International Journal of Comparative Psychology, 2011, 24, 99-111.

Waiting for Grapes: Expectancy and Delayed Gratification in Bonobos

Jeffrey R. Stevens
Max Planck Institute for Human Development, Germany

Alexandra G. Rosati and Sarah R. Heilbronner
Duke University, U.S.A.

Nelly Mühlhoff
Max Planck Institute for Human Development, Germany

Responses to delayed rewards vary widely across individuals and have important implications for personality and temperament. Animals may avoid delayed rewards because the future is uncertain. Therefore, expectations about receiving a future reward should influence the response to delayed payoffs. Here, we offered bonobos (*Pan paniscus*) a delayed gratification task in which food accumulated over time. Once subjects chose to consume the reward, food stopped accumulating. We tested their willingness to wait with a reliable and an unreliable experimenter to vary the subjects' expectations that they would receive the food. Subjects waited less often with the unreliable experimenter but showed individual differences in the degree to which reliability generalized across experimental tasks. These data suggest that the expectations generated about the likelihood of receiving future rewards influence how individuals balance current and future needs.

All organisms face tradeoffs between immediate and future rewards. Accept the available mate or seek out a more attractive one? Raise offspring in the current environment or wait for more prosperous conditions? Stop to eat this piece of food or keep searching for a richer patch? The way in which individuals choose to balance rewards over time, or make intertemporal choices (Read, 2004; Stevens, 2010), can therefore influence many biologically important functions, ranging from reproductive success to foraging efficiency. One of the central paradoxes in the study of intertemporal choice is why most humans and animals seem to prefer immediacy to more delayed outcomes with higher value (Stevens & Stephens, 2009). This may occur because the future is uncertain: delayed outcomes may not come to fruition. Indeed, uncertainty relates to intertemporal choice both in terms of the psychology underlying preferences (Green & Myerson, 2004; Prelec & Loewenstein, 1991; Rachlin, Raineri, & Cross, 1991) and evolutionary models of how decision makers should choose (Kacelnik, 2003; Stevens & Stephens, 2009).

We are indebted to our test experimenters: Juliet Conlin and Verena Köstner. We would like to thank the staff of Zoo Berlin, Bernhard Blaszkiewitz, Peter Rahn, and André Schüle for granting permission for the project and especially Ruben Gralki for invaluable help in testing the bonobos. This work is in accordance with the American Psychological Association's Guidelines for the Ethical Conduct in the Care and Use of Animals, the Animal Behavior Society's Guidelines for the Treatment of Animals in Behavioral Research and Teaching, and the animal welfare laws of Germany. Correspondence concerning this article should be addressed to Jeffrey R. Stevens, Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany. (jstevens@mpib-berlin.mpg.de or jeffrey.r.stevens@gmail.com).

However, few studies have clearly demonstrated a causal relationship between the probability of receiving rewards and intertemporal choice. Consequently, understanding the role of uncertainty in temporal preferences—as well as why individuals differ in their response to uncertainty—will advance our understanding of intertemporal choice at both the proximate and ultimate level.

In one of the first studies to relate uncertainty and intertemporal choice, Mahrer (1956) examined the role of expectancy on delayed reinforcement in children. Here, expectancy refers to building a likelihood of receiving a reward. To generate expectancy, Mahrer had children interact with one of three experimenters that varied in his/her reliability, or probability of delivering a reward. In the 'high-reliability' condition, over four sessions, the experimenter fulfilled the promise to bring a toy to the child on the following day. In the 'low-' and 'medium-reliability' conditions, the experimenter fulfilled the promise in zero or two of the four sessions. Following this expectancy-generating phase, the experimenter then offered each child the choice between an immediate reward and a more preferred but delayed reward. The children generalized the information about the reliability of the experimenter to this choice phase. As reliability decreased, the children chose the delayed reward in 76%, 42%, and 25% of the trials. Thus, the children learned about the reliability of the experimenters and used this information about uncertainty when making choices regarding the future.

Although Mahrer's (1956) work suggests a role for uncertainty in intertemporal choice, whether these effects extend beyond human verbal interactions remains unclear. Here we tested whether an expectation about uncertainty influences intertemporal choice in bonobos (*Pan paniscus*), a species whose temporal preferences (Rosati, Stevens, Hare, & Hauser, 2007) and responses to uncertainty in payoffs (Heilbronner, Rosati, Stevens, Hare, & Hauser, 2008) have both been studied. Subjects participated in a standard delayed gratification task (also called a delay maintenance or food accumulation task, Pelé, Dufour, Micheletta, & Thierry, 2010; Toner & Smith, 1977) previously used by Beran and colleagues (Beran, 2002; Beran & Evans, 2006; Beran, Savage-Rumbaugh, Pate, & Rumbaugh, 1999). In their paradigm, an experimenter typically placed a bowl in front of subjects and then placed single rewards in the bowl one-by-one until either the subject reached into the bowl or all rewards had been placed into the bowl. At that point, the subject could access the bowl of rewards. Once the subject reached into the bowl, however, the experimenter stopped distributing rewards. Therefore, with increased wait times, subjects could obtain larger quantities of reward, so the amount of time subjects waited represented a measure of their willingness to delay gratification. By introducing into this paradigm a potential risk of not receiving the rewards, we examine how uncertainty plays a role in intertemporal choice.

To manipulate expectancy, we varied the reliability of experimenters in an introductory period before offering the bonobos any intertemporal choices. The low-reliability experimenter began a session by offering the subject a highly valued food item but then pulling it away from the subject, thereby preventing the subject from consuming the food. After this pre-feeding period, the low-reliability experimenter conducted a standard experimental delayed gratification session with the bonobo. In contrast, the high-reliability experimenter began a session by

offering the subject food but followed through in providing the food before conducting a standard experimental session. With this manipulation, we can assess whether the bonobos generalize reliability from the initial phase to the subsequent delayed gratification phase, preferring more immediate options in the presence of unreliable experimenters.

Method

Subjects

Four bonobos from Zoo Berlin participated in this experiment: Simon (28 year old male), Santi (26 year old male), Vifijo (13 year-old male) and Opala (9 year-old female). These four bonobos were housed together with the alpha female and her infant in a series of adjacent enclosures. Simon was hand reared by zookeepers, whereas the other subjects were reared by bonobos. None had experienced similar experiments before, but they all experienced a "clicker" training regime conducted by a zookeeper in which they received food for behavioral responses needed for veterinary purposes. The bonobos received fresh fruits and vegetables after the experimental task and had constant access to water.

Materials

We tested subjects individually in different compartments of their home cage. The testing apparatus was similar to Beran (2002). It consisted of one bowl (containing food rewards) and one circular mark on the floor (15 cm in circumference and 41 cm from the subject's cage) in which food rewards were placed in front of the subjects. The bonobos could easily reach their hands between the metal bars of their cage to access the circle (Figure 1). The bonobos received grape halves (about 0.5 g) as rewards because they were a highly preferred food item. The experimenters used a stopwatch to measure time intervals and videotaped all trials for behavioral analysis. Data were analyzed using R statistical software version 2.12.0 (R Development Core Team, 2010) and the plyr (Wickham 2010) and lattice (Sarkar 2008) packages. The original document for this paper used Sweave (Leisch 2002) to embed the R code into the document, ensuring reproducible research (de Leeuw 2001). Data and R code are available as supplementary materials.

Training phase

All sessions occurred between 0915 and 1130 h about five days a week, with one session per subject per day. A 'control' experimenter tested subjects immediately after their clicker training in the morning. In these trials the experimenter sat about 1 m from the subject's cage and placed single grape halves from the bowl into the subject's circle (Figure 1). Initially there was no waiting period between placing the rewards, so the experimenter placed the rewards in the circle at a rate of about one per second (interreward interval = 1 s). A daily session consisted of 8 trials with a maximum reward of 10 grape halves per trial and an intertrial interval of 30 s. Once the subject's hand crossed the line of the circle or after the 10 grapes had been placed, the experimenter covered the bowl with a lid and placed it behind her back, stopping the flow of rewards and initiating the 30 s intertrial interval. If the subject did not place all received rewards into his/her mouth within the 30 s intertrial interval, the experimenter waited until 10 s after the subject placed the last reward into his/her mouth before beginning a new trial.

After a subject completed two consecutive sessions of waiting for all of the 10 rewards in 6 out of 8 trials, we incremented the interreward interval by 1 s for the next session. We continued this procedure for each subject until he/she achieved an interreward interval of 5 s, which we used for the remainder of the experiment. Two subjects (Vifijo and Opala) did not pass this criterion after 43 and 47 sessions, respectively, so they did not advance to the experimental phase. The other two subjects required 7 (Santi) and 15 (Simon) sessions to pass the criterion and begin the experimental phase. During the trials, the experimenter kept verbal remarks to a minimum and avoided eye contact with

the apes both when placing the grapes and during the intertrial intervals.

Figure 1. Experimental apparatus. Experimenters placed grapes in a circle within the bonobo's reach. Once the bonobo's hand crossed the circle, the experimenter stopped placing grapes.

Acclimation phase

After the subjects reached the training criteria for five days on the 5 s interreward interval, we introduced two new test experimenters—the low- and high-reliability experimenters—to the subjects by having them observe two training sessions. After the initial encounters, the test experimenters participated in three pre-feeding sessions before subjects experienced their daily training session. Only one experimenter conducted the pre-feeding session each day and the two test experimenters alternated, with the control experimenter testing in between the test experimenters (e.g., low-reliability, control, high-reliability, control, low-reliability, control, etc.) for a total of 12 sessions (3 sessions for each test experimenter). In these pre-feeding sessions, each experimenter placed a banana slice in her hand about 0.5 m above the feeding circle for about 10 s. During this time, the experimenter alternated her gaze between the banana and the subject five times. After 10 s, the control and high-reliability experimenters dropped the banana slice into the feeding circle. In one randomly chosen trial the low-reliability experimenter gave the subject the banana; however, in the other five trials she took the banana away and placed it in another bowl. Thus, the unreliable behavior of this experimenter introduced uncertainty into the decision-making context. The high-reliability experimenter, in contrast, was as reliable as the control experimenter, but like the low-reliability experimenter was a novel person, thereby controlling for familiarity. After waiting 10 s after the subject put the banana piece into his/her mouth, the test experimenter began another feeding trial until she completed six trials. Following the pre-feeding period, the control experimenter conducted a standard delayed gratification task for all sessions. A standard session consisted of 8 trials with a 5 s interreward interval and a 30 s intertrial interval. We counterbalanced the assignment of the test experimenter as a low- or high-reliability experimenter across subjects.

Experimental phase 1: Generalized expectation

In the experimental phase 1, we examined whether previous experience with reliable or unreliable experimenters would influence subject's expectations about receiving the food in the delay

of gratification task. To do so, we altered which of the three experimenters (control, low-reliability, or high-reliability) conducted the entire test session. Each test experimenter conducted one session with each subject. One control session was conducted between each experimental session, and the experimental session order alternated between test experimenters, with the initial session counterbalanced across subjects. The control experimenter was present in the building for all sessions, but the test experimenters were only present for the experimental sessions.

All experimental sessions began with a pre-feeding period as described above. In the control and high-reliability sessions, the experimenters gave the subject six pieces of banana. In the low-reliability sessions, the low-reliability experimenter gave the subject only one of six pieces of banana. Afterwards, the control experimenter gave the subjects the remaining five pieces of banana to equate consumption across conditions. Following this pre-feeding period, the experimenters conducted a regular delayed gratification session.

Experimental phase 2: Social competition

Next, we examined whether introducing a competitive element would alter the bonobo's expectations about uncertainty in the delay of gratification task. This phase was similar to phase 1, except the low-reliability experimenter gave food to a bonobo in a neighboring cage (in visual contact about 1-1.5 m away), instead of placing it in another bowl. After this pre-feeding period, the experimenter conducted a delayed gratification session (with no visual contact with a neighbor). Both test experimenters conducted one session with each subject in a similar order as in phase 1.

Experimental phase 3: Direct experience

Finally, we examined whether direct experience with unreliable experimenters in the delayed gratification task would influence intertemporal choice. Instead of a pre-feeding period, the control and high-reliability experimenters conducted standard delayed gratification session in which waiting to the end resulted in receiving 10 grape halves. In the low-reliability session, the experimenter interrupted all trials by removing 5 grapes instead of placing the 8th, 9th or 10th grape (randomly selected) into the feeding circle. After the removal, the experimenter waited for the intertrial interval and continued to the next placement up to 10 placements. If the subject waited until the end of a trial, he received a maximum of four rewards (for expected values, see Table 1). Both test experimenters conducted three sessions with each subject. Test experimenters alternated the order of testing with the control sessions in between (as in the acclimation phase).

Table 1

Expected values for planning to stop after each food placement

Number of placements	1	2	3	4	5	6	7	8	9	10
Expected value of grapes (*EV*)	1	2	3	4	5	6	7	5.63*	4.85*	4

*$EV_8 = (1-p_8)\square 8 + p_8\square 2 = 5.63$ and $EV_9 = (1-p_8)(1-p_9)\square 9 + (1-p_8)p_9\square 3 + p_8\square\ 2 = 4.85$, where p_8 and p_9 represent the probabilities of being interrupted in the eighth and ninth grape, respectively (p_8 = 0.4 and p_9 = 0.38).

Results

Phase 1: Generalized expectation

In phase 1, we tested whether uncertainty in a pre-feeding period generalized to a subsequent delayed gratification task. When aggregating the number of grapes before stopping across subjects (Table 2 and left panel in Figure 2), we do not see a strong difference across experimenters. When separated by trials and subject (left panel in Figure 3), however, experimenter condition did influence choice for one of the subjects. Simon waited for 10 grapes in all trials with the control experimenter and in all but the first trial with the high-reliability experimenter. Yet, in the last half of the session with the low-reliability experimenter, Simon stopped after only six or seven grapes. Santi, in contrast, waited for all 10 grapes in all conditions with the exception of the first trials. The median waiting responses did not differ between the control and high-expectancy experimenters. Individual and aggregated descriptive statistics are shown in Table 2.

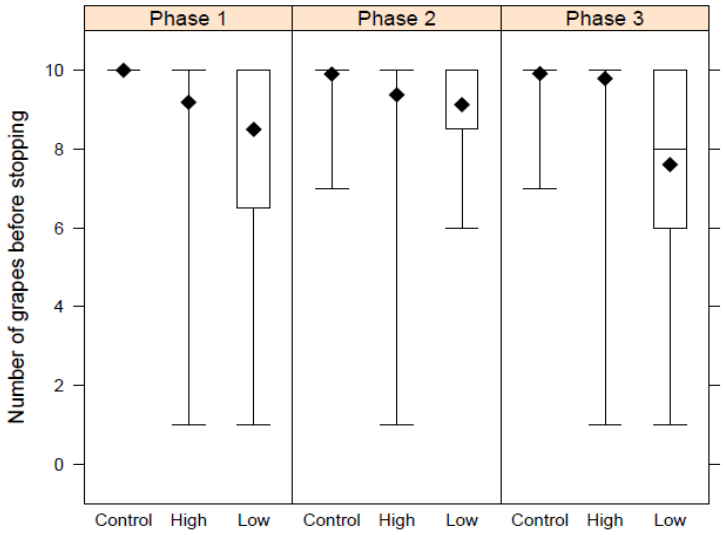

Figure 2. Choices aggregated over trials and subjects. The number of grapes placed before stopping did not differ across the control, high-reliability, and low-reliability conditions in phases 1 (Generalized Expectation) or 2 (Social Competition). In phase 3 (Direct Experience), however, subjects stopped earlier in the low-reliability condition. Lines in boxes represent medians, diamonds represent means, boxes represent 25-75% quantiles (interquartile range), and whiskers represent 1.5 times the interquartile range.

Table 2
Descriptive statistics for number of grapes before stopping in each phase

Subject	Phase	Condition	Median	Mean	Std. Dev.	95% CI	Trials
Simon	1	Control	10	10	0	0	16
	1	High	10	8.88	3.18	2.66	8
	1	Low	8.5	8.12	2.03	1.7	8
	2	Control	10	9.81	0.75	0.4	16
	2	High	10	9.88	0.35	0.3	8
	2	Low	8.5	8.25	1.58	1.32	8
	3	Control	10	9.92	0.45	0.13	48
	3	High	10	9.62	1.84	0.78	24
	3	Low	7.5	7.21	2.59	1.09	24
Santi	1	Control	10	10.06	0.25	0.13	16
	1	High	10	9.5	1.41	1.18	8
	1	Low	10	8.88	3.18	2.66	8
	2	Control	10	10	0	0	16
	2	High	10	8.88	3.18	2.66	8
	2	Low	10	10	0	0	8
	3	Control	10	9.92	0.45	0.13	48
	3	High	10	9.96	0.2	0.09	24
	3	Low	8.5	8	1.77	0.75	24
Combined	1	Control	10	10.03	0.18	0.06	32
	1	High	10	9.19	2.4	1.28	16
	1	Low	10	8.5	2.61	1.39	16
	2	Control	10	9.91	0.53	0.19	32
	2	High	10	9.38	2.25	1.2	16
	2	Low	10	9.12	1.41	0.75	16
	3	Control	10	9.92	0.45	0.09	96
	3	High	10	9.79	1.3	0.38	48
	3	Low	8	7.6	2.23	0.65	48

Phase 2: Social competition

Phase 2 replicated phase 1 with one difference: in the pre-feeding phase, the low-reliability experimenter gave undelivered food to a nearby group mate. As in phase 1, experimenter condition did not influence choice at the aggregate level (Table 2 and center panel in Figure 2), but one subject stopped early with the low-reliability experimenter (center panel in Figure 3). This is the same subject that stopped early in phase 1, and again the effect seemed to increase over the session. The median waiting responses did not differ between the control and high-expectancy experimenters. The other subject treated all experimenters similarly.

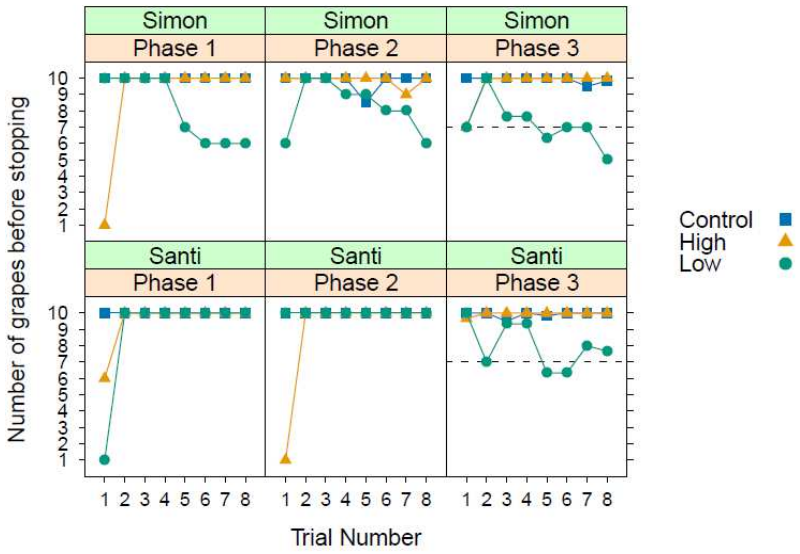

Figure 3. Choices separated by subject and trial number. Though Simon stopped early in all phases, Santi did not stop early until the direct experience of uncertainty in phase 3. Phases 1 and 2 include one session, and phase 3 represents the mean over three sessions. Dashed line represents optimal stopping point in phase 3.

Phase 3: Direct experience

In phase 3, we dispensed with the pre-feeding phase, and had subjects directly experience uncertainty in the delayed gratification task. Here, in both the aggregate data (Table 2 and right panel in Figure 2) and in the individual data (right panel in Figure 3), subjects stopped earlier for the low-reliability experimenter than the other two experimenters. In this phase, both subjects responded to uncertainty by stopping early with the low-reliability experimenter but not with the control or high-expectancy experimenter. The mean (± 95% confidence intervals) number of grapes before stopping dropped from 8.9 ± 0.8 to 6.1 ± 1.3 and 7.8 ± 0.9 across the three sessions.

Given the levels of uncertainty faced with the low-reliability experimenter, we can calculate the expected value of stopping after each grape. For instance, in this phase, the expected value of waiting for the tenth grape was no longer 10; instead, it was 4 because a subject did not receive 1 grape on the interrupted trial and then lost 5 grapes. Table 1 gives the expected values of stopping after each grape, using the actual overall interruption probabilities experienced. From these calculations, we see that stopping after seven grapes yields the highest expected value. This approximates the number of grapes at which the bonobos stopped (right panel in Figure 3).

Discussion

Our data highlight two main points: uncertainty influences delayed

gratification and individuals differ in their responses to uncertainty. To summarize the experimental results, one of the two bonobos generalized about an experimenter's unreliability in the pre-feeding period with bananas to the delayed gratification task with grapes. Namely, this bonobo became more impulsive (that is, began consuming grapes before all were distributed, thereby foregoing larger rewards) with an unreliable experimenter. A similar pattern emerged in the second phase when we added social competition to unreliability: the same bonobo exhibited similar levels of impulsivity towards the low-reliability experimenter as he exhibited in the first phase, thus generalizing about uncertainty across tasks. In the final phase, both subjects responded to the direct experience of uncertainty in the delayed gratification task. When interrupted during a trial, both bonobos chose more impulsively, matching optimal choices for maximizing their expected food intake.

This work provides evidence of uncertainty influencing intertemporal choices in animals. Though theoretical and empirical work relates uncertainty and intertemporal choice (Green & Myerson, 2004; Prelec & Loewenstein, 1991; Rachlin et al., 1991), previous direct tests have failed to find a connection. Notably, Henly et al. (2008) offered blue jays (*Cyanocitta cristata*) a standard 'self-control' task in which the birds chose between a smaller, sooner and a larger, later reward. Henley and colleagues added a constant probability of interruption per unit of time delay such that, as delays increased, the probability of losing the reward increased. Choosing longer delays increased the risk of not receiving the reward, so uncertainty positively correlated with time delay. Despite time acting as a proxy for uncertainty, the blue jays did not choose the smaller, sooner option more as the interruption risk increased. Rather, they either did not change their preferences across interruption rates or even slightly increased their preference to choose the larger, later option at higher interruption rates.

In contrast to the Henly study, our data demonstrate that uncertainty does influence intertemporal choices in bonobos. When uncertainty is embedded directly into the temporal task, both bonobos responded by reducing delayed gratification under uncertainty. Interestingly, one of the bonobos even responded to the expectation of uncertainty by generalizing about an experimenter's reliability from a non-temporal, non-choice task to an intertemporal choice task. Both the temporal element and the choice element differed across the pre-feeding and delayed gratification tasks. Thus, like the children in Mahrer's (1956) task, one bonobo passively acquired information about uncertainty associated with an individual in one context and applied it to a novel context in which information about uncertainty was potentially relevant. Surprisingly, rather than starting off impulsive, this bonobo tended to start sessions in both phase 1 and phase 2 by waiting for the full amount of grapes and only became impulsive in the second half of trials with the low-expectancy experimenter. This may have resulted from the sheer strength of previous reinforcement in this task. Both subjects experience between 20-27 practice and acclimation sessions before beginning the test phases, so the impulsive response had to compete with an overlearned waiting response.

To increase attention to the uncertainty associated with the pre-feeding task, we included a phase in which the low-reliability experimenter gave non-

received food items to a social competitor (phase 2). Surprisingly, this manipulation did not influence the choices of either subject: Simon stopped early at the same rate as in phase 1 and Santi continued to disregard the unreliability of this experimenter. There are at least two possible reasons for social competition not affecting choices. First, bonobos are a very tolerant species, showing food sharing and other cooperative behaviors (Fruth & Hohmann, 2002; Hare & Kwetuenda, 2010; Hare, Melis, Woods, Hastings, & Wrangham, 2007), so perhaps the social competition was not as salient as it might be for other species such as chimpanzees. Second, the social competition does not directly bear on the delayed gratification task. It could be that both subjects already attended to the pre-feeding task in the first phase, so social competition did not increase attention further.

The primary weakness of this experiment is, of course, the small sample size. Though the sample size is not uncommon in experimental studies of bonobos, we would like to see more bonobos tested. Nevertheless, we do see differences between the individuals in this experiment. First, we began the experiment with four subjects, but only two passed the training criterion. The two that could not wait 5 s between grapes were the youngest and lowest ranking male and female of the group. Second, we saw differences between the two subjects that completed the experiment. One subject became more impulsive in the delayed gratification task in response to generalized unreliability, whereas the other subject showed no effect of unreliability until it directly applied to the delayed gratification task.

One potentially interesting way to think about individual differences comes from the psychology of expertise literature (Cokely & Kelley, 2009; Ericsson & Lehmann, 1996) in which researchers distinguish between capacity and strategy. A *capacity* typically refers to the maximum ability to perform a mental task, whereas a *strategy* is a set of actions used to accomplish a task (VandenBos, 2006). Much work in psychology attributes individual differences to differences in capacity, where it is treated as a fixed, rather hard-wired trait. However, many individual differences may result from applying different strategies. For instance, in our experiment, one might postulate that one bonobo did not respond to unreliability in phases 1 and 2 because he could not generalize across contexts (he lacked the capacity). Yet, he may have had the capacity to generalize but simply did not employ the strategy of becoming more impulsive because the contexts differed (namely, there was actually no unreliability in the delayed gratification task). In our case, this strategy actually yielded a higher payoff.

Another application of the capacity/strategy distinction is important for species organized in dominance hierarchies. In many cases, most individuals of a species may have similar psychological capacities. However, the nature of dominance hierarchies results in pressure to employ different strategies depending on where an individual falls on the hierarchy. Subordinates, for instance, may have to demonstrate strong inhibitory control (Amici, Aureli, & Call, 2008; Aureli et al., 2008) and patience (Stevens & Stephens, 2008) in the face of dominants because grabbing food in front of the dominant can be quite costly. Nevertheless, in the absence of dominants, subordinates may act extremely impulsively to take advantage of the opportunity to feed without harassment (Stevens & Stephens, 2002). Thus, depending on the context, individuals will employ different

strategies. This may explain why our subordinate bonobos did not delay gratification enough to pass our training criterion. That being said, humans show enormous variation in delayed gratification (Frederick, Loewenstein, & O'Donoghue, 2002; Mischel, Shoda, & Rodriguez, 1989). In fact, decades of study have revealed that an individual's tendency toward impulsivity or patience predicts such diverse life variables as likelihood of substance dependence, marital success, academic achievement, and cooperativeness (Ayduk et al., 2000; Duckworth & Seligman, 2005; Harris & Madden, 2002; Mischel et al., 1989). Though not well studied yet, there likely is variation in delayed gratification in other animals that cannot be attributed to dominance status, sex, age, etc. Recent work in the animal personality literature proposes adaptive reasons for individual differences (Dall, Houston, & McNamara, 2004; Sih, Bell, Johnson, & Ziemba, 2004; Wolf, van Doorn, Leimar, & Weissing, 2007), but much work in this area remains.

References

Amici, F., Aureli, F., & Call, J. (2008). Fission-fusion dynamics, behavioral flexibility, and inhibitory control in primates. *Current Biology*, *18*(18), 1415–1419.

Aureli, F., Shaffner, C., Boesch, C., Bearder, S., Call, J., Chapman, C., ...Van Schaik, C. P. (2008). Fission-fusion dynamics: New research frameworks. *Current Anthropology*, *49*(4), 627–654.

Ayduk, O., Mendoza-Denton, R., Mischel, W., Downey, G., Peake, P. K., & Rodriguez, M. (2000). Regulating the interpersonal self: Strategic self-regulation for coping with rejection sensitivity. *Journal of Personality and Social Psychology*, *79*(5), 776-792.

Beran, M. J. (2002). Maintenance of self-imposed delay of gratification by four chimpanzees *(Pan troglodytes*) and an orangutan *(Pongo pygmaeus*). *Journal of General Psychology*, *129*(1), 49–66.

Beran, M. J., & Evans, T. A. (2006). Maintenance of delay of gratification by four chimpanzees *(Pan troglodytes*): The effects of delayed reward visibility, experimenter presence, and extended delay intervals. *Behavioural Processes*, *73*(3), 315–324.

Beran, M. J., Savage-Rumbaugh, E., Pate, J. L., & Rumbaugh, D. M. (1999). Delay of gratification in chimpanzees *(Pan troglodytes*). *Developmental Psychobiology*, *34*(2), 119–127.

Cokely, E. T., & Kelley, C. M. (2009). Cognitive abilities and superior decision making under risk: A protocol analysis and process model evaluation. *Judgment and Decision Making*, *4*(1), 20–33.

Dall, S. R. X., Houston, A. I., & McNamara, J. M. (2004). The behavioural ecology of personality: Consistent individual differences from an adaptive perspective. *Ecology Letters*, *7*(8), 734–739.

Duckworth, A. L., & Seligman, M. E. P. (2005). Self-discipline outdoes IQ in predicting academic performance of adolescents. *Psychological Science*, *16*(12), 939-944.

Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Reviews in Psychology*, *47*, 273–305.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, *40*(2), 351–401.

Fruth, B., & Hohmann, G. (2002). How bonobos handle hunts and harvests: Why share food? In C. Boesch, G. Hohmann, & L. F. Marchant (Eds.), *Behavioural diversity in chimpanzees and bonobos* (pp. 231–243). Cambridge, UK: Cambridge University Press.

Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin*, *130*(5), 769–792.

Hare, B., & Kwetuenda, S. (2010). Bonobos voluntarily share their own food with others. *Current Biology*, *20*(5), R230-R231.

Hare, B., Melis, A. P., Woods, V., Hastings, S., & Wrangham, R. W. (2007). Tolerance allows bonobos to outperform chimpanzees on a cooperative task. *Current Biology*, *17*(7), 619–623.

Harris, A. C., & Madden, G. J. (2002). Delay discounting and performance on the prisoner's dilemma game. *Psychological Record*, *52*(4), 429-440.

Heilbronner, S. R., Rosati, A. G., Stevens, J. R., Hare, B., & Hauser, M. D. (2008). A fruit in the hand or two in the bush? Divergent risk preferences in chimpanzees and bonobos. *Biology Letters*, *4*(3), 246–249.

Henly, S., Ostdiek, A., Blackwell, E., Knutie, S., Dunlap, A. S., & Stephens, D. W. (2008). The discounting-by-interruptions hypothesis: Model and experiment. *Behavioral Ecology*, *19*(1), 154–162.

Kacelnik, A. (2003). The evolution of patience. In G. Loewenstein, D. Read, & R. F. Baumeister (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (p. 115-138). New York: Russell Sage Foundation.

Leeuw, J. de. (2001). *Reproducible research: The bottom line* (Department of Statistics Papers Series). Los Angeles: University of California-Los Angeles.

Leisch, F. (2002). Sweave: Dynamic generation of statistical reports using literate data analysis. In W. Härdle & B. Rönz (Eds.), *Compstat 2002—Proceedings in computational statistics* (p. 575–580). Heidelberg: Physica Verlag.

Mahrer, A. R. (1956). The role of expectancy in delayed reinforcement. *Journal of Experimental Psychology*, *52*(2), 101–106.

Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science*, *244*(4907), 933–938.

Pelé, M., Dufour, V., Micheletta, J., & Thierry, B. (2010). Long-tailed macaques display unexpected waiting abilities in exchange tasks. *Animal Cognition*, *13*(2), 263-271.

Prelec, D., & Loewenstein, G. (1991). Decision-making over time and under uncertainty: A common approach. *Management Science*, *37*(7), 770–786.

R Development Core Team. (2010). R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria: http://www.R-project.org

Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective probability and delay. *Journal of the Experimental Analysis of Behavior, 55*(2), 233–244.

Read, D. (2004). Intertemporal choice. In D. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 424–443). Oxford: Blackwell.

Rosati, A. G., Stevens, J. R., Hare, B., & Hauser, M. D. (2007). The evolutionary origins of human patience: Temporal preferences in chimpanzees, bonobos, and adult humans. *Current Biology*, *17*(19), 1663–1668.

Sarkar, D. (2008). *Lattice: Multivariate data visualization with R*. New York: Springer.

Sih, A., Bell, A. M., Johnson, J. C., & Ziemba, R. E. (2004). Behavioral syndromes: An integrative overview. *Quarterly Review of Biology, 79*(3), 241–277.

Stevens, J. R. (2010). Intertemporal choice. In M. D. Breed & J. Moore (Eds.), *Encyclopedia of animal behavior.* Oxford: Academic Press.

Stevens, J. R., & Stephens, D. W. (2009). The adaptive nature of impulsivity. In G. J. Madden & W. K. Bickel (Eds.), *Impulsivity: The behavioral and neurological science of discounting* (p. 361-387). Washington, DC: American Psychological Association.

Stevens, J. R., & Stephens, D. W. (2008). Patience. *Current Biology*, *18*(1), R11–12.

Stevens, J. R., & Stephens, D. W. (2002). Food sharing: A model of manipulation by harassment. *Behavioral Ecology*, *13*(2), 393–400.

Toner, I. J., & Smith, R. A. (1977). Age and overt verbalization in delay-maintenance behavior in children. *Journal of Experimental Child Psychology*, *24*(1), 123-128.

VandenBos, G. R. (2006). *APA dictionary of psychology*. Washington, DC: American Psychological Association.

Wickham, H. (2010). *plyr: Tools for splitting, applying and combining data*. (R package version 1.2.1).

Wolf, M., van Doorn, G. S., Leimar, O., & Weissing, F. J. (2007). Life-history trade-offs favour the evolution of animal personalities. *Nature*, *447*(7144), 581–584.