

Running head: IMPULSIVITY AND CHOICE

Taking More, Now:
The Optimality of Impulsive Choice Hinges on Environment Structure

A. Ross Otto, Arthur B. Markman, and Bradley C. Love
Department of Psychology
University of Texas at Austin

In Press at Social Psychological and Personality Science

Impulsivity is a stable personality trait associated with myopic choice behavior that favors immediate rewards over larger, delayed rewards and is often characterized as maladaptive inside and outside of the laboratory. An alternative view suggests that the consequences of trait impulsivity depend on the nature of the task environment. On this view, the optimal level of impulsivity varies across task payoff structures. This hypothesis is tested in two dynamic decision-making tasks that differ in the relative payoffs of delayed and immediate rewards. In a task that favors delayed rewards to immediate rewards, high-impulsive participants perform worse than low-impulsive participants. In contrast, in a task that favors immediate rewards over delayed rewards, high-impulsive participants outperform low-impulsive participants. These results suggest a more nuanced conceptualization of trait impulsivity as it applies to rewards-related decision-making that may help explain the variability observed in this trait across individuals.

Keywords: Impulsivity, Choice, Decision-Making, Delay of Gratification

Introduction

The ability to forego immediate rewards in the service of receiving larger future rewards is often described as a hallmark of effective self-control in both humans and animals (Logue, 1988; Rachlin & Green, 1972). In this context, impulsive behavior is defined as taking an immediate reward that prevents one from obtaining more valuable future rewards (Ainslie, 1975; Evenden, 1999). In laboratory delay-of-gratification procedures, individuals from impulsive populations are more likely to choose smaller immediate rewards over larger delayed rewards (Cherek & Lane, 1999). Indeed, self-report measures suggest that impulsivity is a stable personality trait associated with a range of pathological behaviors such as substance abuse, gambling and binge eating (De Wit, 2008; Patton et al., 1995; Perry & Carroll, 2008) as well as laboratory-assessed risk-taking (Lejuez et al., 2002).

Much of contemporary research casts impulsivity as a maladaptive trait that predicts negative behaviors and consequences inside and outside of the laboratory (e.g., Mischel et al., 1989). Further, a number of investigations present converging evidence supporting the heritability of the personality trait (Congdon & Canli, 2008). From an evolutionary standpoint, it is puzzling why humans would exhibit continued variability in a trait if it appears to result only in maladaptive behavior. Are there benefits, under any set of circumstances, to impulsive decision-making behavior? In this report, we consider the notion that the behavioral consequences of trait impulsivity hinge on the structure of the environment in which the decision-maker is placed. A highly impulsive person who makes choices as if each day is his or her last is unlikely to exhibit adaptive behavior in modern society, which rewards sacrifices in the short term (such as working rather than playing). Nonetheless, a person who focuses only on the far term may not live long enough to see her sacrifices rewarded. Thus, for any given environment, there is an ideal tradeoff between attention to the short-term and long-term. In this study, we find that self-assessed trait impulsivity predicts how individuals from a non-clinical population adapt to changing rewards in the environment, which, depending on the structure of the environment, benefit or hinder overall task performance.

As a broad example, consider the conflict between short- and long-term goals facing the executive of an expanding company. At each choice point, she can decide to invest in new equipment and training, thus increasing the firm's future output and increasing long-term profits, or she can instead cut costs and receive an immediate boost in short-term profits. Her choices effectively influence the *state* of the company, which in turn affect her future returns. Assuming that the market for the company's products remains stable over a long period of time, the long-term advantageous option is to forgo immediate profits and instead invest in equipment and training because doing so brings about greater profits in the future. By contrast, if the market for the company's product will soon disappear, then investing in equipment and training is suboptimal, because it entails a large opportunity cost—forgoing immediate profits—which will not likely be recovered within the duration of the company's existence. Because the company's time horizon is short, improving the state of the company's production capacity will not afford as much profit as taking the immediately resulting profits from cutting costs. In this case, it is optimal for the executive to maximize short-term rewards and continually cut costs. These two situations illustrate how the long-term advantage of “impulsive” choice depends on the payoff structure of the environment.

In Experiment 1, we adopt a dynamic choice task that instantiates the conflict facing the executive with a long time horizon: the option resulting in larger immediate rewards adversely affects rewards in the long term, while the less immediately attractive option leads to larger

rewards in the long term. The reward functions of the two options are depicted in Figure 1A. The vertical axis represents the immediate rewards resulting from selections to the two options as a function of environment state. In all states, the Long-Term Decreasing option (henceforth LT-Decreasing; solid line) always yields a higher immediate reward than the Long-Term Increasing option (henceforth LT-Increasing; dashed line). The horizontal axis represents the state of the task environment, defined as the number of LT-Increasing choices made over the last fourteen trials. Making LT-Increasing choices—analagous to investing in equipment in training in the above example—moves the state rightward and increases the immediate rewards for both options. Making LT-Decreasing choices—analagous to cutting costs—moves the state leftward and decreases both options' rewards. Because the maximum of the LT-Increasing option (that is, its resulting reward in the rightmost state) is larger than the minimum of the LT-Decreasing option (its resulting reward in the leftmost state), long-term optimal choice requires forgoing the larger immediate rewards resulting from the LT-Decreasing option and continually making LT-Increasing choices. In effect, choices that yield larger immediate rewards negatively affect future rewards, whereas options that are less immediately attractive lead to larger future rewards. Thus, the two options are in conflict with each other: the LT-Decreasing option is locally superior but globally inferior, while the LT-Increasing option is locally inferior but globally superior (cf. Herrnstein et al., 1993; Otto & Love, 2010).

Unlike previous intertemporal choice paradigms in which the delivery schedule and magnitudes of rewards for each option are explicitly presented to decision-makers (e.g. Figner et al., 2010, McClure, Laibson, Loewenstein, & Cohen, 2003; Mischel et al., 1989), the consequences of the options in the present task are unknown to participants and must be learned experientially (Herrnstein, Loewenstein, Prelec, & Vaughan, 1993; Yarkoni, Braver, Gray, & Green, 2005; Warry, Remington, & Sonuga-Bark, 1999). In the present task, participants are not provided with information about the structure of the task environment or the short- and long-term characteristics of the options. Because the setting is novel, participants can bring little task-specific knowledge to bear. Thus, we assume that a decision-maker's choice behavior reflects their default orientation towards reward-related decisions, which, in turn, should be moderated in part by their level of trait impulsivity. To ensure that decision-makers primarily rely on local, trial-by-trial feedback to gauge their choice performance in the task, we did not provide any global performance measures (e.g., history of responses, cumulative payoffs, etc.)

To measure trait impulsivity, we chose the Barratt Impulsiveness Scale (BIS-11), which has been demonstrated to have high internal consistency (Cronbach's $\alpha=.82$) in normal populations and test validity in a number of psychiatric populations (Patton, Stanford, & Barratt, 1995). We hypothesized that high-impulsive individuals—following the conceptualization of impulsive choice in the literature (Herrnstein, 1997; Logue, 1988)—would make decisions informed more by immediate reward differences. More specifically, we predicted that high-impulsive participants should be disinclined to choose the LT-Increasing option repeatedly after they observe the large difference between the immediate rewards of the options, resulting in fewer suboptimal LT-Increasing choices overall. Conversely, we predicted immediate reward differences should not exert a strong influence on the choices of low-impulsives, and thus, these participants would exhibit more optimal, LT-Increasing choices.

Of course, to fully demonstrate that high levels of trait impulsivity can yield more optimal patterns of choice in some decision environments, we devised a second environment in which high-impulsive participants should outperform low-impulsive participants, akin to the example above in which the executive's company has a limited time horizon. Experiment 2

examines the consequences of impulsive choice in an environment where the dynamics of the reward structure are preserved—in that the locally superior option worsens future rewards for both options and the locally inferior option improves future rewards for both options—but the long-term consequences of the options are reversed (see Figure 1B). Because the minimum reward of the LT-Decreasing choice is greater than the reward of the LT-Increasing choice at *every* point, the optimal long-term reward maximizing strategy is to continuously make LT-Decreasing choices (Otto et al., 2009). We predict that high-impulsive participants, upon observing the large immediate reward difference between the Short- and LT-Increasing options, should strongly disfavor the LT-Increasing option and instead make choices to the LT-Decreasing option. Because the LT-Decreasing option is *optimal* in this environment, we expect high-impulsives to perform better than low-impulsives—who we do not predict will be deterred by immediate reward differences from choosing the LT-Increasing option.

Experiment 1

Method

Participants

A total of 45 undergraduates enrolled in an introductory psychology course at a major Southwestern university participated in this experiment in exchange for course credit and a small cash bonus tied to performance. The sample from which our sample was drawn is 54.3% female, 42.5% male, with 3.2% who declined to report their gender. The reported ethnicities of the participant pool were as follows: Hispanic/Latino: 15.5%, African American: 6.8%, Asian: 23.3%, Caucasian: 46.5%, Native American: <0.1%, others: 1.5%. The ages of participants in this pool ranged from 17 to 55, ($M=19.08$, $SD=1.76$).

Materials and Procedure

Participants were administered the BIS-11 questionnaire (Patton et al., 1995) which consists of 30 statements, such as “I do things without thinking” and “I am more interested in the present than the future” with which participants stated their level of agreement on a four-point scale. Higher summed scores indicate higher levels of impulsivity.

Following the questionnaire, participants played the “Farming on Mars” game (see Gureckis & Love, 2009), an adaptation of Herrnstein et al.’s (1993) choice task. Participants read a story about NASA scientists on Mars attempting to extract oxygen from its atmosphere in order to create breathable air for use in a human colony. They were informed that, as members of the project, their job was to extract as much oxygen as possible from the atmosphere. To do this, they needed to repeatedly choose between two “oxygen-extraction robots” with different properties. Beyond this information, participants were only told that the specific oxygen-extracting properties of the two robots were unknown and that their cash bonus was related to the total amount of oxygen extracted (i.e., points earned).

The game was administered on a computer and consisted of 250 trials, and the points earned on each trial were governed by the reward structure depicted in Figure 1A. On each trial, participants were shown a panel with two response buttons labeled “Robot 1” and “Robot 2,” and a display between the two buttons that read “Choose.” The task interface conveyed no information about the long- and short-term properties of the options. The assignment of the response buttons to the LT-Increasing and LT-Decreasing options were randomized across participants.

The reward obtained by choosing the Long-Term option on a given trial was defined by:

$$5 + 70 \times \frac{h}{14}, \quad (1)$$

and the reward obtained by choosing the LT-Decreasing option on a given trial was defined by:

$$65 + 70 \times \frac{h}{14}, \quad (2)$$

where h represents the number of LT-Increasing choices over the last 14 trials. A small amount of Gaussian noise ($\sigma=4$) was added to the reward on each trial. After the participant clicked one of the two response options (i.e., robots) using a computer mouse, the amount of oxygen “points” earned was visually depicted using an 11 by 11 grid of blue dots. The number of dots that were active in this grid indicated the amount of points earned on the current trial (i.e., more dots meant a larger amount of points in the current trial). No information about the cumulative points generated across trials was provided. At the end of the experiment, participants were paid at a rate of 1 cent (.01 USD) per 100 points earned.

Results and Discussion

BIS-11 Questionnaire and Overall Choice Behavior

Total BIS-11 scores ranged from 44 to 82 ($M=62.35$, $SD=9.04$). It is important to note that under the reward structure used in this experiment, trial-by-trial rewards depended on the proportion of LT-Increasing responses made over a moving window of the previous 14 trials, but our analyses concerned participants’ choices over all 250 trials of the experiment. Overall, participants made less than half of their responses ($M=.29$, $SD=.24$) to the LT-Increasing response, but participants exhibited substantial variability in their choice behavior. To assess the extent to which participants’ responses changed over the course of the experiment, we divided the 250 trials into five blocks of 50 trials each. A one-way ANOVA conducted on proportions of LT-Increasing responses over the five blocks revealed a significant effect of block, $F(4,44)=8.41$, $p<.0001$. In other words, participants’ overall tendency to make optimal, LT-Increasing responses increased significantly over the 250 trials. Participants’ evident difficulty in finding and applying the long-term optimal choice strategy mirrors previous work investigating human performance in this task (Gureckis & Love, 2009; Herrnstein et al., 1993; Warry et al., 1999).

Relationship between BIS-11 Scores and Choice Behavior

We were interested primarily in the extent to which self-reported trait impulsivity predicted optimal performance in this task, defined as the proportion of optimal LT-Increasing responses made over the course of the experiment. Figure 2A depicts the relationship between participants’ BIS-11 scores and overall choice behavior. As expected, more impulsive participants exhibited lower proportions of optimal LT-Increasing responses, and this negative correlation was significant, $r(43)=-.31$, $p<.05$. In other words, more impulsive participants were less likely to utilize the optimal, long-term rewards maximizing strategy.

We also examined optimal performance in terms of points per trial averaged over all 250 trials. Impulsive choices also had monetary consequences: we found a suggestive negative correlation between BIS-11 scores and average rewards per trial, $r(43)=-.29$, $p=.054$. Because average rewards have limited sensitivity as an index of performance—due to the underlying task dynamics and stochastic rewards—we found that a quartile analysis, depicted in Figure 2B, provided a more compelling demonstration of the relationship between trait impulsivity and monetary performance. On average, participants in the top quartile of BIS-11 scores earned significantly fewer points per trial ($M=66.28$, $SD=1.01$, $N=10$) than participants in the bottom quartile ($M=67.94$, $SD=2.07$, $N=10$) [$t(18)=2.38$, $p<.05$, $d=1.02$].

As an explanation for this aggregate pattern of choice, we hypothesized that high-impulsive participants were more sensitive to the decrease in immediate rewards observed when switching from the LT-Decreasing to LT-Increasing option (see Figure 1A), and as a result, observing such a decrease would be more likely to bring about a response switch. To test this hypothesis, we used a mixed-effects general linear model (random effects over subjects) and regressed response switches on trial t against each participant's BIS-11 score and a binary indicator specifying whether observed rewards decreased between trials $t-1$ and $t-2$ (constrained to situations where an LT-Decreasing response was made at $t-1$ and an LT-Increasing response was made at $t-2$). The coefficient estimates are reported in Table 1. Crucially, we found a significant positive interaction between BIS-11 scores and observed reward decreases, suggesting that impulsivity was associated with an increased tendency to change responses after observing immediate reward differences between the two options.

Experiment 2

Experiment 2 modified the reward structure of Experiment 1 so that the LT-Decreasing option became the long-term optimal option (Otto et al., 2009). This situation is analogous to the example above in which the executive's time horizon does not warrant improving the company's production capacity. In the structure depicted in Figure 1B, rewards only depended on the last 10 choices instead of the last 14 choices. Critically, because the range of possible states has been truncated, the minimum reward given by the LT-Decreasing option was *larger* than the maximum possible reward given by the LT-Increasing option. Note that this situation requires the opposite pattern of choice as in Experiment 1: choosing the option with the larger immediate rewards (the LT-Decreasing option) is actually the optimal long-term pattern of choice.

Method

Participants

A total of 43 students at a major Southwestern university participated in this experiment, drawn from the same participant pool described in Experiment 1.

Materials and Procedure

The same questionnaire, stimuli and instructions were used as in Experiment 1 with the exception of the modified reward structure. The resultant slopes of the reward functions and the distance between the LT-Increasing and LT-Decreasing reward functions remained the same, but the number of possible states was changed from 14 to 10 states. The rewards for the LT-Increasing and LT-Decreasing options were defined by:

$$5 + 50 \times \frac{h}{10}, \quad (3)$$

and

$$65 + 50 \times \frac{h}{10}, \quad (4)$$

respectively, where h represents the number of LT-Increasing choices over the last 10 trials. A small amount of Gaussian noise ($\sigma=4$) was added to the reward on each trial.

Results and Discussion

BIS-11 Questionnaire and Overall Choice Behavior

Total BIS-11 scores ranged from 49 to 91 ($M=64.65$, $SD=10.62$). It is important to note that under the reward structure used in this experiment, trial-by-trial rewards depended on the proportion of LT-Increasing responses made over a moving window of the previous 10 trials, but our analyses concerned participants' choices over all 250 trials of the experiment. Participants' average proportion of LT-Increasing responses over the entire experiment was .24 ($SD=.23$), and

a one-way ANOVA conducted on choice proportions over five blocks revealed a significant effect of trial, $F(4,42)=3.18$, $p<.05$. That is, participants' overall tendency to make suboptimal LT-Increasing responses increased significantly over the course of the experiment.

Relationship between BIS-11 Scores and Choice Behavior

Figure 2C reveals that, as expected, increasing impulsivity was associated with a smaller proportion of LT-Increasing responses, $r(41)= -.41$, $p<.01$. Thus, in this study more impulsive participants tended to make more optimal responses than did less impulsive participants. In both Experiment 1 and Experiment 2, of course, the impulsive participants were selecting the LT-Decreasing option which results in larger immediate rewards. However, in Experiment 2, the option with larger immediate rewards is also best in the long-term.

We found a significant positive relationship between impulsivity level and rewards-per-trial averaged over 250 trial, $r(41)=.44$, $p<.01$. Figure 2D reveals that the consequences of these choice patterns were opposite Experiment 1: participants in the top quartile of BIS-11 scores earned significantly more points per trial ($M=111.59$, $SD=3.46$, $N=10$) than participants in the bottom quartile ($M=105.41$, $SD=6.55$, $N=11$) [$t(19)=2.63$, $p<.05$, $d=1.18$].

As in Experiment 1, we also hypothesized that high-impulsive participants would be more likely to switch responses after observing a decrease in immediate rewards observed after having sampled the LT-Decreasing and LT-Increasing options. Employing the same mixed-effects regression as above, we found—as in Experiment 1—a significant negative interaction between BIS-11 scores and observed reward decreases (see Table 1) suggesting that impulsivity was associated with an increased tendency to change responses after observing immediate reward differences. Of course, optimal behavior in this environment requires choosing the option with larger immediate rewards.

General Discussion

While impulsivity is often discussed as a maladaptive trait associated with myopic decision-making and a myriad of pathological behaviors (Patton et al., 2003; Perry & Carroll, 2008; Petry, 2001), the present set of results lend credence to the notion that impulsivity is not a purely maladaptive trait but one whose consequences hinge on the structure of the decision-making environment. We found that low- and high-impulsive participants exhibited consistent trial-to-trial choice behavior across the two experiments: impulsive participants were more likely to choose the option with larger immediate rewards—based on their direct experience from sampling the two options—whereas less impulsive participants were more likely to pass up larger immediate gains and opt for the option associated with increasing rewards over time.

Crucially, whether each tendency was advantageous or disadvantageous depended not on any endogenous factors, but solely on the environment. In one environment (Experiment 1), impulsivity was disadvantageous, but in another (Experiment 2), impulsivity led to greater overall returns. Indeed, the observed interaction between personality trait and situation supports the widely held view that this interaction is possibly a more important determinant of behavior than either factor alone (Bowers, 1973).

It is important to note that the present task does not allow us to completely rule out the possibility that low- and high-impulsives may differ not in their basic orientation towards immediate rewards, but rather, in a more cognitive or strategic capacity. For example, the choices of low-impulsive subjects may have been motivated by curiosity about that task structure—that is, repeatedly choosing the less immediately rewarding option in order to uncover the underlying payoff dynamics. Alternately, low- and high-impulsives may differ in their fundamental ability to perceive and understand the underlying task structure. On this view, low-

impulsives would outperform high-impulsives in both Experiments 1 and 2 as understanding the task structure would facilitate optimal choice in both structures. However, Experiment 2—which differed from Experiment 1 only in the number of states—revealed that high-impulsives outperformed low-impulsives, suggesting against the possibility that low-impulsives formed more veridical representations of the task structure than high-impulsives. Future work should investigate the role that strategic differences play in the choice behavior of low- and high-impulsive decision-makers.

The present results provide a more nuanced account of the relationship between impulsive choice and the reward structure in which individuals are making decisions. The idea that the optimality of a particular personality trait hinges on the decision-maker's current environment is not entirely novel (cf. Dickman, 1990; Evenden, 1999). Intuitively, it may be adaptive in the long run for a species to exhibit consistent personality trait variability in the face of a potentially changing environment. Indeed, behavioral ecologists have entertained the idea that fluctuations in environment factors indirectly maintain genetic variation in animal personality (Dingemanse, Both, Drent, & Tinbergen, 2004). A body of theoretical work has also called into question the suboptimality of “impulsive” choice—that is, taking immediate rewards over larger delayed ones—arguing that what appears to be impulsive choice in animal delay-of-gratification procedures can be explained by reward rate maximization (Kacelnik, 1997; Stephens & Anderson, 2001) or a Bayesian assessment of the uncertainty in the hazard rate underlying long-term rewards delivery (Sozou, 1998). Relatedly, substance abusers, who exhibit pathologically risky and impulsive choice behavior (Petry, 2001), have been shown to exhibit—under certain circumstances—more advantageous choice behavior in gambling tasks (Shiv, Loewenstein, & Bechara, 2005).

The literature reveals an inconsistent predictive relationship between self-assessed trait impulsivity and reward-related choice behavior. Richards and colleagues (1999) report a marginally significant relationship between self-report measures of impulsiveness and propensity to choose smaller immediate rewards over larger delayed rewards in explicit tradeoff paradigms. Other studies have reported difficulty correlating self-report measures with delay-of-gratification behavior (Reynolds, Ortengren, Richards, & de Wit, 2006). Self-report measures of impulsivity have been shown to predict delay-of-gratification behavior in populations composed of violent and non-violent parolees (Cherek, Moeller, Dougherty, & Rhoades, 1997), parents of adolescents with disruptive behavioral disorders (Swann, Bjork, Moeller, & Dougherty, 2002), and heroin addicts (Kirby, Petry, & Bickel, 1999)—however, these studies were conducted using populations expected to exhibit wide ranges of psychopathology. The present set of results highlights the behavioral sensitivity of the present choice task to levels of trait variability found in normal, non-clinical populations.

References

- Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82(4), 463-96.
- Bowers, K. S. (1973). Situationism in psychology: An analysis and a critique. *Psychological Review*, 80, 307-336.
- Cherek, D. R., Moeller, F. G., Dougherty, D. M., & Rhoades, H. (1997). Studies of violent and nonviolent male parolees: II. Laboratory and psychometric measurements of impulsivity. *Biological Psychiatry*, 41(5), 523-529.
- Congdon, E., & Canli, T. (2008). A Neurogenetic Approach to Impulsivity. *Journal of Personality*, 76(6), 1447-1484.
- De Wit, H. (2009). Review: Impulsivity as a determinant and consequence of drug use: a review of underlying processes. *Addiction Biology*, 14(1), 22-31.
- De Wit, H., Flory, J. D., Acheson, A., McCloskey, M., & Manuck, S. B. (2007). IQ and nonplanning impulsivity are independently associated with delay discounting in middle-aged adults. *Personality and Individual Differences*, 42(1), 111-121.
- Dickman, S. (1985). Impulsivity and perception: Individual differences in the processing of the local and global dimensions of stimuli. *Journal of Personality and Social Psychology*, 48(1), 133-149.
- Dingemanse, N. J., Both, C., Drent, P. J., & Tinbergen, J. M. (2004). Fitness consequences of avian personalities in a fluctuating environment. *Proceedings of the Royal Society B: Biological Sciences*, 271(1541), 847-852.
- Figner, B., Knoch, D., Johnson, E. J., Krosch, A. R., Lisanby, S. H., Fehr, E., & Weber, E. U. (2010). Lateral prefrontal cortex and self-control in intertemporal choice. *Nature Neuroscience*, 13(5), 538-539.
- Gureckis, T. M., & Love, B. C. (2009). Short-term gains, long-term pains: How cues about state aid learning in dynamic environments. *Cognition*, 113(3), 293-313.
- Herrnstein, R. J., Loewenstein, G. F., Prelec, D., & Vaughan, W. (1993). Utility maximization and melioration: Internalities in individual choice. *Journal of Behavioral Decision Making*, 6(3), 149-185.
- Herrnstein, R. J. (1997). *The matching law*. H. Rachlin & D. I. Laibson (Eds.). Cambridge, MA: The Harvard University Press.
- Kacelnik, A. (1997). Normative and descriptive models of decision making: time discounting and risk sensitivity. In G.R. Bock & G. Cardew (Eds.) *Characterizing human psychological adaptations*. Chichester, UK: Wiley (Ciba Foundation Symposium 208).
- Kirby, K. N., Petry, N. M., & Bickel, W. K. (1999). Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental Psychology: General*, 128(1), 78-87.
- Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., Strong, D. R., et al. (2002). Evaluation of a Behavioral Measure of Risk Taking: The Balloon Analogue Risk Task (BART). *Journal of Experimental Psychology: Applied*, 8(2), 75-84.
- Logue, A. W. (1988). Research on Self-Control: An Integrating Framework. *Behavioral and Brain Sciences*, 11(4), 665-679.
- McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate Neural Systems Value Immediate and Delayed Monetary Rewards. *Science*, 306(5695), 503-507.
- Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior*, 64(3), 263-276.

- Patton, J. H., Stanford, M. S., & Barratt, E. S. (1995). Factor structure of the Barratt impulsiveness scale. *Journal of clinical psychology, 51*(6), 768-74.
- Perry, J., & Carroll, M. (2008). The role of impulsive behavior in drug abuse. *Psychopharmacology, 200*(1), 1-26.
- Petry, N. M. (2001). Substance abuse, pathological gambling, and impulsiveness. *Drug and Alcohol Dependence, 63*(1), 29-38.
- Otto, A. R., Gureckis, T. M., Markman, A. B., & Love, B. C. (2009). Navigating through abstract decision spaces: Evaluating the role of state generalization in a dynamic decision-making task. *Psychonomic Bulletin & Review, 16*(5), 957-963.
- Otto, A. R., & Love, B. C. (2010). You don't want to know what you're missing: When information about forgone rewards impedes dynamic decision-making. *Judgment and Decision Making, 5*(1), 1-10.
- Rachlin, H. (1995). Self-control: Beyond commitment. *Behavioral and Brain Sciences, 18*(1), 109-159.
- Reynolds, B., Ortengren, A., Richards, J. B., & de Wit, H. (2006). Dimensions of impulsive behavior: Personality and behavioral measures. *Personality and Individual Differences, 40*(2), 305-315.
- Richards, J. B., Zhang, L., Mitchell, S. H., & Wit, H. de. (1999). Delay or probability discounting in a model of impulsive behavior: effect of alcohol. *Journal of the Experimental Analysis of Behavior, 71*(2).
- Shiv, B., Loewenstein, G., & Bechara, A. (2005). The dark side of emotion in decision-making: When individuals with decreased emotional reactions make more advantageous decisions. *Cognitive Brain Research, 23*(1), 85-92.
- Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society B: Biological Sciences, 265*(1409), 2015-2020.
- Stephens, D. W., & Anderson, D. (2001). The adaptive value of preference for immediacy: when shortsighted rules have farsighted consequences. *Behavioral Ecology, 12*(3), 330 -339.
- Swann, A. C., Bjork, J. M., Moeller, F. G., & Dougherty, D. M. (2002). Two models of impulsivity: relationship to personality traits and psychopathology. *Biological Psychiatry, 51*(12), 988-994.
- Richards, J. B., Zhang, L., Mitchell, S. H., & Wit, H. D. (1999). Delay or probability discounting in a model of impulsive behavior: effect of alcohol. *Journal of the Experimental Analysis of Behavior, 71*(2).
- Warry, C. J., Remington, B., & Sonuga-Barke, E. J. S. (1999). When More Means Less: Factors Affecting Human Self-Control in a Local versus Global Choice Paradigm, *Learning and Motivation, 30*(1), 53-73.
- Yarkoni, T., Braver, T. S., Gray, J. R., & Green, L. (2005). Prefrontal Brain Activity Predicts Temporally Extended Decision-Making Behavior. *Journal of the Experimental Analysis of Behavior, 84*(3), 537-554.

Table 1

Logistic regression coefficients, indicating the influence of observed reward decreases when switching from the LT-Decreasing to the LT-Increasing option ($decrease_{t-2, t-1}$) and trait impulsivity (BIS-11) on response switching behavior. Note that a positive interaction term between these two quantities suggests that level of trait impulsivity increased the influence of observed reward decreases on subsequent response switches. Asterisks denote significance at the .05 level.

	Coefficient	Estimate (SE)	<i>p</i> -value
Experiment 1	$decrease_{t-2, t-1}$	0.80 (0.04)	< 0.0001*
	BIS-11	-0.08 (0.14)	0.53
	BIS-11 \times $decrease_{t-2, t-1}$	0.21 (0.06)	0.003*
Experiment 2	$decrease_{t-2, t-1}$	0.83 (0.04)	< 0.0001*
	BIS-11	-0.17 (0.16)	0.28
	BIS-11 \times $decrease_{t-2, t-1}$	0.30 (0.009)	< 0.0001*

Figure Captions

Figure 1. Task reward functions used in the two decision environments. Panel A depicts the rewards in Experiment 1 as a function of decision-makers' last 14 responses. Of particular interest is the fact that highest point of the LT-Increasing reward curve is higher than the lowest point of the LT-Decreasing curve. Thus, the long-term reward-maximizing strategy is to choose the LT-Increasing option on every trial. Panel B depicts the rewards in Experiment 2 as a function of decision-makers' last 10 responses—which is truncated from the 14 responses used in Experiment 1. Note that the LT-Decreasing choice always generates higher immediate rewards than the LT-Increasing choice. In contrast to the reward functions in Experiment 1, the global minimum of the LT-Decreasing reward curve is greater than the global maximum of the LT-Increasing reward curve. Therefore, the reward-maximizing strategy is to consistently choose the LT-Decreasing option, the reverse pattern of behavior as Experiment 1.

Figure 2. Relationships between impulsivity, choice behavior, and performance across Experiments 1 (top two panels) and 2 (bottom two panels). Panel A plots participants' proportions of optimal LT-Increasing choices against their BIS-11 scores. Note that consistent LT-Increasing choice was the optimal long-term strategy in this environment. Panel B plots the average points earned, per trial, by participants in the bottom and top quartiles of BIS-11 scores. In Experiment 1, low-impulsive participants significantly outperformed high-impulsive participants. Panel C plots participants' proportions of suboptimal LT-Increasing choices against their BIS-11 scores. In this environment, consistently choosing the LT-Decreasing option is long-term optimal choice strategy. Panel D plots the same performance analysis as in Panel B. In Experiment 2, high-impulsive participants significantly outperformed low-impulsive participants.

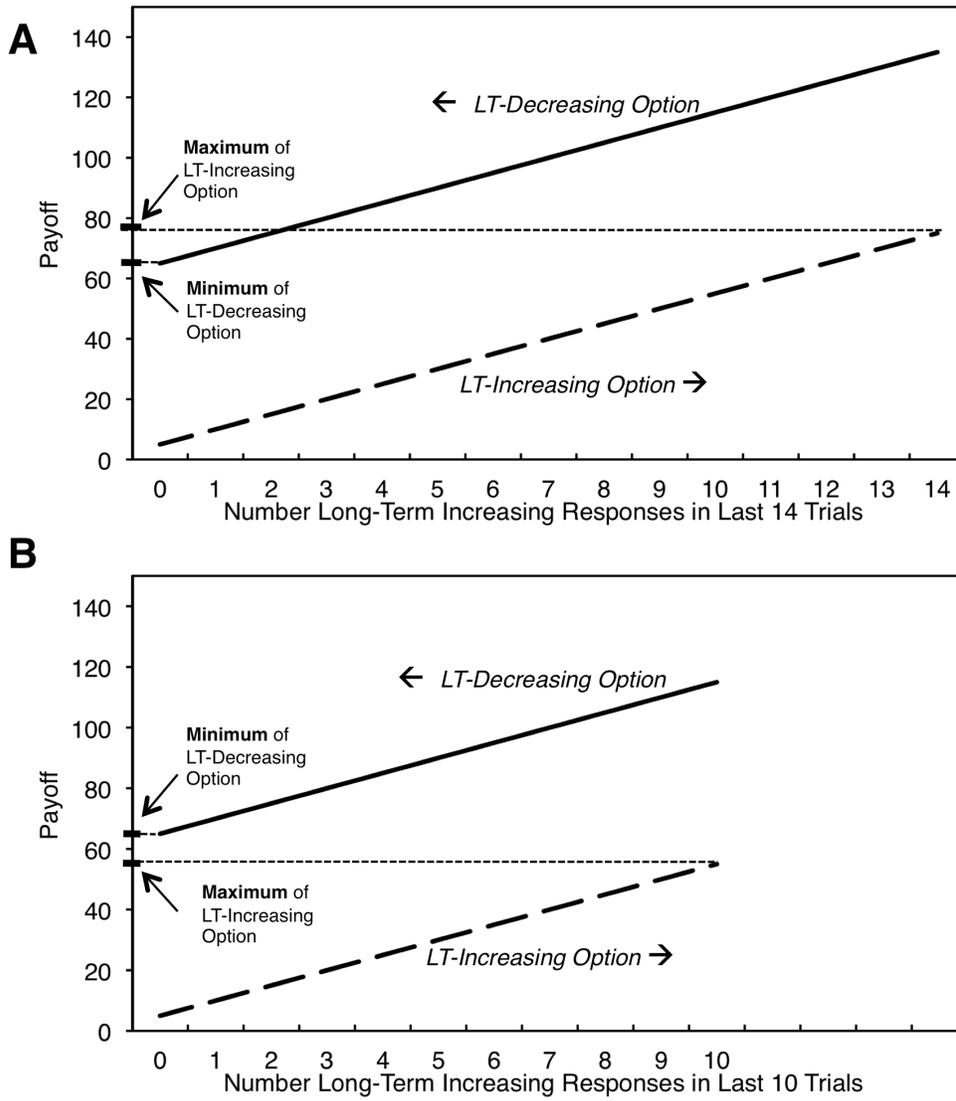


Figure 1.

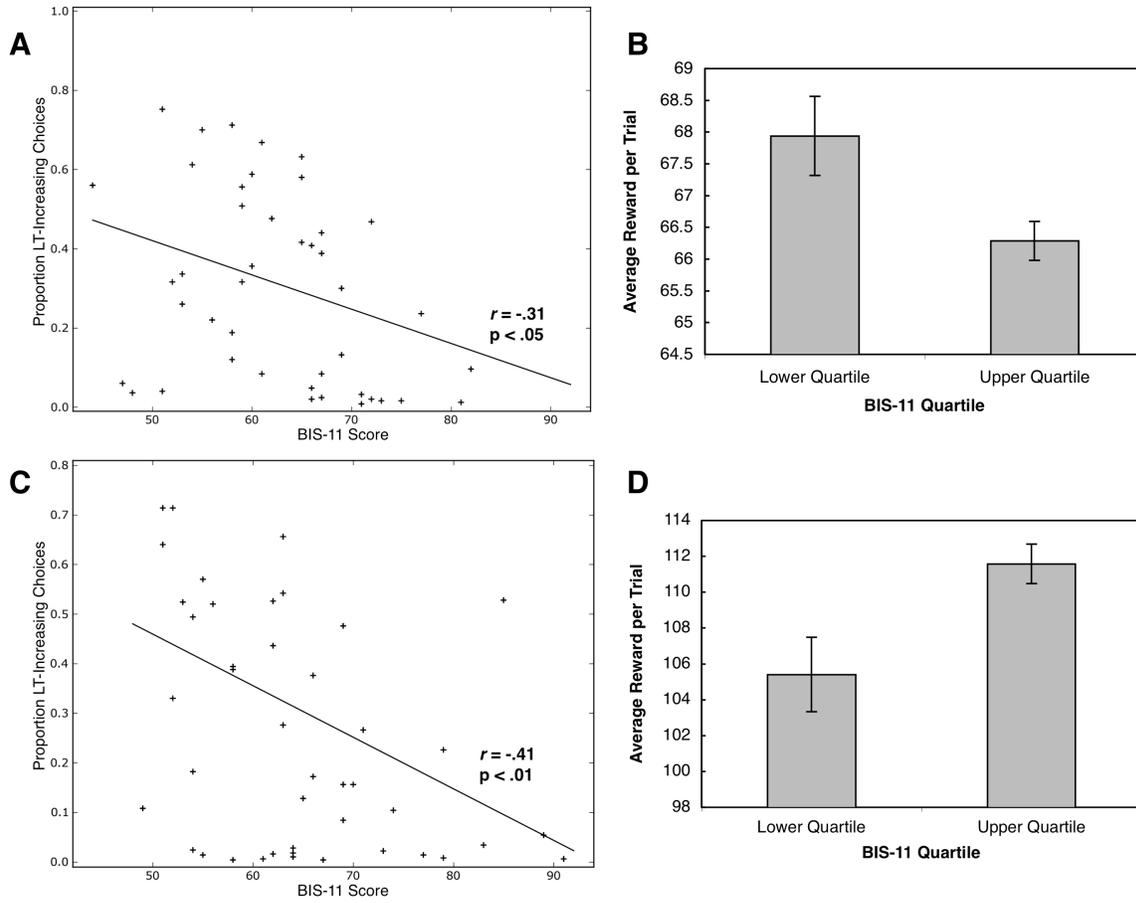


Figure 2.