I want it now so you cannot have it later: The role of impulsive choices in competitive environments

by

Brian Howatt

B.S., DePaul University, 2014
M.S., Arizona State University, 2017

AN ABSTRACT OF A DISSERTATION

submitted in partial fulfillment of the requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Psychological Sciences
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2023

Abstract

People commonly prefer smaller-sooner rewards over larger-later rewards, referred to as impulsive choice. Although impulsive choices are strongly associated with negative behaviors, such as poor diet and exercise habits, substance abuse and gambling, larger delayed rewards are more uncertain than smaller immediate rewards regarding if and when they are delivered. The threat of losing out on delayed rewards may then motivate individuals to shift their preferences towards more immediately available goals. One source of uncertainty for delayed rewards is competition, where multiple individuals exclusively pursue limited resources. The impact of competitors may result in individuals selecting more impulsive choices to ensure a relatively greater acquisition of important resources compared to competitors, despite the environment originally incentivizing waiting. Therefore, the purpose of the current dissertation is to study the degree to which competition-based uncertainty increases impulsive choices above and beyond environmental uncertainty. To test this hypothesis, I conducted two experience-based decision-making experiments that manipulated various dimensions of competition that could potentially control participants' waiting behaviors and performance, such as the visibility of a competitor and the visibility of the competitor's cumulative rewards. Overall, the results provided weak evidence that competition influenced participants' impulsive choices above environmental uncertainty, specifically by slightly increasing their likelihood of obtaining rewards rather than their tendency to select impulsive choices. These findings are more consistent with an economic perspective of choice behavior than an evolutionary perspective, and I discuss the theoretical and methodological implications of this novel line of research.

I want it now so you cannot have it later: The role of impulsive choices in competitive environments

by

Brian Howatt

B.S., DePaul University, 2014
M.S., Arizona State University, 2017

A DISSERTATION

submitted in partial fulfillment of the requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Psychological Sciences
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2023

Approved by:

Major Professor
Michael E. Young

Copyright

© Brian Howatt 2023.

Abstract

People commonly prefer smaller-sooner rewards over larger-later rewards, referred to as impulsive choice. Although impulsive choices are strongly associated with negative behaviors, such as poor diet and exercise habits, substance abuse and gambling, larger delayed rewards are more uncertain than smaller immediate rewards regarding if and when they are delivered. The threat of losing out on delayed rewards may then motivate individuals to shift their preferences towards more immediately available goals. One source of uncertainty for delayed rewards is competition, where multiple individuals exclusively pursue limited resources. The impact of competitors may result in individuals selecting more impulsive choices to ensure a relatively greater acquisition of important resources compared to competitors, despite the environment originally incentivizing waiting. Therefore, the purpose of the current dissertation is to study the degree to which competition-based uncertainty increases impulsive choices above and beyond environmental uncertainty. To test this hypothesis, I conducted two experience-based decision-making experiments that manipulated various dimensions of competition that could potentially control participants' waiting behaviors and performance, such as the visibility of a competitor and the visibility of the competitor's cumulative rewards. Overall, the results provided weak evidence that competition influenced participants' impulsive choices above environmental uncertainty, specifically by slightly increasing their likelihood of obtaining rewards rather than their tendency to select impulsive choices. These findings are more consistent with an economic perspective of choice behavior than an evolutionary perspective, and I discuss the theoretical and methodological implications of this novel line of research.

Table of Contents

List of Figures

List of Tables

Acknowledgements

Unsurprisingly, there are too many people to thank for everything I've accomplished during my graduate studies, and particularly this dissertation work. None of this would have been possible without their assistance and contributions, and I apologize for anyone I was not able to include here.

First, I would like to thank my dissertation committee, Drs. Mike Young, Kim Kirkpatrick, Gary Brase, and Sam McClure.

Mike, I knew even from when we first met at the perspective student's visit that I wanted to come to K-State (even if I took my time actually making that final decision) and have never once doubted whether I made the correct choice. I feel so incredibly fortunate to have worked with and learned from you these last five years, and it's reassuring that I'll be able to continue leaning on you afterwards. I know it hasn't always been easy deciphering my thoughts and editing my writing, but I feel like I'm now marginally more comprehensible, and am extremely grateful for everything we've been through.

Kim, thank you for generously allowing me to participate in your lab to learn even further about impulsive choice from a different perspective and growing important professional developments like science communication and grant writing. Gary, your insight on and dogged pursuit of strong theory certainly left its mark on how I approach psychological research, even if behaviorism is still the root of everything. Sam, taking your classes at ASU shaped my research interests so strongly that I can largely attribute them to the development of my research programs on reinforcement learning and competition, and how they factor into intertemporal choices.

Kari, thank you for your support, discussions and listening to me complain endlessly about this project over the last two years. You, along with Mia and Mo, helped keep me

relatively sane and grounded. Mia and Mo also provided important case studies that cats can be patient to a fault when it comes to waiting for their own treats and impulsive to a fault when it comes to stealing each other's treats.

I'd also like to thank my family for their support and interest in my work. Even if you didn't always understand what I was doing or why I was doing it, hopefully my path onward from K-State shows why this was all worth it. I may not have always wanted to talk about my research over calls or when visiting back home, but doing so forced me to make sure I could meaningfully describe it to people outside of the academia bubble.

Finally, I'd like to thank other members of the lab, Patrick Hancock, Aidan Swall and Jordan Toothaker, for discussing ideas and collecting data that I wouldn't have been able to do on my own. And all other members of the Psychology department who provided important professional, personal, and financial support and feedback during these last five years.

Chapter 1 - Introduction

People face many conflicting intertemporal choices – decisions between smaller short-term and larger long-term rewards – throughout our lives. Some of these choices may be routine – driving aggressively to reach a destination more quickly, watching TV instead of going to the gym to exercise and improve our health, or immediately consuming food versus storing it for later. Other choices are less routine – high school graduates considering whether to immediately start working in a trade or waiting to obtain a college degree, or an alcoholic choosing between drinking at the bar after work or abstaining to save money for rent the next month. While acting on short-term rewards is immediately gratifying, repeatedly doing so may result in lower long-term welfare or even delayed negative consequences (e.g., health issues, eviction, or poor grades). However, unnecessarily waiting for long-term rewards can result in forgoing more profitable, immediate outcomes (e.g., making multiple short-term investments or working in a trade to immediately earn money and avoid student debt) or even losing them completely (losing a highly desirable job offer to another applicant or an animal stealing stored food from a groupmate). Therefore, knowing when to act on short-term rewards and when to wait for long-term rewards is crucial for maximizing our well-being.

Decades of research across the fields of psychology, neuroscience and economics has concluded that organisms tend to prefer acting on smaller sooner rewards over waiting for larger later rewards, referred to as *impulsive choice*[1]. Moreover, although impulsive choices themselves

---

[1] Impulsive choices are often operationalized as requiring 'internal tension' when attempting to delay gratification or 'preference reversals' that leads to succumbing to the temptations of immediately available rewards despite originally preferring patience (Berns, Laibson & Loewenstein, 2007), but for the purpose of the current dissertation I am operationalizing impulsive choices simply as the preference for smaller sooner (or more immediately available) rewards.

are not necessarily irrational (as I am operationalizing them here), they are commonly associated with many individually maladaptive behaviors such as substance abuse, compulsive gambling and risky sexual choices (Gebru et al., 2022; Reynolds, 2006), as well as related to global issues like overfishing and climate change, and thus represent an important behavior to research with practical implications for improving people's lives. Theories of intertemporal choice have proposed numerous explanations for preferences for smaller sooner rewards, ranging from deficits in reward valuation, anticipation, timing and impulse control to evolutionary pressures and the uncertainty of delayed rewards (Ainslie, 1975; Berns, et al., 2007; Marshall, Smith & Kirkpatrick, 2014; McGuire & Kable, 2013; Odum et al., 2020; Stevens et al., 2004; Sozou, 1998; Takahashi, 2009). The present dissertation will broadly focus on testing the effects of one particular form of environmental uncertainty on intertemporal choices: competition.

There are many causes of environmental uncertainty, but one universal case of particular interest is competition, where multiple decision makers exclusively pursue limited and contested resources (Fehr & Schmidt, 1999; van den Bos, Talwar & McClure, 2013) and outcomes are mutually contingent on other organisms' choices (Martin, 2016; Schelling, 1960). Humans and other animals must be able to outperform competitors for critical resources, such as food, water and mates or risk an untimely death and not reproducing. But, doing so requires selecting actions that are the best response to the actions of other decision makers as well as the environmental contingencies. The impact of competitors may result in organisms selecting more impulsive choices to ensure a relatively greater acquisition of these resources compared to competitors (i.e., relative fitness; Darwin, 1859) despite the environment originally incentivizing waiting. Despite the vast literature on competitive decision-making (e.g., Camerer & Ho, 1999; Goldstone, Ashpole & Roberts, 2005; Lee & Seo, 2016; Mishra, Barclay & Sparks, 2016; Murayama &

Elliot, 2012; van den Bos, Golka et al., 2013; von Neumann & Morgenstern, 1947), little research has empirically examined its direct impact on intertemporal choices (e.g., Amita, Kawamori & Matsushima, 2010; Stevens et al., 2011) other than brief theoretical considerations regarding foraging and mate selection (Fawcett, McNamara & Houston, 2012; Hayden, 2016; 2019; Stevens & Stephens, 2008; 2010; Takahashi, 2006).

Therefore, the specific purpose of this dissertation is to experimentally evaluate the degree to which competition increases preferences for more immediately available rewards in a dynamic decision-making task. The paper begins with a selective review of intertemporal choice research and discusses the roles of uncertainty and competition as state variables affecting the relative advantages of acting on smaller sooner rewards. The introduction closes with a theoretical integration of competition and intertemporal choice based on some general principles of evolutionary psychology (e.g., natural selection and relative fitness) that leads into the proposal's specific aims and competing hypotheses.

Intertemporal Choice

Historical Background

The study of intertemporal choice was originally conceived by the economist John Rae (1834) to understand wealth accumulation and its disparities (Frederick et al., 2002). Rae theorized four psychological motives that influenced the enhancement or suppression of wealth accumulation over time, three of which were 1) the appeal of immediate consumption and the discomfort of delayed gratification, 2) the ability to exert self-control while delaying gratification, and 3) the uncertainty of human life[2]. Although Rae primarily theorized about

---

[2] The fourth factor "bequest motives", or the desire to transfer intergenerational wealth, is not considered in the current paper.

wealth accumulation regarding nations, his proposed factors are also clearly relevant to individual decision-making, where every day choices encompass tradeoffs between immediate desires and long-term goals. And the field of psychology (and economics) has a long history of experimentally investigating the roles of Rae's factors in intertemporal choices, and more specifically preferences for smaller, more immediate outcomes.

Intertemporal choice research and the selection of impulsive choices can be divided into two broad categories depending on the nature of the reinforcement delay, delay discounting and delayed gratification (Paglieri, 2013; Reynolds & Schiffbauer, 2005). Delay discounting refers to the rate at which rewards decline in value as a function of their delay (Madden & Bickel, 2010; Madden & Johnson, 2010; Odum et al., 2020), where steeper discounting reflects stronger preferences for smaller sooner rewards relative to larger delayed rewards. In discounting-type tasks, subjects make a series of choices between smaller sooner rewards and larger delayed rewards and reinforcement is presented following an irreversible choice-delay. For example, a subject choosing between a smaller food reward presented two seconds after the response and a larger food reward presented 10 seconds after the response; once a choice is made, the subject cannot alter it and must wait the programmed delay for the reinforcement.

In contrast, delayed gratification involves subjects continuously waiting for the presentation of a larger reward by default, but with the opportunity to defect at any point and choose the smaller amount immediately (McGuire & Kable, 2012; Mischel & Metzner, 1962; Young, Webb & Jacobs, 2011). The inability to wait for the larger reward (or defecting more quickly to the smaller reward) corresponds to being poorer at delaying gratification and more impulsive. For example, consider a person driving home choosing between waiting to eat dinner once they arrive home or stopping sooner to eat fast food. The person may choose to wait until arriving home to eat but is not beholden to this decision because they can change their mind at any point while driving.

Delay Discounting. Delay discounting initially began from behavioral learning researchers using concurrent response procedures to study the degree to which reinforcement delays decreased relative response rates (Chung, 1965; Chung & Herrnstein, 1967) and increased preferences for smaller sooner rewards (Ainslie, 1974; Ainslie & Herrnstein, 1981; Davenport, 1962; Grosch & Neuringer, 1981; Logan, 1965; McDiarmid & Rilling, 1965; Richards, 1974; Williams & Fantino, 1978). These early results established the psychological phenomenon of devaluing delayed rewards and subsequent work extended this notion by examining the multiple control of different variables on subjects' preferences for smaller sooner rewards, such as the relative and absolute values of reinforcement delays and amounts (Gentry & Marr, 1980; Green & Snyderman, 1980; Logue & Chavarro, 1987; Snyderman, 1983) and choice commitment (Rachlin & Green, 1972; Green & Rachlin, 1996). This research also investigated the descriptive and predictive validity of various choice models for waiting behaviors (e.g., the matching law; Baum, 1974; Herrnstein, 1974; Logue, Forzano & Torbin, 1992) under different contingencies. Despite concurrent response procedures' initial success in analyzing intertemporal choices, there were limitations related to conducting human research (Logue et al., 1986) and precisely measuring degrees of choice preference or indifference points (i.e., the point at which subjects reverse their preference for smaller sooner or larger later rewards; Mazur, 1987, pg. 55-56). Later behavioral learning work, in response to some of these limitations, developed *titration* procedures (Mazur. 1987; Reynolds & Schiffbauer, 2004; Richards et al., 1997) and *monetary choice questionnaires* (MCQs; Kirby, Petry & Bickel, 1999; Kirby & Marakovic, 1996; Rachlin et al., 1991) that spawned the study of delay discounting that continues to dominate intertemporal choice research and the classification of impulsivity today. These tasks improved the identifiability of subjects' indifference points (see Mazur, 1987 and Rachlin et al., 1991) for various reinforcement contingencies (e.g., reward amounts, delays and probabilities), which could then be modeled with discount functions that quantitatively describe subjects' psychological sensitivity to delayed outcomes (i.e., the rate at which rewards lose value to them over time; Green & Myerson, 1995; Mazur, 1987; Odum, 2011a; Young, 2017).

Delay discounting research in psychology, neuroscience and economics has consistently revealed that humans and other animals discount future outcomes more strongly than prescribed by rational choice theory[3] (Samuelson, 1937; e.g., Green & Myerson, 1995; Loewenstein & Prelec, 1992; Madden & Bickel, 2010; Mazur, 1987; McClure et al., 2004), motivating them to engage in immediately gratifying behaviors that can conflict with long-term goals. Moreover, delay discounting has stimulated a plethora of clinical research associating steeper discounting (i.e., more impulsive choices) with compulsive gambling, substance use, poor diet and other dysfunctional behaviors (Businelle et al., 2010; Ciccarelli et al., 2019; Odum, 2011b; Steele et al., 2021) and neuroscientific research dissociating the neural mechanisms of intertemporal choices (Frost & McNaughton, 2017; Kable & Glimcher, 2007; Hwang, Kim & Lee 2009; McClure et al., 2004; Tanaka et al., 2020; van den Bos & McClure, 2013).

Delayed Gratification. Research on delayed gratification was popularized by Walter Mischel and his famous *marshmallow test* (Mischel & Ebbesen, 1970; Mischel & Metzner, 1962), where subjects (children or adolescents) were given a single treat (e.g., a marshmallow)

---

[3] Note that Samuelson (1937; cited from Frederick, 2002) did not actually endorse his simplified model as a normative theory (p. 161), nor did he claim it held sufficient descriptive validity (p. 159). Instead, subsequent researchers co-opted it as a framework for optimal decision-making, which persists today as a benchmark for observed behaviors and testing competing theories (Loewenstein & Prelec, 1992).

and told the researcher would return with a second one after a short time (e.g., 15 minutes). If the child could resist eating the first treat by the time the researcher returned, then they would also be given the second one. Otherwise, the child was not given the second treat. In delayed gratification tasks, impulsivity is the inability to wait for the second treat and measured by noting the time at defection.

Across the decades of research using Michel's task (or its variants), researchers have observed that some subjects are unwilling or unable to delay gratification and defect fairly quickly to the smaller immediately available reward (Tobin & Graziano, 2010), similar to delay discounting results. Moreover, longitudinal studies have shown that subjects poorer at delaying gratification were more likely to have behavioral problems later in life, such as attentional deficits, lower academic performance and life satisfaction and substance use (Kluwe-Schiavon et al., 2020; Mischel & Ebbesen, 1970; Mischel, Shoda & Rodriguez, 1989). However, Mischel's delayed gratification tasks have almost exclusively been presented to children or adolescents and only recently has a similar task been developed for adult subjects with multiple decisions over longer choice-delay intervals and different types of reinforcers[4] (Forstmeier et al., 2011; Knolle-Veentjer et al., 2008). Interestingly, these studies have shown that adults modestly-to-strongly preferred larger delayed rewards on average, even when the delivery dates were a month or two away (Forstmeier, Drobetz & Maercker, 2011; Göllner et al., 2018; Knolle-Veentjer et al., 2008).

Overall, there seems to be compelling evidence from the historical intertemporal choice paradigms that organisms are biased towards making impulsive choices, struggle with self-restraint while waiting for delayed rewards and suboptimally discount delayed rewards.

---

[4] Closer inspection of this task reveals that the presented choices more closely resemble those of a discounting paradigm, albeit in a simplified manner. But, subjects were still preferring the delayed rewards on average.

Moreover, impulsive choices are strongly associated with various maladaptive behaviors, both cross-sectionally and longitudinally, often casting preferences for smaller sooner rewards in a negative light with recommendations that they must be avoided at all costs. However, the validity of these findings and interpretations has not gone unchallenged (Green & Myerson, 2019). For example, some studies have revealed cases where humans and nonhuman animals discount rationally in concurrent response tasks (Logue et al., 1986), MCQs (Hwang, Kim & Lee 2009; Kable & Glimcher, 2007) and foraging tasks with real-time constraints (Schweighofer et al., 2006). Additionally, human participants are more appropriately patient in dynamic delayed gratification tasks (McGuire & Kable, 2012, 2015; Young et al, 2011; Young et al., 2013), and some rat subjects are even overly patient in an intertemporal foraging task (Wikenheiser, Stephens & Redish, 2013). These findings suggest a reframing of impulsive choices regarding the ecological benefits of short-sighted behavior.

Reinterpreting and Recasting Impulsive Choices.

Reinterpreting the presumed imperfections of subjects' intertemporal choices is primarily based on two factors: 1) the validity of the experimental methods under which impulsive choices are observed, and 2) the degree to which impulsive choices are suboptimal given the environmental conditions. In sum, correctly cataloging organisms as irrationally shortsighted requires observing behaviors resulting in deviations from reinforcement maximization, particularly in ecologically relevant situations.

The first factor considers the ecological validity of experiments in which impulsive choices are traditionally observed. Interestingly, impulsive choices may not necessarily emerge from subjects' own decision biases, but rather methodological limitations in behavioral tasks. For example, the task designs may not appropriately reflect the naturalistic situations in which organisms have adapted to learn (Hayden, 2016; 2019; Kacelnik, 2003; Schweighofer et al., 2006; Stevens, Hallinan & Hauser, 2005; Stevens, Rosati et al., 2005) or rely too heavily on descriptive contingencies and hypothetical decisions in some titration tasks and MCQs (Dai et al., 2019b; Paglieri, 2013; Wulff, Mergenthaler-Conseco & Hertwig, 2018). Consequently, subjects may not fully understand the task structure of concurrent response procedures, or misjudge the hedonic value of hypothetical outcomes until experiencing them (Steele et al., 2019) or learning visual cues signaling reward contingencies (Hwang et al., 2009). Therefore, seemingly maladaptive impulsive choices could result from misapplications of evolved decision mechanisms to biased experimental contingencies that typically only incentivize waiting (Hayden, 2016; Stevens et al., 2005).

Moreover, a recent review (Bailey, Romeu & Finn, 2021) has roundly criticized the field of Psychology for its hyper-emphasis on the delay discounting framework, calling it an artificial psychological construct lacking clinical validity and generalizability to our everyday choices. In essence, most people do not make such simple choices between static monetary amounts over different time-intervals (e.g., $1000 today or $2000 in one month), or choose whether to consume different amounts of a drug (e.g., one cigarette now or two cigarettes in 10 minutes). This same criticism could also apply to delayed gratification tasks given that they often do not use multiple trials (Mischel’s marshmallow tasks) or enable subjects to experience the delays and change their preferences based on experience (Forstmeier et al, 2011). Additionally, there exist a wide variety of delayed gratification scenarios and simply choosing between one snack now and a second one after some delay is ecologically limited, further suggesting that the conventional experimental methodologies are insufficient to properly study complex manifestations of intertemporal choices.

The second factor considers whether there are environmental conditions (i.e., state variables) under which impulsive choices may actually maximize rewards. The prevailing assumption across most psychological research is that patience is ideal and impulsivity may be a negative trait associated with maladaptive behaviors (Odum, 2011b). However, certain state variables can justifiably incentivize impulsive choices even when patience is theoretically optimal. For example, people may take out high interest loans or cash advances that could financially cripple them in the long run because they do not have any other options available to put food on the table that week. Or, hunger can motivate preferences for smaller sooner food rewards to sate cravings (Skrynka & Vincent, 2019) for which waiting too long will result in severe negative consequences like starvation.

State variables in experiments can also be used to construct more realistic choice scenarios in which either acting now or waiting differentially maximizes rewards. Researchers can then test whether subjects learn optimal waiting times given the appropriate (and discriminable) environmental conditions. And experience-based decision-making (EBDM) research has indeed demonstrated that experimental tasks with ecologically relevant contingencies engender more prudent intertemporal choices (Carter, Pedersen & McCullough, 2015; McGuire & Kable, 2012, 2013, 2015; Otto, Markman & Love, 2012; Schweighofer et al., 2006; Stevens & Stephens, 2010; Stevens et al., 2005; Stevens, Rosati et al., 2005; Young et al., 2011; Young & Howatt, 2021) For example, subjects learn to act quickly on smaller rewards or wait longer for larger rewards when the reward's growth rate (Young et al., 2011) or expected arrival time (McGuire & Kable, 2012) incentivizes doing so. Thus, while patience is a virtue, impulsive choices may also be, in part, a crucial behavior that helps organisms satisfy immediate physiological needs necessary for longer-term survival and reproductive success in dynamic and harsh environments (Green & Myerson, 2019; Fawcett et al., 2012; Fenneman & Frankenhuis, 2020; Fenneman, Frankenhuis & Todd, 2022; Sánchez-Amaro et al., 2021; Stephens, Kerr & Fernández-Juricic, 2004; Stevens et al., 2005; Stevens, Rosati et al., 2005). And EBDM experiments better capture the complex interactions and reward-maximization strategies between environmental variables and impulsive (or waiting) behaviors compared to conventional tasks.

One environmental variable that individuals may be particularly sensitive to when deciding whether to act or wait is the uncertainty in the delivery or arrival time of larger delayed rewards (Rae, 1834; Green & Myerson, 2004; Rachlin et al., 1991). There are many instances of uncertainty decision makers face – e.g., competitors, weather, illness or even general environmental uncertainty without any known cause – and individuals may rationally prefer acting on smaller rewards to secure at least some nominal benefit given that many larger later rewards are not always guaranteed.

Intertemporal Choice and Uncertainty

Larger delayed rewards are frequently more uncertain than smaller immediate rewards regarding if and when they are delivered (Green & Myerson, 2004; Rachlin et al., 1991). Smaller, short-term investments are more likely to be paid back in full and sooner than larger long-term investments, larger prey are rarer and more difficult to catch while also potentially being targeted by competing organisms, and a job applicant does not know if or when a better job offer will arrive (nor are they aware of competing candidates’ likelihood being offered the job instead). Uncertainty may thus motivate preferences for immediacy because waiting too long for larger rewards can be unproductive and even detrimental to our well-being (Green & Myerson, 2004; McGuire & Kable, 2012).

Although the current literature directly integrating uncertainty and intertemporal choice is limited (Dai et al., 2019a; Prelec & Loewenstein, 1991), prior research shows that greater uncertainty in the delivery of larger later rewards does increase impulsive choices in concurrent response and delayed gratification procedures (Glass, 1969; Kidd, Palmeri & Aslin, 2013; Mahrer, 1956; Moffett, Flanagan & Shah, 2020; Michaelson et al., 2013; Mischel & Staub, 1965; Mischel & Grusec, 1967; Navarick, 1987; Stevens et al., 2011). Moreover, in MCQ and titration-based delay discounting tasks, greater uncertainty in the delivery of larger delayed rewards results in even stronger devaluation, and certain rewards are discounted more heavily over time than uncertain rewards (Cox & Dallery, 2016; Vanderveldt, Green & Myerson, 2015; Stevenson, 1992).

This class of studies is again methodologically limited, echoing many of the same concerns and criticisms discussed in the previous section, such as the task structure not reflecting realistic contingencies and delayed rewards' arrival times and probabilities being fixed and descriptive (e.g., choosing between $100 for certain today or $200 with 80% probability next month). In real-world situations, decision makers do not always know the exact arrival time or probability of receiving larger later rewards, nor are these values routinely held constant. Instead, uncertainty commonly takes the form of a probability distribution rather than a single value. For example, an apple grows more nutritious as it ripens but eventually it decays and becomes inedible. Or, the longer the apple remains in a tree the more likely a competitor is to consume it first. But, the rate at which these processes occur is variable, open-ended and unknown, and organisms do not know for certain when an apple will reach maximal ripeness, how long it will remain fresh, or when a competitor is likely to pass by and take it for themselves. Decision makers therefore must be able to learn the value of acting or waiting based on incomplete information and dynamically changing uncertainty, although the conventional research paradigms tend to gloss over this notion for experimental and analytical simplicity.

Alternative perspectives on intertemporal choices have thus studied situations in which the chance of obtaining larger rewards continuously changes over time according to probability distributions (e.g., McGuire & Kable, 2012; Young et al., 2011). Some authors have even theorized that decision-makers learn expectations about larger rewards that approximate their probability distributions rather than just point estimates (e.g., hazard rates, Dasgupta & Maskin, 2005; McGuire & Kable, 2012; Sozou, 1998; see Sherman & Thomas, 1968, for preferences for variable-schedules of reinforcement over fixed-schedules). Theoretical frameworks and empirical research involving the probability distributions of larger rewards provide several advantages compared to those using only static and descriptive contingencies. One advantage is that researchers can calculate normative (optimal) waiting times based on the simulated environment's reward structure and test the degree to which subjects' behaviors align with them. Researchers can then identify which environmental (state) variables most strongly affect subjects' waiting decisions and potentially use this information to improve future choices.

An additional advantage of using probability distributions is that they can be easily embedded, and then manipulated, in EBDM tasks[5] to simulate a wide variety of realistic ecologies (McGuire & Kable, 2012; Young et al., 2014) where acting or waiting differentially maximize reward rates. For example, some animals live in faster-paced and more uncertain

---

[5] Note that a delayed reward's *hazard rate* (i.e., the rate at which the delayed reward becomes successively more or less likely to arrive over time) can also be incorporated into delay discounting paradigms (Dasgupta & Maskin, 2005; Green & Myerson, 1996; Kagel, Green & Caraco, 1986; Sozou, 1998). Although these analyses are theoretically relevant because they can potentially explain hyperbolic discounting, such tasks still tend to retain the static/descriptive contingencies and preclude learning through immediate experience (cf. Bailey et al., 2021).

environments where the availability of delayed rewards rapidly diminishes (high risk due to greater competition for resources like food or mates). Here, impulsive choices may be most beneficial for survival because waiting too long can result in zero consumption. Conversely, other animals live in slower-paced and more certain environments where delayed rewards remain readily available (low risk, such as marmoset monkeys feasting on slow-growing tree sap; Stevens et al., 2005a; 2005b). Here, patience is perfectly suitable and may even be necessary to avoid negative consequences like exhausting resource pools (Young & Howatt, submitted).

In reviewing the literature there emerged two classes of EBDM experiments that study how the uncertainty of larger rewards change over time, by manipulating either the larger reward's arrival time (Fung et al., 2017; Lempert et al., 2018; McGuire & Kable, 2012; 2015) or its probability of being delivered (Webb & Young, 2015; Young, Webb & Jacobs, 2011; Young et al., 2014). Interestingly, both tasks show that individuals can adjust their waiting times to match each contingencies' reward-maximizing (optimal) strategies, but to different degrees.

The first class of tasks are Temporal Persistence tasks (TP; e.g., McGuire & Kable, 2012), where participants are repeatedly faced with a decision in which the default behavior is to wait for a larger reward but with the opportunity to change their mind at any point and choose the smaller reward immediately. If participants continue to wait, at some point in the future the larger reward is automatically delivered, and a new trial begins. Note here that the larger reward is always guaranteed to arrive if participants wait long enough, but its delay to receipt on any given trial is unknown. TP tasks are similar to Mischel's original delayed gratification task, but the arrival times (i.e., delay durations) of the larger rewards are now randomly sampled from different probability distributions that differentially incentivize waiting the full time-interval or quickly defecting (after about 2 s) to the immediately available smaller reward if the larger reward has not yet arrived.

The second class of tasks are Escalating Interest tasks (EI; e.g., Webb & Young, 2015; Young et al., 2011; Young et al., 2014), where participants encounter rewards that are continuously changing in their magnitude and/or delivery probability over a 10 s interval. For example, in one preparation (Experiment 2, Young et al., 2014), the reward magnitudes increased over time while their probabilities decreased. Participants were thus faced with the decision of whether they preferred smaller sooner rewards with higher probabilities or waiting for larger rewards with lower probabilities. Across these various tasks, Young and colleagues then manipulated the way in which the reward dynamics change over time to vary the basic conditions under which responding rapidly or slowly produce a greater relative reward rate. Note here that in contrast to TP tasks, the reward's arrival time always immediately follows a response, but the reward's delivery is no longer guaranteed.

Interestingly, participants' optimality in waiting times differed between these two classes of tasks. In TP tasks, subjects correctly calibrated their waiting times to match the optimal times predicted by the probability distributions, suggesting that they formed reasonable expectations about the larger reward's possible delay durations. And subsequent work extended these findings by testing how state variables such as caloric intake and stress modulated wait times (Fung et al., 2017; Lempert et al., 2018). However, in EI tasks, subjects were overly sensitive to the reward's probability, suggesting that changes in a delayed reward's uncertainty can "overshadow" its amount (Webb & Young, 2015). For example, when the reward's magnitude increased over time and its probability decreased over time, subjects tended to actually act faster than optimal when rapid responding was incentivized and wait longer than optimal when patience was incentivized.

Thus, not only does reward uncertainty play a major role in organisms' waiting times, but different kinds of uncertainty can differentially influence the degree to which organisms adjust their behaviors.

Sources of Uncertainty for Delayed Rewards

Although EBDM tasks have shed considerable light on the behavioral control of continuously changing reward uncertainty on individuals' waiting decisions, it is unknown whether knowledge about the cause of uncertainty further affects these behaviors. The uncertainty in delayed rewards' arrival times for TP tasks was not explicitly attributed to any particular cause, even though open-ended delays are frequently encountered in daily life (e.g., waiting for it to rain or to be seated at a busy restaurant). The uncertainty in obtaining rewards in EI tasks was attributed to the game's weapon operating correctly, but whether this knowledge impacted subjects' wait times was not explored. Given the specific research questions posed by the authors (and subsequent authors) in these two series of studies, it is unsurprising that the influence of this knowledge was not examined. As such, the impact of knowledge-based uncertainty on waiting decisions has yet to be formally investigated.

Why might the source of a delayed reward's uncertainty matter? One plausible explanation is that this knowledge may reasonably inform decision makers about the environment's underlying probability distributions (Griffiths & Tenenbaum, 2006), which are typically unknown and must be estimated from experience. This knowledge can then help decision-makers choose between acting or waiting based on learned expectations. Consider an example of a driver deciding between waiting in a traffic jam or taking an alternate route: are they more likely to act and take side-streets or wait for traffic to clear up if they know whether road construction or a collision caused the delay? Maybe they know from experience that construction takes less time to pass through on average than switching to side streets, so they wait. But, a collision takes a longer time to pass through on average, so they take the side streets if they see an accident up ahead. In either case, the driver's decision depends on knowing the source of uncertainty and relying on a reasonable expectation of the possible delay durations to make a good decision.

There are various sources of uncertainty affecting the delivery of delayed rewards, but one important and understudied cause is *competition*, where multiple individuals exclusively pursue limited and contested resources (van den Bos, Talwar & McClure, 2013). Because obtaining a sufficient share of limited resources in competitive environments can be quite difficult, a second plausible explanation for why knowledge about the source of uncertainty matters is that preferences for smaller sooner rewards could be further motivated by the advantages of outperforming competitors (Stevens & Stephens, 2010). Here, not only can a decision-maker estimate how long a larger reward may remain available given the presence of a competitor, but knowledge about the competitor may also motivate them to acquire the reward even sooner (and at a smaller amount) than expected because otherwise they would also be worse off than the competitor if they do not claim the reward.

Intertemporal Choice and Competition

Research investigating the relationship between intertemporal choice and competition has primarily revolved around Ainslie's (1992) and Rachlin's (2000) proposals that preferences for smaller sooner rewards in discounting tasks are similar to competitive choices in iterated social dilemmas (e.g., repeatedly defecting in the Prisoner's Dilemma; Rapoport & Chammah, 1965). Social dilemmas arise when decision makers must choose between self-interested payoffs (competitive actions) and joint payoffs (cooperative actions), and rationally maximizing self-interested payoffs actually results in lower joint payoffs. Thus, repeatedly competing (i.e., behaving impulsively) in social dilemmas initially produces a large immediate reward but leads to lower cumulative rewards, while sustained cooperation (i.e., delaying gratification) results in larger cumulative rewards.

Empirical studies have largely supported Ainslie's and Rachlin's proposals, finding that steeper delay discounting rates (i.e., more impulsive choices) are associated with more competitive choices in various social dilemmas, such as Prisoner's Dilemma games (Green, Price & Hamburger, 1995; Harris & Madden, 2002; Stephens, McLinn & Stevens., 2002; Yi, Johnson & Bickel., 2005; Yi et al., 2007), Public Goods Games (Curry, Price & Price, 2008; Espin et al., 2012; Ledyard, 1995; but see Jones & Rachlin, 2009 and Rosati, DiNicola & Buckholtz, 2018) and Ultimatum Games (Crockett et al., 2010; Espin et al., 2015; Karagonlar & Kuhlman, 2013; Guth, Schmittberger & Schwarze, 1982). Giddie (2011) also showed that orienting subjects to short-term outcomes decreased inter-group cooperation. However, delay discounting-type questionnaires and dilemma-based economic games are behaviorally independent tasks, and the supporting evidence linking competitive and impulsive choices is largely correlational. In contrast, real-world decision making often involves acting or waiting in competitive situations where larger rewards are at risk of being pilfered by other organisms as time passes.

Social dilemmas also confound whether subjects' competitiveness is due to maximizing their absolute payoffs (what is best for them individually) or their relative payoffs (how can they be better off than others). In many social dilemmas, the economically rational choices maximize both absolute and relative payoffs, but only preferences for relative payoffs constitute "competitive" behavior. For example, competing (defecting) in the prisoner's dilemma is considered "rational" behavior, but different motivations may underlie this preference (see

Messick & McClintock, 1968; McClintock & McNeel, 1966). Subjects may choose to defect because they can potentially receive the highest possible number of points on that trial (maximizing absolute gains), receive more points than the other player (maximizing relative gains), or avoid receiving fewer points than the other player (minimizing relative losses). The previous research thus fails to adequately determine whether “competitive” choices actually constitute impulsive choices, particularly when said choices correctly follow rational choice theory’s prescribed strategies. Experimental analyses of competitive behaviors must be able to dissociate these preferences and how they relate to competitive encounters in a more dynamic and integrated manner.

In one set of experiments designed to more directly investigate the effects of competition on intertemporal choices, Stevens and colleagues (2011) compared how long bonobo apes would wait in a delayed food accumulation task when competitors were identified as a source of reward uncertainty in a training task. In Stevens et al’s preparation, subjects were first trained to receive food from a researcher at different probabilities (e.g., 10% versus 90% probability). In half of the subjects, the forgone food simply went undelivered, but in the other half the forgone food was visibly delivered to a competitor instead. Subjects then completed a delayed gratification task (with no uncertainty or competition in the reward delivery) to measure the degree to which their willingness to wait for larger rewards was affected by expectations about reward delivery and competitors. If subjects only cared about absolute outcomes, they should be insensitive to the competitor’s presence and waiting times should only be affected by delivery probability. But, if subjects did care about relative outcomes, then they should wait less when the competitor was present in the training phase. Stevens and colleagues found that, although the bonobos waited less in the delayed gratification task when exposed to lower probabilities of reward delivery during the training phase, this behavior was not further affected by the presence of a competitor. But, this study was primarily focused on generalized expectations and the degree to which uncertainty transfers across tasks, and not whether delayed rewards become less certain over time due to a competitor.

A second series of studies by Matsushima and colleagues tested baby domestic chicks' preferences for immediate or delayed reinforcement on a concurrent-choice task when they were either responding in isolation or in the presence of other chicks (Amita et al., 2010; Amita & Matsushima, 2011; Amita & Matsushima, 2014; Mizuyama, Uno, & Matsushima, 2016). The chicks were trained in isolation (with probabilistic reinforcement), competition (shared access to the delivered reinforcement with three other baby chicks) or perceived competition (a glass partition separated the three other chicks who were yoked to the test subject) and the authors measured the degree to which the subjects' preferences and response times varied as a function of their training and test groups. Across these different papers and experimental combinations, the results suggested that both perceived and real competition could produce more impulsive choices (Amita et al., 2010; cf. Amita & Matsushima, 2011; Mizuyama et al., 2016) and faster responses to the smaller-sooner option (Amita & Matsushima, 2011; 2014). Moreover, variable rates of reinforcement produced more impulsive choices, regardless of whether the smaller-sooner or larger-later option were more variable, suggesting that greater reward variance can select for more impulsive choices (Mizuyama et al., 2016). These studies provide a good foundation using an animal model and concurrent response procedure while directly testing whether competition can increase impulsive choices above and beyond environmental uncertainty.

Finally, other research using wild animals has also found evidence that competition can increase impulsive choices by shifting preferences towards smaller-sooner rewards and altering caching strategies, the storage of fat deposits, and patch revisitation in animals (Carrascal & Moreno, 1993; Emery et al, 2003; Hopewell et al., 2008).

So, overall, there seems to be correlational and experimental evidence linking competition to stronger preferences for more immediate payoffs, but studies with causal evidence have been limited[6] to animal subjects. There has also yet to be work providing a stronger theoretical account for why impulsive choices may be functionally important in dynamic situations, within the context of human behavior, and establishing a functional relationship using more representative designs (Brunswik, 1955) between these two factors.

Therefore, the primary purpose of the current dissertation is to experimentally investigate the degree to which competitor-based uncertainty increases impulsive choices in a dynamic delayed gratification task in which competitors are directly inserted into the game and can obtain points at the expense of the participant. More specifically, the dissertation's experiments measure participants' waiting decisions and reward rates under various conditions that modify the reward's uncertainty over time, and test whether knowledge about competitors increases participants' propensity to respond more quickly and collect more rewards than environmental uncertainty. A secondary goal of this study is to extend previous research into the human domain

---

[6] Michaelson et al (2013) showed that the trustworthiness of a hypothetical co-player offering smaller sooner or larger later rewards influences intertemporal choices. Specifically, participants significantly preferred smaller sooner rewards (were more impulsive) when co-players were described as less trustworthy. But this design confounds whether the subjects were sensitive to the general uncertainty of larger later rewards being delivered or the competitive motivations of an untrustworthy person.

in a more dynamic manner and lay the groundwork for future research on the psychometrics of competition and theory building.

To set the stage for addressing these open questions, the next chapter provides a theoretical integration of competition and intertemporal choices based on evolutionary theory, where impulsive choices could have been selected for over time by the fitness benefits (i.e., reproductive success) of maximizing relative payoffs and outperforming competitors, even in environments for which waiting was originally incentivized. I then propose competing hypotheses about the relationship between impulsive choices and competitive environments that the experiments are designed to test.

Chapter 2 - A Theoretical Integration of Competition and Intertemporal Choice

Competition and Relative Performance

Competition is ubiquitous. Humans and other animals clash over mutually exclusive and finite resources like food, mates, jobs or wealth, sovereign nations engage in geopolitical struggles for power and influence, and corporations vie for economic superiority in their respective markets. Decision-making under conflict is fundamental to myriad economic and social behaviors where organisms interact with, and compete against, others to obtain rewards that are contingent on each other's actions. Thus, the ability for decision makers to select appropriate actions in competitive environments is necessary for maintaining and improving their well-being. And because decisions are interdependent (Martin, 2016; Rilling & Sanfey, 2011), strategizing requires learning about successes and failures in an environment as well as anticipating the actions of competitors.

Theoretical and empirical research on competitive behaviors is predominately conducted through the lens of *Game Theory*[7] (Nash, 1950; von Neumann & Morgenstern, 1947), the study of strategic interactions between rational agents (see Luce & Raiffa, 1989 for a non-technical introduction). Game Theory emerged from pioneering work on the development of *Expected Utility Theory* (von Neumann and Morgenstern, 1947) and *Subjective Expected Utility Theory* (Savage, 1954), the mathematical analyses of decision-making under risk and uncertainty. These

---

[7] But see Social Comparison Theory (Festinger, 1954; Garcia, Reese & Tor, 2019; Garcia, Torr & Schiff, 2013) for a social psychological approach to studying competitive attitudes and interactions.

theoretical frameworks proved that agents satisfying particular axioms of rational behavior[8] will select actions as if they are maximizing the expected value of probabilistic outcomes. Game theorists apply these rationality principles to strategic social interactions (i.e., games) and the field achieved even greater prominence after John Nash (1950) proved that for all competitive games, each player has a set of strategies that is guaranteed to maximize their expected payoffs regardless of the other players' choices. When all players behave "rationally" by selecting actions from the strategy set that maximizes their expected payoffs, this is referred to as a Nash equilibrium. Nash equilibria were an equally important discovery in game theory and demonstrated that there exist quantitatively precise benchmarks by which decision-making agents can maximize their expected payoffs against other players in economic encounters. Researchers studying economic behavior are then interested in comparing game-theoretic optimal behaviors against empirical behaviors to measure the degree to which they align. Game Theory is an extremely popular research field widely applicable to psychology, biology, economics, political science and computer science, and its expected utility-based models are a powerful analytic tool of strategic reasoning used to study numerous competitive situations such as organizational decision-making, conflict disputes and bargaining, military strategy, and artificial intelligence, at both the individual and group level (Colman, 1995; Parkes & Wellman, 2015; Schelling, 1960).

However, models of competitive behavior based on expected utility make strong assumptions about the environment and decision-makers to satisfy the tenets of game-theoretic rationality. One assumption is that the environment is stationary, where the probability of

---

[8] Example axioms include: 1) transitivity, 2) completeness, 3) independence of irrelevant alternatives and 4) continuity (see Leyton-Brown, 2007 or Luce & Raiffa, 1989).

success for each action is constant throughout the interaction. A second set of assumptions presumes that decision-makers form accurate beliefs about their environments and opponents, perfectly compute the expected utility of all possible actions from these beliefs, and then always select actions according to a 'minimax' strategy (i.e., minimization of expected losses) in zero-sum interactions. A third assumption asserts that decision-makers are self-interested, which is commonly interpreted as decision-makers seeking to maximize their individual utility (i.e., absolute gains; Murphy, Ackermann & Handgraaf, 2011; see also Camerer, 2003 and Fehr & Schmidt, 1999) and ignoring social utility, such as relative utility (i.e., maximizing gains and minimizing losses relative to others) or joint utility (i.e., maximizing collective gains).

Perhaps unsurprisingly then, expected utility models suffer from poor descriptive and predictive validity. Natural environments are usually non-stationary, the concept of utility is difficult to operationalize[9], which muddles the classification of "utility maximizing" behaviors, and living organisms' behavior rarely comports with the strict theoretical principles governing rationality (Kacelnik, 2003; Herrnstein, 1990), in part due to the natural constraints of cognition (i.e., *bounded rationality*; Camerer, 2003; Selten 1990, 1998; Shafir & Tversky, 1992; Simon, 1957). For example, individuals are often overly short-sighted (impulsive), form erroneous beliefs about others, make decision mistakes, and exhibit predictable and exploitable patterns in their choices (Camerer, Ho & Chong, 2004, 2015; Dyson, 2021; Golman, Bhatia & Kane, 2019; Zhang, Moisan & Gonzalez, 2021) that do not follow a minimax strategy (cf. Martin et al., 2014) But, decision makers can compensate for many of these shortcomings by learning to exploit the

---

[9] For example, utility can take different forms that are not quantifiably similar (e.g., health, wealth, happiness, status, opportunity), is operationally circular (i.e., a commodity's utility motivates a person to buy it, and a person buying that commodity shows it has utility), and potentially every decision can be justified as utility maximizing (see Mishra, 2014, p. 282).

underlying characteristics of their environments (Camerer & Ho, 1999; Dyson, 2021; Gigerenzer & Selten, 2001; Lee et al., 2004; Lee, McGreevy & Barraclough, 2005; Lee & Seo, 2018; Martin, 2016; Schelling, 1960; Zhang et al., 2021) and develop optimal solutions that might initially appear irrational (e.g., melioration; Sims et al., 2013).

There is also a lack of consensus about whether decision agents are required to be narrowly self-interested and ignore social utility. Theoretically, narrow self-interest is a simplifying assumption that underlies many expected utility models and analyses of behavior (Camerer, 2011; Fehr & Schmidt, 1999; Murphy et al., 2011). However, it has been argued that utility can contain both selfish and unselfish motives if that is a description of the world and its states that the organism prefers and is motivated to pursue (Harsanyi, 1966; Leyton-Brown, 2007), and that no axioms of game-theoretic rational behavior explicitly require the property of narrow self-interest (Ross, 2019). Empirically, many decision makers do indeed value relative payoffs (Fehr & Schmidt, 1999; Loewenstein et al., 1989). For example, many players in economic games prefer higher gains relative to other players, and these preferences increase when displaying other players' cumulative gains and labeling them as an "opponent" rather than a "player" (Crockett et al., 2010; Espin et al., 2012; Espin et al., 2015; Messick & McClintock, 1968; Messick & Thorngate, 1967; McClintock & McNeel, 1966; McClintock & Nuttin, 1969; Murphy et al., 2011). Knowledge of competitors can also increase risk-taking (Rosati & Hare, 2012), particularly when individuals perceive themselves as socially disadvantaged (Mishra, Barclay & Lalumiere, 2014) and attempt to improve their relative position (Gonzalez, Mishra & Camp, 2016; Hill & Buss, 2010), or when men think they are being observed by someone of equal relative status (Ermer, Cosmides & Tooby, 2008).

Individuals' sensitivity to relative outcomes and competitiveness can have large economic and social consequences (Garcia, Tor & Schiff, 2013; Murayama & Elliot, 2012), which makes clarifying its role in strategic interactions important. People will differentially value financial (in)equity in disputes based on the nature of their relationship with a co-disputant (Loewenstein, Thompson & Bazerman, 1989), prefer lower salaries so long as they earn more than co-workers (Tversky & Griffin, 1991) and rate their financial positions more favorably on average if they are better off than others (Frederick & Loewenstein, 1999; Luttmer, 2005; Tversky & Griffin, 1991; but see Ifcher et al., 2020). Additionally, competitors can negatively affect judgments of risk-reward structures (Pleskac, Conradt, Leuker & Hertwig, 2021) and alter task performance both positively (Cagiltay, Ozcelikb & Ozcelik, 2015; see Subhash & Cudney, 2018 for a review on digital game-based learning paradigms) and negatively (Chen & Chang, 2020; Chen, Liu, & Shou, 2018), depending on the number of competitors (Garcia & Tor, 2009) and the individual's achievement goals (Murayama & Elliot, 2012). Individuals will also over-bid in competitive auctions to ensure winning a prize and improving their social status (Adam, Kraemer & Mueller, 2015; Ku, Malhotra & Murnighan, 2005; Malhotra, 2010; van den Bos et al., 2008; van den Bos et al., 2013), and suboptimally distribute their time across resource patches in foraging tasks to avoid competitors (Goldstone et al., 2005; Gillis & Kramer, 1987; Madden, Peden & Yamaguchi, 2002; Roberts & Goldstone, 2006; Silston et al., 2021).

Because acquiring crucial resources and capital (e.g., food, water, mates, money, social status, grant funding, jobs) is necessary for a decision maker's (or group's) well-being, they must be able to not just do well, but also outperform competitors. To do so, decision makers could prefer smaller sooner rewards to a degree that conflicts with traditional utility-based models because these models eschew important social dynamics like relative payoffs. And yet, preferences for higher relative gains and decision rules prioritizing smaller-sooner rewards could actually be ecologically rational given the environmental conditions in which the organism exists. This suggests that both relative performance and impulsive choices may reflect biologically important, evolved behaviors in competitive environments. Therefore, choice theories couched within an evolutionary framework could provide a superior representation of organisms' waiting behaviors in competitive ecologies (Stevens et al., 2005; Stevens, Rosati et al., 2005), particularly towards understanding these behaviors as adaptations that helped solve the myriad challenges related to survival and reproduction in the wild.

Evolutionary Perspectives on Competition and Relative Performance

Evolutionary perspectives on the behavioral sciences apply the general principles of natural selection (Darwin, 1859; Fisher, 1930; Wright, 1931) and fitness (Orr, 2009) to human behavior and cognition. Broadly, natural selection refers to the differential survival and reproduction of organisms based on differences in their behavioral and cognitive traits. It is an evolutionary mechanism in which heritable traits become relatively more or less frequent in a population over time depending on how well or poorly, respectively, they aid survival and reproduction. Fitness can assume many definitions (Barker, 2009) related to "reproductive success", but here it will generally refer to 1) the degree to which traits are passed down from individual organisms to their (grand)offspring, and 2) the relative frequency of those traits in a population. There are various components theorized to affect both short-term and long-term fitness (Orr, 2009; Sebens, Sarà & Carrington, 2018), such as energy gains (Broom, Cressman & Krivan, 2019; Charnov, 1976; Krebs, Stephens & Sutherland, 1983), utility (Mishra, 2014; Orr, 2007) and reward (Silver et al., 2021; but see Vamplew et al., 2021), and a key link between natural selection and fitness is the importance of an organism's relative achievement of these components (Orr, 2007; Wilson, 2004). Ultimately, an organism's survival and reproductive success is most important in comparison to conspecifics bearing alternative traits (Wilson, 2004). Thus, fitness differences among organisms are necessary for natural selection to lead to changes in traits expressed in a population (Barker, 2009; Orr, 2009), because over time, traits better suited for an environment than others will tend to be expressed more strongly (i.e., selected for) in subsequent generations whereas traits more poorly suited for an environment than others will tend to be filtered out (i.e., selected against).

Competition for resources and fitness components is thus clearly a cornerstone of evolutionary theory (Darwin, 1859; Hintze, Phillips & Hertwig, 2015), and achieving relatively greater amounts of resources or offspring than others is particularly important. Organisms able to obtain more food than competitors, find better shelter than competitors, and propagate more fruitfully than competitors will be more likely to have their traits expressed in the population via their (grand)offspring and genetic relatives. Without competition then, organisms would be less likely to evolve functional traits (Hintze et al., 2015) associated with greater relative fitness (Orr, 2007), because all extant traits would contribute equally to an organism's survival and reproduction and are then equally likely to be passed down to subsequent generations. Therefore, organisms' pursuit of relative outcomes may not be so simply regarded as irrational (as they are in many expected utility models), because this preference tends to positively correlate with the acquisition of crucial primary resources (e.g., food, water, mates) that better ensures their survival and fitness at a higher rate than others (Buss, 1995). And preferences for relative gains, perhaps selected for due to their functional adaptiveness in dynamic environments rife with uncertainty and limited resources, have remained intact today, (mis)guiding many people's pursuits of more contemporary ambitions (e.g., money, social status, grant funding, and jobs).

There are multiple distinct evolutionary approaches to studying the human behavioral sciences, such as human behavioral ecology and evolutionary psychology. These perspectives similarly apply the general principles of evolution to behavior and cognition that are relevant to competition and preferences for relative gains, but differ in various ways in how these principles are applied (see Durrant & Ward, 2015 for a comparative review). This current dissertation is not geared towards experimentally comparing the different approaches, but rather leveraging their biologically informed perspectives towards understanding humans' preferences for relative gains (a potential psychological mechanism) and various waiting behaviors (which can adapt to different situations) in competitive environments. The benefits of doing so seem patently clear. Recasting the traditional notion of economic rationality (via expected utility) into ecological rationality (via survival and fitness) has led to remarkable contributions in understanding other strategic behaviors that conflict with game theory, such as altruism (Trivers, 1971), cooperation and relative payoff maximization (Hammerstein & Selten, 1994), *limited war* as an evolutionarily stable strategy (Maynard Smith, 1982; Maynard Smith & Price, 1973), the *handicap principle* (Mesterton-Gibbons & Adams, 1998), foraging and the matching law (Seth, 2007), kin selection (Hamilton, 1964), status seeking and punishment of free riders (Tooby & Cosmides, 2016), life history, marriage and offspring rearing, and pursuits of individual and collective interests (Mulder & Schacht, 2012). Modeling analyses (Fawcett et al., 2012; Stephens, et al., 2004) have also revealed that fitness maximization can lead to economically "irrational" choices (Houston, McNamara & Steer, 2007; Trimmer, 2013) under dynamic and uncertain environmental conditions and physiological stress (e.g., the organism's current energy state; Marsh et al., 2004; Trimmer, 2013). Thus, it also seems clear that integrating relative payoffs, impulsive choices and competition, guided by evolutionary research programs can help provide a more grounded psychological examination into how and why these preferences and behaviors occur (i.e., the *ultimate* and *proximate* levels of analysis, respectively; Bergman & Beehner, 2022; Tinbergen, 1963).

The Evolutionary Adaptedness of Impulsive Choices under Competition

Evolutionary theory presents an appealing framework for reconciling the discrepancies between the seemingly irrational impulsive behaviors and game-theoretic utility models by means of competition. Stevens and King (2013) argue that Darwin's seminal thesis on organisms' "struggle for existence" in which "natural selection acts by competition" (1859, p. 472; cited from Stevens & King, 2013, p. 410) relates to both the physical and social environment. Competitors are also a frequent example cited when describing the unpredictability of an environment because they impede organisms' delayed consumption either by interrupting foraging or pilfering cached resources (Fawcett et al., 2012; Fenneman & Frankenhuis, 2020; Fenneman et al., 2022; Henly et al., 2008), and this social competition can generate greater uncertainty in reinforcement than the physical environment (Stevens & King, 2013). Stevens and Stephens (2010) also argue that there is clear evidence of competition interfering with organisms' waiting behaviors (see the section above on *Intertemporal Choice and Competition* for relevant citations).

However, to my knowledge, there is currently no formal model that directly bridges intertemporal choices and competition (see Fenneman & Frankenhuis, 2020, p. 271 and Fenneman et al., 2022 for reviews and their brief remarks on this point), for example describing how organisms' need for a relatively greater abundance of resources and social status drive preferences for immediacy (Fenneman & Frankenhuis, 2020). Here, I briefly turn to

Evolutionary Psychology to provide the conceptual tools for proposing a relationship between them.

Buss's *hierarchy of levels of analysis* (see Figure 1 in Buss, 1995) provides a simplified roadmap for how researchers can conduct evolution-based research and develop strong theoretical accounts of psychological mechanisms and behavior in relation to the environment. The first level of this hierarchy is 'Evolution by Natural Selection', a *General Evolutionary Theory* so deeply entrenched in the biological sciences that it is simply assumed to be correct[10]. Researchers interested in testing evolutionary propositions about behavior are not testing this general theory per se, but rather more specific properties that are derived from its outline. The next level down from the general theory of evolution comprises *middle-level theories*, which tend to broadly encompass "entire domains of functioning", for example reciprocal altruism and parental investment theory. These are theories that researchers are most interested in validating as evolutionary propositions through empirical means, such as cross-cultural surveys and rigorous experimentation. And if these theories are unsupported, their falsification calls for their removal or replacement with a better explanation. From these middle-level theories, researchers can make specific hypotheses about the relationships between organisms and their environment. And then these hypotheses generate specific, testable predictions regarding how organisms actually behave or how their cognitive systems evolved to produce specific behaviors.

Proposing a rigorously defined and appropriately constrained middle-level theory of intertemporal choices under competition is beyond the scope of this dissertation. Yet, the general principles derived from these levels of analysis can stimulate a possible explanation on their

---

[10] Note that this statement does not mean evolutionary theory lacks falsifiability, but rather there has yet to be any evidence of its falsehood.

relationship and preliminary hypotheses to be tested. Evolutionarily, humans' sensitivity to relative outcomes could be an evolved psychological mechanism that functionally outputs behaviors associated with greater resource acquisition and reproductive success than conspecifics. One behavior particularly relevant to an organism's fitness goals or solving adaptive problems is acting on smaller sooner rewards, which can be optimal under numerous plausible conditions that organisms routinely encounter in the wild (e.g., environmental harshness and uncertainty, and greater resource collection risks and opportunity costs; Fawcett et al., 2012; Fenneman & Frankenhuis, 2020; Modelling Animal Decisions Group et al., 2014). More concretely said, impulsive choices can be a behavioral solution to environmental uncertainty that aids greater resource acquisition and reproductive success, particularly in competitive environments with limited and contested resources in which decision speed is prioritized over accuracy (Hintze et al., 2015).

Because long-term survival depends on short-term survival, impulsive choices can repeatedly extend an organism's immediate lifespan long enough to reproduce and rear offspring. The psychological preferences that construct reproductively advantageous decision rules prioritizing smaller sooner rewards are then passed down through generations, solidifying their presence in our cognitive systems and behaviors today, just in different contexts that may or may not align with the original conditions from which they arose (Hayden, 2016; 2019; Fawcett et al., 2012; Fenneman & Frankenhuis, 2020; Tooby & Cosmides, 2016). For example, individuals may accept mediocre jobs or high-interest loans because they don't have time to wait for better offers (especially if other people are also waiting for them) due to their immediate need for financial support. Furthermore, businesses may enact short-term survival operations to get through difficult times before considering long-term plans, or act shortsightedly to gain the upper hand against competitors (e.g., the *first-mover advantage*, where companies release unfinished products early to beat competitors to market, such as video games) that allows them to be appropriately far-sighted down the line during more stable conditions. Preferences for short-term actions could have spawned from our genetic ancestors' fight for survival to suboptimally consume resources in the short-term that extended their lives and selected for those behaviors over time.

So, repeatedly taking smaller sooner payoffs may be suboptimal in the long run in many circumstances, but doing so to outperform or outlast competitors can permit one to focus on long-term goals once (some degree of) normality has resumed. This suggests that in the presence of competitors, individuals will be less likely to wait for larger delayed rewards compared to situations where there is no known source of risk in the delayed rewards because it leads to relatively higher payoffs than competitors and produces relative advantages over them in other ways (e.g., survival, reproduction, social status). If an organism is waiting for a larger amount of food to forage but a competitor takes it first, then the competitor is now better off and more likely to survive, regardless of whether the food would have been more nutritious later. Psychological preferences for relatively larger payoffs and decision rules prioritizing smaller sooner rewards will then be evolutionarily selected for in this ecology, given the potentially intense competition for limited resources and fitness benefits of outperforming other organisms because smaller sooner fitness improvements can be more appealing than larger delayed fitness improvements. Moreover, organisms with greater reproductive success due to (or at least being correlated with) more impulsive choices would reproduce more and conserve this behavioral trait over time through its propagating down subsequent generations.

Therefore, this study tests whether impulsive choices are more prevalent in competitive situations than non-competitive ones, holding the amount of uncertainty in participants' reward rates constant across both conditions. In the next chapter I lay out the dissertation's general design and propose competing hypotheses about the relationship between intertemporal choices and competition that the experiments are designed to test.

Chapter 3 - The Current Dissertation

General Design

An Escalating Interest task (EI; Young et al., 2011; Young & Howatt, 2021) was used to measure participants' propensity to respond rapidly (impulsive) or slowly (patient) for rewards in competitive environments. In EI tasks, participants repeatedly encounter virtual targets and respond to them to obtain rewards. The amount of reward delivered following a response to a specific target is indicated by a visual cue that gradually increases in size over a 10 second interval. Each response resets the cue size, interval time and reward amount back to zero and participants can respond immediately after a previous response. Participants can therefore respond rapidly to obtain multiple smaller rewards in quick succession or respond slowly to obtain larger rewards less frequently.

The amount of reward delivered for a given response time is controlled by a mathematical parameter, *power*, of a superellipsoid function:

$$Reward\ Magnitude = MaxMagnitude * \left(1 - \left(\frac{10-WT}{10}\right)^{power}\right)^{\frac{1}{power}} \qquad (1)$$

where the reward magnitude is a function of *MaxMagnitude*, the maximum possible reward attainable after 10 s have elapsed, *WT*, the participant's waiting time[11] (i.e., the amount of time elapsed in the 10 s interval), and *power*, which modifies the function's growth rate and curvature.

---

[11] Note that EI tasks traditionally use participants' *interresponse times* (IRTs), which is the amount of time that has elapsed since their previous response. In the current design, though, IRTs are fundamentally different than wait times (WTs) because the competitors' responses reset the reward magnitude to zero (see the section *Competitors in the Escalating Interest Task* below for more information), so WTs and IRTs can become misaligned. For example, a participant with an IRT of 12 seconds (which would be considered 'patient') could actually have a wait time of 3 s (which would be considered 'impulsive') if the participant makes an initial response, the competitor responds nine seconds later, and then the participant responds again three seconds later.

Power typically ranges from 0.5-1.5, and Figure 1A illustrates how the reward magnitude grows over time for three distinct power values. When power is equal to 1.0 the reward grows linearly in magnitude up to 10 s (i.e., constant growth). When power is less than 1.0 the reward grows slowly in magnitude early in the 10 s interval, but more rapidly as it approaches 10 s (i.e., positively accelerating growth). When power is greater than 1.0 the reward grows rapidly in magnitude early in the 10 s interval, but more slowly as it approaches 10 s (i.e., negatively accelerating growth). Figure 1B illustrates the reward rate (reward magnitude per unit time) for WTs up to 10 s for the same three power values. For example, a WT of 5 s when the power value is 1.0 would produce 50% of maximal reward, or 10% per second (= 50%/5 s); when the power value is 0.75, a 5 s WT would produce 30% of maximal reward, or 6% per second, and when the power value is 1.25, a 5 s WT would produce 65% of maximal reward, or 13% per second.

Figure 1

Reward Magnitudes and Reward Rates in the Escalating Interest Task

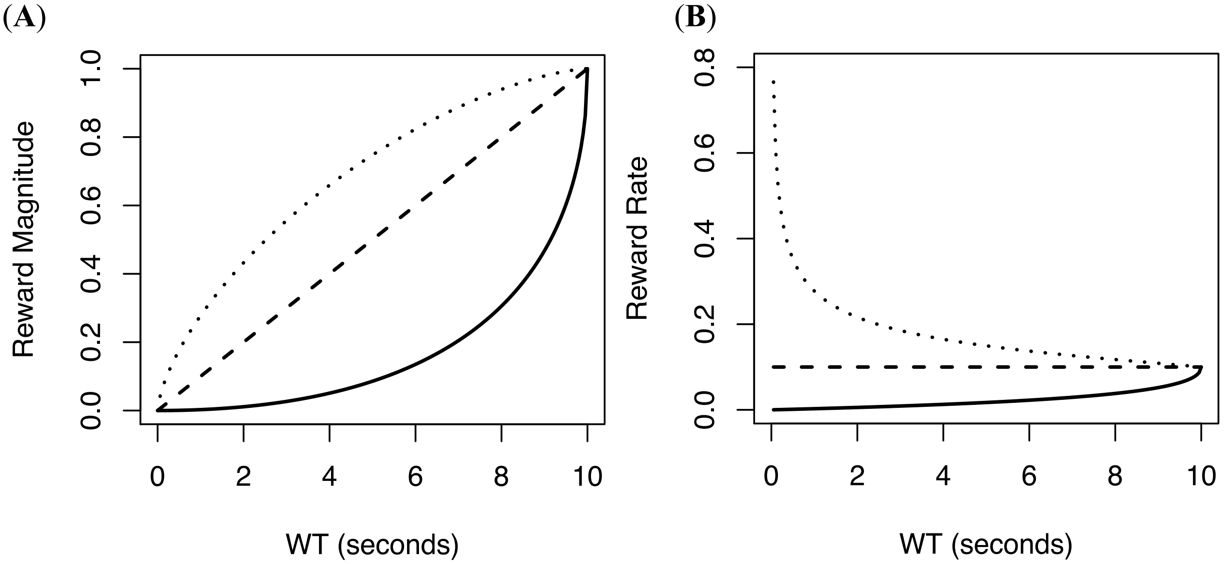

Note: Reward functions for three different power values as a function of a participant's waiting time (WT) in seconds. (**A**) The reward magnitude that can be achieved as the percent of maximum value of the possible reward obtained. (**B**) The reward rate for each waiting time. The top dotted lines represent a power of 1.5, the middle dashed lines a power value of 1.0, and the bottom solid lines a power value of 0.5.

Reward rate can also be thought of as a measure of efficiency in obtaining rewards from the target, of which participants are ideally trying to maximize. When power is equal to 1.0, the reward rate is constant and every WT up to 10 s is equally efficient. For power values less than 1.0, the reward rates increase over time and waiting the full 10 s maximizes efficiency. And for power values greater than 1.0, the reward rates decrease over time and responding as quickly as possible maximizes efficiency. For example, when power is equal to 1.0, two responses of 5 s would produce the same reward amount achievable by waiting 10 s (= 50% × 2), making them equally efficient. But, when power is 0.75, two responses of 5 s would produce only 60% (= 30% × 2) of the reward amount achievable by waiting 10 s, making this a less efficient strategy than waiting. And when power is 1.25, two responses of 5 s would produce 130% (= 65% × 2) of the reward amount achievable by waiting 10 s, making this a more efficient strategy than waiting.

Thus, in these types of EI tasks, the optimal strategy is waiting exactly 10 s when power is less than 1.0 and responding as rapidly as possible when power is greater than 1.0.

However, there are barriers to individuals behaving optimally (Young & Howatt, submitted). First, participants are not informed about the nature of the reward growth function and must learn these dynamics through experience. Second, participants need to discriminate between relatively similar reward rates given by power values that incentivize opposing strategies (e.g., a power value of 0.95 incentivizes waiting 10 s to respond, but a power value of 1.05 incentivizes rapid responding). Third, participants need to precisely estimate how long they have waited since their last response. Fourth, participants need to respond optimally with little-to-no variability in their WTs. Deviations from optimal WTs are thus common in EI tasks and prior research shows that participants are only modestly sensitive to power values (Young & Howatt, 2021). Nonetheless, EI tasks and its variants have been used to model waiting times under various naturalistic conditions related to uncertainty in the reward's delivery and growth rate, choice commitment and defecting, and resource limitations (e.g., Webb & Young, 2015; Young et al., 2011; Young & Howatt, 2021; Young & Howatt, submitted; Young & McCoy, 2015; Young, Vangsness & McCoy, 2018). Therefore, inserting competitors as a specific source of reward uncertainty will continue this fruitful line of research studying real-world contingencies and their impact on intertemporal choices involving relatively short delays.

Even though EI tasks typically use power values ranging from 0.5-1.5, the current dissertation will simplify the complexities of the EI task by assigning all reward targets a power value of 0.75 (see Figure 2 for reward magnitudes and rates). While this simplification will limit the immediate generalizability of the experimental findings, removing power as a predictor variable will greatly reduce the complexity of the statistical models and enable a purer assessment of the effects of competition and competitors' behavior on waiting behaviors. Because uncertainty can incentivize impulsive choices even when patience originally maximizes reward rates, a power value of 0.75 will be assigned to each reward target to initially encourage waiting the full 10 s.

Figure 2

Reward Functions for a Power Value of 0.75

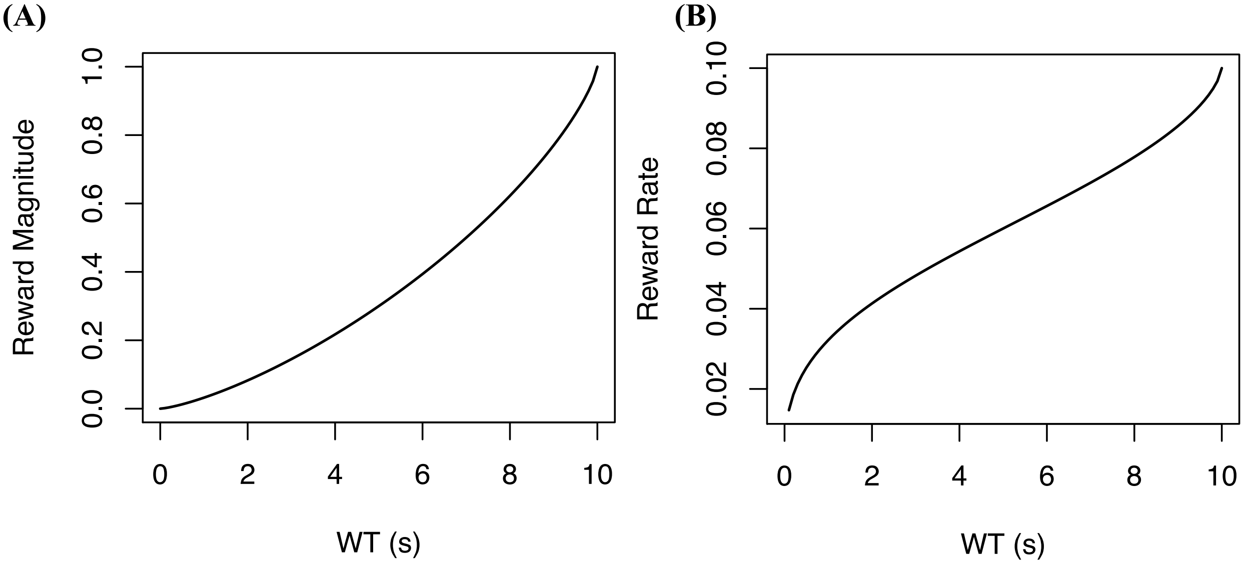

Note: (**A**) Reward magnitudes. (**B**) Reward rates.

A power value of 0.75 is specifically being used here because lower values (e.g., 0.5) may be overly punishing when rapid responding is necessary to collect rewards against impulsive competitors, but higher values (e.g., 0.95) may not punish rapid responding enough. Ideally, participants should be faced with the dilemma of desiring to respond rapidly to collect rewards in the face of uncertainty, but also be able to learn that too rapid of responding is detrimental to maximizing their reward rate. Future research could systematically sample power values across a broader range to assess individuals' sensitivity to competition-based uncertainty as a function of the reward's growth rate.

Competitors in the Escalating Interest Task

To study the effects of competition on participants' waiting behaviors, one computerized avatar was placed next to each target and responded to it to collect rewards. Competitors' waiting times (WTs) were randomly sampled from a 3-parameter Weibull probability distribution function:

$$f(WT \mid k, \lambda, c) = \begin{cases} \frac{k}{\lambda}\left(\frac{WT-c}{\lambda}\right)^{k-1} e^{-\left(\frac{WT-c}{\lambda}\right)^{k}}, \ x \geq c \\ 0, \ x < c \end{cases} \qquad (2)$$

The shape parameter, *k*, controls whether the competitor's WT distribution is right-skewed ($0 < k < 3$), approximately normal ($3 < k < 4$) or left-skewed ($k > 4$). The scale parameter, λ, controls the variance of the competitor's WTs, with larger values corresponding to more variability (i.e., more unpredictable WTs). Lastly, the threshold parameter, *c*, defines the competitors' minimum possible WT.

The probability of a competitor having made a response by a particular time given its assigned parameters is determined by the cumulative distribution function:

$$F(WT \mid k, \lambda, c) = 1 - e^{-\left(\frac{WT-c}{\lambda}\right)^{k}} \qquad (3)$$

Weibull functions are methodologically advantageous to use as a model of the competitors' waiting times because they are highly flexible and can describe a wide array of possible distributions[12] (e.g., right- or left-skewed, approximately normal, bimodal, and low or high variance) The function's parameters can be easily manipulated to generate competitors who

---

[12] Different shape parameters of Weibull functions also produce hazard rates (Dasgupta & Maskin, 2005; McGuire & Kable, 2012; Sozou, 1998) that can incentivize different waiting strategies, but hazard rates are not considered further here.

are impulsive, patient or anywhere in between, on average, with different rates of predictability in their behavior (see Figure 3 for examples). It is important to note that participants' WTs in prior studies are usually bimodal with peaks early in the interval and at 10 s, but for the purpose of this current dissertation the competitors' WTs were Weibull distributed with only one mode. Future studies could use mixtures of Weibull functions to generate bimodal distributions more reminiscent of empirically observed WTs.

Figure 3

Example Competitors

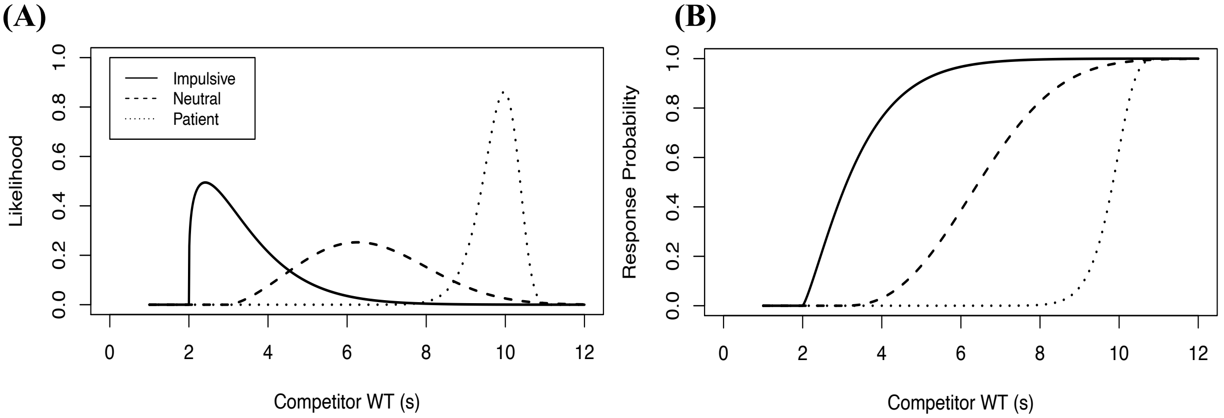

Note: Example Weibull distributions illustrating hypothetical competitors that are impulsive, neutral or patient based on their programmed wait times (WT) in seconds. (**A**) Probability density functions that show the theoretical distributions of different competitors' possible WTs. Higher values on the curves denote the WTs more likely to be observed. (**B**) Cumulative distribution functions that show each competitor's probability of having made a response by that time.

However, the flexibility of Weibull distributions complicates selecting an appropriate set of parameters and cleanly interpreting people's behavior as a function of them (e.g., small changes in the parameters can simultaneously change the distribution's mean, variance, skew and kurtosis). In order to constrain the competitors' behavior, I first fixed all of their WT distributions to be approximately normal (i.e., a shape parameter equal to 3.5). Then, I randomly assigned their mean WTs to be either 3 s or 9 s and the variance of their WTs to be either 0.14

(*sd* = 0.02) or 0.85 (*sd* = 0.73). The mean and variance of the competitors' WTs were thus be manipulated in a 2×2 factorial manner to create four different competitors: 1) shorter average WTs with lower variance, 2) shorter average WTs with higher variance, 3) longer average WTs with lower variance, and 4) longer average WTs with higher variance. Competitors with shorter average WTs (3 s) are considered more impulsive, whereas competitors with longer average WTs (9 s) are considered more patient. Competitors with lower variance (0.14) in their WTs are considered more predictable, whereas competitors with higher variance (0.85) in their WTs are considered less predictable.

Table 1 below provides the Weibull function's parameter values and WT descriptions for the four different competitors, and Figure 4 illustrates their WT distributions (A) and cumulative response probabilities (B) based on the assigned parameters.

Table 1

Weibull parameters and waiting time (WT) attributes for the four different competitors in the current study.

Competitor	**Parameter Values**			**WT (s)**	
	Shape (k)	Scale (λ)	Threshold (c)	Average	Variance
Impulsive-low variance	3.5	0.5	2.55	3	0.14
Impulsive-high variance	3.5	3.0	0.31	3	0.85
Patient-low variance	3.5	0.5	8.55	9	0.14
Patient-high variance	3.5	3.0	6.31	9	0.85

Note: each competitor's threshold parameter is different, but this is only to shift their distributions along the x-axis.

Figure 4

Competitor Wait Times

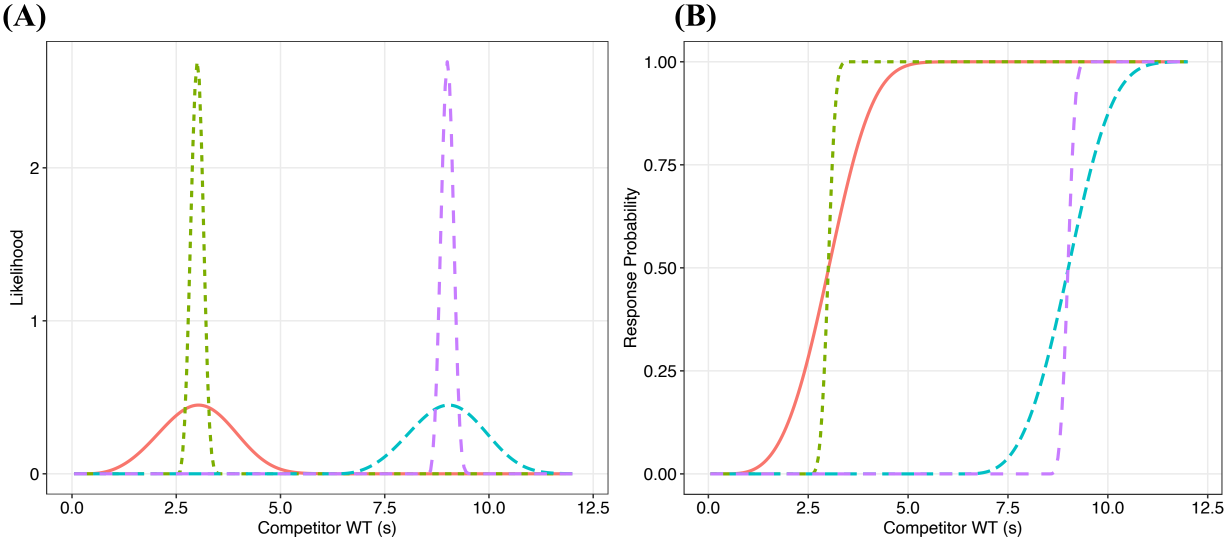

Note: Weibull distributions illustrating the four competitors' response time behavior. (**A**) The competitors' waiting time probability density functions. (**B**) The competitors' waiting time cumulative distribution functions. The red solid line is the impulsive-high variance competitor, the green dashed line is the impulsive-low variance competitor, the blue dashed line is the patient-high variance competitor, and the purple dashed line is the patient-low variance competitor.

When either the participant or computerized competitor responded to a target, the visual cue size, interval time and reward amount reset to zero for both players. Therefore, the amount of reward each player could possibly achieve with each response was mutually contingent upon the time elapsed since either's last response. This is an oppressive form of competition but is analogous to Young et al's study (Experiment 2, 2014) where the magnitude of the reward increased over time while the probability of receiving the reward following a response decreased over time. Unlike Young and colleagues' preparation, in the current study a response always delivered the reward, but the participant could lose out on larger rewards simply by waiting too long to respond. The purpose of the current study's competition-based uncertainty paradigm was to elicit maximal sensitivity to the task's contingencies and relate to naturalistic scenarios in which waiting too long for a larger reward is directly associated with a reduction in its availability (e.g., waiting to long for an apple to ripen results in a groupmate taking it first).

The current paradigm is also akin to a balloon analogue risk task (BART, Lejuez et al., 2002) in which the balloons automatically inflate over time (Young & McCoy, 2019). The key difference between the current paradigm and the BART is that the traditional BART design confounds participants' sensitivity to the increasing magnitude of the reward with their sensitivity to the increasing uncertainty of receiving the reward. In contrast, EI tasks decouple participants' sensitivity to reward magnitude and uncertainty by independently changing the degree to which the rewards grow over time and the degree to which they become less likely to be delivered, enabling a cleaner assessment of which contingency more strongly influences participants' responding.

The primary consequence of the computerized avatar at each target is that participants waiting too long to respond for increasingly larger rewards increase the risk of losing them, even under conditions in which patience originally optimizes reward rates (power ≤ 1.0). Because losing rewards makes participants worse off than their competitor, they may begin to respond more rapidly to increase their chances of obtaining rewards and making them better off than their competitors.

Reward Maximizing Behavior

How long should participants optimally wait to respond to targets in this competitive EI task? Reward maximizing behavior for a power value of 0.75 and against the four competitors was determined by simulating ideal decision-makers who produced constant WTs for the entire experiment. Figure 5 illustrates the expected reward rate (efficiency) for WTs across the four competitors and the reward maximizing WTs. The expected reward rates were computed by multiplying the objective reward rates for a power value of 0.75 (Figure 2B) by the competitor's response probability function (Figure 3B) and reflect a simple expected utility model of optimal behavior against each of the four competitors. These simulations revealed that participants maximize their expected reward rates by responding every 1.8 s against the impulsive-high variance competitor, 2.65 s against the impulsive-low variance competitor, 7.35 s against the patient-high variance competitor and 8.65 s against the patient-low variance competitor.

Figure 5

Reward Maximizing Wait Times

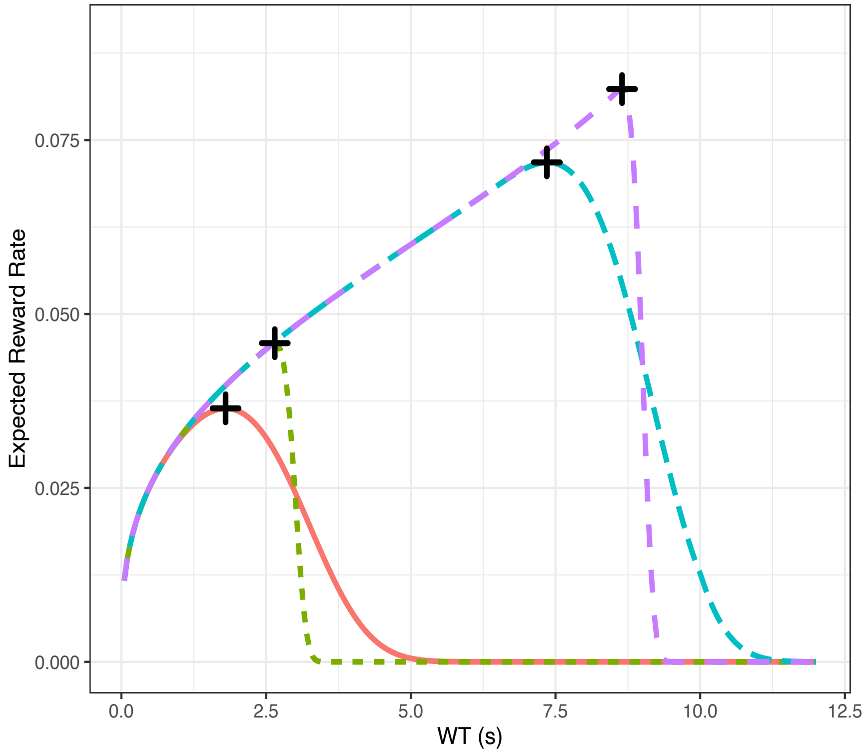

Note: Expected reward rates for various participant WTs against the four different competitors. The black crosses denote participants' reward-maximizing WTs against each competitor. The red solid line is the impulsive-high variance competitor, the green dashed line is the impulsive-low variance competitor, the blue dashed line is the patient-high variance competitor, and the purple dashed line is the patient-low variance competitor.

Participants' waiting times should theoretically match the reward-maximizing times presented in Figure 5, but their empirical behaviors are expected to exhibit considerable variability given the barriers to individuals behaving optimally. Although the differentiation between reward rates is smaller than in the typical EI task in which there is no competition (Figure 1A), the participants will still be able to collect rewards more effectively when they respond more rapidly (but not as fast as possible) than the full 10 s when the power value is below 1.0. Previous research on the EI task without competition, including research with the same gaming environment used here (Young & Howatt, 2021), has shown that participants are moderately sensitive to power, although with considerable individual differences in both their sensitivity to power and overall tendency to wait versus to respond quickly.

One obstacle to participants engaging in reward-maximizing behavior is that when power values are less than 1.0, rapid responding produces a low reward rate coupled with a long target-depletion time. This situation is likely to be highly aversive to human participants because it can substantially prolong the task's duration. Normally, shorter WTs for low power values are uncommon, but this combination is expected to be more frequent given that two of the competitors will be relatively impulsive with average WTs of 3 s. The aversiveness of an extended task duration due to frequent low-reward rates is not a problem for a computer avatar, who will dutifully respond as programmed until the target's rewards have been exhausted. Participants must therefore balance the low power value, competition, and their likely delay aversion in how they select practical waiting times, although it is certainly possible that participants will reject these situations and default to simpler behaviors or even disengage from the task entirely.

General Hypotheses

The factors hypothesized to impact individuals' waiting times in the current EI task are as follows: 1) the power value of the reward targets, 2) the presence of competition versus its absence, 3) the competitor's mean wait time, and 4) the competitor's wait time variance. Power represents the state variables incentivizing different waiting times based on the reward targets' programmed growth rate. Under ideal circumstances a power value greater than 1.0 should shorten participants' average WTs (more impulsive choices) and a power value less than 1.0 should lengthen their average WTs (more patient choices); a power value equal to 1.0 does not provide any relative advantage for either faster or slower average WTs. Because the power values for all targets in the current study will be 0.75, participants should lengthen their average WTs and wait the full 10 s to maximize their reward rate when there is no competitor. The competitors' mean and variance of their WT distributions represent their average rates of responding and the amount of uncertainty around this average, respectively. As a competitor's mean WT increases, the participants' WT should also increase, up to the maximum values shown in Figure 5. As a competitor's WT variance increases, participants' WT should decrease (see Figure 5) because the likelihood that the reward is still available diminishes more rapidly.

Competition refers to a known source of reward uncertainty, but the degree to which its presence strengthens the effect of uncertainty and further shortens individuals' average WTs is still unresolved. Two competing perspectives are considered: *simple subjective utility* and *evolutionary theory*. The simple subjective utility perspective makes the simplifying assumption that rational agents are materially self-interested and ignore social factors (Camerer, 2003; Murphy et al., 2011), such that decision-makers should be insensitive to competition as a specific source of uncertainty. This perspective therefore predicts that there will be no difference in participants' waiting times between conditions of environmental uncertainty and competition-based uncertainty, because the expected reward rate of various WTs (e.g., Figure 5) is unaffected by *why* there is uncertainty in the environment. In contrast, the evolutionary theory perspective would expect decision-makers to be sensitive to social factors and should be influenced by competition. This perspective therefore predicts that participants will have shorter waiting times under competition-based uncertainty because they desire to maximize their outcomes relative to the competitor. These competition-based hypotheses are only broadly defined here, and the methods section of each experiment will provide more specific hypotheses given their respective manipulations.

Chapter 4 - Experiment 1

Experiment 1 provided an initial test of the relationship between competition-based uncertainty and impulsive choices by manipulating the visibility of the competitor. Participants were randomly assigned to one of two groups: 1) a *visible competitor* group, where participants could see the competitor and were explicitly aware that they were the source of reward uncertainty, and 2) an *invisible competitor* group, where participants could not see the competitor and were naïve to the competitor as the source of reward uncertainty.

In addition to the broader hypotheses described above, there are two specific hypotheses relevant to this experimental manipulation. The first hypothesis regards participants' selection of impulsive choices. The simple subjective utility hypothesis, Hypothesis 1a, predicts that participants in the competition-visible group will be equally likely to wait to respond for rewards as participants in the competitor-invisible group. In contrast, the evolutionary hypothesis, Hypothesis 1b, predicts that participants in the competitor-visible group will be less likely to wait to respond for rewards (i.e., more impulsive choices) than participants in the competitor-invisible group.

The second hypothesis regards participants' probability of obtaining rewards. Organisms that make more impulsive choices will likely also collect more rewards, yet similar rates of impulsive choice could produce different reward rates, whereas different rates of impulsive choice could produce similar reward rates. Thus, these two outcomes are different from each other and can reveal unique information about peoples' behavior and performance. The simple subjective utility hypothesis, Hypothesis 2a, predicts that participants in the competition-visible group will be equally likely to obtain rewards as participants in the competitor-invisible group. In contrast, the evolutionary hypothesis, Hypothesis 2b, predicts that participants in the competitor-visible group will be more likely to collect rewards than participants in the competitor-invisible group.

Methods

Participants

A total of 56 people originally participated in this experiment for research credit as part of a course requirement. The study was approved by Kansas State University's IRB, and participants provided written informed consent prior to their participation. Five participants were removed for poor quality data (see Results below) for a final sample size of 51 adults (32 female, 18 male, 1 did not respond) aged 18-27 years old ($M = 19$).

Design

Participants completed a modified escalating interest task, a first person open-world video game programmed in the Unity3D game engine (www.unity.com). The EI task consisted of five visually identical octagonal levels with eight separate chambers, each containing a reward target and a tall, colored pole (see Figure 6) to help participants find the targets more easily. The inner walls forced participants to come into closer contact with each target rather than staying in the middle of the level and clicking on them with minimal delay.

Figure 6

The Escalating Interest Task

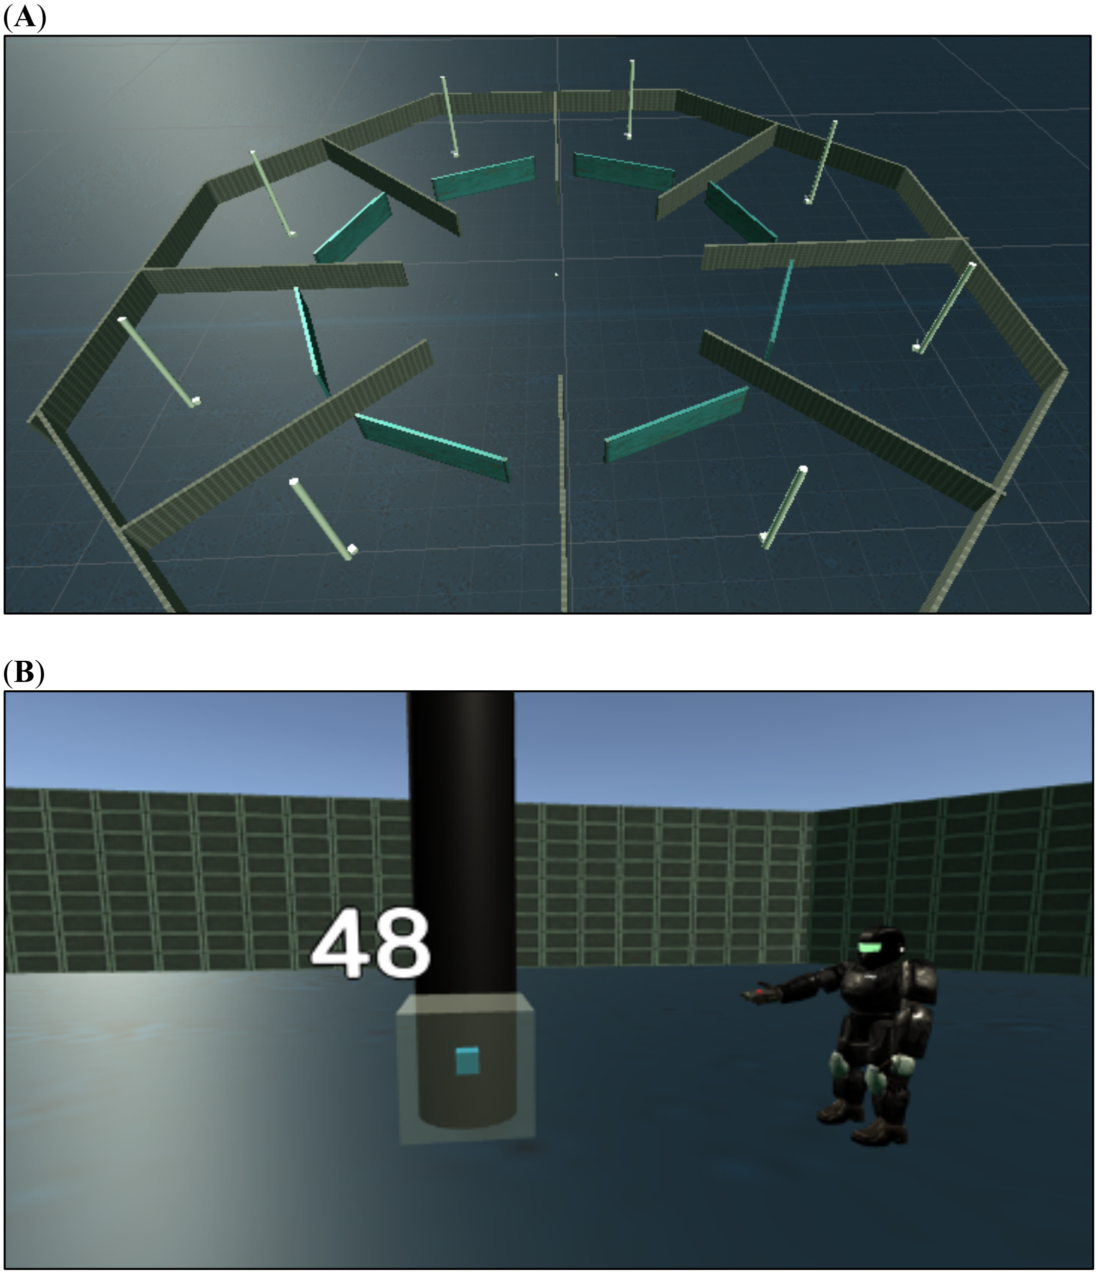

(**A**) An ariel view of the video game environment. (**B**) The player's first-person view of a target (the transparent cube enclosing the blue-colored opaque cube), the remaining rewards for the target (48 points in this case), its black-colored pole that visibly marks its location from a distance, and the computerized competitor (the black-colored robot avatar).

Participants started every level of the game in the center of the octagon and were tasked with collecting as many reward points as possible from the targets in each level; with five levels and eight targets per level, there were a total of 40 targets in the game from which to collect points. Each target was initially supplied with 100 points and a power value that determined the nature of its increasing reward magnitude over a 10 s interval. Reward magnitude was indicated by the size of a small, solid-colored box inside a larger transparent box that began at a size of zero and grew in physical size (according to Eq. 1) until it reached the size of the outer transparent box after 10 s. Clicking on the target at maximum size (after at least 10 s has passed) delivered 25 points and clicking on the target before its full size had been realized resulted in fewer points being disbursed; the exact number of points obtained was determined by the size of the inner, solid-colored box. Clicking on the target also reduced the target's total points remaining, and the target disappeared when its point-value reached zero. A level ended when all eight targets were destroyed.

The power value of the target's box determined the change in reward potential over a 10 s interval. The power parameter within the superellipsoid function determined the percent of maximum reward at any given time during the interval (Eq. 1), where the maximum amount of reward that could be collected per click is 25. To simplify the complexity of the task, the power value for all targets in all five levels were set to 0.75. Under normal circumstances with a power of 0.75, waiting the full 10 seconds maximizes the reward rate (see Figure 2). But under uncertainty, participants should now shorten their waiting times (Figure 5).

Competition in this EI task was generated by placing a single virtual avatar next to each of the eight targets (Figure 6B and Figure 7) that responded at a rate specified by its programmed distribution (Table 1). A competitor would not start responding to its assigned target until the participant had made the first click on that target. Whenever the competitor responded to a target, the target reset in magnitude and physical size for both the competitor and participant. A competitor's response was visually signaled by a yellow-colored line sent from a remote control placed in the competitor's hand to the center of the target for 333 milliseconds (see Figure 7). Whenever the participant responded to a target, the competitor's WT restarted at zero – i.e., the competitors' WTs are based on the amount of time that has passed since the target was previously clicked on by either the competitor or the participant. The competitors were removed in Level 1 to allow participants to acclimate themselves to the task and learn the power contingency before playing against the computerized avatars. Using Level 1 for training also allowed me to analyze whether behavior between the two competitor groups was statistically similar as a control check. The competitors were present in Levels 2-5 and the data in those levels were used in the analyses to test the hypotheses.

Figure 7

Screenshot of a Competitor Response

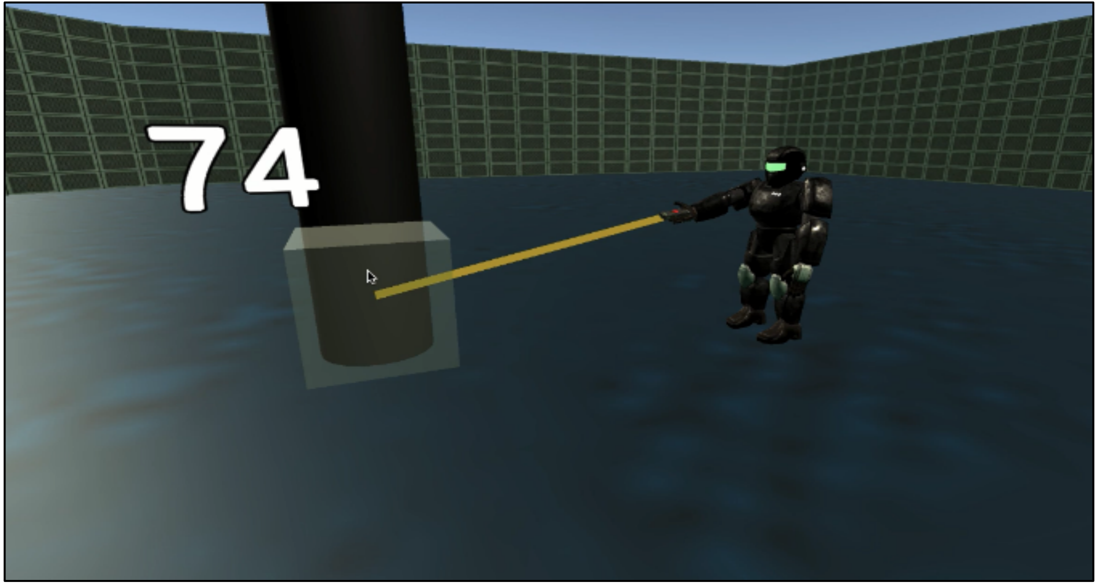

Participants experienced all four competitors across game levels 2-5, making Experiment 1 a 2×2×2 between-within-subject design, where the visibility of the competitor was a between-subject variable and the mean and variance of the competitors' WT distributions were fully within-subject variables. The competitor-types were identical within each level but changed between each level. For example, one competitor-type (e.g., impulsive-low variance) was assigned to all targets in Level 2, a different competitor-type (e.g., patient-high variance) was assigned to all targets in Level 3, and so forth for Levels 4 and 5. The presentation order of the competitor-types encountered between the four levels was Williams Latin square counterbalanced to account for order effects (see Table 2). The competitors were also assigned different colors[13] to clearly signal to participants that they were competing against a different opponent. Because the competitor-invisible group could not see the competitors, the color of the background poles also changed in the same way so that all participants had access to, and were aware of, the same information about the uncertainty contingencies.

Table 2

Williams Latin square groupings of the four competitors across the four game levels

Group	Level 2	Level 3	Level 4	Level 5
1	Impulsive-low	Impulsive-high	Patient-high	Patient-low
2	Impulsive-high	Patient-low	Impulsive-low	Patient-high
3	Patient-low	Patient-high	Impulsive-high	Impulsive-low
4	Patient-high	Impulsive-low	Patient-low	Impulsive-high

Note: 'Impulsive' and 'patient' refer to the means of the wait time distributions, and 'low' and 'high' refer to their variance. Level 1 is absent because there was no uncertainty/competition contingency.

---

[13] Impulsive-low variance = white; Impulsive-high variance = black; Patient-low variance = orange; Patient-high variance = green.

Procedure

Participants were first seated at a computer station in a well-lit lab room. They then read and signed the written consent form and were given verbal instructions related to the game and only broadly informed about the nature of responding to the targets at various WTs and color of the poles (e.g., 10 s interval, clicking earns points based on how long they wait). To preserve participants' naivete about the competition-based experimental manipulation, they were not initially informed about uncertainty in the task. Once participants finished the training in Level 1, they received written details that revealed the subsequent levels would have uncertainty in the target's growth. This information was presented as text on the computer screen so that no other participant could see it. Participants randomly assigned to the competitor-visible group (*n* = 25) were told that the competitor would also be responding to the target with WTs sampled from participants from previous studies. This cover story ideally created a stronger incentive for participants to outperform previous human players, consistent with prior competition research (van den Bos et al., 2008). Participants randomly assigned to the competitor-invisible group (*n* = 26) were simply told that the targets would occasionally malfunction and reset in size and magnitude. Participants could start Level 2 once they had read these details and were ready to continue. After Levels 3 and 4, participants saw another similar-looking screen that simply said, "You have finished this level, please click continue when you are ready to begin the next one". After Level 5, participants were told the game was over and their final point value was displayed. Participants in the competitor-visible group saw the points obtained by the competitor while participants in the competitor-invisible group instead saw the points that they had lost (without knowledge that the competitor had obtained them). Participants who obtained more points than the competitor won a prize of a small candy bar. They were broadly informed about this goal at the beginning of the study to incentivize task engagement but with no mention of competition; only that if they scored enough points they would win. Participants were then free to leave and were granted credit for participating.

Data Analysis

Data Rejection. There are several behaviors in EI tasks that warrant the rejection of individual responses or even an entire participant. First, previous research using EI tasks removes IRTs (WTs) greater than 20 seconds (Young et al., 2011; Young & Howatt, 2022) because this duration typically reflects travel times or inattention to the task since the rewards stops growing after 10 seconds (and participants are informed about this upper limit). Second, although IRTs (WTs) have a theoretical lower bound of zero, they tend to have an empirical lower bound of around 100 milliseconds due to the physical limitations of responding via a computer mouse. However, in the current EI task, participants can technically respond at the same time a competitor does or even more quickly than the usual 100 ms boundary. It is unlikely that these responses reflect true waiting times and instead are more likely to be responses that were prepared before the competitor responded but executed after the competitor responded. Therefore, I removed any WTs greater than 20 s (2.1%) and less than 100 ms (2.8%).

Third, participants' WTs for their first response to each target are uninformative about their preferences because the target will almost always be at its maximum size (Young & Howatt, 2023). Participants' responses were coded as "first", "second", "third" or "other" to represent their ordering (where "other" means any fourth consecutive click or greater), and all responses coded as 'first' were removed (10%). Importantly, the ordering of the responses would reset if a participant switched targets so that participants could have multiple instances of "first", "second" or even "third" click order responses recorded if they switched between active targets.

Fourth, recent work on the EI task (Young & Howatt, 2022) has identified instances in which participants shift between reward targets that are still active, where they click on one target and move to another while the first target resets in size and begins growing again, and so forth. Switching targets removes the need to wait for the current target to reach maximal size and participants can thus use that time to allow the previous target to regrow while traveling to another target, completing the task more quickly for cases in which the target's power value is less than 1.0 (and waiting is optimal). The downside of this strategy in this EI task is that clicking on a target and then switching to another will cause the competitor to begin responding to that target and depleting its points. Although this strategy actually results in participants finishing the game faster, it also causes them to lose the game. Moreover, this strategy results in participants not consecutively interacting with the same target, so that it is more difficult to model how they learn waiting times on the same target because they produce more "first" clicks that are eventually removed. I used visual inspection of participants' raw data to observe approximately how often participants were switching and generated usable data (i.e., responses with click orders greater than "first") and implemented a cutoff of about 100 total responses across Levels 2-5 that were "second", "third" or "other" to determine if a participant should be removed.

The percentage of data removed did not differ between the two competitor-visibility groups, suggesting that there was no differential attrition in participants based on the experimental conditions.

Statistical Models. All analyses were conducted within a Bayesian framework using the *brms* package (Bürkner, 2017) in R (version 4.0.0). Because Bayesian analyses do not provide *p*-values for the regression weights, I measured the degree of evidence in favor of each parameter's estimated effect size using *Evidence Ratios* (ERs) with the *hypothesis* function. Evidence ratios correspond to the ratio of the proportion of the posterior distribution that is of the parameter estimate's sign to the proportion that is not. For example, if a regression weight is estimated to be 2.0 with an estimated error of 1.0, then 97.75% of the posterior distribution is greater than zero, 2.25% of the posterior distribution is less than zero, and the evidence ratio would be about 43 (97.75 / 2.25). If a regression weight is estimated to be -1.0 with an estimated error of 3.0, then 63.05% of the posterior distribution is less than zero, 37.95% of the posterior is greater than zero, and the evidence ratio would be 1.66 (0.6305 / 0.3795). Evidence ratios of 1.0 correspond to a regression weight centered around the hypothesized value, and if the hypothesized value is zero, then this indicates little if any evidence that the predictor is meaningfully associated with the outcome variable. I interpret evidence ratios greater than 3 to indicate positive evidence in favor of a hypothesis and ratios greater than 20 to indicate strong evidence, similar to Bayes factors (Jarosz & Wiley, 2014).

Estimates of the marginal response means and contrast tests were conducted using the *emmeans* package (Lenth et al., 2018). Model comparisons were conducted using WAIC values, which provides estimates of out-of-sample predictions (Vehtari, Gelman & Gabry, 2017). Models with lower values are predicted to fit better and simpler models are usually preferred when WAIC values are similar.

Probability of Waiting. Participants' waiting times in EI tasks tend to be highly bimodal due to the power values incentivizing responding rapidly or waiting 10 s (particularly for power values less than 1.0). However, bimodal data are analytically difficult to work with because they cannot be modeled with traditional statistical distributions (e.g., normal, gamma, Weibull). Specifying a mixture model is one possible solution, but their psychological interpretability is limited. Therefore, participants' waiting times were dichotomized (Young et al., 2011; Young & Howatt, 2021) so that values less than or equal to 5 seconds were classified as 'did not wait (i.e., impulsive choice) and values greater than 5 seconds were classified as 'waited' (i.e., patient choice).

Using the dichotomized waiting times (WTs), I conducted a binomial logistic regression that predicted participants' log odds (i.e., relative likelihood) of waiting to respond as a function of the competitor's visibility (visible or invisible), mean WT (impulsive vs. patient), and WT variance (lower variance vs. higher variance). Positive parameter estimates correspond to higher likelihoods of observing participants waiting to respond and negative parameter estimates correspond to lower likelihoods. The log-odds parameter estimates can then be combined and back-transformed into outcome probabilities to make them more interpretable. The competitor's visibility, mean WT, and WT variance were all treated as categorical variables and then effect-coded to reduce multicollinearity and make the model intercept represent participants' grand-mean probability of waiting to respond. To preserve the regression model's degrees of freedom and statistical power, only the main effects and 2-way interactions of the predictors were tested. Additional models that included the 3-way interactions, game level, ordering of the competitors, and time spent in each level may be used in exploratory analyses to generate hypotheses in future studies.

Because participants repeatedly responded to targets in the task and experienced competitors with multiple levels of the mean and variance of their WTs, their responses are expected to be correlated, which violates the statistical assumption of independent observations. To account for the dependencies in participants' responses, I ran a multilevel (repeated measures) logistic regression (Bolker, 2015). I tested various random-effect structures to assess whether participants exhibited differential sensitivity to the competitor's behavior, with model complexity and convergence, WAIC values, and theoretical implications guiding the final model. The best fitting random effect structure was one that allowed the model intercept and slopes of the competitors' mean WT and WT variance to vary across participants as main effects. These random effects in essence allow the model to estimate participants' differences in their average probability of waiting to respond and how they further change depending on the mean and variance of the competitors' WTs.

For the Bayesian components of the regression model, I set a burn-in period of 2500 iterations with an additional 2500 saved iterations to estimate the posterior distribution of each parameter; four chains were run to derive the parameter estimates. I specified weakly informed priors of *N*(0, 2) for the intercept and all other regression weights, and priors of *Cauchy*(0, 0.5) for the standard deviations (i.e., random effects for individual differences) and *LKJ*(2) for the random effect correlations. The intercept prior initially assumes that participants will be equally likely to wait to respond as not wait, but allows for considerable deviation from this estimate (95% of the means will range from 2% to 98% probability of waiting). The priors for the fixed effect regression weights initially assume that there is no difference in participants' probability of responding as a function of the competitor's visibility, average WT, and WT variance, but still allows for differences in either directions between the groups. To test for sensitivity to these priors, a variety of similar values were evaluated, such as being informed by the simulations presented in Figure 4.

Probability of Reward. Participants' probability of obtaining rewards was estimated using a multilevel binomial logistic regression as a function of the competitor's visibility, average WT, and WT variance and their 2-way interactions. I used the same random effect structure as the analysis predicting participants' probability of waiting to respond, even though a more complex structure that allowed the interaction between the competitor's mean WT and WT variance to vary across participants fit better. This more complex random effect structure is unlikely to generalize to future samples, and model parsimony and consistency with the first analysis was preferred. The competitor's visibility, average WT, and WT variance were also effect-coded to reduce multicollinearity and make the model intercept represent participants' grand-mean probability of obtaining rewards. The Bayesian components of this model were identical to the analysis measuring participants' probability of waiting.

Results and Discussion

Five participants were removed (competitor-visible = 2, competitor-invisible = 3) from the sample due to excessive responding or switching between the reward targets. Excessive responding and switching prevents the model from determining these participants' sensitivity to the experimental contingencies and reflects either a lack of task understanding of or disengagement from the task (or having ulterior goals of finishing the game as quickly as possible without consideration of the task itself). For example, one participant generated 2,500 clicks in Level 1 and 8,500 clicks in Levels 2-4 which resulted in them not finishing the game because they were consistently responding so rapidly. Consequently, this participant never waited to respond, even in Level 1 without competitors or in levels against patient competitors. Moreover, this participant's behavior greatly influenced the regression models for Levels 2-5 and was responsible for artifactual effects between the competitor-visibility groups. The other four participants engaged in excessive switching behaviors, only generating about 15-100 “real” data points that were usable in the analysis. In fact, one participant always switched and did not generate any consecutive responses on any single target, so their data was completely absent.

Three participants did not finish the game, but their behaviors were consistent with the task’s goals up to that point so they were retained because multilevel regressions can appropriately model incomplete and imbalanced data.

All Bayesian regression models presented below converged and visual inspection of the posterior distributions, R-hat values (all below 1.05), sampling chains and ESSs did not reveal any concerns about the reliability of the parameter estimates.

Descriptive Statistics and Control Check

As expected, participants’ raw WTs were highly bimodal across all five levels, even against impulsive competitors (Figure 8). However, the dichotomization of WTs as ‘waited’ (WTs > 5 s) and “did not wait’ (WTs ≤ 5 s) resulted in participants almost never waiting against impulsive competitors. Such a strong effect of the impulsive competitor could possibly create analytical difficulties and are discussed further in Appendix C. Out of the 51 participants in the final sample, 42 won the game by scoring more points than the competitor. In the competitor-visible group 22/25 participants won the game, and in the competitor-invisible group 20/26 participants won the game. A Pearson’s Chi-squared test indicated that these proportions were not significantly different (X^2 (1) = 1.07, $p = 0.29$), suggesting that the competitor's visibility did not differentially impact participants' likelihood of beating the competitor.

The maximum number of points that participants could obtain across Levels 2-5 was 3200 (eight targets, each worth 100 points, in each of the four levels). The median number of points obtained by participants in the competitor-visible group was 2660 (*Median Absolute Deviation* = 398) and the median number of points obtained by participants in the competitor-invisible group was 2281 (*Median Absolute Deviation* = 616). To test whether the competitor-visibility groups differed in their total points obtained, I conducted a non-parametric Wilcoxon rank sum exact test and found that the difference between the groups' rank scores just reached statistical significance ($W = 220$, $p = 0.048$), suggesting that participants in the competitor-visible group obtained more total points than the competitor-invisible group. Figure A.1 in Appendix A illustrates these distributions.

Figure 8

Participants' Wait Time Distributions

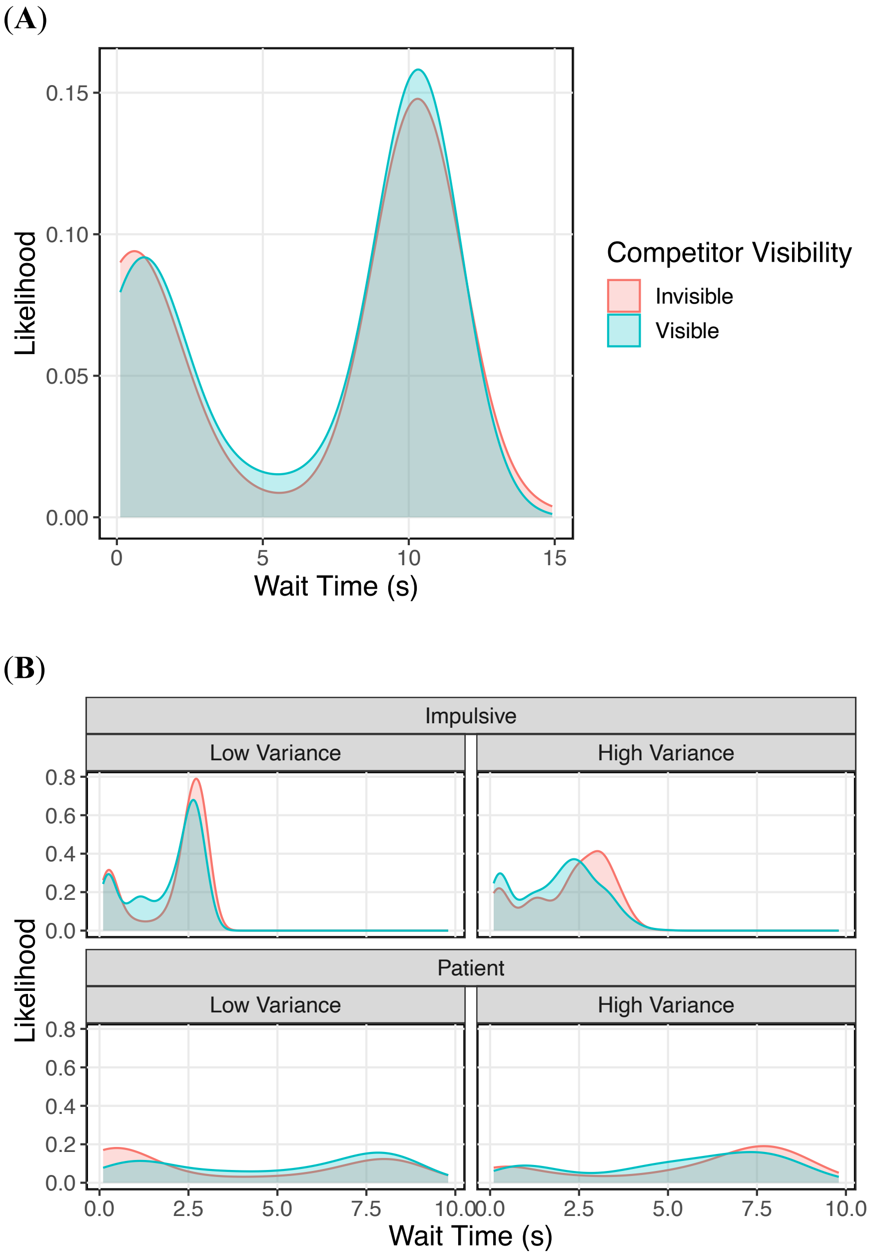

Note: Empirical density plots of participants' wait times (WTs) in Experiment 1 as a function of the competitor's visibility. The data have been downsampled to provide more equal weighting across the participants. (**A**) Participants' WTs in Level 1, with a downsampled frequency of 10 responses. (**B**) WT distributions split by the competitors' mean WT (impulsive vs. patient) and WT variance for Levels 2-5, with a downsampled frequency of 100 responses.

To ensure that any effects observed in the main analyses would not be driven by inherent differences in behaviors across the two competitor-visibility groups rather than the experimental manipulation, I conducted a multilevel logistic regression that predicted participants' probability of waiting to respond to reward targets only in Level 1 when the competitor was absent. Results revealed no statistical difference in participants' waiting behaviors across the two competitor-visibility groups (B = -0.16, SE = 0.55, z = -0.29, p = 0.76). Participants in the competitor-visible group were waiting about 72% of the time on average whereas participants in the competitor-invisible group were waiting about 65% of the time on average. I also explored the degree to which participants' waiting behaviors changed over the course of Level 1 and observed that they became more patient as their time in Level 1 increased (B = 4.76, SE = 0.57, z = 8.31, p < .001) and that there was no statistical difference in this change across the two competitor-visibility groups (B = 0.51, SE = 0.57, z = 0.89, p = 0.37). Figure B.1 in Appendix B illustrates these results.

Probability of Waiting

Figure 9 illustrates participants' model-estimated log odds (relative likelihoods) of waiting to respond against the competitors in Levels 2-5. Figure 9A shows participants' log odds of waiting as a function of the competitor's visibility and there does not appear to be any meaningful difference between these two groups. Figure 9B shows how participants' log odds of waiting across the two competitor-visibility groups changed based on the competitor's mean WT and WT variance, and there is a clear difference based on whether the competitor they faced was impulsive or patient, but not necessarily whether the competitor had lower or higher variance in their WTs. Note that I am keeping the plots of the model fits on the log-odds scale for visualization because back-transforming them into probabilities would compress them too closely to zero due to the large floor effect. The estimated log odds have been back-transformed into probabilities and provided in the main text for comprehension.

Figure 9

Participants' Likelihood of Waiting to Respond

(A)

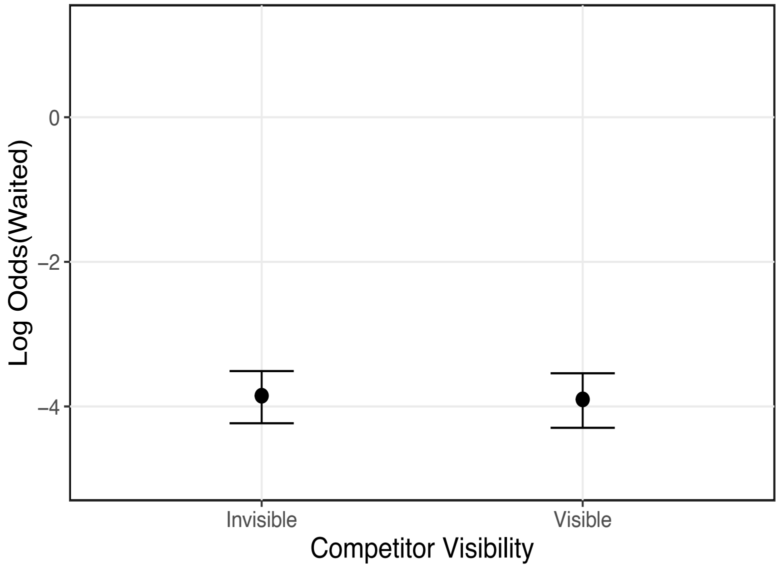

(B)

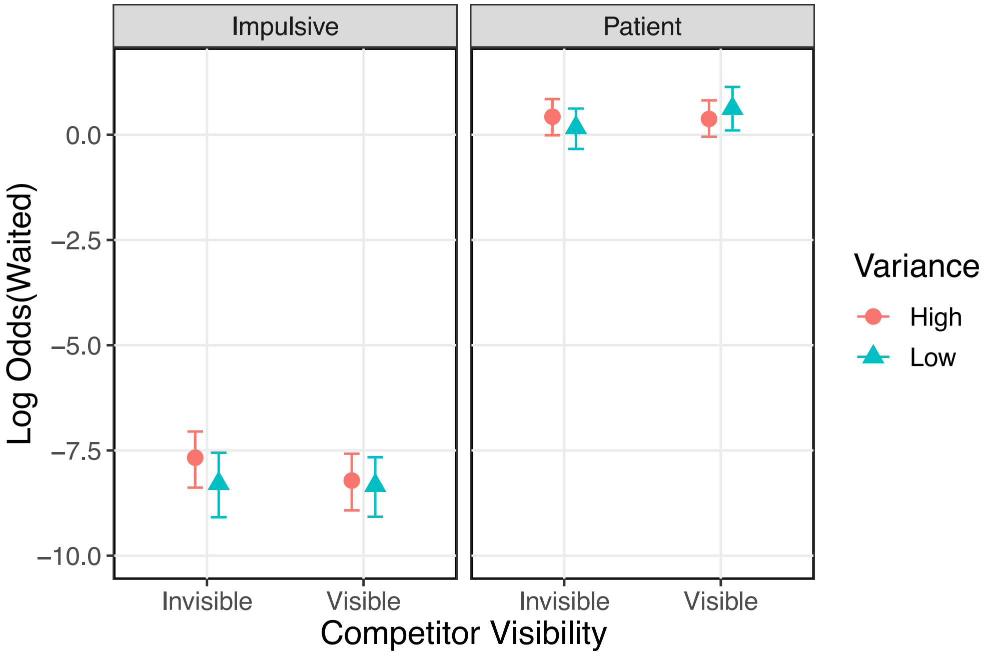

Note: Participants' model-estimated log odds of waiting to respond to the reward target. (**A**) The log odds of waiting split across the competitor-visibility groups. (**B**) Log odds estimates further split across the competitors' mean WT (impulsive vs. patient) and WT variance (low vs. high variance). Error bars represent ±1 SE of the mean.

Results from the regression model confirmed these observations (see Table 3), with weak evidence (*ER*s < 3) for the main effects of competitor visibility and WT variance, and decisive evidence (*ER* > 1000) that participants were less likely to wait against impulsive competitors (*M* < 1%) than patient competitors (*M* = 60%). There was positive evidence of a competitor-visibility × variance interaction (*ER* = 7.8), such that participants were slightly more likely to wait to respond when the low variance competitors were visible (*M* = 2.0%) compared to when they were invisible (*M* = 1.7%), but participants were slightly less likely to wait when the high variance competitors were visible (*M* = 1.9%) compared to when they were invisible (*M* = 2.6%); these differences are fairly small, though. Finally, there was only weak evidence (*ER*s < 3) for a competitor-visibility × mean WT interaction and a mean WT × WT variance interaction.

Table 3

Experiment 1's parameter estimates, 95% credible intervals and evidence ratios for the multilevel logistic regression predicting participants' log odds of waiting to respond.

	B	95% CI	*ER*
Intercept	-3.88	[-4.41, -3.35]	> 1000
Visibility	0.02	[-0.51, 0.52]	1.2
Mean WT	-4.28	[-4.84, -3.74]	> 1000
WT Variance	0.09	[-0.32, 0.52]	2.1
Visibility × Mean WT	0.12	[-0.33, 0.60]	2.4
Visibility × WT Variance	0.12	[-0.08, 0.33]	7.8
Mean WT × WT Variance	0.09	[-0.31, 0.49]	2.0

Note: Visibility was effect-coded as [Invisible = +1; Visible = -1], Mean WT was effect-coded as [Impulsive = +1; Patient = -1], and WT Variance was effect-coded as [High Variance = +1, Low Variance = -1].

Unexpectedly, model comparisons using WAIC did not support the model reported above (WAIC = 3819) over a reduced model that only included the main effect of the competitor-visibility predictor (WAIC = 3808), nor was a reduced model that only included the competitor's mean WT (WAIC = 3118) better supported. This result was unexpected given that the effect of the competitor's visibility on participants' waiting behaviors was meager, but the effect of the competitor's mean WT on waiting was extremely strong. It seems clear that a competitor's mean WT exerted stronger control over participants' behavior than its visibility, yet the model fits were inconsistent with this finding. The reasons for this discrepancy are likely due to the random effect structure and are discussed further in Appendix C, but ultimately do not change the weak evidence found for the competitor's visibility increasing participants' impulsive choices.

Probability of Reward

Figure 10 illustrates participants' model-estimated log odds of obtaining rewards in Levels 2-5. Figure 10A shows participants' reward probabilities as a function of the competitor's visibility and there does appear to be a slight difference between these two groups, specifically that participants were more likely to obtain rewards against visible competitors than invisible competitors. Figure 10B shows how participants' reward probabilities across the two competitor-visibility groups changed based on the competitor's mean WT and WT variance, and the effect of the competitor's visibility seems to be stronger against more impulsive competitors than patient ones. Moreover, there appears to be an interaction effect between competitors' mean WT and WT variance, such that participants were more likely to obtain rewards against competitors with higher variance, but only when those competitors were also patient.

Figure 10

Participants' Likelihood of Reward

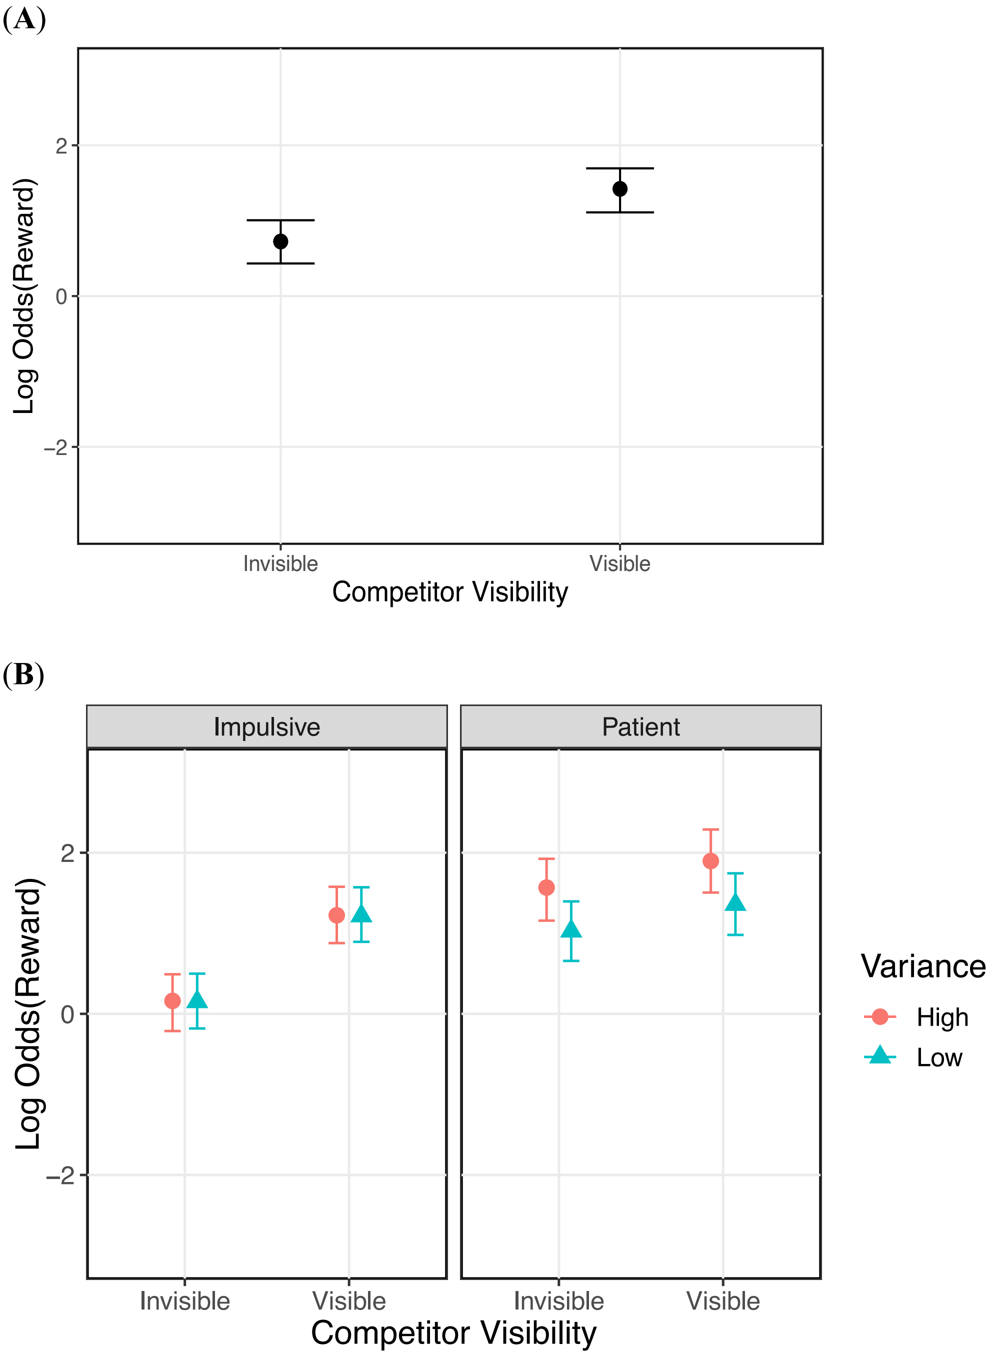

Note: Participants' model-estimated log odds of obtaining the reward. (**A**) The log odds of reward split across the competitor-visibility groups. (**B**) Log odds estimates further split across the competitors' mean WT (impulsive vs. patient) and WT variance (low vs. high variance). Error bars represent ±1 SE of the mean.

Results from the regression model (see Table 4) mostly confirmed these observations, with positive-to-strong evidence (*ER* = 19) that participants were more likely to obtain rewards against visible competitors (*M* = 81%) than invisible competitors (*M* = 67%), positive evidence (*ER* = 7) that participants were more likely to obtain rewards against competitors with higher WT variance (*M* = 77%) than lower WT variance (*M* = 71%), and decisive evidence (*ER* > 1000) that participants were more likely to obtain rewards against patient competitors (*M* = 81%) than impulsive competitors (*M* = 66%). Additionally, there was good evidence of a competitor-visibility × mean WT interaction (*ER* = 23), such that participants were more likely to obtain rewards against impulsive competitors that were visible (*M* = 77%) than invisible (*M* = 53%), but participants were equally likely to obtain rewards against patient competitors regardless of whether they were visible (*M* = 83%) or invisible (*M* = 79%). Finally, there was also decisive evidence of a mean WT × WT variance interaction (*ER* > 1000), such that participants were more likely to obtain rewards against patient competitors with higher variance (*M* = 85%) than lower variance (*M* = 76%), but participants were equally likely to obtain rewards when playing against impulsive competitors regardless of whether their WT distributions had lower variance (*M* = 66%) or higher variance (*M* = 66%).

Table 4

Experiment 1's parameter estimates, 95% credible intervals and evidence ratios for the multilevel logistic regression predicting participants' log odds of obtaining the reward.

	B	95% CI	*ER*
Intercept	1.07	[0.66, 1.49]	> 1000
Visibility	-0.34	[-0.77, 0.06]	19
Mean WT	-0.39	[-0.59, -0.18]	> 1000
WT Variance	0.13	[-0.10, 0.36]	7
Visibility × Mean WT	-0.18	[-0.39, 0.02]	23
Visibility × WT Variance	-0.01	[-0.23, 0.23]	1
Mean WT × WT Variance	-0.13	[-0.18, -0.08]	> 1000

Note: Visibility was effect-coded as [Invisible = +1; Visible = -1], Mean WT was effect-coded as [Impulsive = +1; Patient = -1], and WT Variance was effect-coded as [High Variance = +1, Low Variance = -1].

Even though there was good evidence that the competitor's visibility affected participants' probability of obtaining rewards, model comparisons still supported a reduced model that only included the competitor's mean WT and WT variance as predictors (WAIC = 23,521) over the more complex model reported above that included the competitor's visibility (WAIC = 23,520) due to being simpler. This suggests that the effects of the competitor's visibility may not be strong enough to warrant good out-of-sample generalizability.

Overall, these results provide mild evidence that the visibility of competitors influenced participants' waiting behaviors, mostly by increasing participants' probability of obtaining rewards. Specifically, there was little if any evidence that participants who could see the competitors were less likely to wait to respond, but stronger evidence that these participants were more likely to actually obtain the rewards. These results are more consistent with the simple subjective utility perspective that predicts individuals ignore social factors and are only sensitive to the reward uncertainty. But, I believe the results are also partially consistent with the evolutionary hypothesis that competition increases impulsive choices because the recognition of competitors as a source of uncertainty (social uncertainty) sensitizes decision-makers to relative performance and motivates them to increase their reward rate against competitors.

Nonetheless, the discrepancy between the two analyses could present an interesting interpretation of how competition may differentially influence waiting behaviors and actual performance. Statistically small differences in waiting behaviors could lead to statistically larger outcomes, particularly when behaviors on a continuum like WT are dichotomized compared to a naturally dichotomous outcome of participants obtaining a reward or not. For example, consider two participants playing against an impulsive competitor (mean WT = 3 s): Participant A, who typically responds every two seconds, and Participant B, who typically responds every four seconds. Even though both participants will be estimated to have a nearly identical probability of waiting to respond, Participant A will almost certainly have a higher rate of obtaining rewards. Ultimately, organisms' survival and fitness depend on how well they perform in challenging environments, and performance may be a better metric in this particular case. Modeling participants' raw WTs may be necessary to fully reveal participants' waiting behaviors, but this goal remains elusive until better statistical theory is worked out to improve accessibility to mixture models in regressions.

Although the visibility of an adversary directly signals competition, this knowledge may not be strong enough to reliably control individuals' waiting behaviors and performance. Therefore, Experiment 2 examined the effects of another dimension of competition-based relative processes, knowledge about how many resources the competitor has in relation to the organism – i.e., the competitor’s current point total.

Chapter 5 - Experiment 2

Experiment 2 tested a second manipulation of competition and relative performance by manipulating the visibility of the competitor's cumulative rewards (Messick & McClintock, 1968). Here, all competitors were visible, but now a scoreboard was present at the bottom of the computer monitor screen that displayed the participant's and competitor's current point total. Participants were randomly assigned to one of two groups: 1) a *score-visible* group, where participants could see the competitor's point total, and 2) a *score-invisible* group, where participants could not see the competitor's point total.

I again generated two specific hypotheses relevant to the experimental manipulation. The first hypothesis regards participants' rate of impulsive choice. The simple subjective utility hypothesis, Hypothesis 1a, predicts that participants in the score-visible group will be equally likely to wait to respond as participants in the score-invisible group. In contrast, the evolutionary hypothesis, Hypothesis 1b predicts that participants in the score-visible group will be less likely to wait to respond than participants in the score-visible group.

The second hypothesis regards participants' rate of obtaining rewards. The simple subjective utility hypothesis, Hypothesis 2a, predicts that participants in the scoreboard-visible group will be equally likely to obtain rewards as participants in the score-invisible group. In contrast, the evolutionary hypothesis, Hypothesis 2b predicts that participants in the score-visible group will be more likely to collect rewards than participants in the score-visible group, in part due to behaving more impulsively.

Methods

Participants

A total of 52 individuals originally participated in this experiment for research credit as part of a course requirement. The study was approved by Kansas State University’s IRB, and participants provided written informed consent prior to their participation. Three participants were removed for excessive switching and one participant was removed for being underage. The final sample was therefore 48 adults (22 female, 25 male, 1 did not respond) aged 18-30 years old (*M* = 20).

Design

Experiment 2’s general design and procedure were almost identical to Experiment 1, except that competitors were now visible to all participants and a scoreboard displayed the participant’s and competitor’s cumulative points at the bottom of the computer screen (see Figure 13). The visibility of the competitor’s score was manipulated in a between-subject manner. Half of the participants (*n* = 24) were shown a scoreboard displaying both their own accumulated points and the competitor’s accumulated points as slider bars, whereas the other half of participants (*n* = 24) could not see the competitor’s accumulated points. The slider bars were used to ease participants’ comprehension of the points and their difference without evaluating the symbolic numbers. All participants were again told that the competitors’ WTs were drawn from participants in previous studies. This cover story ideally created a stronger incentive for participants to outperform previous human players when they can see the number of points previous players would be receiving from their response times, consistent with some prior competition research in economic games (Messick & McClintock, 1968).

Figure 11

EI Task Scoreboard

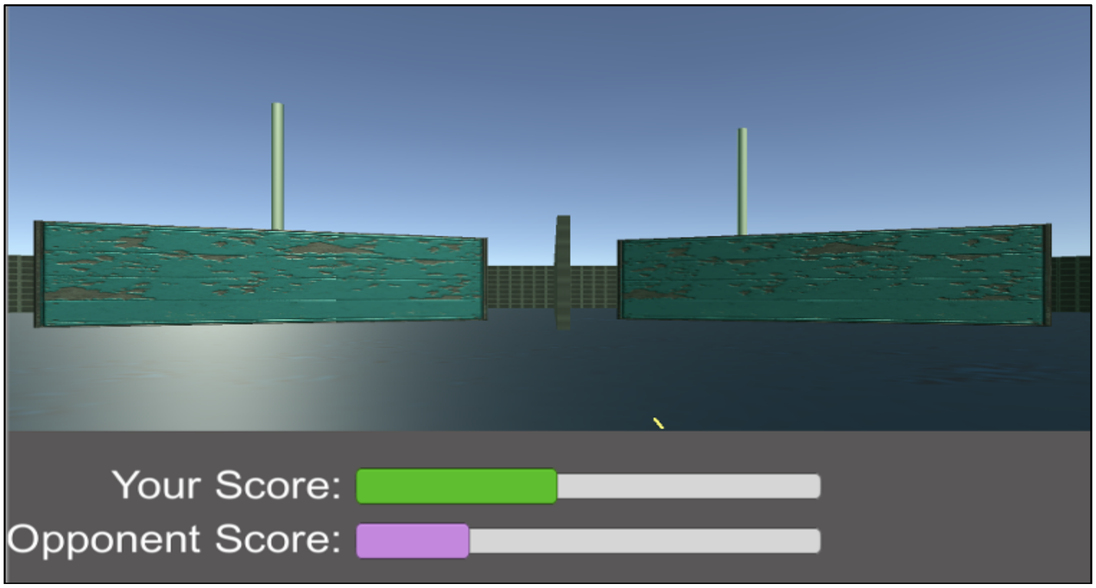

Note: Participant’s perspective of the game with the scoreboard. For half of the participants, the ‘Opponent Score’ text and points slider were invisible.

Data Analyses

The data preprocessing procedure and statistical analyses were identical to Experiment 1. Because the experimental manipulations in Experiment 2 are conceptually similar but methodologically different than Experiment 1, the priors on the parameter estimates in Experiment 2 were informed by Experiment 1’s results, but with some additional uncertainty. Specifically, for the analysis estimating participants’ probability of waiting to respond, I set more informed priors of $N(3.5, 1)$ on the intercept, $N(-4, 0.5)$ on the mean WT, $N(0, 0.5)$ for the competitor’s score visibility, $N(0.1, 0.5)$ for all other regression weights. For the analysis estimating participants’ probability of obtaining the reward, I set more informed priors of $N(1.0, 0.5)$ on the intercept, $N(-0.4, 0.5)$ on the mean WT, $N(-0.3, 0.5)$ for the competitor’s score visibility, $N(0, 0.5)$ for the score-visibility by WT variance interaction, and $N(0.1, 0.5)$ for all other regression weights. I used priors of *Cauchy*(0, 0.5) for the standard deviations (i.e., random effects for individual differences) and *LKJ*(2) for the random effect correlations in both analyses, the same as in Experiment 1.

I again removed any waiting times greater than 20 s (1.9%) and less than 100 ms (2.1%), and all responses that were coded as ‘first’ (9%). The percentage of data removed did not meaningfully differ between the two score-visibility groups.

Results and Discussion

Four participants were removed (score-visible = 2, score-invisible = 2) from the original sample due to excessive switching between the reward targets. One additional participant did not finish the game, but their behavior was consistent with the task’s goals up to that point, so their data were retained.

All Bayesian regression models presented below converged and visual inspection of the posterior distributions, R-hat values (all below 1.05), sampling chains and ESSs did not reveal any concerns about the reliability of the parameter estimates.

Descriptive Statistics and Control Check

Participants’ raw WTs were again highly bimodal (see Figure 12) across all five levels, so WTs were dichotomized as either “waited” (WT > 5 s) or “did not wait’ (WT <= 5 s). Out of the 48 participants in the final sample, 41 won the game by scoring more points than the competitor. In the scoreboard-visible group 21/24 participants won the game, and in the competitor-invisible group 20/24 participants won the game. A Pearson’s Chi-squared test indicated that these proportions were not significantly different (X^2 (1) = 0.16, p = 0.68), suggesting that the visibility of the competitor’s score did not differentially impact participants’ likelihood of beating the competitor.

The maximum number of points that participants could obtain across Levels 2-5 was again 3200. The median number of points obtained by participants in the score-visible group was 2401 (*Median Absolute Deviation* = 636) and the median number of points obtained by participants in the competitor-invisible group was 2550 (*Median Absolute Deviation* = 495). To test whether the score-visibility groups differed in their total points obtained, I conducted a non-parametric Wilcoxon rank sum exact test and found that the difference between the groups' rank scores was not statistically significant ($W = 278$, $p = 0.846$), suggesting that participants were obtaining an equal number of total points. Figure A.1 in Appendix A illustrates these distributions.

Figure 12

Participants' Wait Time Distributions

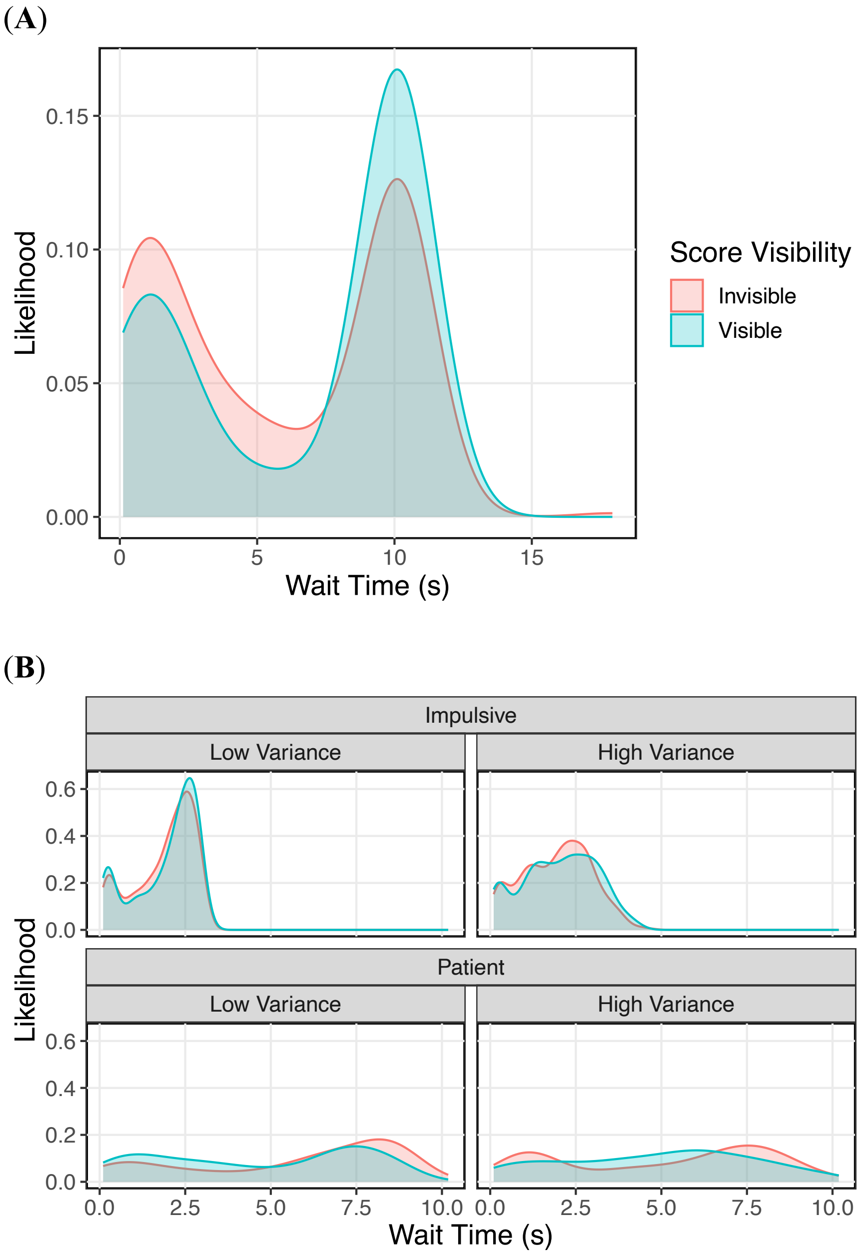

Note: Empirical density plots of participants' WTs in Experiment 2 as a function of the competitor's score visibility. The data have been downsampled to provide more equal weighting across the participants. (**A**) Participants' WTs in Level 1, with a downsampled frequency of 10 responses. (**B**) WT distributions split by the competitors' mean WT (impulsive vs. patient) and WT variance for Levels 2-5, with a downsampled frequency of 100 responses.

To ensure that any effects observed in the main analyses will not be driven by inherent differences in behaviors across the two scoreboard-visibility groups, I conducted a multilevel logistic regression that predicted participants' probability of waiting to respond to reward targets for only Level 1. Results revealed no statistical difference in participants' waiting behaviors across these two groups (*B* = -0.42, *SE* = 0.29, *z* = -1.47, *p* = 0.14). Participants who could see the competitor's score were waiting about 73% of the time whereas participants who could not see the competitor's score were waiting about 55% of the time. Although there is considerable uncertainty in these estimates, their difference still seems fairly large. I also explored the degree to which participants' waiting behaviors changed over Level 1 and observed that they became more patient as their time in Level 1 increased (*B* = 3.06, *SE* = 0.40, *z* = 7.65, *p* < .001) and that there was no statistical difference in this learning across the two competitor-visibility groups (*B* = -0.56, *SE* = 0.39, *z* = -1.41, *p* = 0.16). These results show that participants were sensitive to the power contingency and learned to adjust their behavior over time, but participants in the score-visible group may have been more patient from the start. Figure B.2 in Appendix B illustrates these results.

Probability of Waiting

Figure 13 illustrates participants' model-estimated log odds of waiting against the competitors in Levels 2-5. Figure 13A shows participants' waiting probabilities as a function of the competitor's score visibility and there does not appear to be any meaningful difference between these two groups. Figure 13B shows how participants' waiting probabilities across the two score-visibility groups changed based on the competitor's mean WT and WT variance and there is a clear difference in participants' likelihood of waiting based on whether the competitor they faced was impulsive or patient, but not necessarily whether the competitor had lower or higher variance in their WTs.

Figure 13

Participants' Likelihood of Waiting to Respond

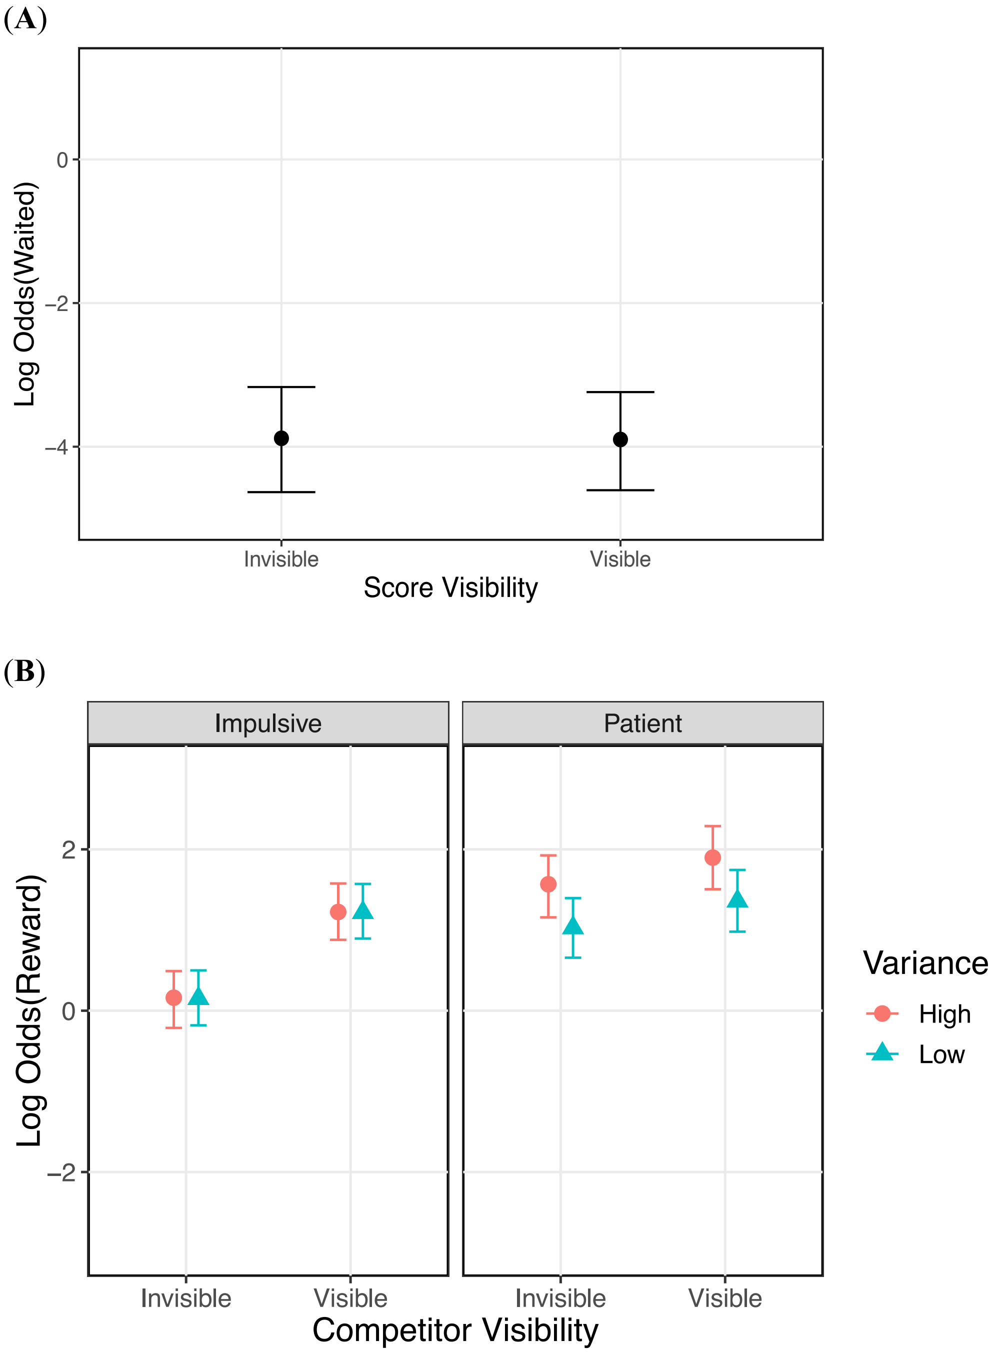

Note: Participants' model-estimated log odds of waiting to respond to the reward target in Experiment 2. (**A**) The log odds of waiting split across the competitors' score-visibility groups. (**B**) Log odds estimates further split across the competitors' mean WT (impulsive vs. patient) and WT variance (low vs. high variance). Error bars represent ±1 SE of the mean.

Results from the regression model (see Table 5) confirmed these observations, with weak evidence (*ER*s < 3) for the main effects of the competitor's score visibility and WT variance, and decisive evidence (*ER* > 1000) that participants were less likely to wait against impulsive competitors (*M* < 1%) than patient competitors (*M* < 59%). There was also only weak evidence (*Er*s < 3.5) for the score-visibility × mean WT interaction, the score-visibility × WT variance interaction, and the mean WT × WT variance interaction.

Table 5

Experiment 2's parameter estimates, 95% credible intervals and evidence ratios for the multilevel logistic regression predicting participants' log odds of waiting to respond.

	B	95% CI	*ER*
Intercept	-3.90	[-4.42, -3.40]	> 1000
Visibility	0.01	[-0.49, 0.50]	1.1
Mean WT	-4.28	[-4.74, -3.82]	> 1000
WT Variance	0.13	[-0.28, 0.55]	2.6
Visibility × Mean WT	-0.16	[-0.60, 0.27]	3.3
Visibility × WT Variance	0.01	[-0.20, 0.22]	1.1
Mean WT × WT Variance	0.06	[-0.32, 0.45]	1.6

Note: The competitor's Score Visibility was effect-coded as [Invisible = +1; Visible = -1], Mean WT was effect-coded as [Impulsive = +1; Patient = -1], and WT Variance was effect-coded as [High Variance = +1, Low Variance = -1].

Similar to Experiment 1, model comparisons using WAIC did not support the more complex model reported above (WAIC = 4747) over a reduced model that only included the main effect of the competitor-visibility predictor (WAIC = 4737), nor was a reduced model with only the competitor's mean WT as a predictor better supported (WAIC = 4746). The competitor’s mean WT clearly exerted stronger control over participants’ behavior than the visibility of their scores, yet the model fits were again inconsistent with this finding. But, this finding seemed be driven by the random effect structure again (see Appendix C) and testing additional models without the random slope effect of the competitor’s mean WT ameliorated this inconsistency without changing the model’s results.

Probability of Reward

Figure 14 illustrates participants’ model-estimated log odds of obtaining rewards in Levels 2-5. Figure 14A shows participants’ reward probabilities as a function of the competitor’s score visibility and there does not appear to be a difference between these two groups, unlike in Experiment 1. Figure 14B shows how participants’ reward probabilities across the two score-visibility groups changed based on the competitor’s mean WT and WT variance, and the effect of the competitor’s score visibility no longer seems dependent on the competitors’ mean WT or WT variance. Participants were again less likely to obtain rewards against impulsive competitors and possibly more likely to obtain rewards against competitors with lower variance (contrary to Experiment 1), but only when they were also patient.

Figure 14

Participants' Likelihood of Reward

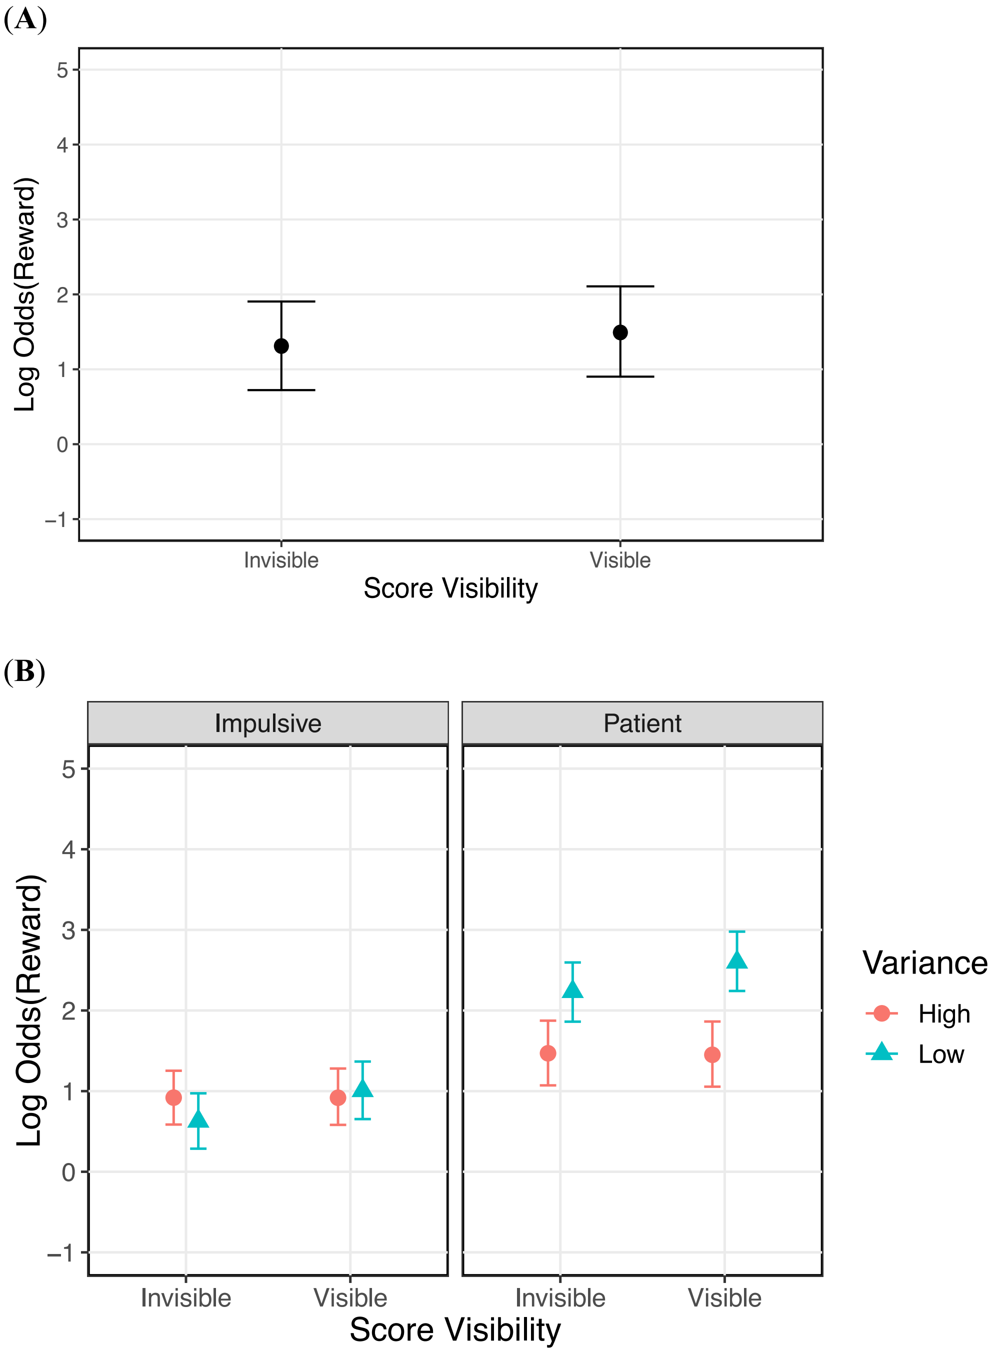

Note: Participants' model-estimated log odds of obtaining the reward in Experiment 2. (**A**) The log odds of reward split across the competitor score-visibility groups. (**B**) Log odds estimates further split across the competitors' mean WT (impulsive vs. patient) and WT variance (low vs. high variance). Error bars represent ±1 SE of the mean.

Results from the regression model (see Table 6) confirmed these observations, with weak evidence (*ER* = 1.6) that participants were more likely to obtain rewards when the competitor's score was visible (*M* = 81%) than invisible (*M* = 79%), positive evidence (*ER* = 28) that participants were more likely to obtain rewards against competitors with lower variance (*M* =84%) than higher variance (*M* = 76%), and decisive evidence (*ER* > 1000) that participants were less likely to obtain rewards against impulsive competitors (*M* = 70%) than patient competitors (*M* < 87%). Additionally, there was weak evidence (*ER* = 1.4) of a score-visibility × mean WT interaction, positive evidence (*ER = 6.8*) of a score-visibility × WT variance interaction, and decisive evidence (*ER* > 1000) of a mean WT × WT variance interaction.

The mean WT × WT variance interaction suggests that participants were equally likely to obtain rewards against impulsive competitors regardless of whether their WT variance was high (*M* = 71%) or low (*M* = 69%), but participants were more likely to obtain rewards against patient competitors when their WT variance was lower (*M* = 92%) than higher (*M* = 81%).

Table 6

Experiment 2's parameter estimates, 95% credible intervals and evidence ratios for the multilevel logistic regression predicting participants' log odds of obtaining the reward.

	B	95% CI	*ER*
Intercept	1.40	[0.96, 1.85]	> 1000
Visibility	-0.07	[-0.52, 0.36]	1.6
Mean WT	-0.54	[-0.78, -0.30]	> 1000
WT Variance	-0.21	[-0.43, 0.02]	28
Visibility × Mean WT	-0.02	[-0.43, 0.02]	1.4
Visibility × WT Variance	0.12	[-0.09, 0.34]	6.5
Mean WT × WT Variance	0.26	[0.21, 0.32]	> 1000

Note: The competitor's score Visibility was effect-coded as [Invisible = +1; Visible = -1], The competitor's Mean WT was effect-coded as [Impulsive = +1; Patient = -1], and WT Variance was effect-coded as [High Variance = +1, Low Variance = -1].

Similar to Experiment 1, model comparisons better supported a reduced model that only included the competitor's mean WT and variance as predictors (WAIC = 22,562) over the more complex model reported above that did include these predictors (WAIC = 22,564), suggesting that the visibility of the competitor's score is likely not impacting participants' performance in a meaningful way.

Overall, these results do not provide any evidence that the visibility of the competitors' cumulative score increases impulsive choices, either on its own or depending on the competitors' mean WT or variance of their WTs, concerning both their probability of waiting and their probability of obtaining rewards. These findings are more consistent with the simple subjective utility hypothesis that predicts individuals will ignore social factors and are only sensitive to the reward uncertainty.

Chapter 6 - General Discussion

Impulsive choices may arise for various reasons, the inability to delay gratification, excessive discounting of future outcomes, distortions of time, reducing effort, or abating uncertainty. Although the vast majority of research studying impulsive choices has focused on delay discounting, there has been an increasing interest on the dynamic interplay between a reward's delay and uncertainty (e.g., Dai et al., 2019a; McGuire & Kable, 2012; Stevens & Stephens, 2010, Fenneman et al., 2022; Webb et al., 2015; Young et al., 2014) since the initial characterization of this relationship (Green & Myerson, 1996; Rachlin et al., 1991). Uncertainty is an important attribute of the everyday choices we make and a contributing factor to preferences for shorter-term goals given the greater collection risks associated with delayed outcomes.

However, in research studying delayed and probabilistic reinforcement, the source of reward uncertainty is often considered environmental (e.g., a tool malfunction, interruptions, or other "extrinsic events") and not explicitly thought to further influence behaviors above just experiencing the uncertainty itself. An important gap in the research literature then is determining whether knowledge about the various sources of delayed reward uncertainty controls behavior to a meaningful degree – i.e., does knowledge of these factors simply serve as a cue of the underlying probability distributions or does it serve additional purposes that further impacts organisms' behavior?

In the current dissertation, I investigated the effects of competition as a specific source of reward uncertainty on participants' preferences for smaller-sooner rewards and their reward rates in a dynamic delayed gratification task (Young et al., 2011; Young & Howatt, 2022). My goal was to lay the foundation for a theoretically informed and ecologically representative experimental method to model the degree to which competition-based uncertainty produces more impulsive choices than environmental uncertainty. Using a 3D video-game environment in which participants freely navigated an open-world and interacted with various reward targets, I conducted two experience-based decision-making experiments that directly tested different dimensions of competition-based uncertainty: 1) the visibility of a competitor, and 2) the visibility of the competitor's score (as a proxy for their accumulated resources), and assessed the degree to which they impacted participants' probability of waiting and their probability of obtaining rewards compared to when this information was withheld.

Competing theoretical perspectives were proposed to predict whether competition-based uncertainty would produce more impulsive choices than environmental uncertainty. The first perspective was simple subjective utility, which is based on the tenets of game-theoretic rationality (i.e., the expected value of waiting times are computed as their reward rate multiplied by the probability that the reward is still available) and the simplifying assumption that individuals ignore relative payoffs and only pursue their own material self-interest. Consequently, a model derived from this perspective would predict that individuals are insensitive to knowledge about competitors and would act similarly regardless of whether they are informed about competitors and their status. The second perspective was evolutionary theory, which is based on the tenets of ecological rationality (e.g., natural selection and fitness) and emphasizes the importance of relative performance. Consequently, a model derived from this perspective would predict that individuals are sensitive to knowledge about competitors and would make more impulsive choices and obtain more rewards when they are informed about competitors and their status (in order to be better off than others).

Results from these two experiments revealed only weak evidence that participants were sensitive to competition-based uncertainty. In Experiment 1, there was little if any evidence that participants were less likely to wait to respond when the competitors were visible, but good evidence that participants were more likely to obtain rewards when the competitors were visible. Furthermore, there was also good evidence that participants' sensitivity to the visibility of the competitor was stronger against impulsive competitors (mean WT = 3 s), suggesting that the severity of competition can further impact participants' tendency to outperform others by being more likely to obtain rewards. However, in Experiment 2, there was no evidence that participants were more likely to either wait to respond or obtain rewards when their competitor's point total was visible.

Overall, the results from both experiments appear to be more consistent with the simple subjective utility perspective. There was mild evidence consistent with the evolutionary perspective that individuals are sensitive to the presence of a competitor, but only when measuring the likelihood with which they obtained rewards. The visibility of the competitor's cumulative rewards did not affect participants' waiting probability or reward probability. It is important to note, though, that although participants' behavior was generally more consistent with the simple subjective utility perspective, the null effects do not necessarily provide confirmatory evidence of it. Bayesian analyses provide quantifiable evidence in support of null effects (unlike standard frequentist analyses), but it is unclear whether the observed null effects reflect the true relationship between competition and behaviors in the EI task or if the study design was not sufficiently powered to detect smaller effects. While smaller effect sizes of competition would be consistent with some previous findings (see the Matsushima et al. and van den Bos et al. studies) the complete lack of an effect of the competitor's score in Experiment 2 is quite different compared to the large effect sizes found in other work (McClintock & McNeel, 1966). Moreover, the lack of an effect would also be inconsistent with the larger body of research documenting the influence of competition and relative gains in economic games. There are a few possible reasons for these findings and how the study could be improved to address its limitations.

Limitations and Potential Solutions

First, if the effects of competition-based uncertainty are simply small in magnitude, the experimental design may not have uncovered them due in part to either the visibility manipulations not being strong enough, the manipulation of the competitors' mean wait times being too strong, or the sample sizes being too small to detect them. Although the visibility of the competitors cannot be manipulated beyond 'visible' and 'invisible', the number of competitors could also be manipulated to strengthen participants' sensitivity to the amount of competition present (Garcia & Tor, 2009). The visibility of the competitor's score is also categorical, but displaying the competitor's point total creates the difference in point totals as another potential source of variance for each participant's behavior. Any visibility effect, or lack thereof, could be confounded with point differences, and more data would need to be collected to model the individual and interactive effects of score visibility and point differences. Additionally, the strength of the competitors' waiting times could be driving similar choice patterns in participants, particularly the impulsive competitors. Moreover, the stark contrast between the impulsive and patient competitors could be contaminating participants' behaviors across the levels, and more data would again be needed to model order effects. But, the massive effect size of the competitor's mean wait time on participants' probability of waiting and obtaining the reward in both experiments also reveals the promise of future work systematically sampling the competitors' mean wait times across the 10 second interval to begin developing a psychometric relationship between competition severity and individuals' waiting behaviors. Lastly, greater variability in participants' behavior due to the open-world nature of this task could be masking the smaller effect sizes, despite the richness of data collected. However, this variability could be advantageous in the long run because the results may replicate and generalize better than under more tightly controlled conditions (Richter et al., 2009). Ultimately, more data will need to be collected to ensure the effect is being accurately measured.

Second, using computerized competitors that were clearly non-human might not have been strong enough motivation for participants to care about outperforming them, even if they were compensated with a candy bar for winning the game. Moreover, the competitors' waiting times being governed by stationary distributions (Figure 3) removes the adaptive nature of opponents who can adjust their behaviors in response to the environment and competitor behavior, even if they are non-human. Most participants did beat the competitor in both experiments (Exp. 1 = 82%; Exp. 2 = 85%), and in Experiment 1 participants who could see the competitors did obtain significantly more total points than participants who could not see the competitors, reflecting a reasonable degree of task engagement; anecdotally, several participants also mentioned that they wanted to beat the competitor. However, previous work has documented that participants were more sensitive to relative outcomes and behaved more competitively in auctions when they believed that they were bidding against human players compared to computer players (van den Bos et al., 2008; van den Bos, Talwar et al., 2013), suggesting that using human (or human-like) participants playing against each other may elicit greater sensitivity to the competition contingencies.

Future work could therefore add greater realism to the task and competition simply by having participants actually play against each other in a multiplayer version of the game (e.g., Goldstone et al., 2005). However, pitting experimentally-naïve participants against each other fundamentally changes the experimental design because the competitors' behaviors are no longer under strict control, which complicates analyzing the effects of competition severity (the mean and variance of the competitors' wait times) on participants' own waiting behaviors and performance. So, if the research question only concerns the determining the influence of competitor visibility and eliciting maximal sensitivity to their presence, and not competition severity or predictability, then utilizing a multiplayer version of the game should be suitable to answer it.

Alternatively, future work could use more believable computerized competitors to help strengthen the impact of playing against them. For example, creating a more plausible cover story so that participants truly think they are playing against each other, rather than just stating the competitors' wait times are sampled from "participants in a previous study", or using a research assistant as a confederate to mimic real competitors. Another potential solution would be developing an adaptive competitor via computational algorithms (e.g., reinforcement learning; Sutton & Barto, 2018) that can calibrate their wait times to maximize their reward rates in response to a competitor. An adaptive algorithm may encourage participants to care more about the task and their relative performance because the task would be more challenging (i.e., a level of desirable difficulty; Bjork, 1994) and engaging, even against a computer model. I am not aware of any dynamic learning algorithm like this, so future research could work on this open problem.

Broader Implications and Future Directions

Experimental and theoretical work on the relationship between intertemporal choices and competition needs to be more thoroughly developed to generate stronger hypotheses and additional lines of research. I believe that EI tasks and the current study can make significant contributions towards achieving this goal.

Impulsive choices are commonly interpreted as maladaptive, and a large number of studies have successfully associated stronger preferences for immediacy with deleterious behaviors like substance abuse and gambling disorders. However, I believe that this interpretation is an incomplete portrayal of intertemporal choices, and a more productive goal may be to instead understand the conditions under which acting on short-term goals and waiting for longer-term goals are beneficial or harmful, for example obtaining resources in harsh and uncertain environments (Fenneman & Frankenhuis, 2020; Fawcett et al., 2012). Using an open-world 3D environment provides a greater level of immersion that is more representative of the naturalistic choices organisms routinely make in the wild, in which acting on more immediately available resources can maximize long-term rates. The current paradigm also enables stronger and richer experimental designs where participants repeatedly interact with the game environment, can learn from feedback, and experience how the environment responds to their actions, all in a highly controlled task. Thus, this current paradigm also underscores the almost unlimited possibilities in how researchers can study intertemporal choices and model the degree to which specific environmental variables impact them using a video game-based task. For example, manipulating state variables like a reward target’s power value or adding uncertainty into the task to incentivize rapid responding demonstrates that there are clearly situations in which impulsive choices are beneficial. Moreover, one could easily imagine a scenario in which game players have health or stamina bars that deplete over time, and repeatedly acquiring smaller resources could help fend off low health until larger resources can be harvested; this contingency could also help raise the in-game stakes that participants need to attend to while playing. The richness of this paradigm could help facilitate psychometric assessments of competition and competitive behaviors through the sampling of numerous variables across a broader range of values.

Next, there does not seem to be a general consensus on how social factors contribute to decision-makers' pursuits of narrow self-interest in game-theoretic utility-based models of choice. This lack of clarity and consistency can make it difficult to analytically derive benchmarks of reward-maximizing behaviors, propose valid hypotheses from these benchmarks, and then conduct experiments to test these hypotheses. Behavioral Game Theory (e.g., Camerer, 2003; Colman, 1995) is one extension of game theory that could be used to improve upon these limitations because it relaxes many of the strict assumptions of rational choice models and explicitly inserts psychological elements of perception, memory, reasoning and learning into game theory. And behavioral game theory has been at the forefront of showing the importance of social utility in economic decision-making and developing cognitive models of strategic interactions (e.g., Camerer et al., 2015; Liu & Halpern, 2023), but has not yet been applied to dynamic intertemporal choices.

One mathematical framework from behavioral game theory that could be relevant to intertemporal choices in strategic encounters is the Fehr-Schmidt model of inequity aversion (Fehr & Schmidt, 1999). The Fehr-Schmidt model is an economic utility function that weights decision-makers' valuation of outcomes by their distaste of disadvantageous inequity (being worse off than others) and their distaste of advantageous inequity (being better off than others), and has been shown to neatly account for the degree to which people value fairness and cooperation in social dilemmas. Moreover, Fehr and Schmidt's model has been integrated into learning models that incorporate how strongly decision-makers pursue advantageous inequity and change their behavior with experience (e.g., Cooper & Stockman, 2002; van den Bos et al., 2013). One limitation of the Fehr-Schmidt model, though, is that it only measures the degree to which people value relative outcomes. It does not predict that decision-makers will be sensitive to relative outcomes, nor can it identify when they should be sensitive to relative outcomes. Thus, this model does not provide a sufficient explanatory model of choices regarding *why* people may be (in)sensitive to relative outcomes.

Despite the lack of convincing evidence consistent with the evolutionary hypotheses across both experiments, I still think that evolution by natural selection is one of the only theoretical perspectives that can accommodate social uncertainty, relative preferences, and intertemporal choices in a way that models both the organism and the environment. However, to my knowledge there is currently no evolutionary theory of how competitive environments select for cognitive processes that prefer relative gain maximization in such a way that produces impulsive choices. Future work therefore could more strongly experimentally test whether sources of uncertainty differentially control behavior and develop theoretically-informed explanations for why this might the case, supplemented with agent-based or mathematical modeling analyses that identify stable behaviors based on environmental conditions, physiological states and competition severity. This work could then lead to the proposal and subsequent testing of a middle-level theory (Buss, 1995) of intertemporal choices under competition that draws from psychology, biology and economics, for example integrating associative learning and the Fehr-Schmidt model with evolutionary processes that can explain the emergence of these cognitive processes and assess the quality of their behavioral output.

As the scientific study of intertemporal choices and their ecological performance continues to advance, so too will the theoretical perspectives and experimental methods that model and explain impulsive behaviors. This work provides a foundation for future work to integrate these psychological phenomenon in a more ecologically valid, and generalizable way.

References

Ainslie, G. W. (1974). Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior*, *21*(3), 485-489.

Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, *82*(4), 463-496.

Ainslie, G. (1992). *Picoeconomics: The strategic interaction of successive motivational states within the person*. Cambridge University Press.

Ainslie, G., & Herrnstein, R. J. (1981). Preference reversal and delayed reinforcement. *Animal Learning & Behavior*, *9*(4), 476-482.

Amita, H., Kawamori, A., & Matsushima, T. (2010). Social influences of competition on impulsive choices in domestic chicks. *Biology Letters*, *6*(2), 183-186.

Amita, H., & Matsushima, T. (2011). Instantaneous and cumulative influences of competition on impulsive choices in domestic chicks. *Frontiers in Neuroscience*, *5*, 1-8.

Amita, H., & Matsushima, T. (2014). Competitor suppresses neuronal representation of food reward in the nucleus accumbens/medial striatum of domestic chicks. *Behavioural Brain Research*, *268*, 139-149.

Bailey, C., Panfil, K., & Kirkpatrick, K. (2021). Hazard function effects on promoting self-control in variable interval time-based interventions in rats. *Journal of the Experimental Analysis of Behavior*, *116*(3), 279-299.

Bailey, A. J., Romeu, R. J., & Finn, P. R. (2021). The problems with delay discounting: a critical review of current practices and clinical applications. *Psychological Medicine*, 1-8.

(Stuart) Barker, J.S. (2009). Defining fitness in natural and domesticated populations. In J. van der Werf, H. U. Graser, R. Frankham, & C. Gondro (Eds.), *Adaptation and Fitness in Animal Populations* (pp 3-14). Springer, Dordrecht. https://doi.org/10.1007/978-1-4020-9005-9_1.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*, 1–48.

Baum, W. M. (1974). On two types of deviation from the matching law: bias and undermatching. *Journal of the Experimental Analysis of Behavior*, *22*(1), 231-242.

Bergman, T. J., & Beehner, J. C. (2022). Leveling with Tinbergen: Four levels simplified to causes and consequences. *Evolutionary Anthropology: Issues, News, and Reviews*, *31*(1), 12-19.

Berns, G. S., Laibson, D., & Loewenstein, G. (2007). Intertemporal choice–toward an integrative framework. *Trends in Cognitive Sciences*, *11*(11), 482-488.

Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe, & A. P. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 185–205). MIT Press.

Bolker, B. M. (2015). Linear and generalized linear mixed models. In G. A. Fox, S. Negrete-Yankelevich & V. J. Sosa (Eds.), *Ecological Statistics: Contemporary Theory and Application* (pp. 309–333). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199672547.003.0014

Broom, M., Cressman, R., & Křivan, V. (2019). Revisiting the “fallacy of averages” in ecology: Expected gain per unit time equals expected gain divided by expected time. *Journal of Theoretical Biology*, *483*, 109993.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, *62*(3), 193-217.

Bürkner, P. (2017). brms: An R package for Bayesian multilevel models using stan. *Journal of Statistical Software, 80*(1), 1–28.

Businelle, M. S., McVay, M. A., Kendzor, D., & Copeland, A. (2010). A comparison of delay discounting among smokers, substance abusers, and non-dependent controls. *Drug and Alcohol Dependence*, *112*(3), 247-250.

Buss, D. M. (1995). Evolutionary psychology: A new paradigm for psychological science. *Psychological Inquiry*, *6*(1), 1-30.

Cagiltay, N. E., Ozcelik, E., & Ozcelik, N. S. (2015). The effect of competition on learning in games. *Computers & Education*, *87*, 35-41.

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. New York, NY: Russell Sage Foundation.

Camerer, C., & Ho, T. H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67*(4), 827-874.

Camerer, C. F., Ho, T. H., & Chong, J. K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, *119*(3), 861-898.

Camerer, C. F., Ho, T. H., & Chong, J. K. (2015). A psychological approach to strategic thinking in games. *Current Opinion in Behavioral Sciences*, *3*, 157-162.

Carrascal, L. M., & Moreno, E. (1993). Food caching versus immediate consumption in the nuthatch: The effect of social context. *Ardea,* 81, 135-141.

Carter, E. C., Pedersen, E. J., & McCullough, M. E. (2015). Reassessing intertemporal choice: human decision-making is more optimal in a foraging task than in a self-control task. *Frontiers in Psychology*, *6*, 95.

Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, *9*(2), 129-136.

Chen, S. Y., & Chang, Y. M. (2020). The impacts of real competition and virtual competition in digital game-based learning. *Computers in Human Behavior*, *104*, 106171.

Chen, C. H., Liu, J. H., & Shou, W. C. (2018). How competition in a game-based science learning environment influences students' learning achievement, flow experience, and learning behavioral patterns. *Journal of Educational Technology & Society*, *21*(2), 164-176.

Chung, S. H. (1965). Effects of delayed reinforcement in a concurrent situation. *Journal of the Experimental Analysis of Behavior*, *8*(6), 439-444.

Chung, S. H., & Herrnstein, R. J. (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior,* 10(1), 67–74. https://doi.org/10.1901/jeab.1967.10-67

Ciccarelli, M., Cosenza, M., D'Olimpio, F., Griffiths, M. D., & Nigro, G. (2019). An experimental investigation of the role of delay discounting and craving in gambling chasing behavior. *Addictive Behaviors*, *93*, 250-256.

Colman, A. M. (1995). *Game theory and its applications: In the social and biological sciences*. London: Psychology Press.

Cooper, D. J., & Stockman, C. K. (2002). Fairness and learning: an experimental examination. *Games and Economic Behavior*, *41*(1), 26-45.

Cox, D. J., & Dallery, J. (2016). Effects of delay and probability combinations on discounting in humans. *Behavioural Processes*, *131*, 15-23.

Crockett, M. J., Clark, L., Lieberman, M. D., Tabibnia, G., & Robbins, T. W. (2010). Impulsive choice and altruistic punishment are correlated and increase in tandem with serotonin depletion. *Emotion*, *10*(6), 855-862.

Curry, O. S., Price, M. E., & Price, J. G. (2008). Patience is a virtue: Cooperative people have lower discount rates. *Personality and Individual Differences*, *44*(3), 780-785.

Dai, J., Pachur, T., Pleskac, T. J., & Hertwig, R. (2019a). Tomorrow never knows: Why and how uncertainty matters in intertemporal choice. In R. Hertwig, T.J. Pleskac, T. Pachur, & the Center for Adaptive Rationality (Eds.), *Taming uncertainty* (pp.175-190).

Dai, J., Pachur, T., Pleskac, T. J., & Hertwig, R. (2019b). What the future holds and when: A description–experience gap in intertemporal choice. *Psychological Science*, *30*(8), 1218-1233.

Darwin, C. (1859). *On the origin of species*. (1859). London: Murray.

Dasgupta, P., & Maskin, E. (2005). Uncertainty and hyperbolic discounting. *American Economic Review*, *95*(4), 1290-1299.

Davenport, J. W. (1962). The interaction of magnitude and delay of reinforcement in spatial discrimination. *Journal of Comparative and Physiological Psychology*, *55*(2), 267.

Durrant, R., & Ward, T. (2015). *Evolutionary criminology: Towards a comprehensive explanation of crime*. San Diego, CA: Academic Press.

Dyson, B. J. (2021). Variability in competitive decision-making speed and quality against exploiting and exploitative opponents. *Scientific Reports*, *11*(1), 1-13.

Emery, N. J., Dally, J. M., & Clayton, N. S. (2004). Western scrub-jays (*Aphelocoma californica*) use cognitive strategies to protect their caches from thieving conspecifics. *Animal Cognition,* 7, 37-43.

Ermer, E., Cosmides, L., & Tooby, J. (2008). Relative status regulates risky decision making about resources in men: Evidence for the co-evolution of motivation and cognition. *Evolution and Human Behavior*, *29*(2), 106-118.

Espín, A. M., Branas-Garza, P., Herrmann, B., & Gamella, J. F. (2012). Patient and impatient punishers of free-riders. *Proceedings of the Royal Society B: Biological Sciences*, *279*(1749), 4923-4928.

Espín, A. M., Exadaktylos, F., Herrmann, B., & Brañas-Garza, P. (2015). Short-and long-run goals in ultimatum bargaining: impatience predicts spite-based behavior. *Frontiers in Behavioral Neuroscience*, *9*, 214.

Espín, A., Martinez Sánchez, A., & Herrmann, B. (2017). Economic preferences 2.0: Connecting competition, cooperation and inter-temporal preferences, CeDEx Discussion Paper Series, No. 2017-04, The University of Nottingham, Centre for Decision Research and Experimental Economics (CeDEx), Nottingham

Fawcett, T. W., McNamara, J. M., & Houston, A. I. (2012). When is it adaptive to be patient? A general framework for evaluating delayed rewards. *Behavioural Processes*, *89*(2), 128-136.

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal Of Economics*, *114*(3), 817-868.

Fenneman, J., & Frankenhuis, W. E. (2020). Is impulsive behavior adaptive in harsh and unpredictable environments? A formal model. *Evolution and Human Behavior*, *41*(4), 261-273.

Fenneman, J., Frankenhuis, W. E., & Todd, P. M. (2022). In which environments is impulsive behavior adaptive? A cross-discipline review and integration of formal models. *Psychological Bulletin*, *148*(7-8), 555.

Festinger, L. (1954). A theory of social comparison processes. *Human relations*, *7*(2), 117-140.

Fisher, R.A. (1930). *The genetical theory of natural selection*. Oxford, England: Clarendon Press.

Forstmeier, S., Drobetz, R., & Maercker, A. (2011). The delay of gratification test for adults: Validating a behavioral measure of self-motivation in a sample of older people. *Motivation and Emotion*, *35*(2), 118-134.

Frederick, S. & Loewenstein, G. (1999). Hedonic Adaptation. In D. Kahneman, E. Diener & N. Schwarz (Eds.), *Well-Being: The Foundations of Hedonic Psychology* (pp. 302-329). New York, NY: Russell Sage Foundation.

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

Frost, R., & McNaughton, N. (2017). The neural basis of delay discounting: A review and preliminary model. *Neuroscience & Biobehavioral Reviews*, *79*, 48-65.

Fung, B. J., Bode, S., & Murawski, C. (2017). High monetary reward rates and caloric rewards decrease temporal persistence. *Proceedings of the Royal Society B: Biological Sciences*, *284*(1849), 20162759.

Garcia, S. M., Reese, Z. A., & Tor, A. (2019). Social comparison before, during, and after the competition. In J. Suls, R. Collins & L. Wheeler (Eds.), *Social comparison, judgment and behavior* (pp. 105-142). Oxford University Press.

Garcia, S. M., & Tor, A. (2009). The N-effect: More competitors, less competition. *Psychological Science, 20*(7), 871-877.

Garcia, S. M., Tor, A., & Schiff, T. M. (2013). The psychology of competition: A social comparison perspective. *Perspectives on Psychological Science*, *8*(6), 634-650.

Gebru, N. M., Kalkat, M., Strickland, J. C., Ansell, M., Leeman, R. F., & Berry, M. S. (2022). Measuring sexual risk-taking: a systematic review of the sexual delay discounting task. *Archives of Sexual Behavior*, 1-22.

Gentry, G. D., & Marr, M. J. (1980). Choice and reinforcement delay. *Journal of the Experimental Analysis of Behavior*, *33*(1), 27-37.

Giddie, L. (2011). *The effect of resource size uncertainty, group size uncertainty, and intertemporal choice on an intergroup choice behavior*. [Doctoral dissertation, The Chicago School of Professional Psychology]. ProQuest Dissertations and Theses Global.

Gigerenzer, G., & Selten, R. (eds.). (2001). *Bounded rationality: The adaptive toolbox*. MIT Press.

Gillis, D. M., & Kramer, D. L. (1987). Ideal interference distributions: population density and patch use by zebrafish. *Animal Behaviour*, *35*(6), 1875-1882.

Glass, T. A. (1969). Effects of Three Situational Variables and Type of Reinforcement on Delay Preference. *The American Journal of Psychology*, *82*(2), 241-247.

Göllner, L. M., Ballhausen, N., Kliegel, M., & Forstmeier, S. (2018). Delay of gratification, delay discounting and their associations with age, episodic future thinking, and future time perspective. *Frontiers in Psychology*, *8*, 2304.

Goldstone, R. L., Ashpole, B. C., & Roberts, M. E. (2005). Knowledge of resources and competitors in human foraging. *Psychonomic Bulletin & Review*, *12*(1), 81-87.

Golman, R., Bhatia, S., & Kane, P. B. (2020). The dual accumulator model of strategic deliberation and decision making. *Psychological Review*, *127*(4), 477-504.

Gonzales, J., Mishra, S., & Camp, R. D. (2017). For the win: Risk-sensitive decision-making in teams. *Journal of Behavioral Decision Making*, *30*(2), 462-472.

Green, L., & Myerson, J. (1996). Exponential versus hyperbolic discounting of delayed outcomes: Risk and waiting time. *American Zoologist*, *36*(4), 496-505.

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

Green, L., & Myerson, J. (2019). On the complexity of discounting, choice situations, and people. *Perspectives on Behavior Science*, *42*(3), 433-443.

Green, L., Price, P. C., & Hamburger, M. E. (1995). Prisoner's dilemma and the pigeon: Control by immediate consequences. *Journal of the Experimental Analysis of Behavior*, *64*(1), 1-17.

Green, L., & Rachlin, H. (1996). Commitment using punishment. *Journal of the Experimental Analysis of Behavior*, *65*(3), 593-601.

Green, L., & Snyderman, M. (1980). Choice between rewards differing in amount and delay: Toward a choice model of self control. *Journal of the Experimental Analysis of Behavior*, *34*(2), 135-147.

Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, *17*(9), 767-773.

Grosch, J., & Neuringer, A. (1981). Self-control in pigeons under the Mischel paradigm. *Journal of the Experimental Analysis of Behavior*, *35*(1), 3-21.

Güth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, *3*(4), 367-388.

Hamilton, W. D. (1964). The genetical evolution of social behaviour. II. *Journal of Theoretical Biology*, *7*(1), 17-52.

Hammerstein, P., & Selten, R. (1994). Game theory and evolutionary biology. In R.J. Aumann & S. Hart (Eds.), *Handbook of game theory with economic applications*, *2* (pp. 929-993). Elsevier: Amsterdam.

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

Harsanyi, J. C. (1966). A general theory of rational behavior in game situations. *Econometrica: Journal of the Econometric Society*, 613-634.

Hayden, B. Y. (2016). Time discounting and time preference in animals: a critical review. *Psychonomic Bulletin & Review*, *23*(1), 39-53.

Hayden, B. Y. (2019). Why has evolution not selected for perfect self-control?. *Philosophical Transactions of the Royal Society B*, *374*(1766), 20180139.

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, 154-162.

Herrnstein, R. J. (1974). Formal properties of the matching law. *Journal of the Experimental Analysis of Behavior*, *21*(1), 159-164.

Herrnstein, R. J. (1990). Rational choice theory: Necessary but not sufficient. *American Psychologist*, *45*(3), 356-367.

Hill, S. E., & Buss, D. M. (2010). Risk and relative social rank: positional concerns and risky shifts in probabilistic decision-making. *Evolution and Human Behavior*, *31*(3), 219-226.

Hintze, A., Phillips, N., & Hertwig, R. (2015). The Janus face of Darwinian competition. *Scientific Reports, 5*(1), 1-7.

Hopewell, L. J., Leaver, L. A., & Lea, S. E. (2008). Effects of competition and food availability on travel time in scatter-hoarding gray squirrels (Sciurus carolinensis). *Behavioral Ecology*, *19*(6), 1143-1149.

Houston, A. I., McNamara, J. M., & Steer, M. D. (2007). Violations of transitivity under fitness maximization. *Biology Letters*, *3*(4), 365-367.

Hwang, J., Kim, S., & Lee, D. (2009). Temporal discounting and inter-temporal choice in rhesus monkeys. *Frontiers in Behavioral Neuroscience*, 3, 1-13.

Ifcher, J., Zarghamee, H., Houser, D., & Diaz, L. (2020). The relative income effect: an experiment. *Experimental Economics*, *23*(4), 1205-1234.

Jarosz, A. F., & Wiley, J. (2014). What are the odds? A practical guide to computing and reporting Bayes factors. *The Journal of Problem Solving, 7*(1), 2-9.

Jones, B. A., & Rachlin, H. (2009). Delay, probability, and social discounting in a public goods game. *Journal of the Experimental Analysis of Behavior*, *91*(1), 61-73.

Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*(12), 1625-1633.

Kacelnik A. 2003. The evolution of patience. In G. Loewenstein, D. Read & R. Baumeister (Eds), *Time and Decision: Economics and Psychological Perspectives on Intertemporal Choice* (pp. 115–38). New York, NY: Russell Sage Foundation.

Kagel, J. H., Green, L., & Caraco, T. (1986). When foragers discount the future: Constraint or adaptation?. *Animal Behaviour*, *34*, 271-283.

Karagonlar, G., & Kuhlman, D. M. (2013). The role of social value orientation in response to an unfair offer in the ultimatum game. *Organizational Behavior and Human Decision Processes*, *120*(2), 228-239.

Kidd, C., Palmeri, H., & Aslin, R. N. (2013). Rational snacking: Young children's decision-making on the marshmallow task is moderated by beliefs about environmental reliability. *Cognition*, *126*(1), 109-114.

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.

Kirby, K. N., & Maraković, N. N. (1996). Delay-discounting probabilistic rewards: Rates decrease as amounts increase. *Psychonomic Bulletin & Review*, *3*(1), 100-104.

Kluwe-Schiavon, B., Viola, T. W., Sanvicente-Vieira, B., Lumertz, F. S., Salum, G. A., Grassi-Oliveira, R., & Quednow, B. B. (2020). Substance related disorders are associated with impaired valuation of delayed gratification and feedback processing: a multilevel meta-analysis and meta-regression. *Neuroscience & Biobehavioral Reviews*, *108*, 295-307.

Knolle-Veentjer, S., Huth, V., Ferstl, R., Aldenhoff, J. B., & Hinze-Selch, D. (2008). Delay of gratification and executive performance in individuals with schizophrenia: Putative role for eating behavior and body weight regulation. *Journal of Psychiatric Research*, *42*(2), 98-105.

Krebs, J. R., Stephens, D. W., & Sutherland, W. J. (1983). *Perspectives in optimal foraging* (pp. 165-216). Cambridge: Cambridge University Press.

Ku, G., Malhotra, D., & Murnighan, J. K. (2005). Towards a competitive arousal model of decision-making: A study of auction fever in live and internet auctions. *Organizational Behavior and Human Decision Processes*, *96*(2), 89-103.

Ledyard, O. (1995). Public goods: some experimental results. In J. Kagel & A. Roth (Eds.), *Handbook of Experimental Economics* (pp. 111-194). Princeton, NJ: Princeton Univ. Press.

Leyton-Brown, K. (2007). Lecture 2: *Self-Interested Agents and Utility Theory* [PowerPoint slides]. The University of British Columbia. Models of Strategic Behaviour. Available at: https://www.cs.ubc.ca/~kevinlb/teaching/cs532l%20-%202007-8/lectures/lect2.pdf

Lee, D., Conroy, M. L., McGreevy, B. P., & Barraclough, D. J. (2004). Reinforcement learning and decision making in monkeys during a competitive game. *Cognitive Brain Research*, *22*(1), 45-58.

Lee, D., McGreevy, B. P., & Barraclough, D. J. (2005). Learning and decision making in monkeys during a rock–paper–scissors game. *Cognitive Brain Research*, *25*(2), 416-430.

Lee, D., & Seo, H. (2016). Neural basis of strategic decision making. *Trends in Neurosciences*, *39*(1), 40-48.

Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., Strong, D. R., & Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: The balloon analogue risk task (BART). *Journal of Experimental Psychology: Applied, 8*, 75–84.

Lempert, K. M., McGuire, J. T., Hazeltine, D. B., Phelps, E. A., & Kable, J. W. (2018). The effects of acute stress on the calibration of persistence. *Neurobiology of Stress*, *8*, 1-9.

Lenth, R., Singmann, H., Love, J., Buerkner, P., & Herve, M. (2018). Emmeans: Estimated marginal means, aka least-squares means. *R package version*, *1*(1), 3.

Liu, X., & Halpern, J. Y. (2023). Strategic Play By Resource-Bounded Agents in Security Games. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems* (pp. 2478-2480).

Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*, *107*(2), 573-597.

Loewenstein, G. F., Thompson, L., & Bazerman, M. H. (1989). Social utility and decision making in interpersonal contexts. *Journal of Personality and Social psychology*, *57*(3), 426.

Logan, F. A. (1965). Decision making by rats: Delay versus amount of reward. *Journal of Comparative and Physiological Psychology*, *59*(1), 1-12.

Logue, A. W., & Chavarro, A. (1987). Effect on choice of absolute and relative values of reinforcer delay, amount, and frequency. *Journal of Experimental Psychology: Animal Behavior Processes*, *13*(3), 280-291.

Logue, A. W., Forzano, L. B., & Tobin, H. (1992). Independence of reinforcer amount and delay: The generalized matching law and self-control in humans. *Learning and Motivation*, *23*(3), 326-342.

Logue, A. W., Peña-Correal, T. E., Rodriguez, M. L., & Kabela, E. (1986). Self-control in adult humans: Variation in positive reinforcer amount and delay. *Journal of the Experimental analysis of Behavior*, *46*(2), 159-173.

Luce, R. D., & Raiffa, H. (1989). *Games and decisions: Introduction and critical survey*. Courier Corporation.

Lüdecke, D., Ben-Shachar, M. S., Patil, I., Waggoner, P., & Makowski, D. (2021). performance: An R package for assessment, comparison and testing of statistical models. *Journal of Open Source Software*, *6*(60). 3139. https://doi.org/10.21105/joss.03139.

Luttmer, E. F. (2005). Neighbors as negatives: Relative earnings and well-being. *The Quarterly Journal of Economics*, *120*(3), 963-1002.

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

Madden, G.J. & Johnson, P.S. (2010) A Delay-Discounting Primer. In *Impulsivity: The behavioral and Neurological Science of Discounting* (pp. 11-37). https://doi.org/10.1037/12069-001

Malhotra, D. (2010). The desire to win: The effects of competitive arousal on motivation and behavior. *Organizational Behavior and Human Decision Processes*, *111*(2), 139-146.

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

Marsh, B., Schuck-Paim, C., & Kacelnik, A. (2004). Energetic state during learning affects foraging choices in starlings. *Behavioral Ecology*, *15*(3), 396-399.

Marshall, A. T., Smith, A. P., & Kirkpatrick, K. (2014). Mechanisms of impulsive choice: I. Individual differences in interval timing and reward processing. *Journal of the Experimental Analysis of Behavior*, *102*(1), 86-101.

Martin, N. (2016). Strategy as mutually contingent choice: New behavioral lessons from Thomas Schelling's the strategy of conflict. *SAGE Open*, *6*(2), 2158244016651912.

Martin, C. F., Bhui, R., Bossaerts, P., Matsuzawa, T., & Camerer, C. (2014). Chimpanzee choice rates in competitive games match equilibrium game theory predictions. *Scientific Reports*, *4*(1), 1-6.

Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge University Press, Cambridge.

Maynard Smith, J. & Price, G. R. (1973). The logic of animal conflict. *Nature*, *246*(5427), 15-18.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. *Quantitative Analyses of Behavior*, *5*, 55-73.

McClintock, C. G., & McNeel, S. P. (1966). Reward and score feedback as determinants of cooperative and competitive game behavior. *Journal of Personality and Social Psychology*, *4*(6), 606-613.

McClintock, C. G., & Nuttin Jr, J. M. (1969). Development of competitive game behavior in children across two cultures. *Journal of Experimental Social Psychology*, *5*(2), 203-218.

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.

McDiarmid, C. G., & Rilling, M. E. (1965). Reinforcement delay and reinforcement rate as determinants of schedule preference. *Psychonomic Science*, *2*(1), 195-196.

McGuire, J. T., & Kable, J. W. (2012). Decision makers calibrate behavioral persistence on the basis of time-interval experience. *Cognition*, *124*(2), 216-226.

McGuire, J. T., & Kable, J. W. (2013). Rational temporal predictions can underlie apparent failures to delay gratification. *Psychological Review*, *120*(2), 395-410.

McGuire, J. T., & Kable, J. W. (2015). Medial prefrontal cortical activity reflects dynamic re-evaluation during voluntary persistence. *Nature Neuroscience*, *18*(5), 760-766.

Messick, D. M., & McClintock, C. G. (1968). Motivational bases of choice in experimental games. *Journal of Experimental Social Psychology*, *4*(1), 1-25.

Messick, D. M., & Thorngate, W. B. (1967). Relative gain maximization in experimental games. *Journal of Experimental Social Psychology*, *3*(1), 85-101.

Mesterton-Gibbons, M., & Adams, E. S. (1998). Animal contests as evolutionary games: paradoxical behavior can be understood in the context of evolutionary stable strategies. The trick is to discover which game the animal is playing. *American Scientist*, *86*(4), 334-341.

Michaelson, L., De la Vega, A., Chatham, C. H., & Munakata, Y. (2013). Delaying gratification depends on social trust. *Frontiers in Psychology*, *4*, 1-7.

Mischel, W., & Ebbesen, E. B. (1970). Attention in delay of gratification. *Journal of Personality and Social Psychology*, *16*(2), 329-337.

Mischel, W., & Grusec, J. (1967). Waiting for rewards and punishments: effects of time and probability on choice. *Journal of Personality and Social Psychology*, *5*(1), 24-31.

Mischel, W., & Metzner, R. (1962). Preference for delayed reward as a function of age, intelligence, and length of delay interval. *The Journal of Abnormal and Social Psychology*, *64*(6), 425-431.

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

Mischel, W., & Staub, E. (1965). Effects of expectancy on working and waiting for larger reward. *Journal of Personality and Social Psychology*, *2*(5), 625-633.

Mishra, S. (2014). Decision-making under risk: Integrating perspectives from biology, economics, and psychology. *Personality and Social Psychology Review*, *18*(3), 280-307.

Mishra, S., Barclay, P., & Lalumière, M. L. (2014). Competitive disadvantage facilitates risk taking. *Evolution and Human Behavior*, *35*(2), 126-132.

Mishra, S., Barclay, P., & Sparks, A. (2017). The relative state model: Integrating need-based and ability-based pathways to risk-taking. *Personality and Social Psychology Review*, *21*(2), 176-198.

Mizuyama, R., Uno, L., & Matsushima, T. (2016). Food variance and temporal discounting in socially foraging chicks. *Animal Behaviour*, *120*, 143-151.

Modelling Animal Decisions Group, Fawcett, T.W., Fallenstein, B., Higginson, A.D., Houston, A.I., et al. (2014). The evolution of decision rules in complex environments. *Trends in Cognitive Science, 18*(3), 153–61.

Moffett, L., Flannagan, C., & Shah, P. (2020). The influence of environmental reliability in the marshmallow task: An extension study. *Journal of Experimental Child Psychology*, *194*, 104821. https://doi.org/10.1016/j.jecp.2020.104821

Mulder, M. & Schacht, R. (2012). Human Behavioural Ecology. In: eLS. John Wiley & Sons, Ltd: Chichester. https://doi.org/10.1002/9780470015902.a0003671.pub2.

Murayama, K., & Elliot, A. J. (2012). The competition–performance relation: A meta-analytic review and test of the opposing processes model of competition and performance. *Psychological Bulletin*, *138*(6), 1035-1070.

Murphy, R. O., Ackermann, K. A., & Handgraaf, M. (2011). Measuring social value orientation. *Judgment and Decision Making*, *6*(8), 771-781.

Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior*, *64*(3), 263-276.

Nash Jr, J. F. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, *36*(1), 48-49.

Navarick, D. J. (1987). Reinforcement probability and delay as determinants of human impulsiveness. *The Psychological Record, 37*(2), 219-226.

Odum, A. L. (2011a). Delay discounting: I'm a k, you're a k. *Journal of the Experimental Analysis of Behavior*, *96*(3), 427-439.

Odum, A. L. (2011b). Delay discounting: trait variable?. *Behavioural Processes*, *87*(1), 1-9.

Odum, A. L., Becker, R. J., Haynes, J. M., Galizio, A., Frye, C. C., Downey, H., ... & Perez, D. M. (2020). Delay discounting of different outcomes: Review and theory. *Journal of the Experimental Analysis of Behavior*, *113*(3), 657-679.

Orr, H. A. (2007). Absolute fitness, relative fitness, and utility. *Evolution: International Journal of Organic Evolution*, *61*(12), 2997-3000.

Orr, H. A. (2009). Fitness and its role in evolutionary genetics. *Nature Reviews Genetics*, *10*(8), 531-539.

Otto, A. R., Markman, A. B., & Love, B. C. (2012). Taking more, now: The optimality of impulsive choice hinges on environment structure. *Social Psychological and Personality Science*, *3*(2), 131-138.

Paglieri, F. (2013). The costs of delay: Waiting versus postponing in intertemporal choice. *Journal of the Experimental Analysis of Behavior*, *99*(3), 362-377.

Parkes, D. C., & Wellman, M. P. (2015). Economic reasoning and artificial intelligence. *Science, 349*(6245), 267-272.

Pleskac, T. J., Conradt, L., Leuker, C., & Hertwig, R. (2021). The ecology of competition: A theory of risk–reward environments in adaptive decision making. *Psychological Review*, *128*(2), 315-335.

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

Rachlin, H. (2000). *The science of self-control*. Harvard University Press.

Rachlin, H., & Green, L. (1972). Commitment, choice and self-control 1. *Journal of the Experimental Analysis of Behavior*, *17*(1), 15-22.

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

Rae, J. (1834). *The Sociological Theory of Capital* (reprint 1834 ed.). London: Macmillan

Rapoport, A., & Chammah, A. (1965). *Prisoner's Dilemma: A study in conflict and coooperation*. University of Michigan Press, Ann Arbor.

Reynolds, B. (2006). A review of delay-discounting research with humans: relations to drug use and gambling. *Behavioural Pharmacology*, *17*(8), 651-667.

Reynolds, B. & Schiffbauer, R. (2004). Measuring state changes in human delay discounting: an experiential discounting task. *Behavioural Processes*, *67*(3), 343-356.

Reynolds, B., & Schiffbauer, R. (2005). Delay of gratification and delay discounting: A unifying feedback model of delay-related impulsive behavior. *The Psychological Record*, *55*(3), 439-460.

Richards, J. B., Mitchell, S. H., De Wit, H., & Seiden, L. S. (1997). Determination of discount functions in rats with an adjusting-amount procedure. *Journal of the Experimental Analysis of Behavior*, *67*(3), 353-366.

Richards, R. W. (1974). Inhibitory stimulus control and the magnitude of delayed reinforcement. *Journal of the Experimental Analysis of Behavior*, *21*(3), 501-509.

Richter, S. H., Garner, J. P., & Würbel, H. (2009). Environmental standardization: cure or cause of poor reproducibility in animal experiments?. *Nature Methods, 6*(4), 257-261.

Rilling, J. K., & Sanfey, A. G. (2011). The neuroscience of social decision-making. *Annual Review of Psychology*, *62*(1), 23-48.

Roberts, M. E., & Goldstone, R. L. (2006). EPICURE: Spatial and knowledge limitations in group foraging. *Adaptive Behavior*, *14*(4), 291-313.

Rosati, A. G., DiNicola, L. M., & Buckholtz, J. W. (2018). Chimpanzee cooperation is fast and independent from self-control. *Psychological Science*, *29*(11), 1832-1845.

Rosati, A. G., & Hare, B. (2012). Decision making across social contexts: competition increases preferences for risk in chimpanzees and bonobos. *Animal Behaviour*, *84*(4), 869-879.

Ross, D. (2019). Game Theory. *Stanford encyclopedia of philosophy*. See https://plato.stanford.edu/entries/game-theory/index.html

Samuelson, P. A. (1937). A note on measurement of utility. *The review of Economic Studies*, *4*(2), 155-161.

Sánchez-Amaro, A., Tan, J., Kaufhold, S. P., Fernández-Navarro, S., & Rossano, F. (2021). How environmental unpredictability and harshness affect chimpanzees (Pan troglodytes) in risk-choice and temporal discounting tasks. *Journal of Comparative Psychology, 2,* 232-243.

Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.

Schelling, T. (1960). *The strategy of conflict*. Cambridge, Mass.: Harvard University Press.

Schweighofer, N., Shishida, K., Han, C. E., Okamoto, Y., Tanaka, S. C., Yamawaki, S., & Doya, K. (2006). Humans can adopt optimal discounting strategy under real-time constraints. *PLoS Computational Biology*, *2*(11), e152.

Sebens, K. P., Sarà, G., & Carrington, E. (2018). Estimation of fitness from energetics and life-history data: An example using mussels. *Ecology and Evolution*, *8*(11), 5279-5290.

Selten, R. (1990). Bounded rationality. *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, *146*(4), 649-658.

Selten, R. (1998). Features of experimentally observed bounded rationality. *European Economic Review*, *42*(3-5), 413-436.

Seth, A. K. (2007). The ecology of action selection: insights from artificial life. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1485), 1545-1558.

Shafir, E., & Tversky, A. (1992). Thinking through uncertainty: Nonconsequential reasoning and choice. *Cognitive Psychology*, *24*(4), 449-474.

Sherman, J. A., & Thomas, J. R. (1968). Some factors controlling preference between fixed-ratio and variable ratio schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, *11*(6), 689-702.

Silston, B., Wise, T., Qi, S., Sui, X., Dayan, P., & Mobbs, D. (2021). Neural encoding of perceived patch value during competitive and hazardous virtual foraging. *Nature Communications, 12*(1), 5478. https://doi.org/10.1038/s41467-021-25816-9.

Silver, D., Singh, S., Precup, D., & Sutton, R. S. (2021). Reward is enough. *Artificial Intelligence*, *299*, 103535. https://doi.org/10.1016/j.artint.2021.103535.

Simon, H. A. (1957). *Models of man; social and rational*. New York: John Wiley & Sons.

Sims, C. R., Neth, H., Jacobs, R. A., & Gray, W. D. (2013). Melioration as rational choice: Sequential decision making in uncertain environments. *Psychological Review*, *120*(1), 139-154.

Skrynka, J., & Vincent, B. T. (2019). Hunger increases delay discounting of food and non-food rewards. *Psychonomic Bulletin & Review*, *26*(5), 1729-1737.

Snyderman, M. (1983). Delay and amount of reward in a concurrent chain. *Journal of the Experimental Analysis of Behavior*, *39*(3), 437-447.

Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *265*(1409), 2015-2020.

Steele, C. C., Gwinner, M., Smith, T., Young, M. E., & Kirkpatrick, K. (2019). Experience matters: the effects of hypothetical versus experiential delays and magnitudes on impulsive choice in delay discounting tasks. *Brain Sciences*, *9*(12), 379. https://doi.org/10.3390/brainsci9120379.

Steele, C. C., Steele, T. J., Gwinner, M., Rosenkranz, S. K., & Kirkpatrick, K. (2021). The relationship between dietary fat intake, impulsive choice, and metabolic health. *Appetite*, *165*, 105292. https://doi.org/10.1016/j.appet.2021.105292.

Stephens, D. W., Kerr, B., & Fernández-Juricic, E. (2004). Impulsiveness without discounting: the ecological rationality hypothesis. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *271*(1556), 2459-2465.

Stevens, J. & King, A. (2013). The life of others: Social rationality in animals. In R. Hertwig, U. Hoffrage & the ABC Research Group (Eds.) *Simple heuristics in a social world* (409-431). New York: Oxford University Press.

Stephens, D. W., McLinn, C. M., & Stevens, J. R. (2002). Discounting and reciprocity in an iterated prisoner's dilemma. *Science*, *298*(5601), 2216-2218.

Stevens, J. R., Hallinan, E. V., & Hauser, M. D. (2005). The ecology and evolution of patience in two New World monkeys. *Biology Letters*, *1*(2), 223-226.

Stevens, J. R., Rosati, A. G., Heilbronner, S. R., & Mühlhoff, N. (2011). Waiting for grapes: Expectancy and delayed gratification in bonobos. *The International Journal of Comparative Psychology, 24*, 99-111.

Stevens, J. R., Rosati, A. G., Ross, K. R., & Hauser, M. D. (2005). Will travel for food: spatial discounting in two new world monkeys. *Current Biology*, *15*(20), 1855-1860.

Stevens, J. R., & Stephens, D. W. (2008). Patience. *Current Biology*, *18*(1), R11-R12.

Stevens, J. R., & Stephens, D. W. (2010). The adaptive nature of impulsivity. In G. J. Madden & W. K. Bickel (Eds.), *Impulsivity: The behavioral and neurological science of discounting* (pp. 361–387). American Psychological Association. https://doi.org/10.1037/12069-013

Stevenson, M. K. (1992). The impact of temporal context and risk on the judged value of future outcomes. *Organizational Behavior and Human Decision Processes*, *52*(3), 455-491.

Subhash, S., & Cudney, E. A. (2018). Gamified learning in higher education: A systematic review of the literature. *Computers in Human Behavior*, *87*, 192-206.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Takahashi, T. (2006). Queuing theory under competitive social foraging may explain a mathematical equivalence of delay and probability in impulsive decision-making. *Medical Hypotheses*, *67*(2), 276-279.

Takahashi, T. (2009). Theoretical frameworks for neuroeconomics of intertemporal choice. *Journal of Neuroscience, Psychology, and Economics*, *2*(2), 75-90.

Tanaka, D., Aoki, R., Suzuki, S., Takeda, M., Nakahara, K., & Jimura, K. (2020). Self-controlled choice arises from dynamic prefrontal signals that enable future anticipation. *Journal of Neuroscience*, *40*(50), 9736-9750.

Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift für tierpsychologie*, *20*(4), 410-433.

Tobin, R. M., & Graziano, W. G. (2010). Delay of gratification: A review of fifty years of regulation research. *Handbook of Personality and Self-Regulation*, 47-63.

Tooby, J., & Cosmides, L. (2016). The theoretical foundations of evolutionary psychology. In D. M. Buss (Ed.), *The handbook of evolutionary psychology: Foundations* (pp. 3–87). John Wiley & Sons, Inc.

Trimmer, P. C. (2013). Optimal behaviour can violate the principle of regularity. *Proceedings of the Royal Society B: Biological Sciences*, *280*(1763), 20130858.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *The Quarterly Review of Biology*, *46*(1), 35-57.

Tversky, A., & Griffin, D. (1991). Endowment and contrast in judgments of well-being. *Subjective well-being: An interdisciplinary perspective*, *21*, 101-118.

Vamplew, P., Smith, B. J., Källström, J., Ramos, G., Rădulescu, R., Roijers, D. M., ... & Foale, C. (2022). Scalar reward is not enough: A response to Silver, Singh, Precup and Sutton (2021). *Autonomous Agents and Multi-Agent Systems*, *36*(2), 1-19.

van den Bos, W., Golka, P. J., Effelsberg, D., & McClure, S. M. (2013). Pyrrhic victories: The need for social status drives costly competitive behavior. *Frontiers in Neuroscience*, *7*, 1-11.

van den Bos, W., Li, J., Lau, T., Maskin, E., Cohen, J. D., Montague, P. R., & McClure, S. M. (2008). The value of victory: social origins of the winner's curse in common value auctions. *Judgment and Decision Making*, *3*(7), 483-492.

van den Bos, W., & McClure, S. M. (2013). Towards a general model of temporal discounting. *Journal of the Experimental Analysis of Behavior*, *99*(1), 58-73.

van den Bos, W., Talwar, A., & McClure, S. M. (2013). Neural correlates of reinforcement learning and social preferences in competitive bidding. *Journal of Neuroscience*, *33*(5), 2137-2146.

Vanderveldt, A., Green, L., & Myerson, J. (2015). Discounting of monetary rewards that are both delayed and probabilistic: delay and probability combine multiplicatively, not additively. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(1), 148-162.

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, *27*(5), 1413-1432.

Von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

Webb, T. L., & Young, M. E. (2015). Waiting when both certainty and magnitude are increasing: Certainty overshadows magnitude. *Journal of Behavioral Decision Making*, *28*(3), 294-307.

Wikenheiser, A. M., Stephens, D. W., & Redish, A. D. (2013). Subjective costs drive overly patient foraging strategies in rats on an intertemporal foraging task. *Proceedings of the National Academy of Sciences*, *110*(20), 8308-8313.

Williams, B. A., & Fantino, E. (1978). Effects on choice of reinforcement delay and conditioned reinforcement. *Journal of the Experimental Analysis of Behavior*, *29*(1), 77-86.

Wilson, D. S. (2004). What is wrong with absolute individual fitness? *Trends in Ecology & Evolution*, *19*(5), 245-248.

Wright, S. (1931). Evolution in Mendelian populations. *Genetics*, *16*(2), 97-159.

Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, *144*(2), 140-176.

Yi, R., Buchhalter, A. R., Gatchalian, K. M., & Bickel, W. K. (2007). The relationship between temporal discounting and the prisoner's dilemma game in intranasal abusers of prescription opioids. *Drug and Alcohol Dependence*, *87*(1), 94-97.

Yi, R., Johnson, M. W., & Bickel, W. K. (2005). Relationship between cooperation in an iterated prisoner's dilemma game and the discounting of hypothetical outcomes. *Learning & Behavior*, *33*(3), 324-336.

Young, M. E. (2017). Discounting: A practical guide to multilevel analysis of indifference data. *Journal of the Experimental Analysis of Behavior*, *108*(1), 97-112.

Young, M. E., & Howatt, B. C. (2022). Learning when to wait and when to act. *Learning & Behavior*, *50*(2), 207-221.

Young, M.E. & Howatt, B.C. (2023). Cue transparency leads to a myopic decision habit: Optimal training ratios to avoid overshadowing. *Learning and Motivation, 82*.

Young, M.E. & Howatt, B.C. (submitted). When smaller sooner depletes a pool of resources faster.

Young, M. E., & McCoy, A. W. (2015). A delay discounting task produces a greater likelihood of waiting than a deferred gratification task. *Journal of the Experimental Analysis of Behavior, 103*(1), 180-195.

Young, M. E., & McCoy, A. W. (2019). Variations on the balloon analogue risk task: A censored regression analysis. *Behavior Research Methods*, *51*(6), 2509-2521.

Young, M. E., Vangsness, L., & McCoy, A. W. (2018). The temporal dynamics of waiting when reward is increasing. *Behavioural Processes, 149*, 16-26.

Young, M. E., Webb, T. L., & Jacobs, E. A. (2011). Deciding when to "cash in" when outcomes are continuously improving: An escalating interest task. *Behavioural Processes*, *88*(2), 101-110.

Young, M. E., Webb, T. L., Rung, J. M., & Jacobs, E. A. (2013). Sensitivity to changing contingencies in an impulsivity task. *Journal of the Experimental Analysis of Behavior*, *99*(3), 335-345.

Young, M. E., Webb, T. L., Rung, J. M., & McCoy, A. W. (2014). Outcome probability versus magnitude: When waiting benefits one at the cost of the other. *PloS one*, *9*(6), e98996.

Zhang, H., Moisan, F., & Gonzalez, C. (2021). Rock-paper-scissors play: Beyond the win-stay/lose-change strategy. *Games*, *12*(3), 52. https://doi.org/10.3390/g12030052.

Appendix A - Total Points Obtained

Figure A.1

Total Points Distributions

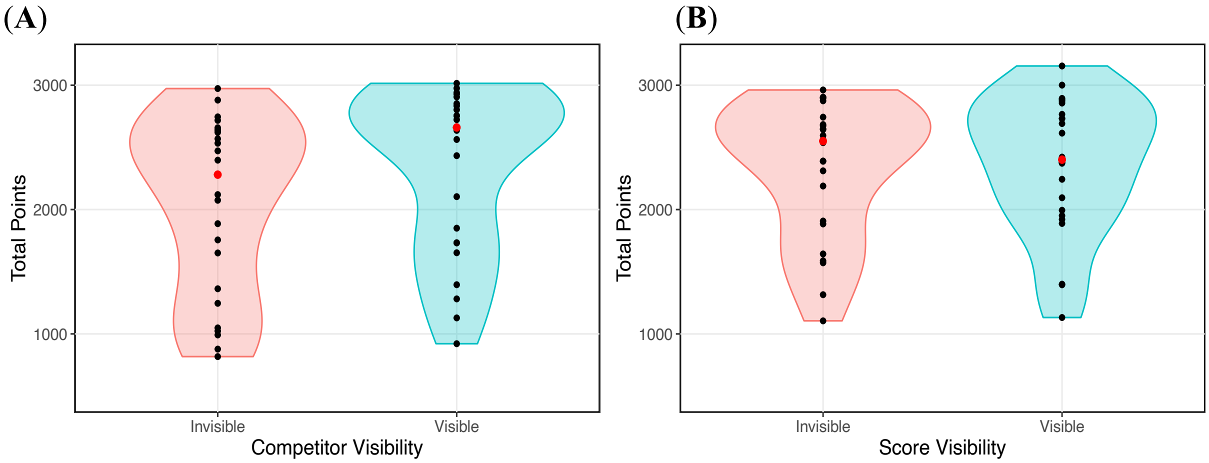

Note: Violin plots showing the distribution of participants' total reward points obtained across Levels 2-5 in both Experiment 1(**A**) and Experiment 2 (**B**). The black dots are the individual participants, and the red dots are the median total points of each visibility group.

Appendix B - Level 1 Responses

Figure B.1

Level 1 Likelihood of Waiting to Respond in Experiment 1

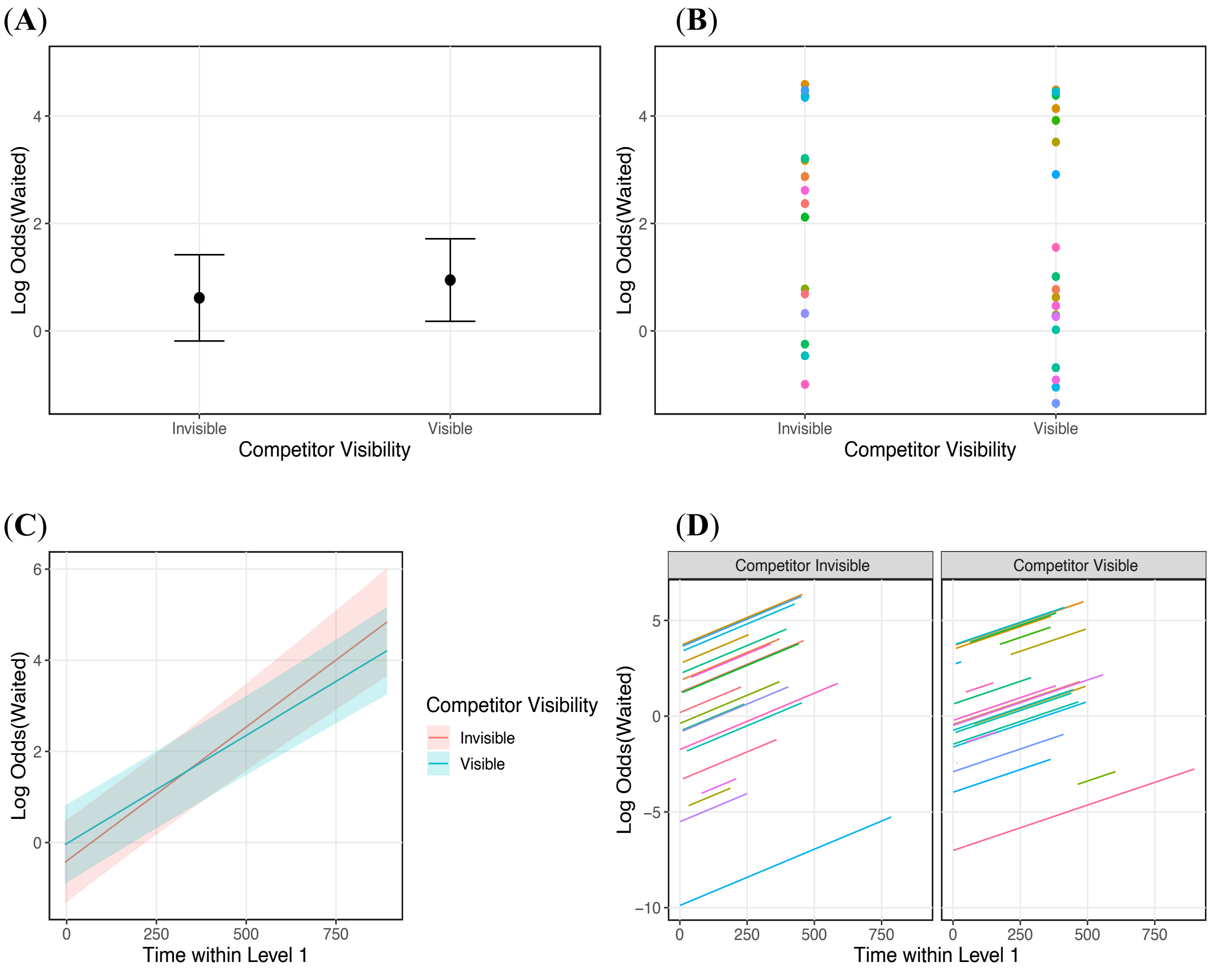

Note: Experiment 1's model-estimated log odds of waiting to respond in Level 1 as a function of the competitor-visibility group. Note that the competitor was removed in Level 1, so these analyses serve as a baseline of participants' behavior. (**A**) The grand-average log odds of waiting for each visibility group. (**B**) Individual participants' estimated average log odds of waiting. (**C**) How each competitor-visibility group's estimated grand-average log odds of waiting changed over time in Level 1. (**D**) Individual participant slopes estimating their change in log odds of waiting to respond over time in Level 1.

Figure B.2

Level 1 Likelihood of Waiting to Respond in Experiment 2

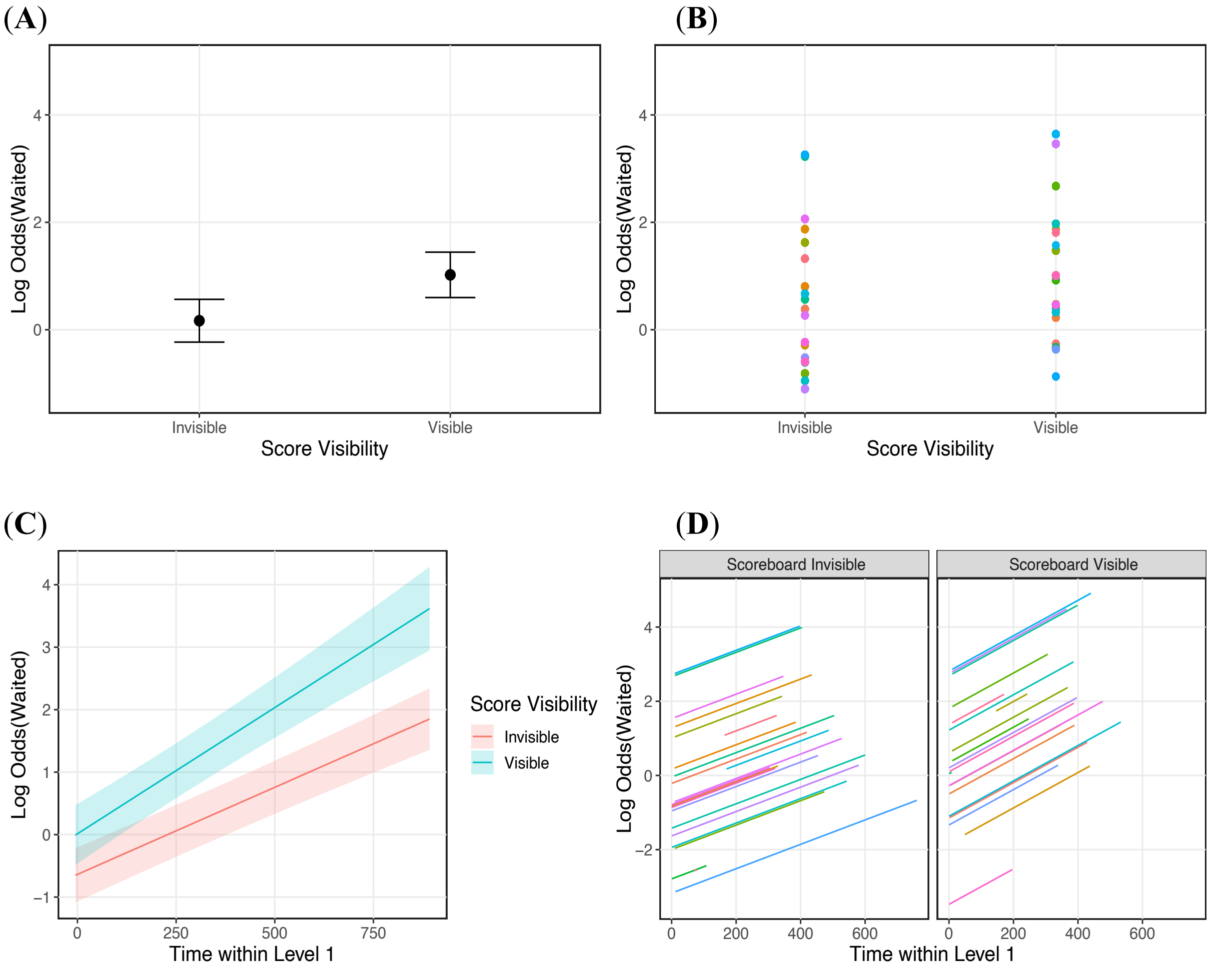

Note: Experiment 2’s model-estimated log odds of waiting to respond in Level 1 as a function of the competitor score-visibility group. The competitor was again removed in Level 1. (**A**) The grand-average log odds of waiting for each visibility group. (**B**) Individual participants’ estimated average log odds of waiting. (**C**) How each score-visibility group’s estimated grand-average log odds of waiting changed over time in Level 1. (**D**) Individual participant slopes estimating their change in log odds of waiting to respond over time in Level 1.

Appendix C - Separability

In Experiment 1's analysis of participants' probability of waiting to respond, there was a disagreement between the magnitude of the effect size for the competitor's mean waiting time (WT) and the model fits. Specifically, there was decisive evidence that the competitors' mean WT greatly influenced participants' probability of waiting, yet WAIC scores did not support models that included this predictor over simpler models that excluded it. This disagreement can likely be attributed to the strength of the manipulation of the competitor's mean WTs and the random effects structures.

First, the strength of the mean WT manipulation resulted in 'separability', which occurs when a predictor is associated with only outcome value. In this experiment, participants (almost) never waited to respond when playing against impulsive competitors, regardless of their visibility or WT variance. Separability can result in a lack of model convergence, unstable parameter estimates, and greater uncertainty in the model's fit because there could be an infinite number of estimates that differentiate peoples' responding against impulsive and patient competitors. In Bayesian analyses, separability can possibly be mitigated by using stronger priors that limits the range of plausible parameter estimates the sampling algorithm evaluates. Therefore, I re-ran the final model reported in the main text several times and tested different priors that more strongly constrained the sampling around zero (e.g., mean WT ~ $N(0, 1)$, mean WT ~ $N(0, 0.5)$) or around the posterior that was initially observed, $B = -4.28$, $SE = 0.30$ (e.g., mean WT ~ $N(-4, 1)$, mean WT ~ $N(-4.5, 0.5)$). Note that setting a prior distribution on a parameter estimate from a posterior distribution derived from the same data is usually considered poor practice, but I was doing this for testing purposes, not inferential purposes. Results across these various models were consistent, suggesting that the model was converging on the same solution and separability was not driving the differences in parameter estimates and model fits. In contrast, I ran the same model as a frequentist-based multilevel logistic regression using the *lme4* library (Bates, Mächler, Bolker & Walker, 2014) and encountered major convergence issues, poor parameter estimates and extremely large standard errors (some *SE*s > 300). Thus, I feel confident that the disagreement is not due to just separability.

Second, I probed the quality of the regression models using the *performance* library (Lüdecke et al., 2021) and discovered that this disagreement is instead likely due to the random effects structure of allowing the slope of the competitor's mean WT to vary across participants. In essence, the random slope effect of the competitor's mean WT was accounting for such a large amount of participants' behavior, that the fixed effect of the competitor's mean WT was not contributing unique information to the predictive accuracy of the model. Figure C.1 illustrates the model fits for each individual participant from the model reported in the main text (see Figure 11), and it is clear that the probabilities of waiting to respond against impulsive competitors (mean WT = 3 s) is strongly homogeneous (which also relates to the separability issue above). I then re-ran the regression models without the random slope effect of the competitor's mean WT and found that the models including the fixed effect of the competitor's mean WT fit participants' choices decisively better than models that did not include this predictor. This result is consistent with what I had originally expected, and similar results were found in Experiment 2.

Figure C.1

Individual Participants' Likelihood of Waiting to Respond

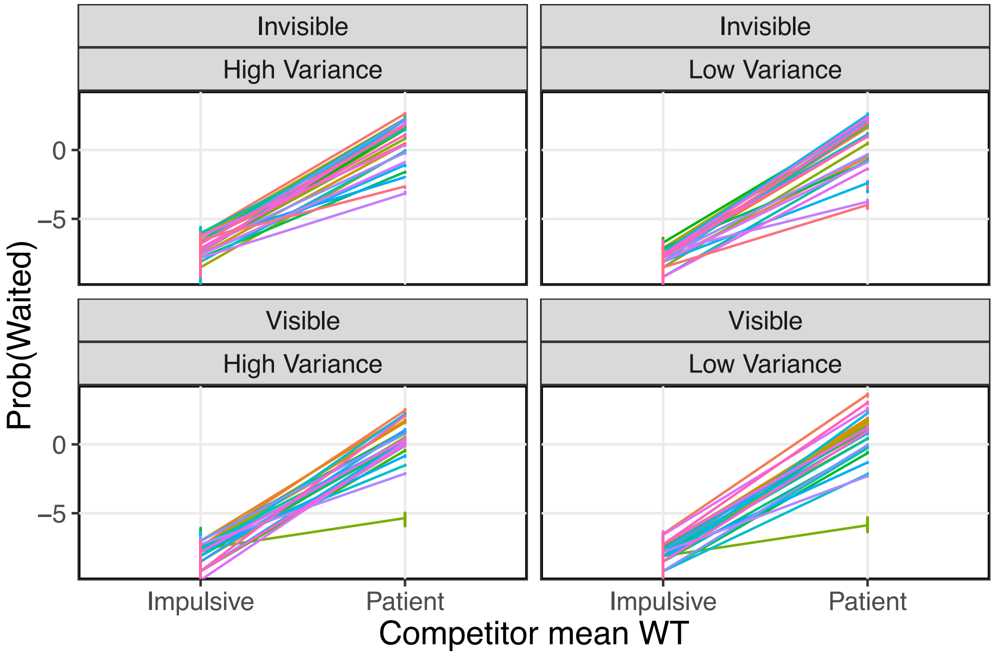

Note: Experiment 1's model-estimated log odds of waiting to respond in Levels 2-5 as a function of the competitor-visibility group, the competitor's mean WT (Impulsive = 3 s; Patient = 9 s) and WT variance. Each line is an individual participant.