\,A^{0.082}}.$$

Now consider the case in which the probability of winning is .50 (and thus *X*, the odds against, equals 1). If the probabilistic amount, *A*, is \$1,000, then the exponent equals $0.295*1000^{0.082}$ or 0.52, whereas if the probabilistic amount is \$1,000,000, the exponent equals 0.92. Given that $X = 1$, the expression in brackets, $[1/(1 + 7.75\,X)]$, is equal to $1/(1 + 7.75)$ or 0.114, which, when raised to the 0.52 power to obtain the weighted probability, is 0.323, but when raised to the 0.92 power is 0.136 (compare the closed and open circles in Figure 7). Thus, when the possible reward is \$1,000, the subjective value of a 50% chance of winning is close to one-third of the actual amount, whereas when the possible reward is \$1,000,000, the subjective value of a 50% chance of winning is less than one-seventh of the actual amount. Although in both cases the subjective value is less than the expected value (i.e., one-half the actual amount), it is as if people give more weight to the fact that they might not win when \$1,000,000 is at stake than when \$1,000 is what is at stake. Simply put, people's behavior is more risk-averse when the rewards are larger.

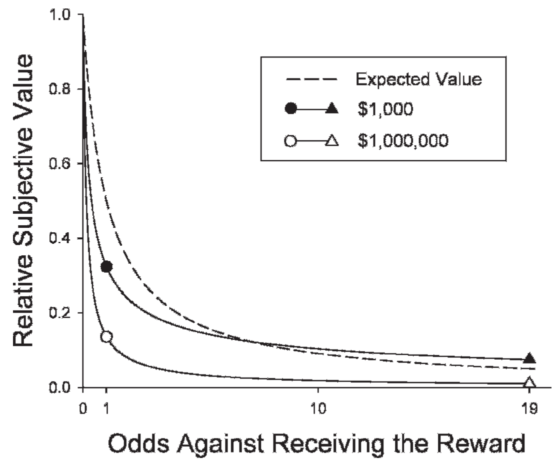

Fig. 7. Relative subjective and expected value as a function of the odds against receiving a reward. The dashed curve represents the relative Expected Value (i.e., $EV / A = P$). The two solid curves represent the predictions of Equation 3, rewritten in terms of relative subjective value, with the parameters obtained from fitting the equation to the group medians for the \$1,000 and \$1,000,000 amount conditions. The solid circle and solid triangle represent the relative subjective value of a \$1,000 reward when the odds against are 1 and 19 ($P = .50$ and $P = .05$), respectively; the open circle and open triangle represent the corresponding relative subjective values for a \$1,000,000 reward.

As the odds of winning get worse, however, behavior does not necessarily remain risk-averse (i.e., the subjective value of a reward is not necessarily lower than its expected value). Whether behavior will be risk-averse or not depends on both the odds and the amount involved. For example, consider the case where the probability of winning is only .05 (and thus $X = 19$). In this case, the expression in brackets in the preceding equation is equal to 0.0067, which when raised to the 0.52 power (in the case of a possible \$1,000 reward) gives a weighted probability of 0.074 (see the closed triangle in Fig. 7). Thus, when the possible reward is $1,000, the subjective value is nearly 7.5% of the actual amount, which is greater than the expected value of 5% of the actual amount, indicating that risk-taking behavior is predicted. In contrast, when the possible reward is $1,000,000, behavior is still predicted to be risk-averse when the probability of winning is 0.05: With this larger amount, the expected value is still 5% of the actual amount, but the subjective value is much less. In fact, the subjective value, calculated by raising 0.0067 to the 0.92 power, is only 1% of the actual amount (see the open triangle in Fig. 7).

These examples suggest that people will switch from being risk-averse to risk-taking as the odds against winning increase when smaller amounts are involved but remain risk-averse when very large amounts are involved. But what about the fact that people buy lottery tickets for very large amounts? Such behavior is definitely not risk-averse because the expected value of a lottery ticket is less than its price. We would note, however, that the preceding examples were based on average parameter values, and the average person does not necessarily buy lottery tickets. Moreover, extrapolating beyond the range of odds for which we have data is itself risky business, and therefore one would want to have further data about choice under situations in which the odds against winning are considerably higher than those studied here before extending these results to choices such as whether or not to buy a lottery ticket. Nevertheless, the general principle captured by the present analysis is that the larger the amount of possible reward, the lower the likelihood of winning will have to be before people switch from being risk-averse to being risk-taking.

The preceding analysis of the effect of reward amount is consistent with Kahneman and Tversky's (1979) idea that the role of reward likelihood in probability discounting (or decision under risk, as they put it) is best thought of in terms of the relative weight given to a reward's likelihood versus its amount. What is new here, and which is not even considered in prospect theory, is the idea that these weights are themselves amount-dependent. That is, people appear to put more weight on reward likelihood (i.e., their behavior is more controlled by the degree to which the larger reward is probabilistic) when the amount of a probabilistic reward is large, and thus are more likely to choose the certain reward (i.e., to be risk-averse), and (perhaps paradoxically) they put less weight on the likelihood of reward (and thus relatively more weight on the amount) when the amount of the probabilistic reward is small.

Conclusions

Taken together, the present results converge on the conclusion that one does not need to fit a two-parameter hyperboloid discounting function to the data for each amount of probabilistic reward separately. Rather, a model in which the b parameter is amount-invariant provides a more efficient description. Moreover, the value of the exponent increases monotonically as the amount increases, and the power model represents an effort to describe that increase even more parsimoniously. It is possible, of course, that another mathematical form might describe the relation between the exponent of the discounting function and the amount of probabilistic reward better than a power function does. Nevertheless, it seems clear that there is a functional relation between the exponent and the reward amount, and that a model is needed that describes that relation. A power function would appear to provide at least a reasonable approximation to that relation.

The change in the exponent of the probability-discounting function with reward amount can be understood in terms of the concept of decision weights introduced by Kahneman and Tversky (1979). People's decisions under risk are basically risk-averse, as Kahneman and Tversky have shown, and in the model proposed here, this is captured by the fact that the value of the b parameter in the discounting function is typically greater than 1.0. However, this tendency towards risk aversion is modulated by the exponent of the discounting function. The lower the exponent, the less risk-averse the behavior, as if the exponent reflects the weight put on the likelihood of reward.

Although this interpretation builds on the concept of decision weights, the weights envisioned by prospect theory are not affected by reward amount. In addition, although Kahneman and Tversky (1979) propose that the weight placed on the likelihood of a reward changes as a function of its probability according to a nonlinear probability-weighting function, their theory does not posit a specific mathematical form for this function. Tversky and Kahneman (1992) later used a one-parameter weighting function, but they deliberately emphasized the qualitative predictions of their theory rather than parameter estimates and goodness of fit. In any case, their one-parameter weighting function does not allow for effects of amount on probability weighting like those reported here. In contrast, our power model (Eq. 3) is the first probability-discounting model to specifically incorporate an amount-dependent probability-weighting function.

The present findings with respect to the effects of amount on the discounting of probabilistic rewards reinforce the need to distinguish between delay discounting and probability discounting. First, increases in the amount of delayed reward lead to shallower discounting, whereas increases in the amount of probabilistic reward lead to steeper discounting. Second, increases in amount of delayed reward are associated with decreases in the rate parameter of the delay discounting function, whereas increases in amount of probabilistic reward are associated with increases in the exponent of the probability-discounting function. Third, the amount-independent exponent of the delay discounting function is consistent with a psychophysical scaling interpretation, whereas the amount-dependent exponent of the probability-discounting function is consistent with a probability-weighting interpretation. Taken together, these findings indicate that despite the fact that the delay and probability-discounting functions are of similar mathematical form, underlying this form are fundamentally different mechanisms.

REFERENCES

Estle, S. J., Green, L., Myerson, J., & Holt, D. D. (2006). Differential effects of amount on temporal and probability discounting of gains and losses. *Memory & Cognition*, *34*, 914–928.

Gonzalez, R., & Wu, G. (1999). On the shape of the probability-weighting function. *Cognitive Psychology*, *38*, 129–166.

Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin*, *130*, 769–792.

Green, L., & Myerson, J. (2010). Experimental and correlational analyses of delay and probability discounting. In G. J. Madden & W. K. Bickel (Eds.), *Impulsivity: The behavioral and neurological science of discounting* (pp. 67–92). Washington, DC: American Psychological Association.

Green, L., Myerson, J., & Ostaszewski, P. (1999). Amount of reward has opposite effects on the discounting of delayed and probabilistic outcomes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 418–427.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.

McKerchar, T. L., Green, L., & Myerson, J. (2010). On the scaling interpretation of exponents in hyperboloid models of delay and probability discounting. *Behavioural Processes*, *84*, 440–444.

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

Myerson, J., Green, L., Hanson, J. S., Holt, D. D., & Estle, S. J. (2003). Discounting of delayed and probabilistic rewards: Processes and traits. *Journal of Economic Psychology*, *24*, 619–635.

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

Prelec, D. (1998). The probability-weighting function. *Econometrica*, *66*, 497–527.

Rachlin, H. (2006). Notes on discounting. *Journal of the Experimental Analysis of Behavior*, *85*, 425–435.

Rachlin, H., Brown, J., & Cross, D. (2000). Discounting in judgments of delay and probability. *Journal of Behavioral Decision Making*, *13*, 145–159.

Stevens, S. S. (1957). On the psychophysical law. *Psychological Review*, *64*, 153–181.

Tversky, A., & Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review*, *102*, 269–283.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.

Wu, G., Zhang, J., & Gonzalez, R. (2004). Decision under risk. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 399–423). Oxford, UK: Blackwell Publishing.

Received: August 25, 2010
Final Acceptance: December 13, 2010

APPENDIX

Psychophysical Scaling of Amount and the Probability-Discounting Function

Here we show why the observed increase in the exponent of the probability-discounting function as amount increases is not inconsistent with psychophysical scaling of amount (even though it is inconsistent with psychophysical scaling of reward likelihood). Assume, as is frequently done, that the scaling of amount is described by a simple power function: $V = A^d$ in which, because this *is* psychophysical scaling, the exponent remains constant as amount varies. As in Equation 1, the value, V_p, of a probabilistic reward of amount A_p would then be given by

$$V_p = A_p^d \ * \ 1/(1+bX)^s, \qquad \text{(A1)}$$

and the value, V_c, of a certain reward (for which $X = 0$) of amount A_c would be given by

$$V_c = A_c^d. \qquad \text{(A2)}$$

At indifference, the value of the certain reward is equal to the value of the probabilistic reward, $V_c = V_p$. Thus,

$$A_c^d = A_p^d \ * \ 1/(1+bX)^s \qquad \text{(A3)}$$

To find A_c, the certain equivalent of the probabilistic reward, one takes the d^{th} root of both sides of the preceding equation, which yields

$$A_c = A_p \ * \ 1/(1+bX)^{s/d}. \qquad \text{(A4)}$$

This is equivalent to Equation 1 in the text (in which V is measured as A at indifference) but reveals that if amount is psychophysically scaled, then the exponent of Equation 1 contains the exponent of the psychophysical power function that describes the scaling of amount.

Psychophysical Scaling of Amount and the Power Model of Probability Discounting

The same argument may be extended to the power model of probability discounting (Eq. 3 in the text). If $s = aA^c$, then substituting into Equation A4 yields

$$A_c = A_p \ * \ 1/(1+bX)^{(aA^c)/d}. \qquad \text{(A5)}$$

Dropping the subscripts, Equation A5 may be rewritten as

$$\begin{aligned} V &= A \ * \ 1/(1+bX)^{(aA^c)/d} \\ &= A/(1+bX)^{(a/d)\ A^c}, \end{aligned} \qquad \text{(A6)}$$

which has the same mathematical form as Equation 3 in the text. That is, the exponent of the hyperboloid discounting function (Eq. 3) is still equal to the product of a constant (now *a/d*, rather than just *a*) and the amount of probabilistic reward raised to a power. Thus, only the interpretation of this constant is affected by assuming psychophysical scaling of amount, not the form of the discounting function.