The Dynamic Nature of Procrastination

by

Pei Yuan Zhang

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Neural Science

New York University

January, 2024

Dr. Wei Ji Ma

Acknowledgements

My Ph.D. journey that, at times, felt more like an adventure. This thesis, more than a mere compendium of my research, marks the culmination of an extraordinary adventure – one teeming with the thrill of discovery, the rigors of challenge, and, admittedly, a fair share of successes and failures. I am proud to walk this path as the first in my family to achieve a college degree, now even venturing into the realms of advanced academia. Here, I am, at its end, immeasurably grateful to those who have guided and supported me through this remarkable adventure.

My adventure took me from Nanjing to Beijing, then to Berlin, and finally to New York City. Along the way, I transitioned from studying Physics to diving into Computational Neuroscience and ultimately found my passion in Computational Cognitive Science.

At the heart of this journey has been my mentor, Wei Ji Ma. He has pulled me out of a really bad situation in another lab and has supported me to study a subject that ignites my deepest curiosities – the dynamic nature of procrastination. Wei Ji has chosen to tread a path less traveled, embracing the inherent risks of a pioneering topic that demands ingenuity and interdisciplinary exploration. The five years working with Wei Ji have been full of joy and excitement.

My collaborators, Yijun Lin and Falk Lieder, have been the co-navigators in this expedition. Their insights and partnership have been nothing short of indispensable. To my circle of scientific comrades – Karolina Lempert, Evgeniya Lukinova, Cate Hartley, Silvia Lopez-Guzman, Paul Glimcher, Joseph Kable, Sophie Arnold, Heiko Schuett, Shenglong Wang, David Bosch, Todd Gureckis, David Amodio, Zhiwei Li, and Peter Gollwitzer – each of you has lent a unique perspective, a critical eye, and a helping hand through the various stages of this research.

Here, especially, I want to thank Brenda Woodford for the arduous work of extracting the students'research participation data from NYU's Sona system for the study in Chapter 3.

This academic trek was more than a pursuit of knowledge; it was enriched by the joy of companionship of my friends. To all of my friends in New York: Yotta, Zhu zhu, Jane, Monica, Shaonan, Sixuan, Dongqi, Qixiu, Sam, Mingyu and 706 community; to my friends in Dali (during COVID): Mai, Yifeng, Huoshi, Way, Boyun, Caodiao, Dengrang, and the subject of my documentary film Tang; to my friends in Beijing and Shanghai: Jianru, Ruiqing, Xiaoyue, liliu and Qizheng. As it is impossible to address everyone, I apologize for any omissions.

I can't forget the incredible people in Wei Ji's lab—who are not just colleagues but friends. Each one of you has played a part in this journey, giving me feedback that sharpened my ideas and helping me do better presentations. Your encouragement, especially on days when things didn't go as planned, meant the world to me. Your support turned the lab into a place I looked forward to being in every day. Thank you for being an essential part of my PhD adventure and for all the laughs and wisdom along the way.

Lastly, a heartfelt thank you to my partner, Jinglong Pang, for the love and stability you've brought into my life during this Ph.D. journey. Your decision to move from China to New York to be with me has touched my heart deeply. The meals you've lovingly prepared and your soothing words during stressful times have been my comfort. You've been a beacon of positivity, gently reminding me of my worth and capabilities, especially when I've been filled with self-doubt.

Your encouragement has quietly lifted me up, reinforcing my steps toward this significant achievement. For all the times you've been both my cheerleader and my solace, I am truly grateful. Your support means more to me than these words can convey. Thank you for being my steadfast support, Pangpang!

Preface

The work described in Chapter 2 is being prepared for journal submission. A short version is available as an abstract submitted to the Cognitive Computational Neuroscience (CCN) conference 2019 (Zhang and Ma, 2019). It can be accessed here. Pei Yuan will be the first author and Wei Ji will be the senior author.

The work presented in Chapter 3 is a preprint (Zhang and Ma, 2023b) and can be accessed here. Pei Yuan is the first author and Wei Ji is the senior author.

The work presented in Chapters 4 and 5, in collaboration with Ph.D. candidate Yijun Lin from the University of Florida and Ph.D. Falk Lieder from The University of California, Los Angeles (UCLA), is being prepared for journal submission. A short version has been submitted as an abstract to CCN 2023 (Zhang and Ma, 2023a) and can be accessed here. Pei Yuan will be the first author, Yijun Lin will the second author, Falk Lieder will be the third author, and Wei Ji will be the senior author.

Abstract

People procrastinate. Procrastinators usually make little progress at the start (start late) and a significant increase in progress shortly before the deadline (rush to complete). The more progress ramps up over time, the greater the level of procrastination. Yet, the cognitive mechanisms underlying this intriguing dynamic feature of procrastination–time course of progress–remain poorly understood.

To investigate this, we started by enhancing the theoretical understanding of the time course of progress. We proposed a normative model (Chapter 2) that considers the time course of progress as the output of sequential decision-making: whether to work now (and, if so, how much) or later. If they decide to work now, they pay the cost of investing mental effort immediately but also make progress, and more work leads to more progress. If they decide not to work, they make no progress and also pay no effort cost. The amount of work on each day is derived from the Bellman equation, which assumes that a person's goal on each day is to maximize the discounted value gained by making progress while minimizing the immediate effort cost. This model predicted three patterns of procrastination: a delay in the beginning and then ramping up, working at the last minute, and not working at all. It also identified several correlates of procrastination, including perfectionism, the shape of the cost function, temporal discounting, and the total given time. Additionally, the model predicted the effect of interventions, such as offering immediate rewards, on reducing procrastination and improving performance.

Next, we asked: What is the key factor contributing to procrastination in the real world?

We tested one key potential correlate of procrastination–temporal discounting–with a long-term real-world behavioral study (Chapter 3). We observed a positive correlation between individual discount rates estimated from an inter-temporal choice task and the levels of procrastination in a research participation task assigned by a psychology course, suggesting that temporal discounting is a cognitive mechanism underlying procrastination.

The observed association between temporal discounting and procrastination suggested that if we brought a future reward temporally closer, then that would help reduce procrastination. So, we tested the causal relationship between reward timing and procrastination (Chapter 4). We asked: Can offering immediate rewards help reduce procrastination regardless of reward rule? To answer the question, we created a novel experimental paradigm named BORE (Boring Online Reading Experiment) that mimics real-world procrastination while still allowing for manipulation. People worked on a self-paced, week-long online reading task consisting of numerous units of work that take about 3 hours to complete. We utilized a between-subject design, crossing two levels of reward timing (either delayed or immediate upon task completion) and three levels of reward rules. Our results revealed that offering an immediate reward upon task completion helped people start the task earlier, helped people who generally procrastinate more complete the task earlier, and helped them complete units of work earlier. Both of which held true regardless of reward rules.

Last but not least, we asked: What is the cognitive process underlying the time course of progress? We proposed two models and fit them into the rich data of the time courses of progress that we collected from the BORE (Chapter 5). The first model was the normative model that we discussed before. The second model is the roll-out model. It shares the normative model's sequential decision-making framework and computational goal. However, the value of making progress is computed instead by simulating future time courses with an anticipation of the average workload in the future. We found that the normative model provided a poor fit to the data, while the roll-out model fit the data quite well. Therefore, we found some evidence against people behaving rationally. We had some evidence for people simulating their future work progress, and they had an idea of how much work they would do on average in the next few days.

Taken together, our findings enhance our understanding of the cognitive mechanism of the dynamic nature of procrastination and offer implications for reducing procrastination. This thesis shows a successful attempt at applying cognitive science to the real world.

Contents

List of Figures

List of Tables

1 Introduction

1.1 The dynamic nature of procrastination

People procrastinate—for example, take me. When writing my undergraduate thesis in physics, ideally, I split work into two months, making some progress in the first month and gradually increasing the workload over time until the deadline day (Fig. 1.1 upper panel). However, in reality, I made little progress until the last week (Fig. 1.1 lower panel).

This little progress in the beginning and rushing in the end is a typical temporal pattern of work progress in procrastinators (Schouwenburg and Groenewoud, 2001; Steel et al., 2018; Konradt et al., 2021; Moon and Illingworth, 2005; Vangsness and Young, 2020; Zhang and Ma, 2023a) (Fig. 1.2). The more progress ramps up over time, the greater the level of procrastination.

As a cognitive scientist, this dynamic feature of procrastination fascinates me. I am curious about the cognitive mechanism underlying it. So, I dove into the literature to find some clues.

It turns out that procrastination is quite a popular topic both inside and outside of academics. Outside academia, authors and researchers have written self-help books offering advice to a general audience on overcoming procrastination (Fiore, 2007; Burka and Yuen, 2007; Ludwig and Schicker, 2018). Additionally, there are lots of discussions about procrastination on social media and in educational media. For instance, blogger Tim Urban gave a TED talk titled "Inside the mind of a master procrastinator," which garnered millions of views (Urban, 2016).

Within academics, there has been a growing number of papers every year studying procras-

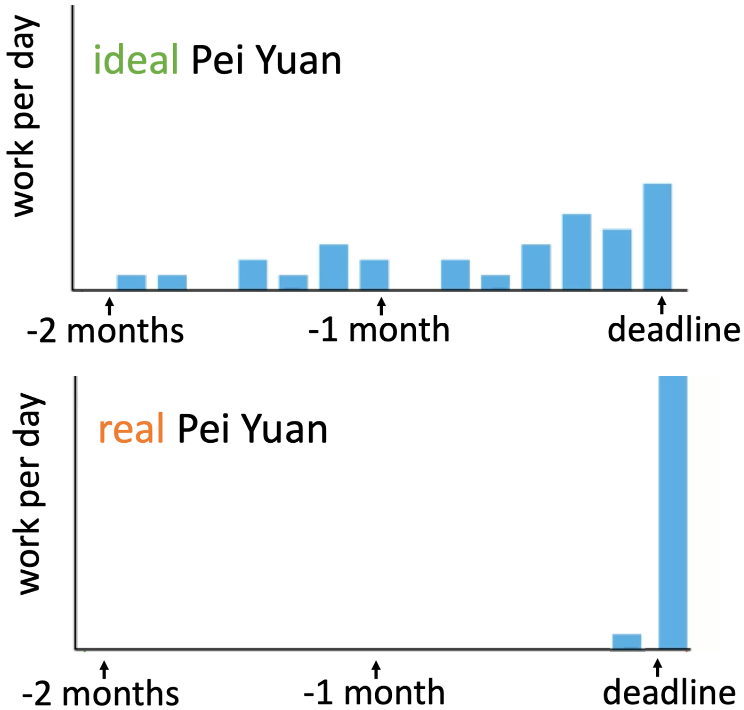

Figure 1.1: How Pei Yuan wrote her undergraduate thesis in Physics: the amount of work per day over two months before the deadline day. Upper panel: The ideal version. Lower panel: The reality.

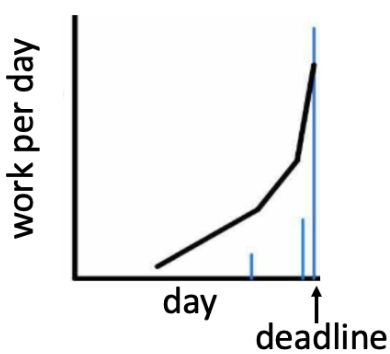

Figure 1.2: A typical temporal pattern of work progress in procrastinators: little progress in the beginning and a significant amount of progress approaching the deadline.

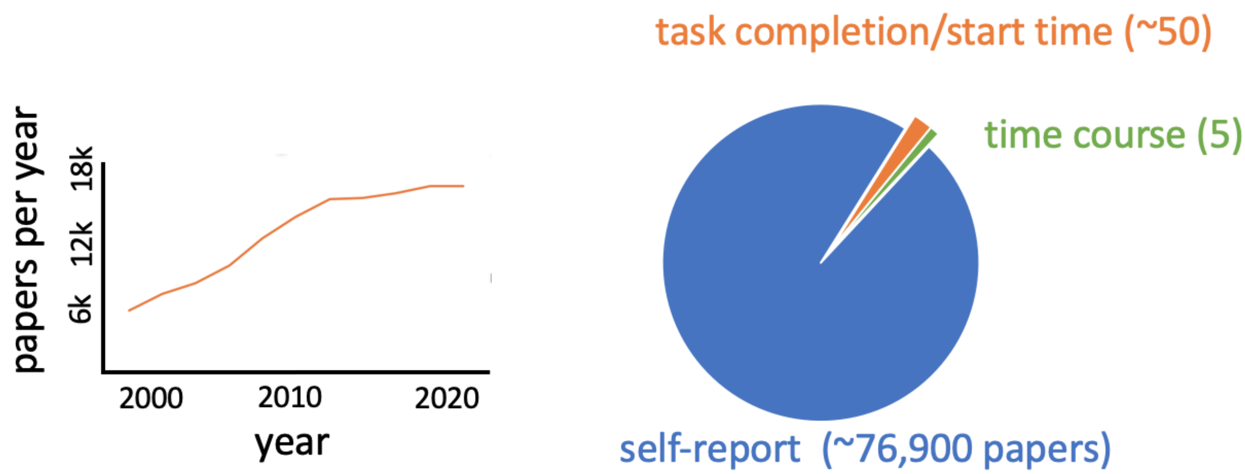

Figure 1.3: Few studies characterized procrastination through the time course of progress, even if there has been a growing number of papers per year studying procrastination over the past two decades. Left panel: An increasing number of studies of procrastination have been published every year since 2000. Right panel: The pie chart demonstrates a few studies measuring procrastination through time courses of progress.

tination over the past two decades (Fig. 1.3 left panel). It is a multi-disciplinary subject studied across various fields such as personality and industrial psychology (Steel and Klingsieck, 2016; Steel, 2007; Schouwenburg and Lay, 1995; Kim et al., 2017), clinical psychology (Scent and Boes, 2014; van Eerde and Klingsieck, 2018); education (Diver and Martinez, 2015; Putnik et al., 2013), behavioral economics (O'Donoghue and Rabin, 1999; O'Donoghue and Rabin, 2001; Fischer, 1999; Fischer, 2001), and neuroscience (Raphaël and Mathias, 2022; Zhang and Ma, 2019).

In terms of research topics, research on procrastination primarily revolves around three main aspects: the correlates of procrastination, the consequences of procrastination, and interventions to reduce procrastination.

Regarding the correlates, studies mainly focus on personality traits, such as perfectionism (Flett et al., 1995; Flett et al., 1992), impulsiveness (Steel and Klingsieck, 2016; Ferrari, 1993), self-efficacy (Klassen et al., 2008), as well as conscientiousness and its facets of self-control, distractibility, organization, and achievement motivation (Steel, 2007). Studies also delve into task characteristics such as task aversiveness (Blunt and Pychyl, 2000; Milgram et al., 1995), task difficulty (Janssen and Carton, 1999; Froese et al., 1984; Ackerman and Gross, 2005), etc. Research also explores other factors such as self-regulation failure (Senécal et al., 1995; Park and Sperling, 2012; Rakes and Dunn, 2010; Ferrari, 2001), fear of failure (Haghbin et al., 2012; Schouwenburg, 1992), self-handicapping (Ferrari and Tice, 2000; Beck et al., 2000; Marshall et al., 2008; Strunk and Steele, 2011), and depression (Stöber, 2001; Constantin et al., 2017).

Regarding the consequences of procrastination, the focus has been on both how people feel and what they achieve. In terms of how people feel, studies discussed stress (Sirois, 2014; Tice and Baumeister, 1997), anxiety (Haycock et al., 1998), guilt (Fee and Tangney, 2000), and low self-esteem (Saleem and Rafique, 2012; Yang et al., 2021). In terms of what they achieve, procrastinators tend to be worse off in terms of academic performance (Kim and Seo, 2015), financial outcome (Martinez et al., 2017; Akerlof, 1991; O'Donoghue and Rabin, 1999), and political decision-making (Farnham, 2021; Kegley, 1989; Holland, 2001).

In terms of interventions, studies explore strategies such as emotion regulation (Eckert et al., 2016), addressing maladaptive beliefs, acceptance strategies (Wohl et al., 2010), commitment or cognitive-behavioral therapy (Scent and Boes, 2014; van Eerde and Klingsieck, 2018), training about time management (C. et al., 2007; Wieber and Gollwitzer, 2010) and process focus instead of an exclusive emphasis on the ultimate goals (Krause and Freund, 2014a).

However, little is known about the underlying cognitive mechanisms of procrastination. Few studies have tried to understand procrastination from a cognitive science perspective. We do not yet know the cognitive components that influence the time course of work progress, nor do we understand how these components interact with each other over time.

What also surprises me is that the dynamic feature of procrastination—the time course of work progress—is rarely studied (Fig. 1.3 right panel). Most studies, numbering approximately 76,900 papers, measure procrastination using self-reported surveys (Steel, 2007; Pychyl and Flett, 2012; Sirois et al., 2003). Only a few studies, approximately 50 papers, have examined behavior, either through observed behavior or based on self-reflection. Among these, most have characterized procrastination by a single time point only: task starting time (Ariely and Wertenbroch, 2002; Ariely and Wertenbroch, 2002; Buehler and Griffin, 2003; Owens et al., 2008; Niermann and Scheres, 2014; Milgram et al., 1992; McElroy and Lubich, 2013; Lubbers et al., 2010; Janssen and Carton, 1999; Hensley, 2014; Diver and Martinez, 2015) or task completion time (Ferrari and Scher, 2000; Roberts et al., 1988; Reuben et al., 2015; Pittman et al., 2008; McElroy and Lubich, 2013; McCrea et al., 2008; Malatincová, 2015; Lim, 2016; Janssen and Carton, 1999; Hensley, 2014; Green, 1982; De Paola and Scoppa, 2015; Cerezo et al., 2017; Buehler and Griffin, 2003; Ariely and Wertenbroch, 2002; Raphaël and Mathias, 2022), or the duration of work (Ji Won You, 2015; Xu et al., 2016; Onji, 2013; Niermann and Scheres, 2014; Lim, 2016; Liborius et al., 2019; Lay, 1987; Krause and Freund, 2014b; Kerdijk et al., 2015; Johnson et al., 2016; Elvers et al., 2003; Ackerman and Gross, 2005).

However, examining, for instance, the completion time alone does not differentiate between two individuals who complete a task at the same time but with different temporal patterns of work progress—one making steady progress toward the end (like the ideal Pei Yuan) and the other rushing to complete at the end (like the real Pei Yuan). Similarly, looking at the task starting time alone cannot differentiate between two people who start the task at the same time but with different work patterns. Also, looking at the duration of work alone fails to capture variations in the temporal patterns of work progress—one works in an early phase, and the other works in a later phase.

To my knowledge, merely five studies have analyzed the time course of progress and tried to categorize its patterns, including both observed behavior and self-reported estimation of effort or work hours (observed behavior: Moon and Illingworth, 2005; Steel et al., 2018; Vangsness and Young, 2020; self-reported effort or work hours: Konradt et al., 2021; Schouwenburg and Groenewoud, 2001).

We assert that unraveling the dynamic nature of procrastination—the time course of work progress—is crucial as it provides insights into the underlying cognitive mechanisms of procrastination.

The significance of studying the temporal patterns of work progress to enhance our understanding of procrastination has been echoed by seasoned researchers in the fields of procrastination and motivation (Steel et al., 2018; Roe, 2014). Roe advocates for a more nuanced research approach, recommending not just using longitudinal designs but also more detailed observations over time. He notes that while longitudinal studies are rare in motivational research, the few available often sample two or three chosen time points. Roe argues that, instead, a more thorough sampling across the entire duration of the task is essential. Such a meticulous approach is vital to understanding the nature of how people manage goal pursuit over time (Lord et al., 2010). Unfortunately, exceedingly few studies have met any of these criteria, which are crucial to having a meaningful understanding of the “temporal footprint of work” (Steel et al., 2018).

In this thesis, we aim to characterize procrastination through the "temporal footprint of work"—time course of work progress—and uncover its underlying cognitive mechanisms.

1.2 The fascinating questions

In this thesis, we aimed to characterize procrastination through the time course of work progress and uncover its underlying cognitive mechanisms.

In Chapter 2, we started by contributing to the understanding of the time course of progress from a theoretical aspect.

To date, even though there are many theoretical works on procrastination (Steel and König, 2006; Akerlof, 1991; O’Donoghue and Rabin, 1999; Zhang and Feng, 2020; Fischer, 1999; Fischer, 2001), they fall short of predicting 1) the temporal dynamics of procrastination in terms of the time course of work progress, 2) the negative consequences in terms of both the performance and the feeling of exhaustion, and 3) the interventions to reduce procrastination and improve performance.

We assumed that the time course of progress arises from a sequential decision-making process. We proposed a normative process model that predicts all aspects of procrastination listed above: the time course of work progress, the consequences of procrastination in terms of performance and the feeling of exhaustion, and the interventions to reduce procrastination.

Next, in Chapter 3, we asked: What is the key factor contributing to procrastination? This question was inspired by the predictions of our model in Chapter 2 and also by other theories suggesting that temporal discounting is a cognitive mechanism underlying procrastination (O'Donoghue and Rabin, 1999; O'Donoghue and Rabin, 2001; Fischer, 1999; Fischer, 2001; Steel and König, 2006; Steel, 2007).

When faced with a task in its initial stages, where the eventual reward is distant, people temporarily discount the value of that future reward. As a consequence, the temporarily discounted future reward fails to provide sufficient motivation for people to start working until the deadline looms near. For instance, in thesis writing, the initially discounted future reward makes other activities, like socializing, more appealing. This leads to procrastination until the looming deadline, which significantly increases the utility of the task. These theories predict that individuals with stronger temporal discounting procrastinate more. However, no empirical evidence for this hypothesis exists. Even though there are a few attempts (Raphaël and Mathias, 2022; Xin et al., 2020), both of which failed to find the association.

In Chapter 3, we tested this association in a real-world task. We provided the first compelling evidence for the long-standing hypothesis that temporal discounting is a mechanism underlying procrastination. Note that we quantified procrastination by a statistical summary of the entire time course of progress: mean unit completion day.

Chapter 3 was a correlational study; next, in Chapter 4, we tested the causal relationship between reward timing and procrastination. We asked: Can we reduce procrastination?

The association between temporal discounting and procrastination we observed in Chapter 3 suggested that if we brought a future reward temporally closer, then that would help reduce procrastination. Often, in real-world situations, like submitting a homework assignment, a reward is given after the deadline. Even if we get it done earlier, we still have to wait until the deadline to be rewarded. One way to bring a future reward closer is to receive an immediate reward upon task completion. So, we asked in Chapter 4: Can offering an immediate reward upon completion reduce procrastination?

Existing studies primarily investigate the influences of added immediate rewards on task persistence rather than procrastination (e.g., when to start, when to complete, or even the dynamic process of procrastination), such as integrating fun into workouts or offering points to enhance task completion (Woolley and Fishbach, 2016; Woolley and Fishbach, 2017; Lieder et al., 2019). Moreover, these rewards are additional benefits, not alterations to the original task rewards. To truly assess the effect of reward timing on procrastination, studies should focus on changing the timing of the original task rewards.

We hypothesized that offering immediate rewards helps reduce procrastination, regardless of reward rules. To test our hypothesis, we created a novel experimental paradigm named BORE (Boring Online Reading Experiment). We found that offering immediate rewards helped people start the task earlier and helped those who generally procrastinate more complete the task and units of work earlier, both of which held true regardless of reward rules.

Last but not least, in Chapter 5, we asked: What is the cognitive process underlying procrastination?

Cognitive neuroscientists have extensively studied the cognitive and neural mechanisms of mental effort (Kool and Botvinick, 2018; Shenhav et al., 2017) in the field of cognitive control. The commonly used experimental paradigm asked people to choose repeatedly between performing a high-demand task for a larger amount of money and performing a low-demand task for a smaller amount. They found that exerting mental effort is costly. However, it is unclear how those studies can inform us about how people make mental efforts in their daily lives for long-term projects such as writing articles or building software that often extends over long periods of time, ranging from days to months. Making sustained efforts during those intervals is essential for people to achieve their goals and for a society to function efficiently.

Neuroscientists have used reinforcement learning to study how animals allocate time between work and leisure over time (Niyogi, Shizgal, et al., 2014; Niyogi, Breton, et al., 2014). Animals, however, tend to maximize instantaneous reward rate, in contrast to humans, who allocate effort over time towards a temporally distant goal.

There is also a growing number of efforts to account for and explain goal pursuit using computational process models (Lieder et al., 2019; Prystawski et al., 2022; Singhi et al., 2023). Still, these theoretical works focus on difficult decision-making problems rather than the simple temporal problems we face every day, like working to finish a project.

To uncover the underlying cognitive processes of procrastination, we proposed two models and fit them into the data that we collected in Chapter 4. The first model is the rational model, which we have discussed intensively in Chapter 2. The second model is the roll-out model. It shares the normative model's sequential decision-making framework and computational goal. However, the value of making progress is computed instead by simulating future time courses with an anticipation of the average workload in the future. We found that the roll-out model fits better than the rational model. We found some evidence against people behaving rationally. We had some evidence for people simulating their future work progress, and they had an idea of how much work they would do on average in the next few days.

1.3 The challenges

We highlight three challenges. The first is to characterize procrastination through the time course of work progress. To accomplish this, a real-world task must meet the following criteria:

1. A clear definition of the work unit.

2. The ability to measure or ascertain when each work unit is completed.

3. The involvement of multiple units of work to establish a time course of work progress.

Many real-world tasks, such as writing or taking an academic course, often lack clearly defined units of work. For instance, in writing, is the work unit a sentence, a paragraph, or even a page? Additionally, there are real-world tasks with clearly defined units of work where the completion of each unit cannot be readily accessed. Also, some tasks, like mailing a letter, require only a single action and, therefore, don't involve multiple units of work.

We found a real-world task that met all the criteria in Chapter 3, and we created a novel task paradigm that met all the criteria as well in Chapter 4.

The second challenge involves designing a task to test if offering immediate rewards upon completion reduces procrastination. An ideal task should mimic real-world procrastination while still allowing for manipulation. It is often the case that real-world behavior is observational and hard to manipulate. A controlled experiment allows for manipulation but is far from the real world. We created a novel task that bridges these two in Chapter 4.

The task is named BORE (Boring Online Reading Experiment). The key attributes that made it a successful design were:

1. Spanning multiple days: Participants were assigned seven days to work on a lengthy reading task, which takes about 3 hours to complete all the reading units.

2. Being online and self-paced, granting the participants the autonomy to decide when to work and where to work.

3. Being intentionally boring: the task was deliberately made to be boring, with the reading difficulty set at a 7-year-old child's level.

The third challenge involves developing computational process models that fit the time course of progress data. The process model that fits the data well will uncover the underlying cognitive process of procrastination.

2 | A normative account of temporal dynamics of procrastination

Let's start by building a normative model
that predicts the time course of work progress:
the dynamic nature of procrastination!

2.1 Introduction

Procrastination permeates people's lives. To some degree, almost everyone procrastinates at certain things, both inside and outside of academia. People delay filing their taxes until the last minute (Martinez et al., 2017). Researchers postpone until the last minute to submit abstracts or turn in papers for academic conferences (Flandrin, 2010). College students rush to complete their work on academic tasks such as self-paced quizzes (Solomon and Rothblum, 1984; Rothblum et al., 1986; Steel et al., 2001; Howell et al., 2006), course assignments (Hensley, 2014; Cerezo et al., 2017; Lim, 2016), as well as administrative tasks (De Paola and Scoppa, 2015; Reuben et al., 2015). They also delay starting their work (Green, 1982; Niermann and Scheres, 2014; Solomon and Rothblum, 1984; Rothblum et al., 1986; Steel et al., 2001; Howell et al., 2006).

If we characterize the dynamic process of procrastination by the time course of work progress, which consists of how much work is done each day, typical patterns of the time course of work progress emerge. Often, there is a delay in working in the beginning (start late) and an increase in the amount of work shortly before the deadline (rush to complete) (Schouwenburg and Groenewoud, 2001; Steel et al., 2018; Konradt et al., 2021; Moon and Illingworth, 2005; Vangsness and Young, 2020; Zhang and Ma, 2023a).

In addition, there are negative consequences associated with procrastination. Procrastinators tend to be worse off regarding their achievements and feelings. Regarding achievements, procrastination can prevent individuals from completing their work on time (Steel, 2007; Kim and Seo, 2015; Solomon and Rothblum, 1984). People often fail to catch up with work by the deadline or even abandon work at the last minute (Wrosch et al., 2003). In terms of feelings, even though some people manage to meet the deadline at the last minute, rushing to do so by staying up late can lead to severe health costs. This often leaves individuals feeling exhausted and overworked (Tice and Baumeister, 1997; Sirois et al., 2003; Sirois, 2007).

Furthermore, several interventions addressed in the literature were suggested to reduce procrastination or improve performance. For example, immediate reward upon task completion reduces procrastination compared to delayed reward (Zhang and Ma, 2023a). Rewarding progress immediately was suggested to improve performance (Lieder et al., 2019; Milkman et al., 2014). Interim deadlines were suggested to reduce procrastination (Wesp, 1986; Burger et al., 2011; Lamwers and Jazwinski, 1989), and the empirical study showed that they increased final performance (proportion completed) (Ariely and Wertenbroch, 2002).

However, to date, existing theories of procrastination largely fall short of predicting 1) the temporal dynamics of procrastination in terms of the time course of work progress, 2) the negative consequences in terms of both performance and the feeling of exhaustion, and 3) the interventions to reduce procrastination and improve performance (Steel and König, 2006; Akerlof, 1991; O'Donoghue and Rabin, 1999; Zhang and Feng, 2020).

Temporal motivation theory proposed by Steel and Konig describes how the utility of working changes over time, which has the potential to describe the temporal dynamics of work progress (Steel and König, 2006). However, since it is a descriptive theory, it cannot generate a sequence of actions to predict procrastination behavior, nor can it evaluate the consequences of procrastination or interventions for procrastination. GA Akerlof, O'Donoghue, and Rabin proposed a process model to study procrastination behavior (Akerlof, 1991; O'Donoghue and Rabin, 1999). They modeled for a particular case where a task requires acting only once rather than the common long-term real-life task, such as writing a Ph.D. thesis, where work must be divided into multiple actions to be completed over time. Similarly, the theory proposed by Zhang and Feng (Zhang and Feng, 2020) predicts the number of days of delay but does not address the temporal dynamics of work progress after people start working nor evaluate the consequence of procrastination or the interventions. The most close attempt is a process model proposed by Fischer (Fischer, 1999; Fischer, 2001). The model assumes that the task requires a certain number of hours to complete, and the remaining time for leisure is an exhaustive resource. The model predicts the optimal time allocation between leisure and work before the deadline, i.e., work hours over time. However, it did not evaluate the consequences of procrastination or interventions for procrastination.

To bridge the gap, the primary goal of this paper is to propose a normative process model that predicts all aspects of procrastination: the temporal dynamics of work progress, the consequences of procrastination in terms of performance and the feeling of exhaustion, and the interventions to reduce procrastination.

To realize that, we developed a computational model using reinforcement learning theory. We consider the time course of progress as the output of sequential decision-making. We assume that each day, people decide whether to work now (and, if so, how much) or later. While cast in the broader framework of reinforcement learning theory, our process model is also inspired by existing theories of procrastination. One central idea of existing theories of procrastination is temporal discounting (O'Donoghue and Rabin, 1999; Fischer, 2001; Steel and König, 2006). Recently, empirical evidence has supported that temporal discounting is one cognitive mechanism of procrastination (Zhang and Ma, 2023b). Because temporal discounting is naturally rooted in reinforcement learning theory, one key parameter of our model is the discount rate. However, temporal discounting is only part of the story. Studies revealed that several other individual differences are correlated with procrastination, task characteristics, and environmental factors (Steel, 2007). In this paper, we highlight the role of perfectionism, the shape of the cost function, and the given total time. We will see how these correlates of procrastination affect the temporal dynamics of work progress, performance, and the total cost of mental effort.

The structure of the paper is as follows: In Part 1, we develop a model framework for procrastination and its consequences. Then, in the results session, we first demonstrate three procrastination patterns in terms of the time course of work progress. Next, we show how individual differences and task characteristics affect the time course of progress, performance, and the total effort cost. Last, we show the empirical support of our model's potential to represent perfectionism in our designed correlational study. In Part 2, we address two interventions: immediate reward (at various immediacy levels) and interim deadlines, and how these interventions affect the time course of work progress, the number of days of delay, and the consequences of procrastination. Last, in the general discussion, we draw connections between our model predictions and empirical evidence in the literature and suggest several future experimental tests to test our model predictions. We also discussed the limitations of our model in the end.

2.2 Part 1: Patterns and Correlates of Procrastination

2.2.1 Model

Before going into the details of the model, we first provide a broader view. We consider the time course of progress as the output of sequential decision-making. According to our model, on each day, people decide whether to work now (and, if so, how much) or later. If they decide to work now, they pay the cost of investing mental effort immediately but also make progress, and more work leads to more progress. If they decide not to work, they make no progress and also pay no effort cost. Therefore, two vital cognitive components affect our sequential decisions: expected reward and costly mental effort, which interact with each other over days.

The optimal amount of work on each day is derived from the Bellman equation (Bellman, 1957), which assumes that a person's goal on each day is to maximize the discounted value gained by making progress while minimizing the immediate effort cost. In the Bellman equation, the value of the future state is exponentially discounted. However, a wealth of empirical literature suggests that people estimate the value function by discounting future rewards hyperbolically instead of exponentially (Frederick et al., 2002; Frederick et al., 2002; Green, 1982; Green and Myerson, 2004). To approximate the hyperbolical discounting, inspired by (Fedus et al., 2019), we showed that we can compute the state-action value according to the hyperbolic discounting factor by an integral of an infinite set of state-action values in a standard exponential discounting manner. The following results qualitatively hold if the value of the future state is exponentially

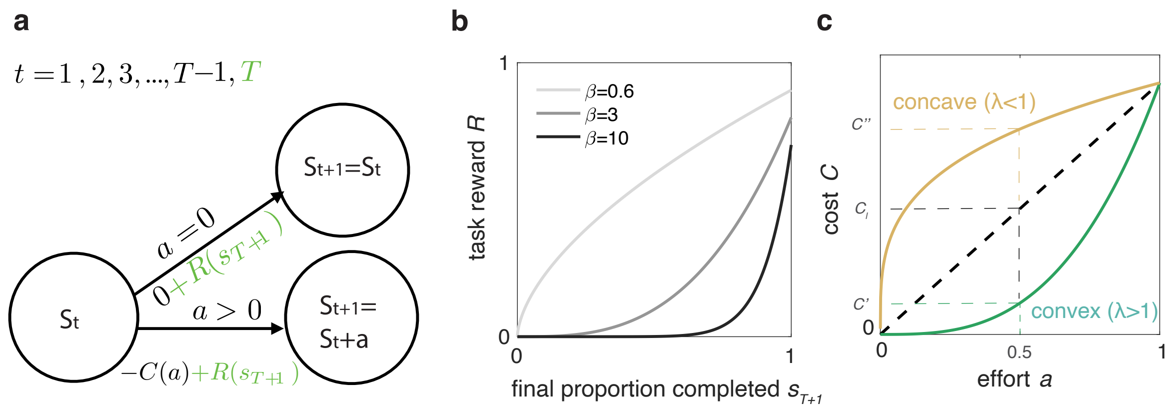

Figure 2.1: Model and functions. (a) Model framework. (b) Task reward as a function of the final proportion completed. (c) Cost as a function of effort.

discounted. For simplicity, we demonstrate the results only from exponential discounting.

We model the reward-performance and cost-effort functions as generally as possible to capture the variety of relationships in the real world and across different individuals. We evaluate the consequence of procrastination by 1) the performance, i.e., the final proportion completed, and 2) the feeling of exhaustion, i.e., the total cost of mental effort.

The details of the model framework are below.

We assumed discrete time; we arbitrarily refer to the unit of time as a day. We denote by T the total number of days provided to complete the task (deadline). A small T suggests a shorter total time. The agent chooses an action every day from 1 to T. We define the task state s as the work progress: the proportion of task completion (between 0 and 1).

2.2.1.1 Reward Schedule: Delayed Reward

The reward schedule that we believe is representative of real-world situations of procrastination is delayed reward, where the agent does not receive any reward until $T + 1$. For example, homework is only graded one day after the deadline. We formalize the reward schedule as follows:

$$r_t = \begin{cases} 0 & \text{when } t < T+1 \\ R(s_{T+1}) & \text{when } t = T+1 \end{cases} \tag{2.1}$$

where $R(s_{T+1})$ is the reward given in the end and $t = 1, 2, .., T, T+1$.

2.2.1.2 Task Reward

We assume a power-law relationship between task reward and final proportion completed (i.e., final performance) (Fig. 3.5b).

$$R(s_{T+1}) = \alpha s_{T+1}^{\beta}. \tag{2.2}$$

α denotes the maximum reward. β is the exponent of the power-law relationship between the maximum reward and final performance.

Despite this simplicity, this relationship has the potential to represent various reward rules in real-world scenarios as well as capture individual levels of perfectionism.

First, β can represent **reward rules**: a relationship between the value of an external reward and final performance, with a higher β indicating a higher stake. One extreme reward rule with the highest stake is having $\beta \to \infty$, which represents the make-or-break case where the external reward is all-or-none. A real-world scenario is work paid only if it is complete. A reward rule with a lower stake could be proportional ($\beta = 1$), in which the payoff is proportional to the proportion completed. A real-world scenario is a wage worker who is paid by the hour. There are also real-world reward rules sitting between these two extremes.

Second, in cases where the relationship between the external reward and final performance is not well defined (e.g., giving a presentation), people value their performance subjectively, in other words, how satisfied and how rewarded they feel about their final performance. In these cases, β has the potential to capture individual **levels of perfectionism**. Perfectionists tend to have all-or-nothing thinking, whereby only total success or total failure exist as outcomes (Flett et al., 1995). They have difficulty appreciating any progress made towards a goal; as long as the task is incomplete, it seems that nothing at all has been accomplished. We formalize this all-or-nothing thinking by a power-law relationship between the subjective reward and the final performance. The larger the exponent β, the higher the level of perfectionism. When the exponent β is extremely large ($\beta \to \infty$), the agent only feels rewarded when the work is nearly perfect, with the satisfaction level being all-or-nothing. We further designed a questionnaire and collected data to verify that this power-law relationship could represent the all-or-nothing aspect of perfectionism (See Results).

In the result session, for the sake of convenience, we interpret this power-law relationship only as the level of perfectionism. Note that the results related to individual levels of perfectionism apply to reward rules as well. In the discussion section, we will discuss both interpretations: the reward rule and the individual level of perfectionism.

2.2.1.3 Effort Cost

Working on a task requires the agent to invest mental effort, and such effort is costly. Empirically, researchers suggest that the cost function of effort could be convex, linear, or concave (Kool and Botvinick, 2018). So to cover all the possible shapes of the cost function (to have cost function to be generally enough), we assumed that the cost function follows a power law (Fig. 3.5c),

$$C(a) = C_{\max} a^{\lambda}, \tag{2.3}$$

where a denotes the amount of effort. $C_{\max}$ denotes the maximum cost–how aversive the task is ($C_{\max} > 0$). When $\lambda > 1$, the cost function is convex: increasing cost is associated with successive increments in effort (Navon and Gopher, 1979; Glimcher and Fehr, 2013). When $\lambda < 1$, the cost function is concave: the cost is relatively sensitive to increases in effort when it is low, but not when it is high (Kool and Botvinick, 2018). From now on, we refer to $C(a) = C_{\max} a^{\lambda}$ with $\lambda > 1$ convex cost function, and we refer to $C(a) = C_{\max} a^{\lambda}$ with $\lambda < 1$ concave cost function.

2.2.1.4 Optimal Policy

In a state s, at time t, the goal of the agent is to maximize the sum of the task reward gained by making progress and the discounted value of the next state. Meanwhile, the agent needs to minimize the effort cost. To reach this goal, the optimal policy is derived from the Bellman equation (Bellman, 1957).

We assume that the agent is in an environment with fun activities as distractions (e.g., surfing social media, eating snacks), so the agent faces two options every day: either working on the task or having fun in alternative activities instead. We denote the utility of alternative activities as J. If the agent chooses to have fun ($a = 0$) (Fig. 3.5a), then they get an immediate small utility of alternative activities J, and they stay in the same state. The value of the state-action pair is

$$Q_t(s, a = 0) = J + r_t + \gamma\, V_{t+1}(s), \tag{2.4}$$

where the value of s is specified as

$$V_t(s) = \max_a Q_t(s, a). \tag{2.5}$$

γ is the discount rate, which determines how much the state value is discounted on the next day ($\gamma \in [0, 1]$). The larger γ, the more the agent takes future value into account.

By contrast, if the agent decides to work ($a > 0$) (Fig. 3.5a), then the agent moves the task forward to a new state s' and pays the effort cost immediately. The value of the state-action pair is

$$Q_t(s, a > 0) = r_t - C(a) + \gamma\, V_{t+1}(s'). \tag{2.6}$$

Here, we assumed a simple deterministic linear relationship between s and s', $s' = s + \eta a$: every additional unit of effort adds a fixed amount of progress scaled by the task difficulty. η represents the level of difficulty of a task. The larger the η, the easier the task for the agent. Here, we let $\eta = 1$ without loss of generality. Therefore, we have $s' = s + a$, or additional progress is equal to the amount of effort $\Delta s = a$. Since s is between 0 and 1, a is between 0 and 1 as well. If an agent chooses to work, the minimum effort an agent can make should be some value above 0. Here, we choose an arbitrary number as $a = 0.01$ for the sake of computation expense in the model simulation.

$Q_t(s, a)$ is derived from solving the Bellman equation using dynamic programming. A policy consists of the amount of effort made each day. The optimal policy maximizes $Q_t(s, a)$ on each day:

$$\pi_t(s) = \underset{a}{\operatorname{argmax}}\, Q_t(s, a). \tag{2.7}$$

2.2.1.5 Total Effort Cost

We evaluate the exhaustion level of working by the total cost of mental effort following the optimal policy: $\sum_a C(a)$.

2.2.2 Results

2.2.2.1 Three Patterns of Procrastination

Though the definitions of procrastination in the literature are different, they all agree on one essential behavioral element: delay in working. So, we define procrastination as a delay in starting work before task completion, i.e., $a = 0$ when $s < 1$.

We first explored the parameter space in our model to look for procrastination behavior in the optimal policies and found three procrastination patterns. Then we further explored the effect of the following on procrastination, which are parameters that might represent individual differences ($\lambda, \beta, \gamma, C_{\max}$), parameters of task characteristics (maximum reward α, total time T and β), and the utility of alternative activities (J) (Table 2.1). So, the results are divided into two major sections: three procrastination patterns and the correlates of procrastination.

Table 2.1: List of parameters and their psychological interpretations.

parameters	psychological interpretations
$\alpha \in [0, 1]$	maximum task reward
$T \in \mathbb{Z}^{+}$	total time
$\beta > 0$	level of perfectionism or reward rule
$\lambda > 0$	shape of cost function
$\gamma \in [0, 1]$	discount rate
$C_{\max} > 0$	task aversion
$J > 0$	utility of alternative activities

For all the results, we provide intuition to ease the understanding of the model. Our results are all verified by Monte Carlo simulations.

Our model produces four patterns of work progress, three of which show procrastination. Four patterns of work progress are: 1) ramping up without delay in the beginning (Fig. 2.2a), 2) a delay in the beginning and then ramping up (Fig. 2.2a), 3) working at the last minute (Fig. 2.2a), and 4) not working at all (Fig. 2.2a). The last three patterns show procrastination (delay in the beginning). Working at the last minute could be viewed as an extreme case of a delay in the beginning and then ramping up. Fig. 2.2b illustrates the overall summary of the parameter space in which we can observe these three procrastination patterns. In the following sections, We will return to Fig. 2.2b to discuss in great detail which parts of the parameter space we get these three procrastination patterns.

A delay in the beginning, and then ramping up. A delay in the beginning and then ramping up is an optimal policy where the agent postpones working on the task in the early days and then gradually increases the amount of effort over the days toward the deadline. We explored the parameter space to determine the necessary conditions for this procrastination pattern. We

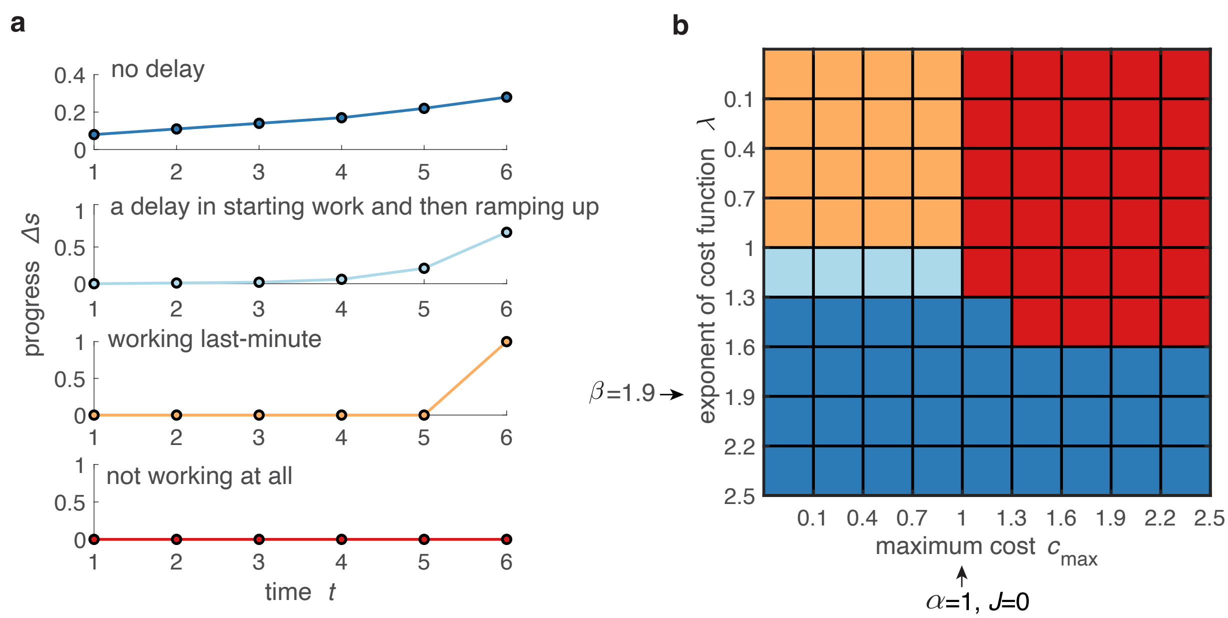

Figure 2.2: Temporal patterns of work progress and its associated parameter space. (a) Patterns of the time course of work progress: (from the top panel to the bottom panel) no delay, a delay in starting work and then ramping up, working last-minute, and not working at all. (b) The parameter space where four temporal patterns of work progress were found in panel a.

first explored the space where an agent has a convex cost function ($\lambda > 1$) and then explored the area where an agent has a concave cost function ($\lambda < 1$).

If an agent has a convex cost function ($\lambda > 1$), a necessary condition to find this procrastination pattern is that the agent discounts the future reward ($\gamma < 1$). The stronger the temporal discounting (having a smaller γ), the more likely the agent delays in the beginning. In contrast, an agent who discounts future rewards to a lesser degree does not procrastinate in the beginning days; they start working from the first day (Fig. 2.2b darker blue cases).

Why does an agent with stronger temporal discounting delay initially and then ramp up later? We provide intuition here. Each day, an agent evaluates the value and cost of making efforts. The value of making efforts is temporarily discounted. For an agent with stronger temporal discounting, the discounted future reward tends to be less than the cost of making efforts, and as a result, it is not worth working on that day. The next question is why procrastination happens at the beginning and not at the end. It is because the discounting factor for a future reward is given by γ^{T-t}, which means that rewards in the distant future are discounted more heavily during earlier days than on later ones. So, in the early days, the agent will tend to choose not to work, and when the deadline is closer, they gradually increase the amount of effort towards the deadline.

If an agent has a concave cost function ($\lambda < 1$), they only work for one day, and that is the last day. Working only on the last day is a special case of a delay in the beginning and then ramping up. We consider this special case as a separate pattern of procrastination, and we discuss it in the next session.

Working at the last minute. When an agent has a concave cost function ($\lambda < 1$), they work only on the last day, regardless of the values of the other parameters (Fig. 2.2b). This is because of the convexity of the cost function.

For a convex cost function ($\lambda > 1$), the total cost of dividing the effort into smaller units across days is lower than the cost of expending the total effort in one day. For example, (Fig. 3.5c) shows the total cost of dividing the effort into two halves (0.5 effort each day). The total cost of making 0.5 effort per day is $2C'$, which is smaller than the total cost of making the full effort in just one day, $C_{\max}$.

However, the opposite case is for an agent with a convex cost function ($\lambda > 1$), where it is more costly to divide the effort into smaller units across days rather than expending the total effort in one day. Fig. 3.5c illustrates that the total cost of making 0.5 effort on each day is $2C''$, which is larger than the total cost of making the full effort in just one day, $C_{\max}$. So, the optimal policy for an agent to maximize the value of making efforts while reducing the total cost is to work for one day only. The effort the agent makes on the last day depends on the combination of parameters.

Not working at all. Not working at all is an extreme pattern of procrastination where an agent does not make any effort for all the days. The necessary condition to find the pattern—not working at all—is when the utility of alternative activities exceeds the net utility of completing the task on one day ($\alpha - C_{\max} < J$); otherwise, if $\alpha - C_{\max} > J$, agents always choose to work on the task.

We studied the additional condition that causes an agent to choose not to work at all. We found that an agent is more likely to choose not to work at all if the parameters change in any of the following ways in terms of task characteristics: 1) when the maximum reward task is smaller, or 2) when the total time is shorter; in terms of individual differences 3) when the exponent of the cost function is smaller (less convex) (Fig. 2.2b), 4) when the discount rate is smaller (stronger temporal discounting) (Fig. 2.2b) or 5) when the level of task aversion if higher, or environmental factor 6) when the utility of alternative activities is greater.

Next, we provide the intuition to understand why an agent chooses not to work at all in the above cases.

Among all the cases, the reason is straightforward: having a smaller task reward, a higher level of task aversion, or a higher utility of alternative activities. We can think of extreme cases where having no task reward ($\alpha = 0$) or having an infinite level of task aversion ($C_{\max} \to \infty$) or infinite utility of alternative activities ($J \to \infty$) is definitely not worth working on.

The fourth case in which we provide intuition is having a smaller exponent of the cost function (less convex). Due to the convexity of the cost function, given the same amount of effort between 0 and 1, the effort cost per day for an agent with a larger λ is lower than the one with a smaller λ. So, having a smaller exponent of the cost function means having a higher effort cost per day. That explains why the agent with a smaller λ is more likely to choose not to work at all.

The fifth case has a shorter total time. Note that agents more likely to choose not to work at all under shorter time only apply to those with convex cost function $\lambda > 1$ because agents with concave cost function $\lambda > 1$ work only on the last day and the work amount on the last day is not affected by the total time (see supplement). We provide intuition here as to why agents are more likely to choose not to work given a shorter total time. It is due to the convexity of the cost function again. Agents with a convex cost function reduce the total effort cost by splitting the work into days. If the total time is shorter, the number of days the agents split the work into is fewer; the total effort cost would increase. Thus, the agents are more likely to choose not to work at all.

The last case has stronger temporal discounting (smaller γ). The intuition behind this is that as stronger temporal discounters discount the future expected reward to a greater degree, they are less likely to work for the total days. Therefore, they have a higher total cost and are more likely to choose not to work at all.

2.2.2.2 The Correlates of Procrastination

Many factors contribute to procrastination. Here, in terms of individual differences, we highlighted the effects of the discount rate, the exponent of the cost function, and the level of perfectionism. In terms of task characteristics, we highlight the given total time. Specifically, we evaluate the effect on 1) the temporal dynamics of work progress, 2) the number of days of delay: the number of days agents choose not to work before the task completion, 3) the performance: the final proportion completed, and 4) the total effort cost. Please refer to the supplement to see the results from those factors that affect procrastination in an obvious way: Agents procrastinate for longer with higher levels of task aversion under a lower maximum task reward; agents procrastinate for longer, given the higher utility of alternative activities.

Stronger temporal discounters procrastinate for longer. We simulated the time course of progress among agents with various discount rates and found that an agent with stronger temporal discounting procrastinates for longer and finishes less work in the end. Even in some cases, an agent with stronger temporal discounting finishes the whole task; they pay a higher total cost (more exhausted). Fig. 2.3 upper panel illustrates that an agent with stronger temporal discounting procrastinates for longer days and has less work done. In one extreme case, the agent does not discount future reward ($\gamma = 1$) at all and does not procrastinate. They work steadily over days and complete the task in the end. Another extreme case is an agent who does not consider the future reward ($\gamma = 0$). They do not start working until the last day and leave

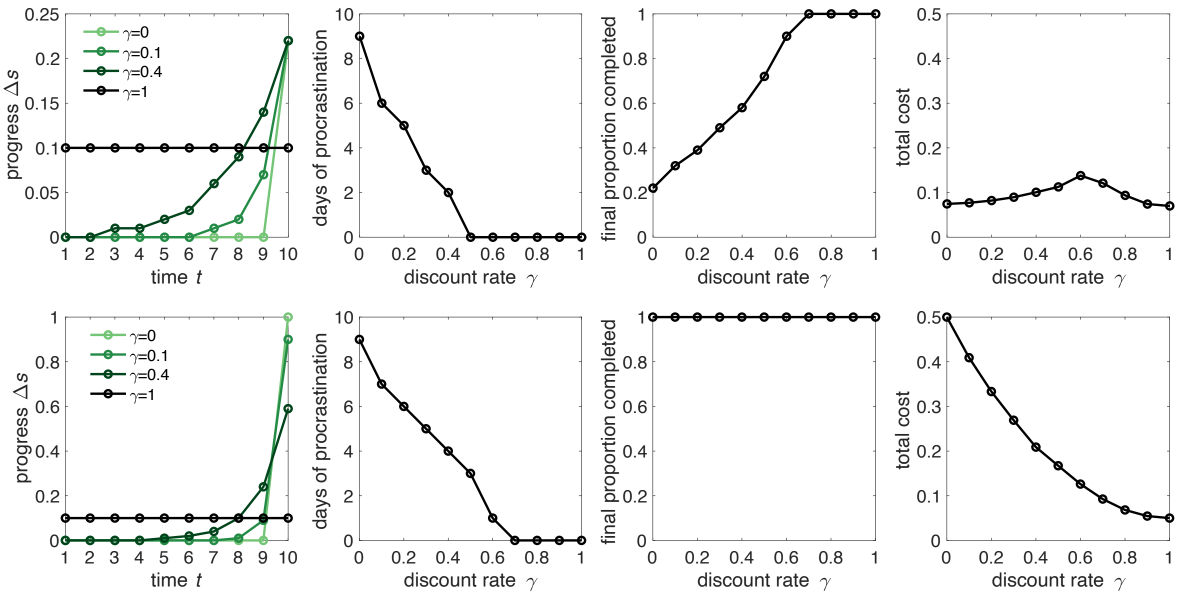

Figure 2.3: Effects of discount rate on the time course of progress, the number of days of delay, the final proportion completed, and the total cost. The upper panel shows that the stronger temporal discounters work less in the end ($C_{\max} = 7$, $\lambda = 3$); the lower panel shows that stronger temporal discounters pay more total effort cost while completing the task ($C_{\max} = 0.5$, $\lambda = 2$).

the task unfinished. Fig. 2.3 lower panel illustrates that another agent with stronger temporal discounting procrastinates more. More importantly, different from Fig. 2.3 upper panel, even though they complete the task, rushing in the end (e.g., $\gamma = 0.1$) is costly. They pay a higher total cost (more exhaustion) to reach the same performance.

Intuitively, stronger temporal discounters are more myopic and discount the future reward to a greater degree. Therefore, they have more days of delay, have less work done, and pay higher total costs.

Agents with convex cost functions procrastinate less than those with concave cost functions, and less convex procrastinate for longer. We simulated the time course of progress among agents with concave or convex cost functions. We found that agents with convex cost functions, compared to those with concave cost functions, have fewer days of delay. Fig. 2.4a illustrates where an agent with a concave cost function procrastinates for six days, whereas an agent with a convex cost function $\lambda = 1.8$ procrastinates for four days. Intuitively, as we know,

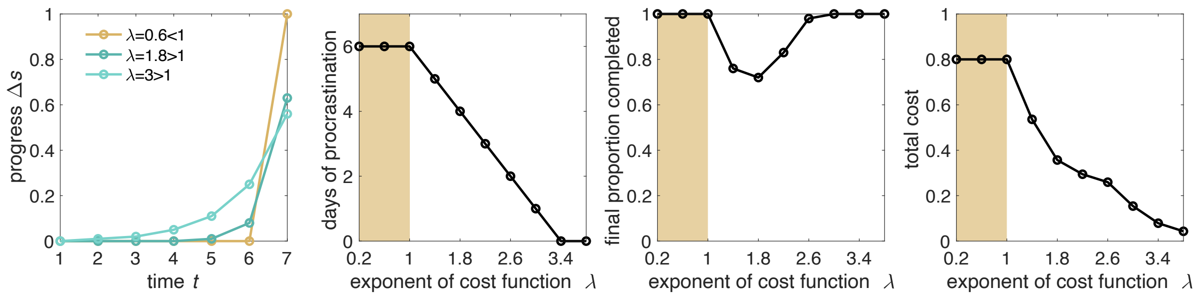

Figure 2.4: Effects of the shape of the cost function on the time course of work progress, the number of days of delay, final proportion completed, and total cost.

agents with concave cost functions, if they work, only work on the last day; namely, they procrastinate for $T-1$ days. In contrast, agents with convex cost functions would divide the effort across days to reduce the total effort cost so that they would procrastinate for fewer days.

In terms of performance, the change is not unidirectional. An agent with a convex cost function can get more or less work done than an agent with a concave cost function.

Among agents with convex cost functions, we simulated the time course of progress among agents with various λ and found that if $J = 0$, an agent with a larger λ procrastinates for fewer days. Fig. 2.4a shows an example of having a larger exponent of the cost function $\lambda = 3$ procrastinates for only one day, three days fewer than the one having $\lambda = 1.8$. Intuitively, due to the convexity of the cost function, agents with larger λ have lower effort costs per day. It would encourage them to work for more days to reduce the total effort cost. If $J > 0$, we did not find the above relationship. In terms of performance, the change is not unidirectional.

A diverse relationship between levels of perfectionism and procrastination. We simulated the time course of progress among agents with various levels of perfectionism. We found a variety of relationships between the level of perfectionism, the number of days of delay, and the final proportion completed. We first illustrate the variety of relationships with some examples and provide the intuition needed to understand the reason.

Fig. 2.5b1 and Fig. 2.5d1 illustrate two opposite trends of the relationship between perfectionism and procrastination. Fig. 2.5b1 shows where an agent with a higher level of perfectionism

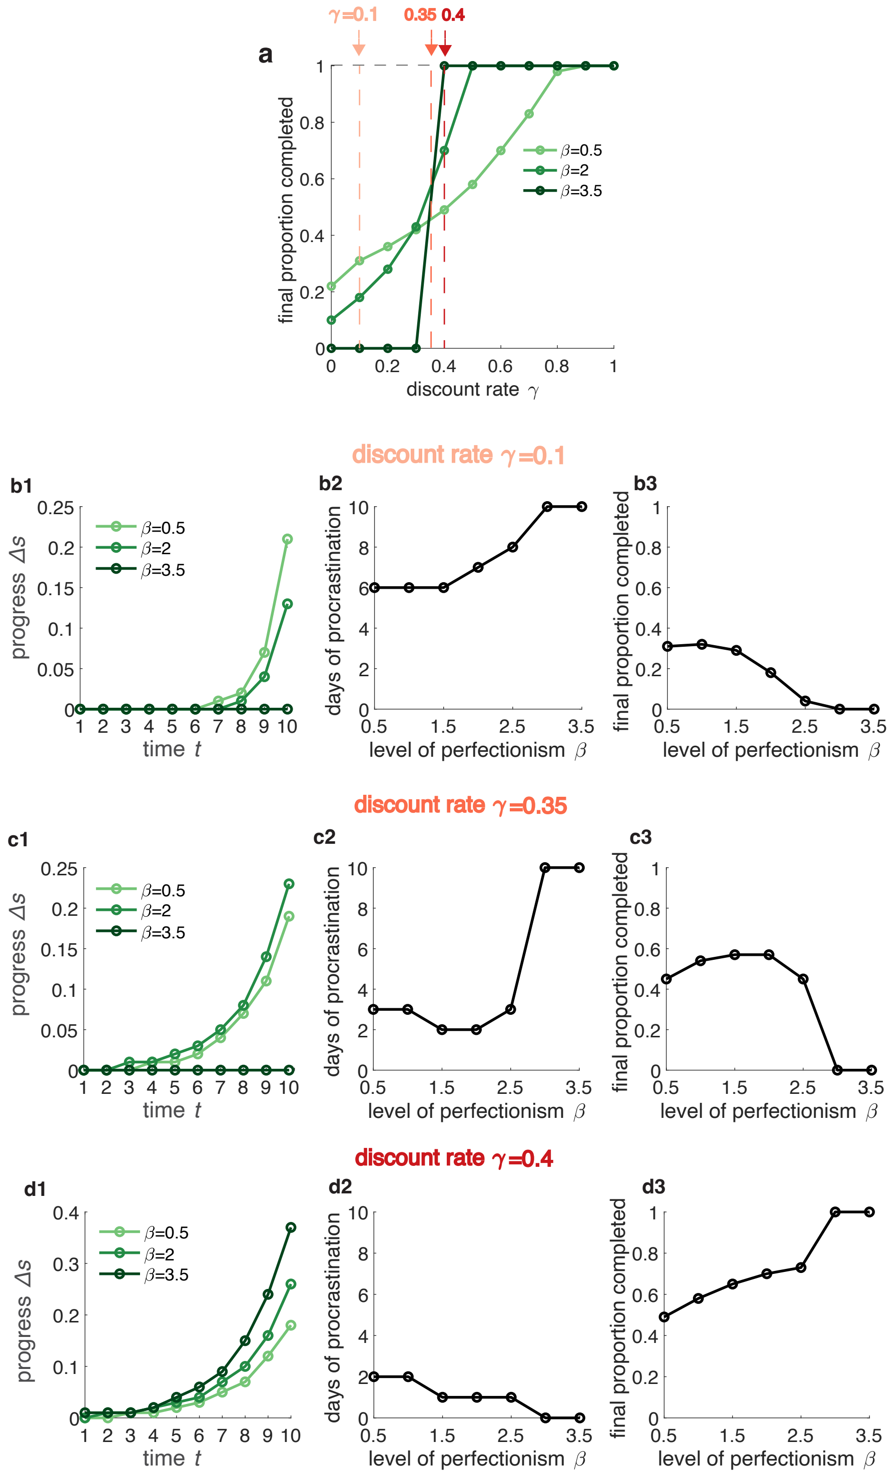

Figure 2.5: Diverse relationship between levels of perfectionism and procrastination. (a) The effect of discount rate on final proportion completed at three levels of perfectionism. (b1,2,3) The effect of the level of perfectionism on the time course of progress, the number of days of delay, and the final proportion completed when the discount rate is at 0.1. (c1,2,3) when the discount rate is at 0.35. (d1,2,3) when the discount rate is at 0.4.

procrastinates for longer and finishes less work. An extreme case in this example is that an agent with $\beta = 3.5$ does not work at all. In contrast, Fig. 2.5d1 shows the opposite trend: an agent with a higher level of perfectionism procrastinates for fewer days and finishes more work in the end. However, the agent with the same level of perfectionism $\beta = 3.5$ in this example is at the opposite extreme. They work all day and complete the task in the end.

There are also cases where a mixed version of these two opposite trends exists. For example, as Fig. 2.5c1 illustrates, in a specific range of levels of perfectionism ($0.5 < \beta < 2$), an agent with a higher level of perfectionism procrastinates for fewer days and finishes more work in the end, whereas in another range of levels of perfectionism ($2 < \beta < 3.5$), an agent with a higher level of perfectionism procrastinates for longer and finishes less work in the end.

The observed mixed trend arises due to the concurrent presence of two opposite trends. So, understanding why the two opposite trends exist is the key to understanding the diverse relationship between perfectionism and procrastination.

The two opposite trends are found at two different discount rates while keeping all the other parameters the same. The trend of having a higher level of perfectionism leading to more procrastination is found in a stronger temporal discounter ($\gamma = 0.1$), and the opposite trend is found in a weaker temporal discounter ($\gamma = 0.4$).

Why does the trend flip from stronger to weaker temporal discounters? To find out why, we simulated the progress of agents with different discount rates. We plotted their performance against various discount rates at each example perfectionism level (Fig. 2.5a). The key finding is that as the discount rate lowers, there is a jump from choosing to complete the task to choosing not to work at all in an agent with high perfectionism (as long as $\beta > \lambda$, here $\lambda = 3$); the final proportion completed jumps from 1 to 0 (Fig. 2.5a darker green curve). As the level of perfectionism increases, the performance function of the discount rate approaches the high perfectionism's function with the jump (Fig. 2.5a lighter green curves). This jump between not working at all and completing the task in an agent with high perfectionism causes the opposite trend. Why is there a jump in an agent with a high level of perfectionism (as long as $\beta > \lambda$)? It is because an agent with a high level of perfectionism has a more curved reward function. In other words, if an agent completes the task only partially, the reward is disproportionate to the performance. For example, very little reward is offered even when the task is almost finished. Therefore, it is not worth it to leave the task partially done. An agent with a high level of perfectionism should either choose to complete the task or choose not to work at all.

We then give an insight into why the trend is flipped from stronger to weaker temporal discounters, i.e., stronger temporal discounters procrastinate more. In comparison, weaker temporal discounters procrastinate less than agents with high perfectionism. Stronger temporal discounters are more myopic and barely look ahead. As a result, at the beginning of the task, they value only the expected reward they gain and overlook the reward they gain in the end. Therefore, if they have a high level of perfectionism, they will not feel like working because the reward for the initial effort is disproportionally low. However, weaker temporal discounters tend to take the far future into account. As a result, at the beginning of the task, they value the reward they gain far into the future. If they have a high level of perfectionism, they work hard to complete the task because they foresee that the reward is decent only if they complete the task.

Last, for an agent with a high level of perfectionism, varying discount rates is not the only way to have this flip of trend. Other ways include varying the maximum task reward, the total time, the exponent of the cost function, the level of task aversion, or the utility of alternative activities.

Agents switch from not working at all to working all days, given longer total time. Since agents with a concave cost function only work for one day, and their procrastination is not affected by the total time, we studied the effect of the total time on procrastination only in agents with convex cost functions. We found that given a longer total time, agents switch from not working at all to working all days, i.e., they procrastinate for fewer days and finish more work. Fig. 2.6 illustrates an example where agents choose not to work at all if $T = 3$, but given

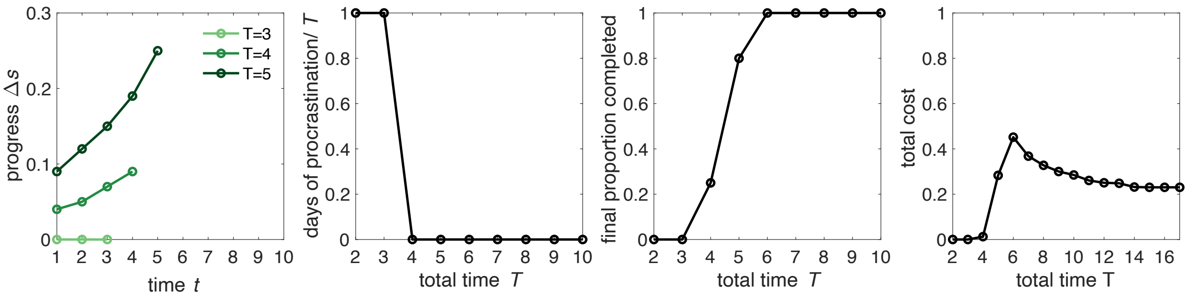

Figure 2.6: Effects of the total given time on the time course of work progress, the number of days of delay, final proportion completed, and total cost.

one more day, the agents work all days. It is due to the convexity of the cost function again. An agent with a convexity cost function splits the work into days to reduce the total effort cost. If the allotted total time is too short, the total effort cost would be more than the task reward; as a result, the agent chooses not to work at all. However, the agent will choose to work once the total time is long enough for the task reward to outweigh the total effort cost. As the task reward barely outweighs the total effort cost, the agent will work every day to reduce the total effort cost to the maximum.

2.2.2.3 Empirical Support of the Potential to Represent Perfectionism.

In the model session, to characterize perfectionism, we proposed the power-law relationship between the utility of the task reward and the proportion of task completion. Specifically, given the same proportion of task completion, a person with a higher level of perfectionism will always have a lower level of satisfaction until the task is complete. In other words, a person with a higher level of perfectionism has a larger exponent β of the power law function.

However, our proposal remains theoretical, and supportive empirical evidence is needed. We designed an experiment to verify that this power-law relationship represents perfectionism. The goal was to examine whether the perfectionism individual score (measured by the commonly used perfectionism scale (Flett et al., 1995) correlates with the exponent β of our proposed power

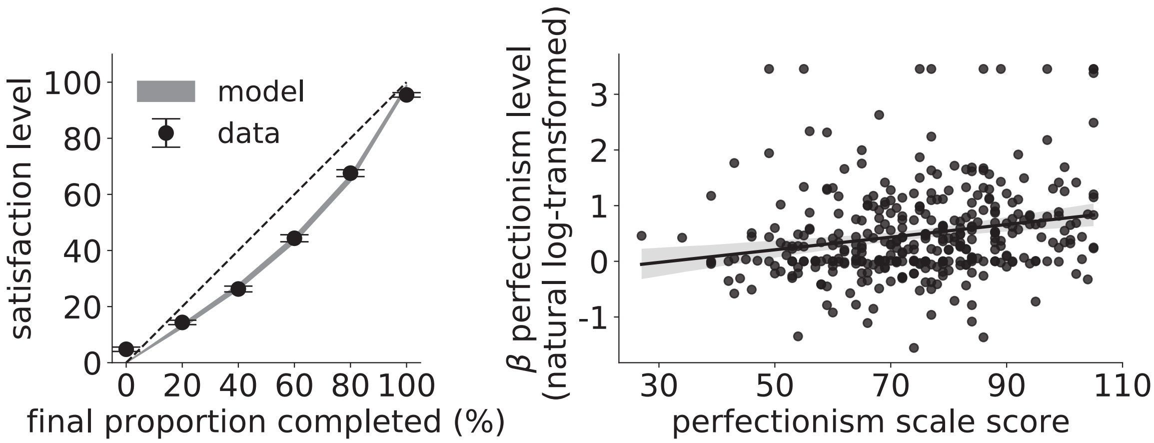

Figure 2.7: Empirical support of the model's potential to represent the level of perfectionism. Left panel: People's satisfaction level as a function of their final performance is well fit by a power law function of exponent β. Right panel: The fitted exponent of the power law function, which indicates our model's level of perfectionism correlated with people's self-reported perfectionism level indicated by their perfectionism scale score.

law function. To compute β at the individual level, we developed a questionnaire. In the questionnaire, subjects were asked to imagine that they were assigned to a task, and based on different proportions of the task they completed in the end, they were asked to indicate their level of satisfaction on a scale from 0 to 100. An example item was "If you complete 60% of the hypothetical task, please indicate your satisfaction level on a scale from 0 to 100." Then, to compute the exponent β, we fit a power law function to each subject's reported level of satisfaction as a function of the proportion of completion in the hypothetical task (Fig. 2.7 left panel).

We found a significant positive correlation between the individual scores in the perfectionism scale and the exponent β of the power-law function to which we fit the data ($r = 0.22, p < 0.0001$) (Fig. 2.7 right panel). We conclude that our proposed power-law relationship between the utility of the task reward and the proportion of task completion has the potential to represent perfectionism.

2.3 Part 2: Interventions

Empirical studies showed that two interventions effectively reduced procrastination or increased final performance. One involves rewarding immediately instead of in a delayed manner (Zhang and Ma, 2023a; Lieder et al., 2019; Milkman et al., 2014; Woolley and Fishbach, 2016; Woolley and Fishbach, 2017). Another involves interim deadlines (Ariely and Wertenbroch, 2002).

Zhang et al. found that by rewarding upon task completion versus rewarding after the deadline, people had fewer days of delay (start working earlier) and even completed the task earlier before the deadline (Zhang and Ma, 2023a). We modeled the intervention, reward upon task completion, in the framework of our normative theory to see if the model prediction is aligned with the empirical results (Fig. 2.8). Besides the intervention reward upon task completion, we also tested the other two immediate reward inventions with higher levels of immediacy than the reward upon task completion. They involve immediate reward while agents are making progress before the task completion. One is an immediate reward at each milestone, e.g., rewarding once completing every third of the task (Fig. 2.8). Another one, and the highest level of immediacy, is rewarding each unit of progress (Fig. 2.8). The model predictions provide interesting hypotheses that are worth testing empirically in the future.

Ariely et al. designed an error correction task to test the effect of interim deadlines on performance. Participants were asked to correct grammar errors in three texts, each with 100 errors. Subjects had up to 21 days to work on the task. The experimental condition is an evenly-spaced deadline with a penalty of missing each interim deadline (Fig. 2.9). In specific, participants were asked to complete one text every seven days. If the text submission is delayed, there is a small monetary penalty (i.e., 3%) for each day of delay. While in control condition, there is a single deadline on the last day. They found that more errors were corrected in the interim deadline condition than in the single deadline condition. We modeled this intervention—interim deadlines—in our theoretical framework and predicted the effect of interim deadlines on the time course of work

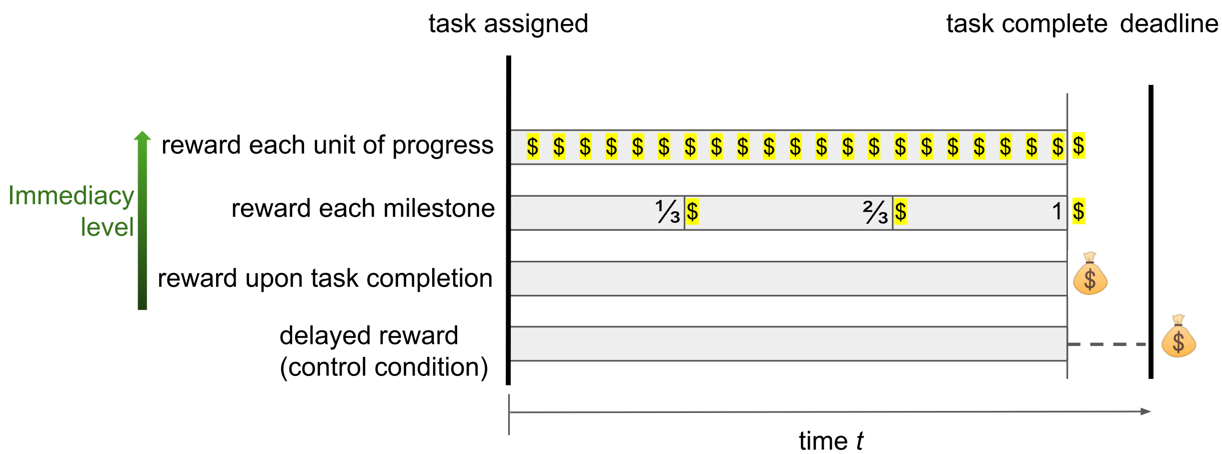

Figure 2.8: Immediate reward interventions with different immediacy levels.

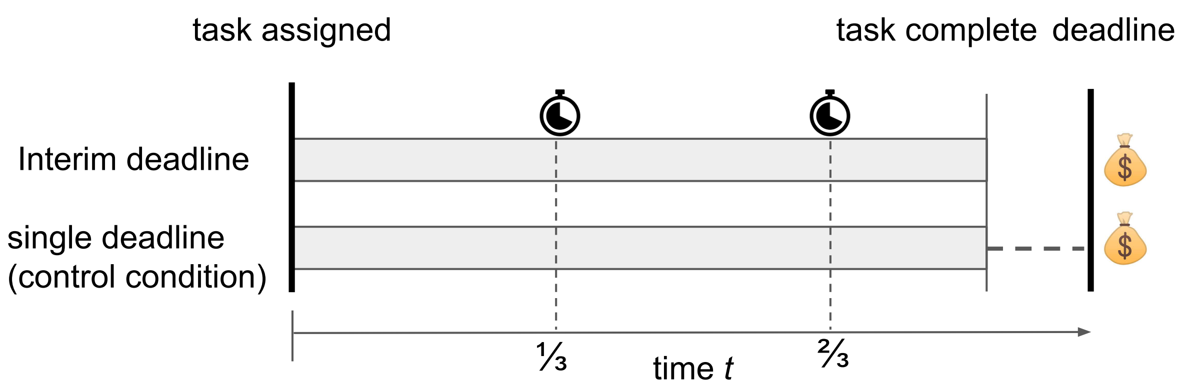

Figure 2.9: interim deadline intervention.

progress and the number of days of delay, which can be tested in future experiments.

2.3.1 Model

2.3.1.1 Immediate Rewards

The control condition –Delayed Reward– is formalized in Part 1. We have repeated the formula here to facilitate the comparison with other interventions.

Control Condition.

$$r_t = \begin{cases} 0 & \text{when } t < T+1 \\ R(s_{T+1}) & \text{when } t = T+1 \end{cases} \tag{2.8}$$

The other interventions are formalized as below.

Reward Upon Task Completion. There is no reward until the task is completed.

$$r_t = \begin{cases} 0 & \text{when } S_t < 1 \\ R(1) & \text{when } S_t = 1 \end{cases} \tag{2.9}$$

Reward Each Milestone. Let's say there are n milestones. $n \in 2, 3, 4, ...$ The reward for each milestone is based on the proportion of completion at the moment of completing the milestone.

$$r_t = \begin{cases} 0 & \text{when } S_t < \frac{1}{n} \\ R(\frac{1}{n}) & \text{when } S_t = \frac{1}{n} \\ 0 & \text{when } \frac{1}{n} < S_t < \frac{2}{n} \\ R(\frac{2}{n}) - R(\frac{1}{n}) & \text{when } S_t = \frac{2}{n} \\ ... & \\ 0 & \text{when } \frac{n-1}{n} < S_t < 1 \\ R(1) - R(\frac{n-1}{n}) & \text{when } S_t = 1 \end{cases} \tag{2.10}$$

Reward Each Unit of Progress.

$$r_t = R(s_{t+1}) - R(s_t). \tag{2.11}$$

2.3.1.2 Interim Deadlines

In our model of Experiment 2 from Ariely et al., we retain the delayed reward condition. In addition, there is a penalty for failing to meet interim deadlines. The policy for penalization is delineated as follows: The total time is set at $T = 21$, with the maximum reward being \$30. Deadlines are evenly spaced, with an interim deadline occurring every seven days. Participants are required to complete an additional $\frac{1}{3}$ of the total work every seven days. Therefore, by the 7th day, $\frac{1}{3}$ of the work should be completed, by the 14th day, $\frac{2}{3}$ of the work should be completed, and by the 21st day, the entire work should be completed. If an intermediate deadline is missed, a daily penalty ensues, where each day's delay results in a \$1 deduction (approximately 3% of the maximum reward); for simplicity, designate this loss as $l_t = 0.03$. It's important to note that penalties are cumulative for each missed interim deadline. For example, if by the 14th day, the progress is less than $\frac{1}{3}$, the penalties for missing both the first (completing less than $\frac{1}{3}$ of the work by the 7th day) and the second interim deadlines (completing less than $\frac{2}{3}$ of the work by the 14th day) are applicable, culminating in a loss of $l_t = 0.06$. There are no additional rewards for submitting work before these interim deadlines.

$$Q_t(s, a = 0) = J + r_t - l_t + \gamma\, V_{t+1}(s), \tag{2.12}$$

$$Q_t(s, a > 0) = r_t - C(a) - l_t + \gamma\, V_{t+1}(s'). \tag{2.13}$$

where l_t is the associated loss at time t due to the policy of penalty. l_t is in both the state-action value function for having fun or for working.

$$
l_t = \begin{cases} 0 & \text{when } t < 7 \\ \\ 0.03 & \text{if } S_t < \frac{1}{3}\text{, when } 7 \le t < 14 \\ 0 & \text{if } S_t \ge \frac{1}{3}\text{, when } 7 \le t < 14 \\ \\ 0.06 & \text{if } S_t < \frac{1}{3}\text{, when } 14 \le t < 21 \\ 0.03 & \text{if } \frac{1}{3} \le S_t < \frac{2}{3}\text{, when } 14 \le t < 21 \\ 0 & \text{if } S_t \ge \frac{2}{3}\text{, when } 14 \le t < 21 \\ \\ 0.09 & \text{if } S_t < \frac{1}{3}\text{, when } t = 21 \\ 0.06 & \text{if } \frac{1}{3} \le S_t < \frac{2}{3}\text{, when } t = 21 \\ 0.03 & \text{if } \frac{2}{3} \le S_t < 1\text{, when } t = 21 \\ 0 & \text{if } S_t = 1\text{, when } t = 21 \end{cases} \tag{2.14}
$$

2.3.2 Results

2.3.2.1 Immediate reward

We simulated the time course of progress across the entire parameter space under three different immediate reward interventions and compared these to the control condition, which was a delayed reward scenario.

Regarding the time course of progress, we present an example in Fig. 2.10d, demonstrating the differences in the time courses of progress under the three immediate reward interventions compared to the delayed reward. With the delayed reward, the agent procrastinates for the first

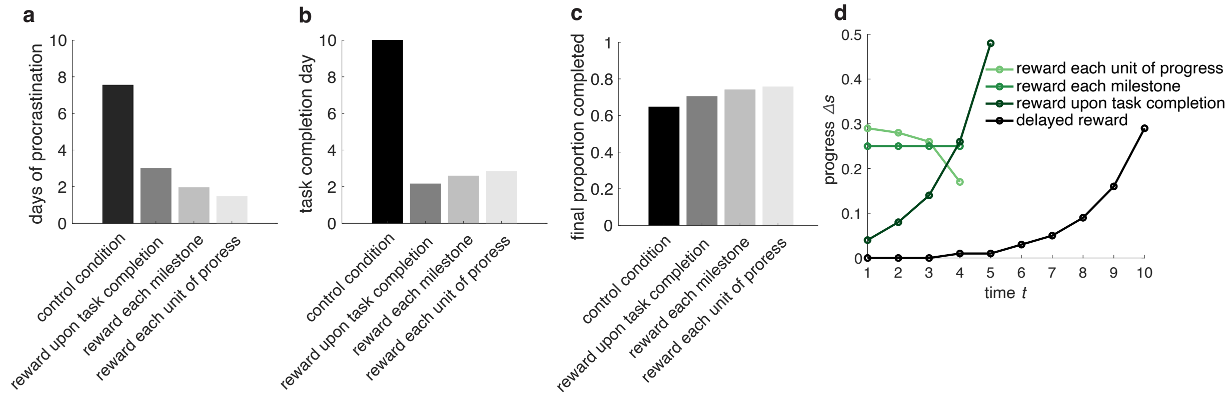

Figure 2.10: The effects of all immediate reward interventions (a) on the number of days of delay, (b) on task completion day, (c) on final proportion completed, and (d) on the time course of progress.

three days but begins work immediately under the immediate reward conditions. Furthermore, the agent completes the task several days ahead of the deadline when immediate rewards are offered, whereas it finishes the task on the final day when only a delayed reward is available. Additionally, the agent completes the task one day sooner when rewards are given for each unit of progress or for reaching each milestone (indicating a higher level of immediacy) than when the reward is given only upon completion of the task (a lower level of immediacy).

In terms of procrastination, we found that offering immediate rewards reduces the number of days of delay (Fig. 2.10a). The higher immediacy of interventions, the fewer days of delay observed.

Moreover, offering immediate rewards helps agents complete the task before the deadline (Fig. 2.10b). This is in contrast to the delayed reward condition, where agents always complete tasks on the last day.

With respect to performance, agents get more work done in rewarding upon completion or in rewarding each unit of progress than in control condition (Fig. 2.10c). On average, a higher level of immediacy in the intervention corresponds to an increase in the amount of work completed by the agents.

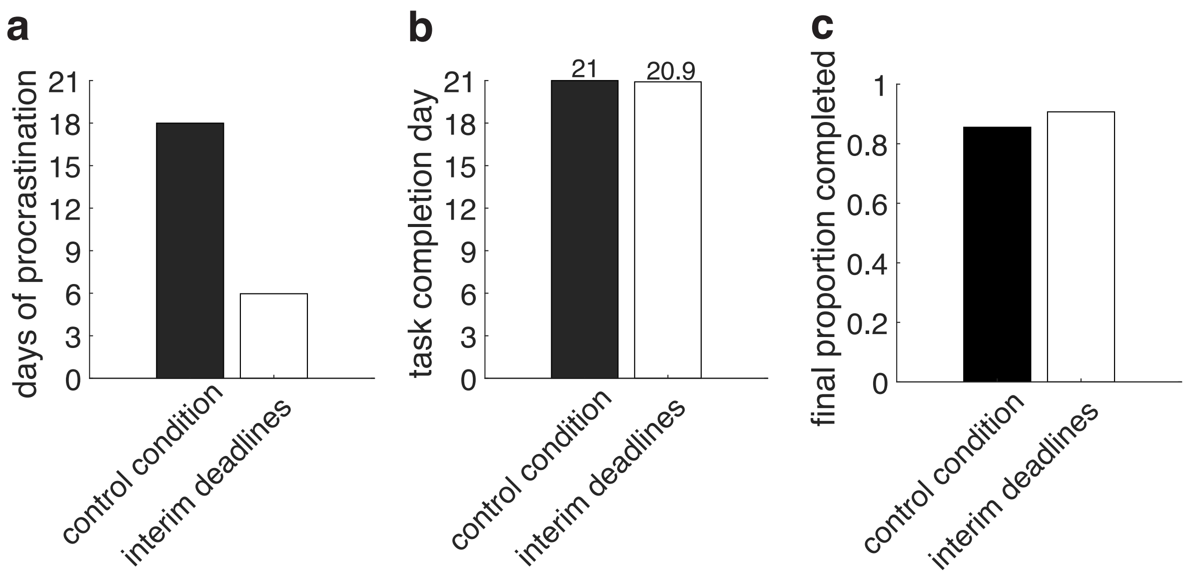

Figure 2.11: The effects of interim deadline interventions (a) on the number of days of delay, (b) on task completion day, (c) on final proportion completed.

2.3.2.2 Interim deadlines

We simulated the time courses of progress under interim deadlines condition and compared the results with those from the control condition, namely the delayed reward or, for the sake of comparison, what we could term the single deadline condition. We found that interim deadlines reduce the number of days of delay, help agents complete tasks earlier than the deadline, and help agents get more work done (Fig. 2.11).

There emerge two distinct progress patterns under the interim deadlines (Fig. 2.12). Fig. 2.12 left panel illustrates the first pattern, which consists of three sequential ramping activities, each aligning with an interim deadline. Fig. 2.12 right panel illustrates the second pattern, where there is a single ramping activity leading up to the final deadline.

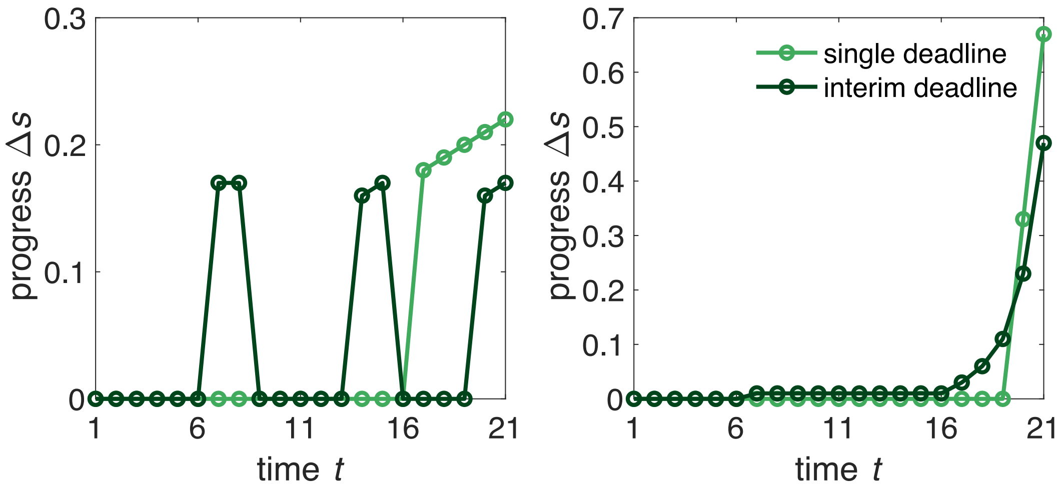

Figure 2.12: The effects of interim deadline interventions on the time course of progress. The left panel and right panel show two distinct patterns of time course of progress.

2.4 Discussion

We propose a normative account of the temporal dynamics of procrastination. This normative theory predicts three patterns of procrastination: a delay in the beginning and then ramping up, working at the last minute, and not working at all. This theory also predicts several correlates of procrastination in terms of personality traits and task characteristics. Last, the theory reproduced the effect of interventions on reducing procrastination or improving performance, as suggested by the empirical literature.

Here, we discussed the empirical literature related to our model predictions and how our model can be improved in several aspects.

Our model predicts three procrastination patterns. A delay in the beginning and then ramping up is an often-seen pattern in the literature (Konradt et al., 2021; Vangsness and Young, 2020; Vangsness and Young, 2020; Putnik et al., 2013; Dewitte and Schouwenburg, 2002). Konradt et al. asked students to self-report their effort allocation over time during exam preparation. They found that most students (41.6%) delay working initially and then ramp up. Similarly, Dewitte and Schouwenburg asked the students to report weekly the hours they spent studying. They found that all students tend to postpone most of their study activities to the last week before an exam, and that a hyperbolic curve could nicely describe this trend. Vangsness and Young observed the real-world behavior of undergraduate students completing their mandatory research credits by participating in research opportunities over a semester. They found that 18.9% of all the students had the temporal pattern of work progress as a delay in the beginning and then ramping up. Punik et al. observed how students worked on an online assignment and observed indirect evidence of ramping up towards the deadline, i.e., the number of views and posts increasing towards the deadline. Working at the last minute (e.g., did not work on the first six days and rushed to complete the task on the last day, (4.4%) and not working at all (18.3%) were observed in Zhang and Ma 2023.)

However, there are patterns of work progress observed in the real world but not predicted by our model. For example, Vangsness and Young categorized the temporal pattern of work progress into three categories. Besides procrastination (18.9%) and steady working (60.3%), they also observed early completion (they call them pre-crastinator): they completed the task ahead of the deadline, 20.8%). Similarly, Konradt et al. found that in students' self-reported time course of work progress, besides dominant procrastination pattern (41.6%) and steady working (33.0%), few students have U-shape (11.7%) or inverted U-shape (13.5%). Zhang and Ma also observed U-shape in a few participants.

Our model predicts that stronger temporal discounters procrastinate for longer, and this is consistent with empirical evidence (Zhang and Ma, 2023a). All the other model predictions remain to be tested empirically, including the effect of the shape of the cost function, the diverse effect of perfectionism depending on personality traits and task characteristics, and the effect of given total time.

Our model can be improved in several other aspects besides predicting other patterns of time course of progress (e.g., early completion). First, we assumed a deterministic state transition for simplicity. This state transition could instead be stochastic.

Second, the mapping from effort to progress sometimes depends on the task state; for instance, when writing a paper, even though we make the same amount of effort, we might make very little progress in the beginning but much more progress later. But this mapping is state-independent in our model.

Third, we often do not face one task only in real-world scenarios. We usually have multiple tasks at hand, and from moment to moment, we need to decide which task we would like to work on. However, so far, our model only targets a single task. To model a multi-task situation, we need to consider switching costs between tasks (Wylie and Allport, 2000; Altmann and Gray, 2002; Gopher et al., 2000). However, our simple model has the potential to be extended to model a multi-task situation.

Last, sequential decision-making is notoriously challenging from a computational perspective. Using the Bellman equation to find the optimal solution depends on backward induction, and several studies show that even for shallow planning depth, humans do not use backward induction (e.g., Hotaling and Busemeyer, 2012; Huys et al., 2015; Zhang and Yu, 2013). So, whether people are using the Bellman equation to perform optimally in our model framework is a question for future experiments. Also, future models need to formulate plausible, bounded rational strategies as alternatives. Nevertheless, this simple model is a basis for quantitatively examining procrastination.

2.5 Supplement

Agents with higher levels of task aversion procrastinate for longer

We simulated the time course of progress among agents with various levels of task aversion. We found that (when J=0), an agent with a higher task aversion procrastinates for longer and finishes less work. When J>0, we did not find the above relationship bidirectional.

Agents procrastinate for longer under lower maximum task reward

We simulated the time course of progress under various maximum task rewards. We found that agents with a smaller maximum reward procrastinate for longer and finish less work in the end.

Agents procrastinate for longer, given the higher utility of alternative activities

We simulated the time course of progress under various utility of alternative activities. We found that agents procrastinate for longer and finish less work, given the higher utility of alternative activities.

3 | Temporal discounting predicts procrastination in a real-world task

People procrastinate, but why?

What is the key factor contributing to procrastination?

Can we find real-world evidence for better understanding?

3.1 Introduction

People procrastinate. For instance, people delay filing their taxes until the last minute (Martinez et al., 2017). Researchers postpone until the last minute registering for academic conferences (Alfi et al., 2007), and submitting abstracts and papers (Flandrin, 2010). College students commonly put off starting self-paced quizzes and find themselves rushing to complete them by the end of the semester (Solomon and Rothblum, 1984; Rothblum et al., 1986; Steel et al., 2001).

The question arises: why do people procrastinate? One long-standing hypothesis suggests that temporal discounting is the mechanism underlying procrastination (O'Donoghue and Rabin, 1999; O'Donoghue and Rabin, 2001; Fischer, 1999; Fischer, 2001; Steel and König, 2006; Steel, 2007). When faced with a task in its initial stages, where the eventual reward is distant, people temporarily discount the value of that future reward. As a consequence, the temporarily discounted future reward fails to provide sufficient motivation for people to start working until the deadline looms near. To illustrate, consider the process of writing a thesis. Initially, the perceived utility of working on the thesis is diminished due to temporal discounting, making it less appealing than alternative activities like socializing. As a result, a student may delay writing the thesis until the utility of working on the thesis outweighs the utility of socializing, which occurs as the deadline approaches.

This hypothesis suggests a positive correlation between the degree to which individuals discount future rewards and the extent of their procrastination. In other words, individuals who discount future rewards to a greater degree procrastinate more.

However, no empirical evidence for this hypothesis exists (Raphaël and Mathias, 2022; Reuben et al., 2015). Raphaël at al. did not find a correlation between temporal discounting and procrastination in a survey completion task. Reuben et al. found a correlation in two real-world tasks, each offering enhanced rewards as incentives for early completion. However, such incentives could be a confound because the actual correlation might be between temporal discounting and achievement motivation (Xin et al., 2020; Lee et al., 2012). Indeed, when early-completion incentives were removed in a third task, the authors found no correlation between temporal discounting and procrastination.

In this study, we tested for a relationship between temporal discounting and procrastination. we looked for a real-world task that satisfied three criteria. First, to rule out the potential confound mentioned above, no incentives should be given for early completion. Second, our task should measure the entire time course of work progress rather than a single endpoint such as task completion time (Raphaël and Mathias, 2022; Reuben et al., 2015). This is because individuals who complete a task at the same time can exhibit very different patterns of work progress, as observed in previous studies (Konradt et al., 2021; Vangsness and Young, 2020): some people maintained steady progress from beginning to end, whereas others made very little progress at the start and rushed to complete their work on the very last day. Thus, to get a fine-grained metric of procrastination, our task should measure the entire time course of work progress. This, in turn, requires that the task a) has an unambiguous definition of a unit of work, b) the completion time of each unit of work is measured, and c) involves multiple units of work to establish a time course of work progress. Real-world tasks such as writing or taking an academic course often lack clearly defined units of work. Finally, we need a task in which an individual's work progress is not affected by others.

The real-world task that satisfied all the above criteria and used in this study was the research participation requirement in the *Introduction to Psychology* course at New York University. To receive course credit, all enrolled students were required to participate in research studies for a total of 7 hours before the end of the semester; the semester lasted a total of 109 days. This task was self-paced, granting students the autonomy to decide when to participate. All three criteria were met in this task. First, since course credit was independent of the time at which the research requirement was completed, no incentives were given for early completion. Second, a unit of work was clearly defined as 0.5 hours because research participation opportunities involved a time commitment of 0.5, 1, 1.5, or 2 hours. The vast majority (91.2%) of participation opportunities took 0.5 hours or 1 hour. The date of each research participation is documented in the NYU Sona System and is accessible to the system administrator. Students needed to participate multiple times to fulfill the 7-hour requirement. In practice, all students participated at least 6 times, with a median of 10 times. Last, research participation opportunities were plentiful: an average of 15 hours per student. Thus, there was no need for students to compete for these opportunities, and each student's work progress could reasonably be assumed to be independent of that of others.

The secondary objective of our study is to examine the relationship between risk attitude and the behavioral level of procrastination. By postponing the research participation until the end of the semester, students face an increased risk of not being able to complete the 7-hour research participation, particularly when considering other competing obligations near the end of the semester, such as final exams. Consequently, procrastination in the research participation task can be viewed as a risk-seeking behavior. A previous study proposed that there was no correlation between risk attitude in active risk-taking and self-reported procrastination (Keinan and Bereby-Meyer, 2012). We examined the relationship between people's risk attitude and procrastination behavior in this real-world research participation task.

3.2 Methods

3.2.1 Procedure

To estimate the students' temporal discount rate and risk attitude at an individual level, we sent email invitations with a link to our online study to all the students enrolled in the Introduction to Psychology course two weeks after the semester ended. In the email, we provided a broad description of the study's aim, investigating the factors that influence student research participation. However, we did not disclose the specific focus of the study on procrastination.

Participants were compensated with $5 for their participation and had the opportunity to earn a bonus of up to $66 based on their choices during the tasks. This study was approved by New York University's Institutional Review Board (IRB-FY2020-4262), and it was pre-registered on Open Science Framework (https://osf.io/4sxrw). Experimental stimuli, anonymized data, and scripts for analysis are available through the Open Science Framework (https://osf.io/z548y/).

3.2.2 Measures

Behavioral indices of procrastination. We pre-registered four indices to quantify procrastination: Mean Unit Completion Day (MUCD), the Day of Halfway Point (Rothblum et al., 1986), Task Completion Day (Raphaël and Mathias, 2022; Reuben et al., 2015; Howell et al., 2006) and Hours of research participation in the Last Third of the semester (Hours in Last Third Semester) (Solomon and Rothblum, 1984).

Mean unit completion day (MUCD) is the average of the completion days of the fourteen half-hour work units (formula below). We consider a general situation in which a person completes N half-hour units of work over T days. We denote the number of units completed per day by $x_1, x_2, \ldots, x_T$. Then $N = \sum_{t=1}^{T} x_t$, $MUCD$ is the average of the completion days of the fourteen work units:

$$MUCD = \frac{1}{N} \sum_{t=1}^{T} t x_t$$

We discuss below two more metrics in the Supplement, the area under the cumulative progress curve (Steel, 2018) and task starting day. The former reduces to MUCD, and the latter is not suitable for our task.

Discount rate. To estimate the degree of reward discounting, we used a delay discounting task. We used a widely used choice set that was designed to capture a broad range of discount rates (Senecal et al., 2012). Participants were asked to indicate their monetary preferences between smaller but sooner rewards and larger but delayed rewards. For example, $30 today or $35

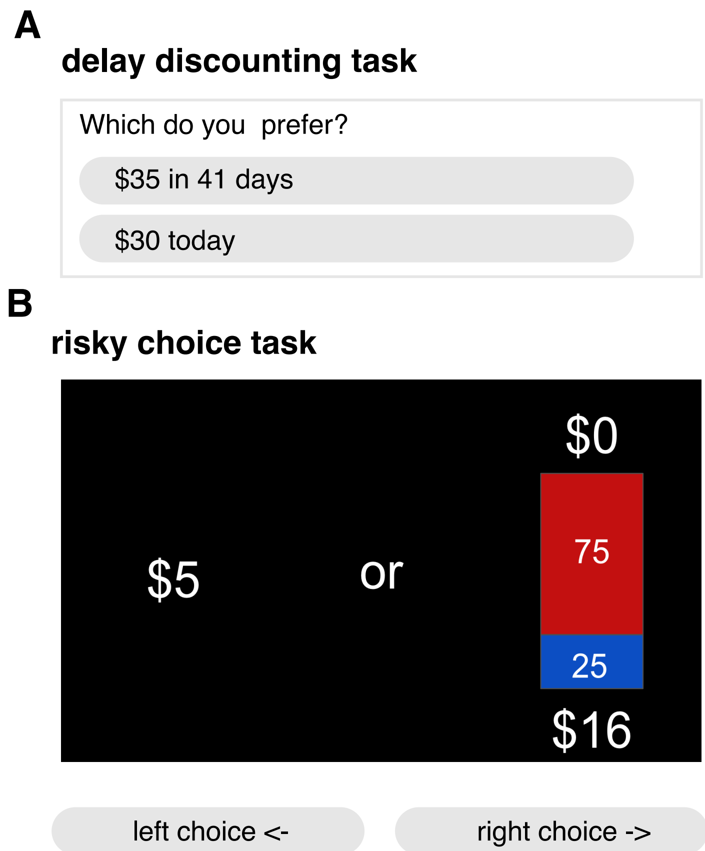

Figure 3.1: Task illustration. (a) One trial in delay discounting task. (b) One trial in risky choice task.

in 41 days (Fig. 4.1A). Moreover, this task was designed to be incentive-compatible, in contrast to the hypothetical nature of rewards in the previous study (Raphaël and Mathias, 2022).

The delay discounting task consisted of 51 self-paced trials in which participants chose between receiving a smaller amount of money immediately or a larger amount after a specific number of days. The immediate reward ranged from \$10 to \$34, while the delayed reward was fixed at \$25, \$30, or \$35, with delays ranging from 1 to 180 days. This choice set was designed to capture a broad range of discount rates evenly distributed in log space within the range of [-1.6, -8.4]. It was adapted from Kirby's choice set (Kirby et al., 1999) and has been widely used in the temporal discounting literature (Senecal et al., 2012; Yu et al., 2017; Parthasarathi et al., 2017; Lempert et al., 2020; Bulley et al., 2022; Batistuzzo et al., 2022). Additionally, we included five attention check trials in which participants were asked to choose between a larger immediate amount of money and a smaller amount with a delay.

We estimated temporal discount rates by fitting a hyperbolic choice model to the choice data of each participant. The utility of each option (immediate or delayed) is given by: $U = \frac{v}{1+kD}$, where U is the subjective discounted value, v is the monetary reward, D is a delay in days, and k is the individual discount rate. We used the softmax function to generate choice probabilities from option values.

$$\mathrm{Pr}_{\mathrm{delayed}} = \frac{1}{1 + e^{-\beta(U_{\mathrm{delayed}} - U_{\mathrm{immediate}})}}$$

where $\mathrm{Pr}_{\mathrm{delayed}}$ is the probability that the participant chose the delayed option on a given trial, and is the inverse temperature which captures the stochasticity of the choice data. We used maximum-likelihood estimation to estimate the model parameters. We calculated the average goodness of fit as 1 minus the ratio between the log-likelihood of the model and that of a random-response model.

Risk attitude. To estimate participants' risk attitude, we employed two methods: the incentive risky choice task and the Domain-Specific risk-taking (DOSPERT) Scale (Blais and Weber, 2006). The risky choice task assesses individual risk attitude primarily with the financial domain, while the DOSPERT scale assesses individual risk attitude across five domains: ethical, financial, health/safety, recreational, and social domain. Additionally, we designed specific questions related to the research participation task to measure participants' risk attitude in delaying research participation until the end of the semester (Supplement).

The risky choice task consisted of 57 trials, each involving a choice between receiving \$5 for sure and participating in a lottery where participants had a chance to win a larger amount with a certain probability, otherwise receiving \$0. For example, one trial presented participants with a choice between \$5 for sure and a 25% chance of winning \$16 or a 75% chance of receiving \$0 (Fig. 3.5B). The larger amounts ranged from \$6 to \$66, and we used three different winning probabilities: 25%, 50%, and 75%. The choice set was adapted from a study by Lopez-Guzman et al., 2018. To minimize any potential biases, we counterbalanced the position of the sure-bet option on the screen (left or right) and the associated color of the larger amount (blue or red). Additionally, we included seven attention check trials that presented participants with a choice between $5 for sure and a certain chance of receiving $4 or $5.

To help participants better understand the probabilities involved, the instructions included a visual representation of the choices. Each lottery image depicted a physical bag containing 100 poker chips, including red and blue chips. The size of the colored area and the number written inside indicated the number of chips of each color in the bag. The process of randomly drawing a chip was referred to as "playing the lottery."

We estimated individual risk attitudes by fitting a power utility model to the trial-by-trial choice data. In this model, the utility of each option (safe or lottery) is given by: $U = pv^{\alpha}$, where v is the dollar amount, p is the probability of winning, and α is the individual's risk attitude. A participant with $\alpha > 1$ is considered risk-seeking, $\alpha = 1$ is considered risk-neutral, or $\alpha < 1$ is considered risk-averse. Like in the delay discounting task, we used the softmax function to generate choice probabilities from option values.

$$\mathrm{Pr}_{\mathrm{lottery}} = \frac{1}{1 + e^{-\gamma(U_{\mathrm{lottery}} - U_{\mathrm{safe}})}}$$

where $\mathrm{Pr}_{\mathrm{lottery}}$ is the probability that the subject chose the lottery on a given trial, and γ is the inverse temperature which captures the stochasticity of the choice data. We used maximum-likelihood estimation to estimate the model parameters.

Incentive compatibility. Both the delay discounting task and the risky choice task were incentive-compatible. Participants were offered a bonus: at the end of the study, their choice from a randomly selected trial in either the delay discounting task or the risky choice task determined the amount of this bonus. The bonus was provided as an electronic Amazon Gift Card.

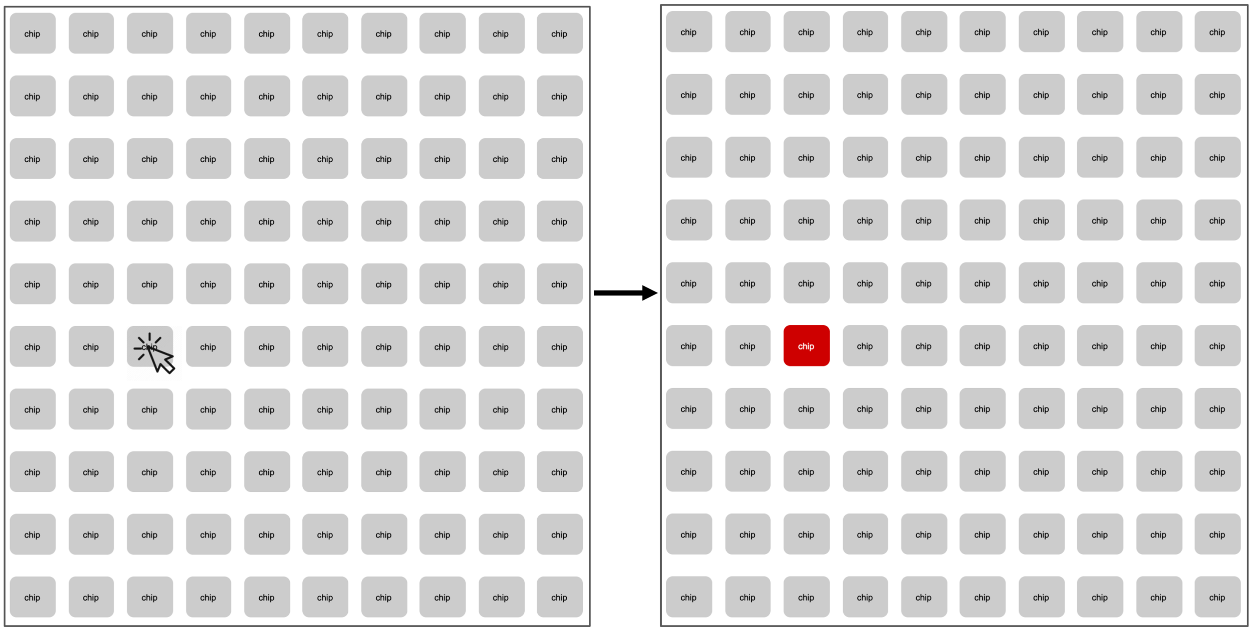

Figure 3.2: Visual aid of "playing the lottery": the process of drawing a chip at random from a set of 100 chips.

If the one randomly selected trial is from the delay discounting task, the timing of receiving the bonus depends on the chosen option. Specifically, for payment today, participants received the gift card on the same day. For delayed payments, participants received the gift card at a time corresponding to the delay associated with their chosen option.

If the one randomly selected trial is from the risky choice task, if participants chose the sure bet in the selected trial, they would receive $5. However, if they chose the lottery, they would engage in the process of drawing a chip at random from a set of 100 chips. As the task was conducted online, we provided participants with a visual aid of the chip-drawing process. We displayed 100 chips (Fig. 3.2 left panel) and instructed participants to click on a chip to simulate the random draw. After clicking, the color of the chip would be revealed (Fig. 3.2 right panel). If a participant drew a red chip, they would win the lottery; otherwise, they would lose the lottery and receive $0.

Questionnaires. To test the convergent validity of our measure of procrastination in this task (i.e., the correlation between procrastination in fulfilling the 7-hour research participation requirement and general procrastination behavior in academic settings), we used a widely-used questionnaire to measure general procrastination tendency in academic settings: Procrastination Assessment Scale-students (PASS) (Solomon and Rothblum, 1984). Participants were asked to report the frequency with which they procrastinated on tasks such as writing term papers, studying for exams, and four other academic scenarios.

For exploratory analyses (See supplement), we included additional surveys associated with procrastination and custom-designed questions that specifically addressed individuals' procrastination in the real-world research participation task, such as questions aimed to assess participants' awareness of their procrastination levels and their level of regret regarding their procrastination in research participation.

3.2.3 Participant inclusion

The sample size of the online study was 194, which was 25.9% of the students who had been enrolled in the *Introduction to Psychology* course. To ensure that our measures of procrastination would not be confounded by the total number of work units completed, we only included the subset of participants who did not continue to do research sessions after they had met their 7-hour requirement. For example, we would include a participant who, after completing 6.5 hours, did a final research session to meet the requirement. However, we would exclude one who, after completing 7 hours, did an additional session that was not required. This resulted in a total of 93 participants. Of the remaining 101 participants, 80 continued to do research sessions beyond the 7-hour requirement, potentially to earn extra credit. The remaining 21 completed fewer than 7 hours; in some cases, this was because they completed an alternative assignment (i.e., writing critique papers).

To test the hypothesis of correlation between temporal discounting and procrastination, out of 93 participants, we excluded 9 who either failed two or more of the five attention check questions or who consistently chose one option, as that would make it impossible to determine their discount rate. To ensure that participants were not responding randomly, we conducted a quality control procedure (Pehlivanova et al., 2018). We verified that participants' responses were influenced by task-relevant variables. This involved fitting to each participant's responses a logistic regression model that included as predictors the immediate amount, the delayed amount, the delay, and the squares of these variables. The goodness of fit of the model was assessed using the coefficient of discrimination, and any participant with a value below 0.2 was considered a random respondent. No participants were excluded as random respondents. This left us with a final sample of 84 participants (53 female, 28 male, 2 non-binary, 1 unknown; 19.4 ± 1.4 years old).

To test the hypothesis of correlation between risk attitude and procrastination, out of 93 participants, we excluded two subjects who either chose the objectively worse option in two or more of seven attention check trials or who consistently chose one option, as that would be impossible to determine their risk attitude. Similarly to the delay discounting task, we conducted a quality control procedure to ensure that participants were not responding randomly. No participants were excluded as random respondents. This left us with a final sample of 91 participants (56 female, 31 male, 3 non-binary, 1 unknown; 19.3 ± 1.8 years old).

3.3 Results

High individual variability in the behavioral level of procrastination. In the research participation task, we found that the time course of work progress differed greatly between individuals, ranging from participants who started and finished early and ones who worked steadily over time to ones who waited until the last two weeks of the 109-day period (Fig. 3.3A). Fig. 3.3B shows all the students' cumulative progress. We see the high individual variability.

MUCD had a wide distribution, ranging from 19.1 to 100.9 (M = 49.6, SD = 18.2), further demonstrating the high level of individual variability in procrastination (Fig. 3.3C).

Convergent validity. Before turning to our main question, we first assessed the convergent

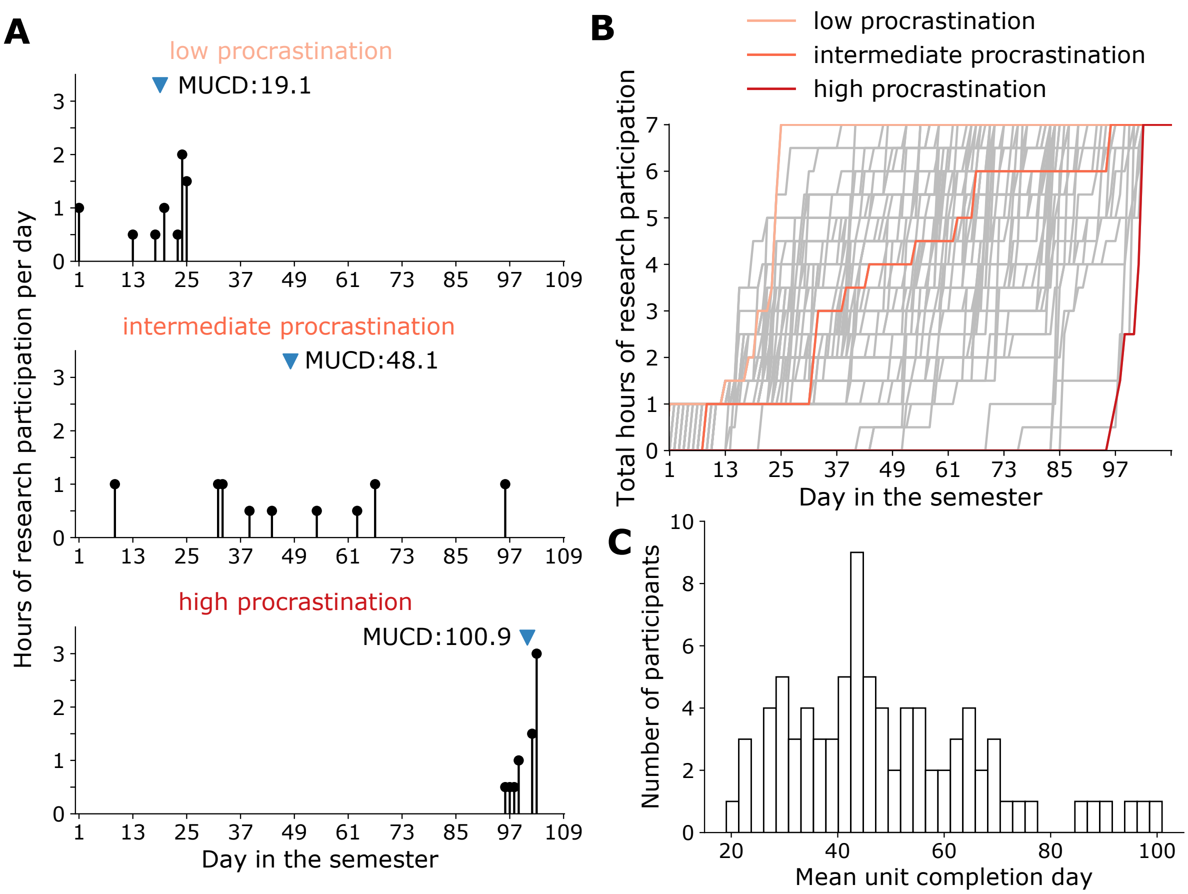

Figure 3.3: Procrastination in a real-world task. (*A*) Example time courses of work progress, with blue triangles marking the Mean Unit Completion Day (MUCD). Top: a low procrastinator who started on the first day and finished early. Middle: an intermediate procrastinator who worked steadily throughout the semester. Bottom: a high procrastinator who rushed to complete the task in the last two weeks of the semester. (*B*) Time courses of cumulative work progress for all the participants, with the three examples from (A) highlighted. (*C*) Histogram of MUCD. (*D*) Histogram of the natural log-transformed discount rate in the delay discounting task.

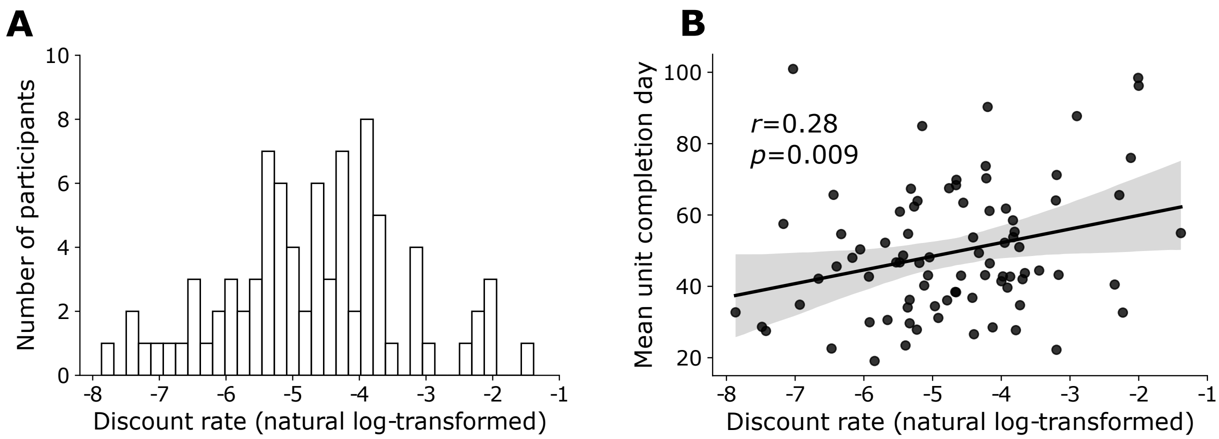

Figure 3.4: Correlation between discount rates and procrastination. (*A*) Histogram of the natural log-transformed discount rate. (*B*) Correlation between MUCD and the natural log-transformed discount rate.

validity of our measure of procrastination, that is, whether MUCD in this research participation task is associated with self-reported procrastination in general academic situations. Our findings revealed a moderate positive correlation between MUCD and PASS score (Pearson $r = 0.42$, $p < 0.001$), which provides support for the convergent validity of our measure.

Discount rates correlate with the behavioral level of procrastination. Turning to our main question, we examined the correlation between temporal discounting and procrastination. Individual discount curves were well characterized by hyperbolic functions (The goodness of fit was on average 0.73, with an SD of 0.14.). Fig. D left shows the distribution of the natural log-transformed discount rate, which ranged from -7.87 (equivalent to a 1.14% discount of reward value after 30 days) to -1.39 (an 88.2% discount of reward value after 30 days). We found a positive correlation between the discount rate and MUCD ($r = 0.28$, $p = 0.009$) (Fig right).

So far, we have only considered MUCD as a metric of procrastination, but other metrics have been used in the literature, including a) day of halfway point (median of time course of work progress) Rothblum et al., 1986; b) hours of research participation in the last third of the semester Solomon and Rothblum, 1984; and c) task completion day Raphaël and Mathias, 2022; Reuben et al., 2015. In our data, all three of these metrics correlated with the PASS score, suggesting the convergent validity of these measures ($r = 0.42$, $p < 0.001$; $r = 0.41$, $p < 0.001$; $r = 0.31$, $p = 0.005$, respectively). Regarding the correlation between temporal discounting and procrastination, metrics (a) and (b) both correlated with the discount rate ($r = 0.28$, $p = 0.009$; $r = 0.24$, $p = 0.030$, respectively), but metric (c) did not ($r = 0.21$, $p = 0.061$).

We found no significant correlation between the discount rate and the PASS score ($r = 0.21$, $p = 0.056$), highlighting the advantage of behavioral measures of procrastination over survey-based measures.

No evidence of a correlation between risk attitude and behavioral level of procrastination. Fig. left shows the distribution of the natural log-transformed risk attitude. Similarly, as the previous study proposed no correlation between risk attitude and self-reported procrastination Keinan and Bereby-Meyer, 2012, we found no evidence of a correlation between individual risk attitude, as estimated from the risky choice task, and their behavioral level of procrastination in this research participation task characterized by Mean Unit Completion Day ($r = 0.03$, $p = 0.749$) (Figure), Day of Halfway Point ($r = 0.10$, $p = 0.353$), Task Completion Day ($r = -0.06$, $p = 0.547$) and Hours in Last Third Semester ($r = 0.07$, $p = 0.525$). Similarly, no significant correlation was observed between the likelihood of participants engaging in risky activities across the five domains measured by the Domain-Specific risk-taking scale Blais and Weber, 2006 and the behavioral level of procrastination. Take Mean unit completion day as an example; no correlation was observed in ethical domain ($r = -0.11$, $p = 1.0$), in financial domain ($r = 0.08$, $p = 1.0$), in health/safety domain ($r = -0.02$ $p = 1.0$), in recreational domain ($r = -0.22$, $p = 0.31$) and in social domain ($r = -0.06$, $p = 1.0$) (corrected by Holm–Bonferroni method).

We also examined the correlation between risk attitude, specifically in relation to postponing research participation until the end of the semester, measured by our designed questionnaire, and behavioral level of procrastination. Interestingly, most participants strongly agreed that postponing research participation until the end of the semester increases the risk of not being able to fulfill the requirement (median of rating from strongly disagree (1) to strongly agree (7) is

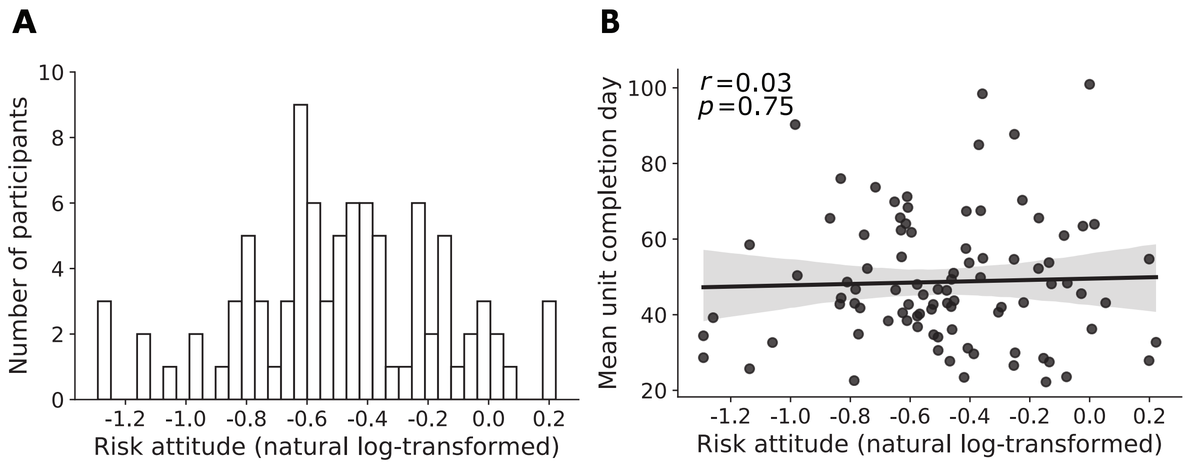

Figure 3.5: Correlation between risk attitude and procrastination. (A) Histogram of the natural log-transformed risk attitude. (B) Correlation between MUCD and the natural log-transformed risk attitude.

strongly agree (7)). However, we did not find evidence of a significant correlation between risk attitude, specifically in relation to postponing research participation and level of procrastination characterized by Mean Unit Completion Day ($r = -0.18$, $p = 0.084$), Day of Halfway Point ($r = -0.20$, $p = 0.059$), Task Completion Day ($r = -0.18$, $p = 0.088$) and Hours in Last Third Semester ($r = -0.12$, $p = 0.267$).

3.4 Discussion

We have presented empirical evidence for an association between reward discounting and procrastination behavior in a long-term real-world task. This suggests that temporal discounting is a cognitive mechanism underlying procrastination. Future work should investigate this association in other real-world tasks, especially those that are non-academic and in more diverse global samples.

Why did prior studies (Raphaël and Mathias, 2022; Reuben et al., 2015) fail to find a correlation between temporal discounting and procrastination, while ours succeeded? One reason might be that the choice sets they used might not have allowed for estimating the discount rate with the same precision as ours. An additional reason might be that their delay discounting task was not incentive-compatible. Finally, the earlier studies only used task completion day as a metric of procrastination, perhaps because they did not measure the entire time course of work progress and therefore had to resort to this metric. In contrast, we measured the entire time course of work progress and computed a fine-grained metric of procrastination, MUCD, that might provide greater statistical power than merely relying on task completion day.

In terms of no evidence of a correlation between risk attitude and behavioral level of procrastination, it is important to note that the risk attitude we measured here is associated with one type of risk-taking: active risky taking, which involves engaging in risky activities (Keinan and Bereby-Meyer, 2012). Another type of risk-taking is passive risk-taking, which arises from inaction or the avoidance of action. Procrastination, characterized by delaying actions, can be seen as a form of passive risk-taking. To gain further insights into the interplay between risk attitudes and procrastination in passive risk-taking contexts, future research should focus on utilizing validated scales such as the Passive Risk Taking scale developed by Keinan and Bereby-Meyer, 2012.

3.5 Supplement

3.5.1 Explorative analyses

Participants were aware of their own level of procrastination in research participation. To examine whether participants were aware of their own level of procrastination in research participation, they were first asked to recall how they allocated their time throughout the semester to fulfill the research participation requirement. Then, they were asked to rate from not at all to an extreme extent their procrastination level to fulfill the research participation requirement. Having the recall question before the rating question is to have participants rate their procrastination level based on their recalled time course of progress in fulfilling the requirement.

We found that the rating of their own procrastination level in research participation significantly correlates with their behavioral level of procrastination ($r = 0.679$, $p < 0.001$, quantified by Mean unit completion day and applies to other indices of procrastination).

Procrastination in research participation did not represent general procrastination tendency beyond academic settings. We tested whether this specific procrastination behavior in the research participation scenario could represent general procrastination even beyond academic situations (measured by the General Procrastination Scale (Lay, 1986 for student populations). We did not find a correlation between the average score in general procrastination and any behavioral index of procrastination in research participation (Mean Unit Completion Day ($r = 0.19$, $p = 0.09$), day of the halfway point ($r = 0.17$, $p = 0.12$), task completion day ($r = 0.20$, $p = 0.08$), and hours in last third semester ($r = 0.21$, $p = 0.06$)).

High procrastinators were less satisfied with (and regret more) the way they allocated their time over the semester to fulfill the research participation requirement (preregistered hypothesis). We found a correlation between how much participants were satisfied with the way they allocated their time over the semester to fulfill the requirement and their behavioral level of procrastination ($r = -0.571$, $p < 0.001$) (quantified by Mean unit completion day and applies to other indices of procrastination). To test whether high procrastinators regret more in both the cognitive domain and affective domain about how they allocate their time throughout the semester to fulfill the research participation requirement, we used the regret element scale (Buchanan et al., 2016). We found a correlation between how much participants regret the way they allocated their time both in affective and cognitive domains and their behavioral level of procrastination (affective regret: $r = 0.514$, $p < 0.001$; cognitive regret: $r = 0.512$, $p < 0.001$, quantified by Mean unit completion day and applies to other indices of procrastination).

Attribution of procrastination and success in fulfilling the requirement (preregistered hypothesis). We used the Causal Dimension scale (Russell, 1982) to test our hypotheses that high procrastinators attribute their high procrastination to more external, temporal, and uncontrollable factors, whereas low procrastinators attribute their low procrastination to more internal, stable, and controllable factors. There was no evidence of a correlation between behavioral level of procrastination quantified by Mean unit completion day and locus of causality (internal or external) ($r = -0.12$, $p = 0.271$), stability ($r = -0.09$, $p = 0.373$) and controllability ($r = -0.17$, $p = 0.105$). We also used the Causal Dimension scale to test another hypothesis that high procrastinators attribute their success in fulfilling the requirement to more external, temporary, and uncontrollable factors, whereas low procrastinators attribute their success to more internal and stable factors. There was no evidence of a correlation between the behavioral level of procrastination quantified by Mean unit completion day and locus of causality (internal or external) ($r = -0.05$, $p = 0.605$), stability ($r = 0.03$, $p = 0.75$) and controllability ($r = -0.05$, $p = 0.644$).

Top-rated procrastination reasons. We tested self-reported reasons for procrastination in fulfilling the research participation requirement by asking the participants to rate not at all reflects why (1) to definitely reflects why (7) on statements that consist of the four reasons: the excitement of doing things at last moment, time management, task aversiveness, and laziness. The statements are adapted from reasons for procrastination (Solomon and Rothblum 1984). The top-rated reason is time management ($M = 2.57$, $SE = 0.10$), followed by task aversiveness ($M = 2.19$, $SE = 0.10$), laziness ($M = 1.90$, $SE = 0.10$) and excitement of doing things at last moment ($M = 1.33$, $SE = 0.06$). But note, in general, ratings for all four reasons are relatively low.

No evidence of a correlation between risk perception and behavioral level of procrastination. We tested the hypothesis that people who perceive risky activities as more risky tend not to procrastinate in fulfilling the research participation requirement. This hypothesis is along a similar line with testing the relationship between risk attitude and level of procrastination but uses risk perception instead. We measured people's risk perception by Domain-Specific Risk-taking Scale (risk-perception scale), which has the same items as the risk-taking scale but with different instructions asking the subjects their gut-level assessment of how risky each situation or behavior is. We found no evidence of a correlation between risk perception in any of the five domains and the level of procrastination in research participation (quantified by Mean unit completion day and same qualitative results if using other indices) in ethical domain ($r = 0.03$, $p = 0.753$), in financial domain ($r = -0.21$, $p = 0.04$), in health/safety domain ($r = -0.03$ $p = 0.764$), in recreational domain ($r = 0.0$, $p = 0.999$) and social domain ($r = 0.01$, $p = 0.937$). None of the correlations reaches the significance level after the Bonferroni correction.

Impulsivity, self-control, and perfectionism (preregistered hypothesis) Impulsivity, self-control, and perfectionism are traits associated with self-reported procrastination (Steel, 2007). We found a positive correlation between impulsivity and behavioral level of procrastination ($r = 0.26$, $p = 0.013$). There is no evidence of a correlation between the behavioral level of procrastination and self-control ($r = -0.14$, $p = 0.184$), maladaptive perfectionism ($r = 0.12$, $p = 0.249$), or adaptive perfectionism ($r = -0.02$, $p = 0.881$). The measurement of maladaptive and adaptive perfection follows (Enns et al., 2001), where maladaptive perfectionism is measured from three subscales: socially prescribed perfectionism from Hewitt (Flett et al., 1995), concern over mistakes and doubts about actions from Frost (Frost and Marten, 1990), and adaptive perfectionism is measured from two subscales: self-oriented perfectionism from Hewitt and personal standard from Frost.

3.5.2 Two more indices of procrastination

Area under the cumulative progress curve (AUC). Steel et al., 2018 proposed the area under the cumulative progress course curve (AUC) as a measure of procrastination. We demonstrate here that AUC is a linear transformation of MUCD. To calculate AUC, we start by calculating cumulative progress as $y_t = \sum_{i=1}^{t} x_i$. AUC is now the sum of the cumulative progress values:

$$\text{AUC} = \sum_{t=1}^{T} y_t$$

which can be evaluated as

$$
\begin{aligned}
\mathrm{AUC} &= \sum_{t=1}^{T} \sum_{i=1}^{t} x_i \\
&= \sum_{i=1}^{T} (T+1-i) x_i \\
&= (T+1) \sum_{i=1}^{T} x_i - \sum_{i=1}^{T} i x_i \\
&= (T+1)N - N \times MUCD
\end{aligned}
$$

Because for us, N is fixed at 14 half-hour units, AUC is a linear transformation of MUCD and the results in this paper involving MUCD would be identical if we used AUC instead. We chose to use MUCD due to its more intuitive psychological interpretation as the mean of the time course of work progress. Additionally, the direction of MUCD aligns with the level of procrastination, where higher values indicate more procrastination. Conversely, AUC follows the opposite direction, with higher values indicating less procrastination.

Task starting day. We did not use task starting day as an index of procrastination in our study because students were asked to do the first research session within the first two weeks of the semester. Thus, the first research session had its own deadline, which was not related to the deadline of the 7-hour research participation task.

3.5.3 Designed questions to measure participants' specific risk attitudes in the research participation task

We designed three questions to measure the participants' risk attitude regarding delaying their research participation until the end of the semester. The first question measures their perception of the risk associated with not being able to fulfill the research participation requirement by postponing it until the end of the semester. Participants were asked to rate their agreement with the statement: "I believe that postponing one's research participation until the end of the semester increases the risk of not being able to fulfill the research participation requirement."

The second and third questions, designed as a pair, aim to measure the extent of aversion to the risk of not fulfilling the requirement due to delaying research participation until the end of the semester. Participants were asked to rate their agreement with the two statements: "The increased risk of not being able to fulfill the research participation requirement due to postponing the research participation was motivating and exciting for me" (second question, with reversed key) and "The increased risk of not being able to fulfill the research participation requirement due to postponing the research participation was stressful or anxiety-inducing for me" (third question).

Individuals who both perceive a higher level of risk and exhibit risk aversion are considered to be more risk-seeking in the research participation task. Therefore, we utilized the averaged score from the aforementioned three questions to measure their specific risk attitude in the research participation task.

4 Offering immediate rewards upon completion reduces procrastination

Great! We now understand why people procrastinate!

Since we know the reason,

can we use what we have learned to beat procrastination?

The association between temporal discounting and procrastination

we observed in chapter 3 suggests that

if we brought a future reward temporally closer,

then that would help reduce procrastination.

Is that true?

4.1 Introduction

People struggle with procrastination and want to reduce it. Procrastination's consequences can be quite harmful (Steel, 2007), and over 95% of procrastinators wish to reduce it (O'Brien, 2000).

Essentially, procrastination leads to two major negative outcomes: poor performance and miserable well-being. In terms of poor performance, in educational settings, students who procrastinate tend to perform worse academically. For example, research has shown that students who delay their work often receive lower grades (Beswick et al., 1988; Steel et al., 2001; Wesley, 1994). Procrastination also affects people financially. Individuals who delay tasks like filing taxes or planning for retirement often end up losing money (Martinez et al., 2017; Akerlof, 1991; O'Donoghue and Rabin, 1999). Even in politics and banking, procrastination in decision-making can lead to significant problems and losses (Farnham, 2021; Kegley, 1989; Holland, 2001).

In terms of personal well-being, procrastinators often face mental health issues such as depression and anxiety. They also tend to struggle with a lack of motivation and low self-esteem (Tice and Baumeister, 1997; Sirois et al., 2003; Sirois, 2007). Overall, procrastination is not just about delaying tasks; it has real, negative effects on our performance and overall happiness.

The negative effects of procrastination call for interventions to reduce procrastination. The interventions discussed mostly in existing literature are either emotion regulation (Eckert et al., 2016), acceptance (Wohl et al., 2010), commitment or cognitive-behavioral therapy (Scent and Boes, 2014; van Eerde and Klingsieck, 2018), or training about time management (C. et al., 2007; Wieber and Gollwitzer, 2010) and focus on the process (Krause and Freund, 2014a. All of the above interventions tell what people should do.

However, few studies in the literature explore how manipulating task characteristics can help people reduce procrastination (Schunk and Swartz, 1993; Howell and Watson, 2007). One key task characteristic is the timing of rewards. While many theories predict that adjusting reward timings can reduce procrastination, empirical evidence supporting this is limited. For instance, theories of temporal discounting suggest that procrastination is driven by the delay in rewards (Steel and König, 2006; Steel, 2007; Akerlof, 1991; O'Donoghue and Rabin, 1999; O'Donoghue and Rabin, 2001; Fischer, 1999; Fischer, 2001; Zhang and Ma, 2019). The essential idea is that at the beginning of a task, when the reward is relatively distant (usually after the deadline), the temporarily discounted reward is too small to motivate people to start working. This theory would predict a causal relationship between reward timing and procrastination: people should procrastinate less if a reward is given immediately. However, while there is theoretical backing, empirical evidence is scant in confirming whether altering reward timings, i.e., providing immediate rewards, effectively reduces procrastination.

Most related studies have focused on the addition of immediate rewards to increase task persistence. For example, some researchers found that adding immediate rewards to a task improved people's persistence (Woolley and Fishbach, 2016; Woolley and Fishbach, 2017; Lieder et al., 2019). Woolley and Fishbach found that adding a fun workout helps people persist longer in their workout. Lieder et al. found that offering points during the task versus no points helped people complete more pieces of writing. However, these studies mainly explored the impact of extra, immediate rewards on task continuation and overall work done rather than specifically focusing on procrastination (e.g., when to start, when to complete, or even the dynamic process of procrastination). Moreover, these rewards are additional benefits, not alterations to the original task rewards. In order to test the effect of reward timing, we expect studies testing the effect of changing the timing of the original task rewards on procrastination.

In this study, we aim to examine if altering the timing of original task rewards influences procrastination. One way is to offer immediate rewards upon completion. Often, in real-world situations, like submitting a homework assignment, a reward is given after the deadline. Even if we get it done well ahead of the deadline, we still have to wait until the deadline to be rewarded. One way to bring the future reward closer is to receive immediate rewards upon task completion.

This study aims to investigate whether changing the timing of task rewards impacts procrastination behaviors. One approach to this is offering immediate rewards upon task completion. Typically, in real-world scenarios such as submitting homework, rewards are only received after the deadline has passed, necessitating a wait even if the task was completed well ahead of the deadline. A strategy to make the reward feel closer and possibly reduce procrastination could be to offer rewards immediately after a task is completed.

To test if offering immediate rewards upon completion reduces procrastination, the task design is challenging because we want a task to mimic real-world procrastination while still allowing it for manipulations. It is often the case that real-world behavior is observational and hard to manipulate. Controlled experiment allows for manipulations but is far from the real world. We created a novel task that bridges these two.

We call the task BORE, which stands for Boring Online Reading Experiment. Participants were assigned 7 days to work on this lengthy reading task, which takes about 3 hours to complete all 120 reading units. Their progress is shown on the web page. The task is online and self-paced and was deliberately made to be boring, with the reading difficulty set at the level of a child. We manipulated reward timing as either an immediate reward upon task completion or a delayed reward. We measured the time course of progress and tested whether offering immediate rewards helps to reduce procrastination.

We also introduced another cognitive factor: the reward rule or the payoff function, embodying the relationship between reward value and final task performance, which is crucial in real-world contexts susceptible to procrastination. One extreme reward rule is Make-or-Break, where the reward is either all if the task is complete or none if the task is incomplete. Make-or-Break has the highest stake, with a huge loss of the exerted effort if the task is left incomplete. This reward rule represents many free-lancer jobs without partial reward in the real world. In contrast, another extreme reward role is Proportional, where the payoff is proportional to the final performance. Proportional has the lowest state since each piece of work will be rewarded, akin to hourly wage jobs in the real world. There are also real-world reward rules sitting between these two extremes, e.g., Proportional Plus Bonus, where the reward is proportional to the final proportion completed but capped at half of the maximum reward.

Existing theories primarily focus on the reward rule's influence over effort allocation across different tasks (Analytis et al., 2019) and its impact on overall performance and the time course of progress (Zhang and Ma, 2019). Empirically, it is unclear how the reward rule affects people's procrastination as well as their final performance.

We hypothesize that

1. The task completion rate is highest in the Make-or-Break condition, lowest in the Proportional condition, and intermediate in the Proportional Plus Bonus condition.

2. Offering immediate rewards upon task completion reduces procrastination regardless of reward rules.

4.2 Methods

4.2.1 Participants

We recruited subjects with a New York University NetID. This includes students, staff, and alumni of New York University. We recruited subjects through the subject pool of the Center for Experimental Social Studies at New York University, through the "Paid SONA" pool at New York University, through flyers, and through direct emails to students in different majors at New York University. A total of 611 participants joined our experiment. All participants receive a base payment of $5 for joining our experiment regardless of their performance.

4.2.2 Task paradigm

Participants were assigned 7 days to work on a lengthy reading task. The task is online. Participants received a link to the task once the experiment started. The task is self-paced, meaning participants can access the reading task anytime, anywhere, during the 7 days. They can continue the task by reopening the link to the task. They will not be able to continue the task after the deadline. The task consists of 120 paragraphs, each followed by four multiple-choice, single-answer questions (Fig. 4.1). Participants have to answer all four questions correctly to "complete the paragraph." This unambiguous definition of a unit of work allows for characterizing procrastination by the time course of progress: number of paragraphs completed per day over days. An average reader takes about 3 hours to complete all 120 paragraphs.

We used the paragraphs along with the questions and answers adapted from MCTest (Richardson et al., 2013) We revised the original MCTest to resolve ambiguities in questions and answers, corrected typos, and selected paragraphs of similar length. The reading comprehension level required for this reading test was that of 7-year-old. Participants (New York University students, staff, and alumni) should have no difficulty answering all questions correctly if they pay attention. We also provide an extra 45 paragraphs to allow them to complete 120 paragraphs in case they fail to complete a few paragraphs.

Participants' cumulative progress (i.e., the total number of paragraphs completed so far) will be shown in the header of the web page, and the number of extra paragraphs they used will be shown in the footer of each page.

4.2.3 Procedure

The experiment has four phases (Fig. 4.2): prescreening, day 1 for instruction phase, day 2 through day 8 for the reading task, and day 9 for post-task survey. The details of each phase is described below.

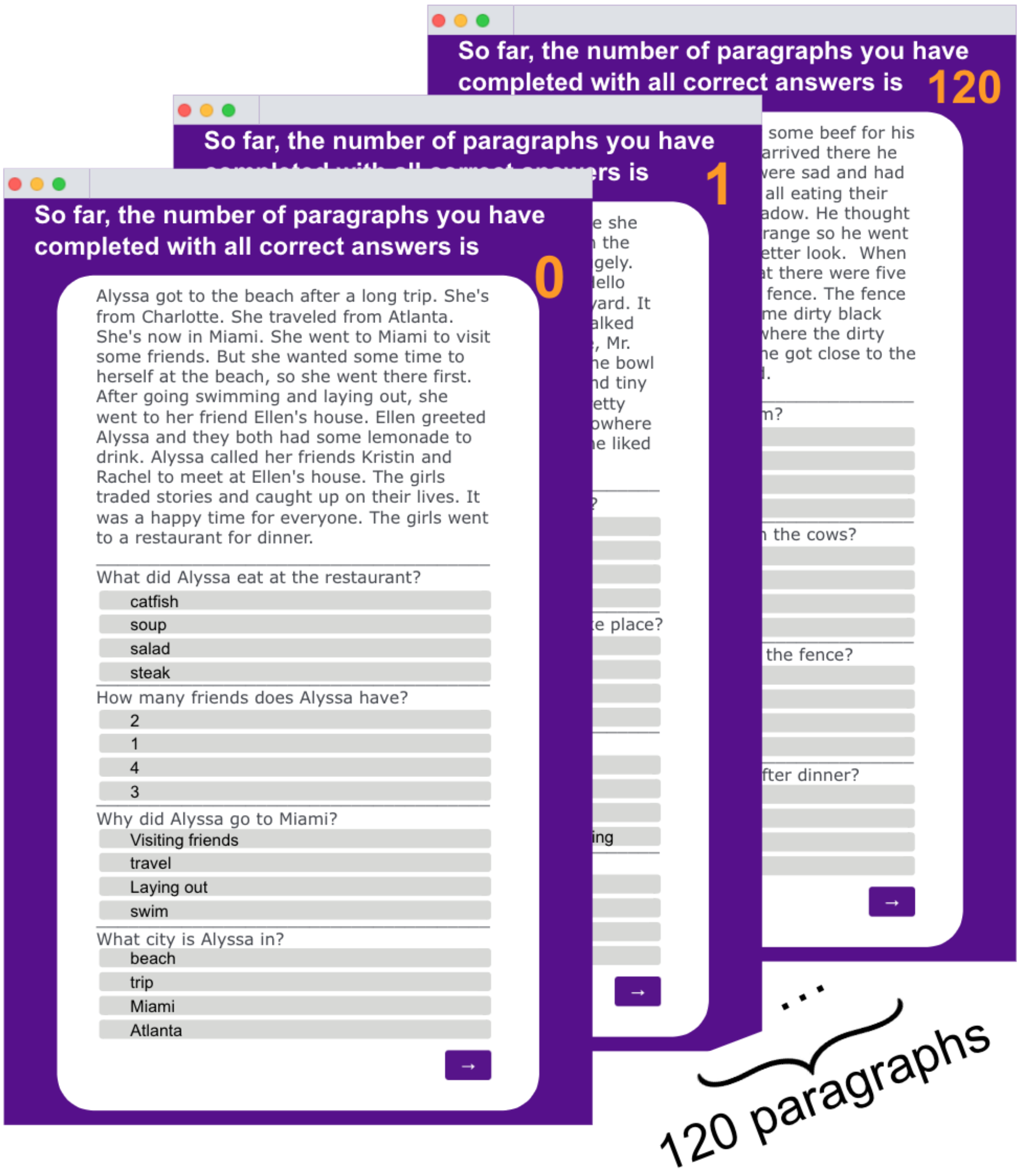

Figure 4.1: Task paradigm. The task consists of 120 paragraphs, each presented on a single page. Occasionally, participants are asked to report their feelings about the task.

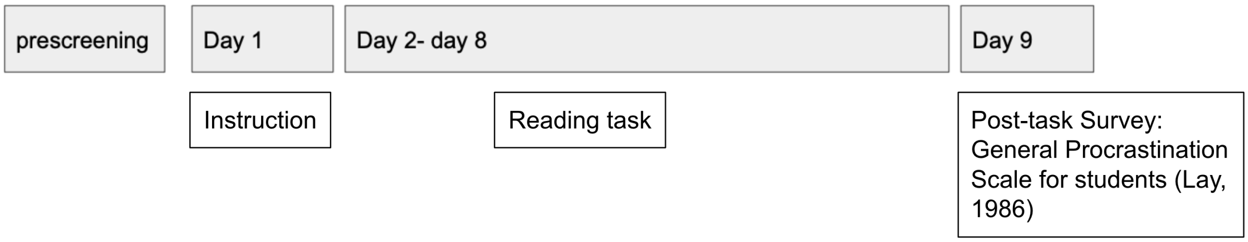

Figure 4.2: Task timeline

Prescreening: Participants must meet certain prerequisites to be eligible for the experiment:

1. having daily access to a laptop or desktop,
2. residing in the Eastern Time zone,
3. and passing a 1-minute reading quiz demonstrating English comprehension.

The participants were invited to join our experiment if they met the first two prerequisites and passed the reading quiz. The criteria for passing the reading quiz is answering all 4 questions correctly within 200 seconds.

Day 1: Invited participants were asked to complete the instruction phase of the reading task. This phase includes understanding the instructions, reporting their motivations to start and complete the task as soon as possible, and practicing with three reading paragraphs. Participants who did not complete the instruction by the end of the day were considered to have dropped out of the experiment.

On days 2 through 8: The link to the task was sent to the participants through NYU email at 1 a.m. on day 2. Participants had access to the reading task from then until 11:59 p.m. on day 8.

Day 9: A post-task survey was sent out one day after participants completed the reading task; if a participant did not complete the reading task before the deadline, the survey would be sent one day after the deadline. The post-task survey included several personality traits questionnaires, our designed questions asking about their feelings about the task and their reflections on their work. Participants received a base payment of $5 for participating in the experiment, irrespective of their task performance or whether they completed the post-task survey.

4.2.4 Experimental design

4.2.4.1 2 by 3 reward manipulations

We utilized a between-subject design, crossing two levels of reward timing and three levels of reward rule. Participants were randomly assigned to each of the six conditions, with around 100 participants in each condition. Specifically, 102 participants in the Delayed Reward and Make-or-Break condition, 102 participants in the Immediate Reward and Make-or-Break condition, 101 participants in the Delayed Reward and Proportional Plus Bonus condition, 101 participants in the Immediate Reward and Proportional Plus Bonus condition, 103 participants in the Delayed Reward and Proportional condition, and 102 participants in the Immediate Reward and Proportional condition.

In the immediate reward condition, subjects receive the payment immediately after they complete the task; in the delayed reward condition, subjects receive the payment one day after the deadline (Fig. 4.3A). If participants did not complete the task by the deadline, the timing of the reward does not differ between immediate reward and delayed reward: they receive the payment one day after the deadline, and the amount of the reward is according to their final proportion completed.

To test if the effect of reward timing on procrastination is consistent among various real-world reward rules, we manipulated three reward rules, which are representative of real-world scenarios (Fig. 4.3B). In all three reward rule conditions, if the participant completes the task, they receive the maximum reward of $36. The amount of money differs if they do not complete the task, as follows: in the Make-or-Break condition (MoB for short), they receive $2 regardless of their final proportion completed; in the Proportional condition (Pro for short), the reward is

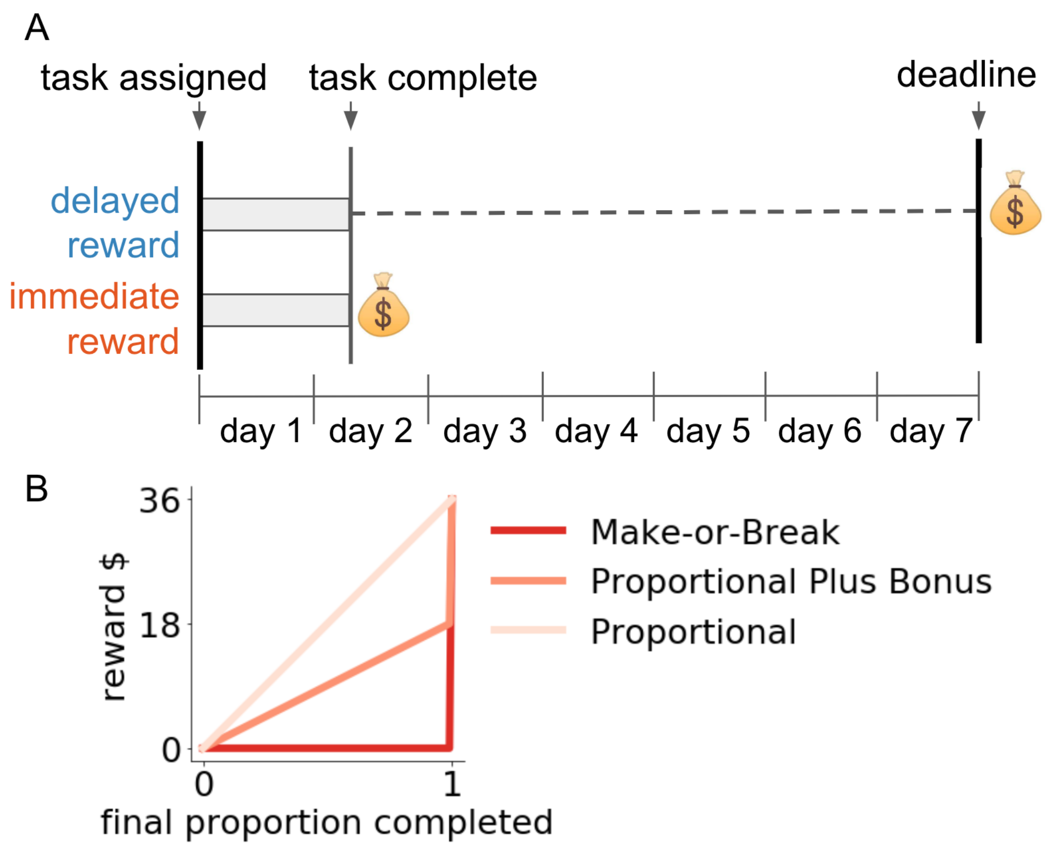

Figure 4.3: Experimental manipulations. (A) Two reward schedules: delayed reward or immediate reward upon task completion. (B) Three levels of reward rule: Make-or-Break (MoB), Proportional Plus Bonus (Pro+B), and Proportional (Pro).

proportional to their final proportion completed, with a maximum of \$36; in the Proportional Plus Bonus condition (Pro+B for short), the reward is proportional to their final proportion completed but with a maximum of \$18 (i.e., half of \$36). Make-or-Break condition is the extreme situation with the highest stake: if the participant completes the task, they earn \$36, but if they leave only one paragraph incomplete, they earn almost nothing. Proportional condition is another extreme with the lowest stake: they earn more if they work more, and each unit of work makes \$36/120. Proportional Plus Bonus condition sits between two extremes.

4.2.4.2 Motivation

To test the effect of reward timing on participants' motivation to start and complete the task as soon as possible, participants were asked to report their level of motivation after they understood the task instructions and before they received the link to the task. The pre-registered questions are: 1) I feel motivated to start working on this reading task as soon as possible, and 2) I feel motivated to complete 120 paragraphs as soon as possible. Participants were asked to rate their motivation on a scale ranging from one (not at all motivated) to five (very motivated).

4.2.4.3 Questionnaires

Post-task, participants filled out questionnaires assessing various personality traits, feelings towards the task, and self-reflection on their behavior. Here, we highlight the General Procrastination Scale for Students and task aversion levels.

4.2.5 Control measures

Since it is a naturalistic task, even if our goal is to test the effect of reward timing, many other factors unrelated to our hypothesis could affect procrastination. We kept those factors in check so that they would not overshadow the impact of reward timing. Those factors are reading speed, the day of the task within a week, the device to use to work on the task, and the time zone.

Reading speed and reading accuracy. We controlled for the reading speed and reading accuracy by the following approaches.

1. **Task difficulty.** Imagine the reading task is difficult (e.g., at the level of the Graduate Record Examinations (GRE)); there would be great individual differences in reading ability (i.e., reading speed and reading accuracy). This individual difference in reading ability will affect the total hours needed to complete 120 paragraphs and finally affect procrastination. It is expected that participants with better reading comprehension will need fewer hours to complete the task, and thus, they complete the task earlier. To control for reading comprehension, we used the paragraphs along with the questions and answers adapted from MCTest (Richardson et al., 2013). The reading comprehension level required for this reading test was that of 7-year-old. Since participants are all New York University students, staff, or alumni, this reading task should be very easy for them.

2. **Fictional stories versus facts.** Another reason we chose MCTest is that all the paragraphs are fictional stories written by Amazon Mechanical Turk workers, so there are no known facts. Participants all have to read the paragraphs to be able to answer the follow-up questions. Pre-known knowledge does not help people to read faster.

3. **Prescreening to exclude extremely slow readers.** In the prescreening, all participants were requested to work on a reading paragraph (a sample from MCTest). They were asked to work on the task without interruptions. They were told that their accuracy and the time they spent answering the questions would be recorded to determine whether they passed the prescreening. In this way, we measure their reading speed and exclude those who are extremely slow in reading. Participants who are invited to our study need to answer all 4 questions correctly within 200 seconds. The criteria of 200 seconds is set according to the 2SD of the reading duration in the pilot data.

Day of the task within a week. We randomized the first day of the task in a week for each participant. So we had equal numbers of participants whose first day of the task was on Mondays, Tuesdays, etc. Specifically, one-seventh of the participants started the task on Monday, another one-seventh on Tuesday, etc. The motivation to control the day of the task within a week is to avoid that business level of the day being confounded with procrastination. For example, people are less busy on weekends, and if we have all the participants' first day of the task on Monday, and the deadline day is on Sunday, and we observe the ramping up progress towards the deadline, it could simply be that people are less busy and have more time to work on the task on weekends, instead of a typical process of procrastination as rushing in the end.

Device access. We control the access availability to the task. One of the prerequisites to joining our experiment is having daily access to a laptop or desktop. Participants were asked to check if they met this prerequisite in the prescreening by answering the following two questions. 1) Do you have your own laptop or desktop computer available for the next two weeks? Yes or no. 2) When do you have access to your laptop or desktop computer for the next two weeks? Three options were on weekdays only, on both weekdays and weekends, and on weekends only. The motivation to have this prerequisite is to make sure participants have full access to the task at any time. This essentially grants participants the automaticity to decide when to work on the task. Otherwise, if a participant has limited access to a laptop or desktop in a week, for example, only for a few hours on a single day, then their time course of progress is not influenced by the task manipulations but instead by the access availability.

Next, we control the device type. People's behavior on Mobile devices could be different from the behavior on desktops or laptops; for example, people could be more easily distracted if they are working on Mobile than on laptops. To control for any confounding factors originating from the device type, we granted access to a laptop or desktop but blocked access to a Mobile device (e.g., mobile phone, iPad). Participants were told in the prescreening that they were allowed to have access to our experiment only through desktops or laptops. The web page of the experiment will be locked if they use Mobile devices instead.

Time zone. Another prerequisite to joining our experiment is that the participant needs to be located in the same time zone as Eastern Time. Since our experiment is time-sensitive, we control the potential confounding factors originating from the different times of the day. An example of a confounding factor could be the energy level associated with different times of the day. Let's say two participants with a 12-hour time-zone difference try to catch up for the last three hours right before the deadline; one is in the daytime, and another is in the nighttime. As both of them are more tired in the nighttime than in the daytime, their final performance could be different due to when they work on a day for the last few hours.

4.2.6 Quantify procrastination levels

To quantify procrastination by taking into account the full dynamic process of procrastination, we quantify procrastination as the following three indices:

1. Task starting day (i.e., on which day the participants start working on the task),

2. Task completion day (i.e., on which day the participants complete the task),

3. Mean unit completion day (short for MUCD).

   Mean unit completion day is calculated as the average day on which each of the 120 paragraphs is completed, where each completed paragraph is considered a unit. A participant who finishes most paragraphs later in the process will have a higher MUCD compared to someone who completes them earlier. For instance, in the extreme case where all paragraphs are completed on the last day, the MUCD would be 7. A higher MUCD indicates a higher level of procrastination, and it takes into consideration the entire dynamic process of procrastination.

4.3 Results

4.3.1 Desired task properties

Success in controlling confounding variables. We chose an extremely simple reading task in order to control the reading speed or accuracy. If the control is successful, we expect neither reading speed nor reading accuracy should correlate with procrastination. The reading speed is measured in the prescreening reading quiz as duration takes the participant to complete the sample reading paragraph without interruptions ($M = 66.07$ seconds, $SD = 24.25$). We found no correlation between reading speed and procrastination quantified by any of the index (for task starting day: $r = 0.002$, $p = 0.956$; for mean unit completion day: Pearson $r = 0.058$, $p = 0.318$; for task completion day: $r = 0.075$, $p = 0.194$). Reading accuracy is quantified as the proportion of the number of paragraphs completed with all answers correctly in the total number of paragraphs worked. As expected, since the task is designed to be extremely easy, the accuracy is high and with little variance ($M = 0.92$, $SD = 0.06$). We found the same thing to reading accuracy: no correlation with procrastination (for task starting day: $r = -0.084$, $p = 0.061$; for mean unit completion day: Pearson $r = 0.092$, $p = 0.115$; for task completion day: $r = 0.068$, $p = 0.239$).

We also controlled the day of the week when participants worked on the task. We had equal numbers of participants whose first day of the task was on Mondays, Tuesdays, etc. If the control is successful, we expect to see no significant difference in the number of paragraphs completed over the days of the week (Mondays, Tuesdays, etc.). We conducted a one-way ANOVA, and there was no significant difference in the number of paragraphs completed over the days of the week ($F(6, 4270) = 1.24$, $p = 0.28$).

Convergent validity. We assessed the convergent validity of our measure of procrastination, that is, whether self-reported procrastination in general situations (measured by the General Procrastination Scale for Students) is correlated with behavioral procrastination in this specific reading task. Our findings revealed a positive correlation between the General Procrastination Scale score and task starting day ($r = 0.16$, $p = 0.001$), mean unit completion day ($r = 0.27$, $p < 0.001$), and task completion day ($r = 0.20$, $p < 0.001$).

Great individual variability in the time course of progress. One key aspect of a good design is to have great individual variability in the levels of procrastination. Figure shows the time course of progress for a total of 611 participants. We did observe great individual variability in the time course of progress, as desired. 48.8% of the participants completed the task, and among them, we observed two extreme time courses of progress. One extreme is that participants completed the task on the first day (5.4%, Fig.4.4 upper panel , MUCD=1), whereas another extreme is that participants did not work at all on the first six days and rushed to complete the task on the last day (4.4%, MUCD=7). In addition, there is a full spectrum between these two extremes (examples in Fig. 4.4 lower panel, $MUCD = 4.30 \pm 1.79$). We observed the "types" of time courses described in previous work (steady workers, pre-crastinators, and procrastinators). Among non-completers, 35.8% were non-starters, and the rest of them started the task but left the task incomplete.

4.3.2 Reward rule affected completion and persistence

Final performance. The distribution of the final proportion completed showed three categorical performances: non-starting ($N = 112$), starting but incomplete ($N = 201$), and complete ($N = 298$) (Fig. 4.5). We further confirmed that there are three distinct categories of performance by comparing models of binomial distribution with peaks at 0 and 1 or both 0 and 1, and the model of binomial distribution with both peaks wins overwhelmingly.

Since there are three categorical performances, instead of testing the effect on the final proportion completed, which combines these three categories, we test the effect of reward manipulations (i.e., reward timing and reward rule) on performance in the following three situations:

1. Whether to or not start the task;

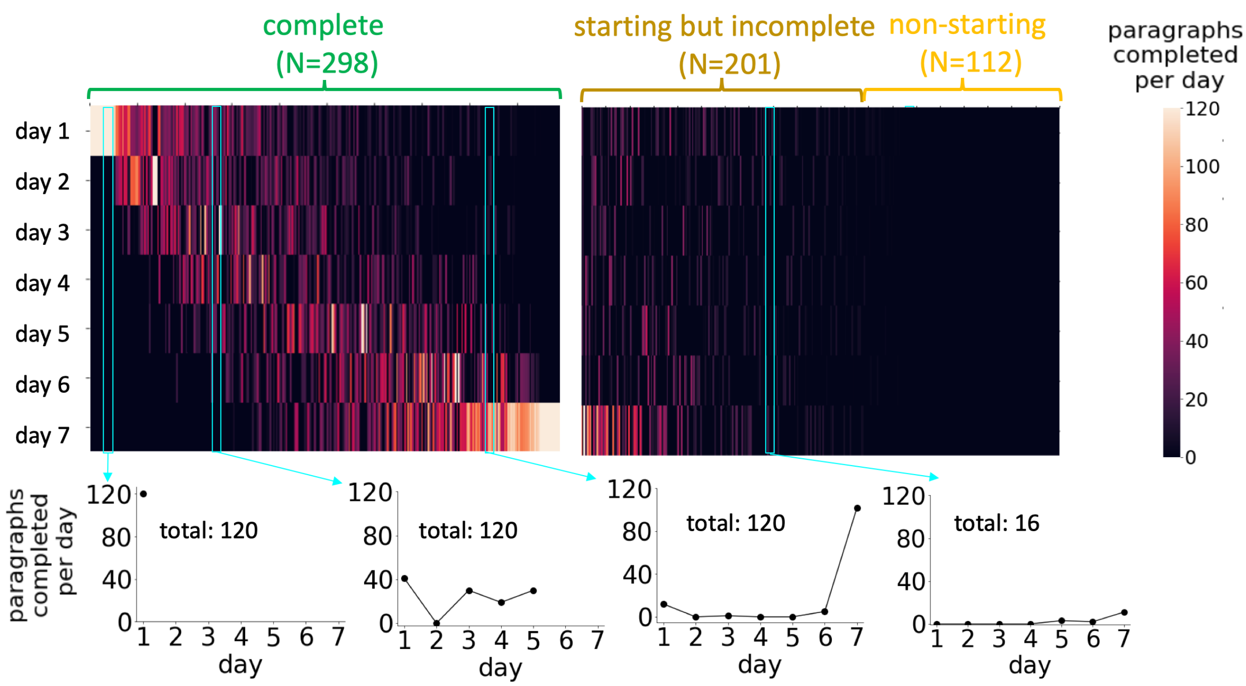

Figure 4.4: Time course of work progress for all participants. Upper panel: each column represents a single participant's time course of work progress. Each row is a day. The color indicates the number of paragraphs completed per day, with lighter the color indicating more paragraphs completed on that day. Among all 611 participants. 298 participants completed the task and we sorted the time courses according to the mean unit completion day (lower on the left and higher on the right). 201 participants started the task but did not complete the task and the rest of the participants did not work on the task (non-starting). We sorted these starting but incomplete and non-starting participants together based on their final proportion completed (higher on the left and lower on the right). Lower panel: examples of time course of progress. From the left to the right, the first person completed the task on the first day; the second person roughly divided his work into 5 days; the third person worked a bit on the first day and rushed to complete the task on the last day; and the last person did not complete the task.

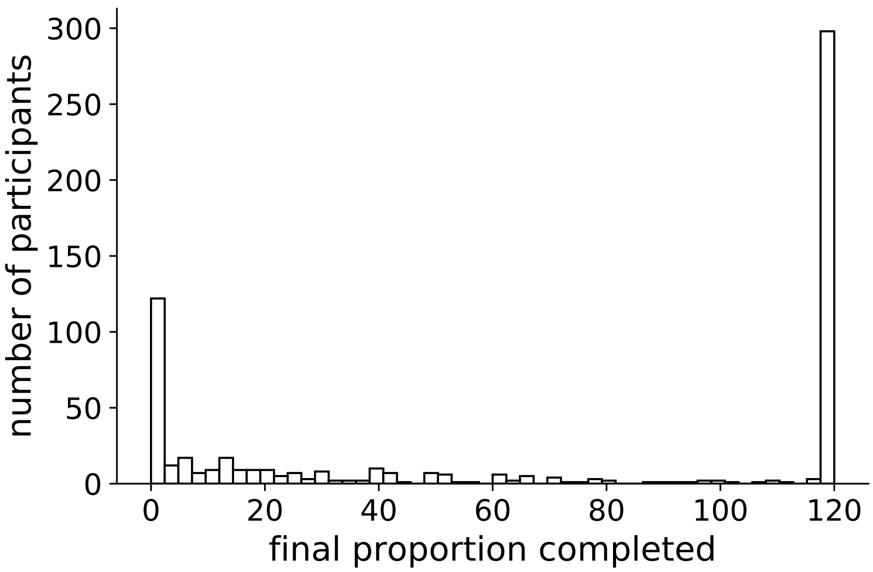

Figure 4.5: The distribution of final proportion completed. It showed three distinct categorical final proportion completed: non-starting, starting but incomplete, and complete.

2. Among starters, whether or not complete the task;

3. Among all the participants, whether or not they complete the task.

For the data analysis, we first test the effect of reward manipulations, and then we test the effect of reward manipulations together with the General Procrastination Scale score to see if the effect depends on participants' procrastination tendencies in general.

First, we conducted logistic regression of whether or not they started the task over reward timing and reward rule and its interaction. We found no main effect of reward timing ($\chi^2(1) = 1.06, p = 0.31$) or reward rule ($\chi^2(2) = 0.64, p = 0.73$) and no interaction either ($\chi^2(2) = 2.74, p = 0.25$). So, reward manipulations did not affect whether participants started the task or not.

Among starters, we conducted logistic regression to see whether or not completing the task is affected by reward timing and reward rule. We found a main effect of the reward rule on whether or not completing the task ($\chi^2(2) = 13.41, p = 0.001$). A posthoc test revealed that participants in Make-or-Break condition ($z = 3.25, p = 0.003$, Cohen's h for proportions= 0.64) or in Proportional Plus Bonus ($z = 3.00, p = 0.008$, Cohen's h for proportions= 0.58) are significantly more likely to

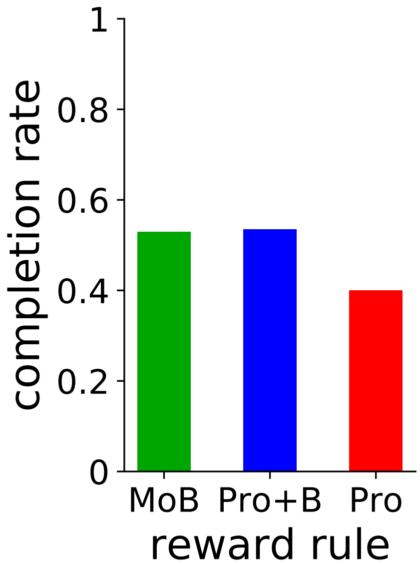

Figure 4.6: Completion rate in three reward rule conditions. Completion rate is significantly different across three rewar rule conditions ($\chi^2(2) = 9.56$, $p = 0.008$). More participants completed the task in MoB than in Prop condition ($z = 2.63$, $p = 0.023$). The same thing is for Pro+B ($z = 2.73$, $p = 0.017$).

complete the task than those in Proportional condition, and no difference between Make-or-Break and Proportional Plus Bonus ($z = -0.28$, $p = 0.96$, $h = -0.06$).

There was no main effect of Reward Timing ($\chi^2(1) = 0.19$, $p = 0.66$) or no interaction ($\chi^2(2) = 2.13$, $p = 0.35$). So reward rule, but not reward timing, affects starters' completion.

Now, among all the participants (i.e., including non-starters), we test whether or not completing the task is affected by reward timing and reward rule. We got the same results as considering starters only (Fig. 4.6): a main effect of reward rule ($\chi^2(2) = 9.56$, $p = 0.008$), and a posthoc test revealed that participants in Make-or-Break condition ($z = 2.63$, $p = 0.023$, Cohen's h for proportions= 0.45) or in Proportional Plus Bonus ($z = 2.73$, $p = 0.017$, Cohen's h for proportions= 0.47) are significantly more likely to complete the task than those in Proportional condition, and no difference between Make-or-Break and Proportional Plus Bonus ($z = 0.1$, $p = 0.99$, $h = 0.02$); no main effect of reward timing ($\chi^2(1) = 0.04$, $p = 0.84$) or no interaction ($\chi^2(2) = 3.89$, $p = 0.14$). So reward rule, but not reward timing, affects all the participants' completion.

Second, we conducted logistic regression of reward timing, reward rule, and General Procrastination Scale and their two-way interactions and three-way interaction on whether participants completed the task or not. We did this regression only among starters; we did not include non-starters because, for a total of 112 non-starters, only 8 participants (0.7%) did the post-task survey and had the General Procrastination Scale Score. We found a main effect of the reward rule again ($\chi^2(2) = 19.22, p < 0.0001$) and a main effect of General Procrastination Scale score ($\chi^2(1) = 9.72$, $p = 0.002$), with no effect of reward timing or any interactions. Participants who tended to procrastinate in general were more likely to leave the task incomplete. This is the consequence of having a higher procrastination tendency that people with high GPS tend to have poor performance.

Persistence. In the above session, we talked about the effect of reward manipulations on the final performance, which is the final status of the work; here, instead, we asked about the process taking to the final status, that is, the effect of reward manipulations on participants' persistence/-continuation in working on the task. To answer this question, we used the Cox Proportional Hazards regression model, a model for survival curve analysis, using the *survival* and *survminer* packages in the R programming environment (R Core Team et al., 2013). Although these packages are more commonly used in studies of health-related outcomes (e.g., to model rates of morbidity after patients receive one of two cancer treatments), these packages have also been used to model persistence in psychological studies (e.g., to predict the probability of children persisting in playing the game) (Rhodes et al. 2019; McGuire and Kable 2012). Here, we used these models to predict the probability of participants choosing to stop working on the task across a maximum of 120 paragraphs.

There was a main effect of the reward rule ($\chi^2(2) = 10.64$, $p = 0.005$). Participants in the Make-or-Break condition ($\beta = -0.26$, SE= 0.14, $z = -1.93$, $p = 0.054$) were less likely to drop out of the task than those in the proportional condition. Same thing for participants in the Proportional Plus Bonus condition ($\beta = -0.29$, SE= 0.14, $z = -2.14$, $p = 0.032$). No difference between

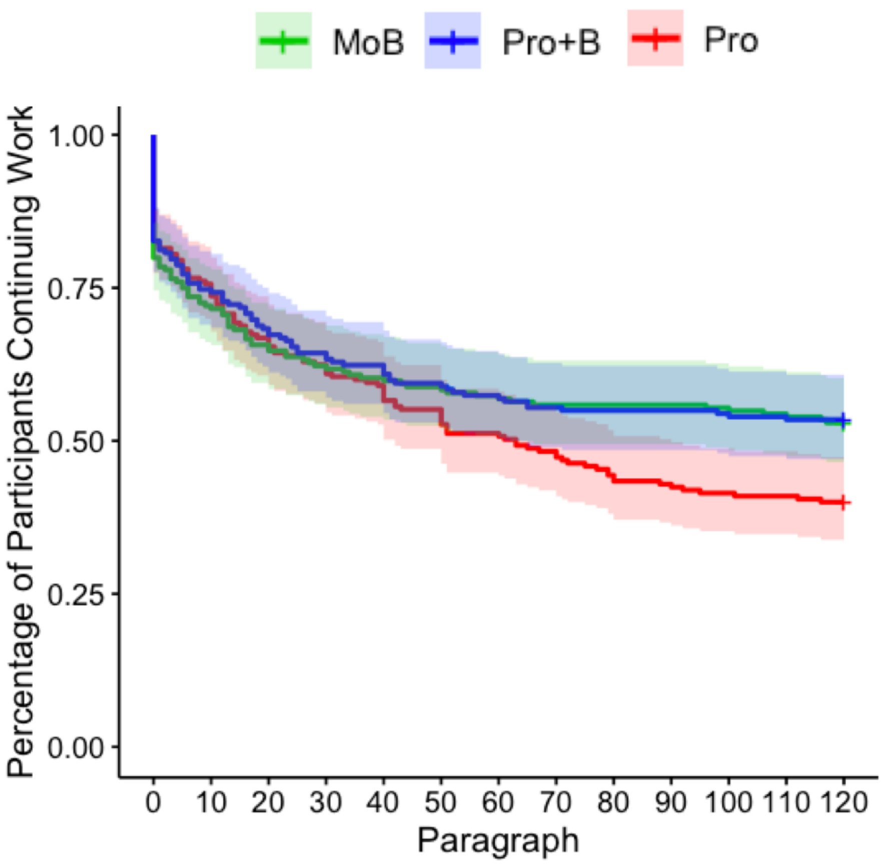

Figure 4.7: Percentage of participants persisting/continuing work. The big drop of percentage in the beginning is a reflection of non-starters. The final percentage of persisting at 120th paragraph is the completion rate.

the Make-or-Break condition and the Proportional Plus Bonus condition ($\beta = 0.03$, SE$= 0.14$, $z = 0.22$, $p = 0.83$).

In summary, the reward rule affects whether participants complete the task or no, and also affects their persistence in working on the task.

4.3.3 Offering immediate rewards motivated people to start and complete earlier

Before we examine the effect of reward manipulations on procrastination behavior, we first examined their effects on participants one key psychological state: their motivation. Aligned with indices of procrastination behavior, when to start and when to complete, we examined their motivation to start as soon as possible and their motivation to complete as soon as possible.

For data analysis, we first did a two-way ANOVA analysis for reward timing and reward rule excluding the General Procrastination Scale Score, and then we did a three-way ANOVA by including the General Procrastination Scale Score. The reason we include the General Procrastination Scale Score is to see if the effect of reward manipulations differs on people with different levels of procrastination tendencies in general.

Motivation to start as soon as possible. There was a main effect of reward timing on participants' motivation levels to start the task as soon as possible ($F(1, 605) = 6.41$, $p = 0.012$) (Fig. 4.8A). Participants reported a higher level of motivation to start the task as soon as possible given immediate reward treatment ($M = 3.30$, $SD = 0.81$) than that given delayed reward treatment ($M = 3.13$, $SD = 0.80$). There was no main effect of the reward rule ($F(2, 605) = 1.76$, $p = 0.26$) or interaction ($F(2, 605) = 0.72$, $p = 0.49$), suggesting that the effect of reward timing on motivation is consistent across all three reward rules.

After we include the General Procrastination Score, there was the main effect of reward timing again ($F(1, 449) = 8.94$, $p = 0.003$), and a main effect of the General Procrastination Score

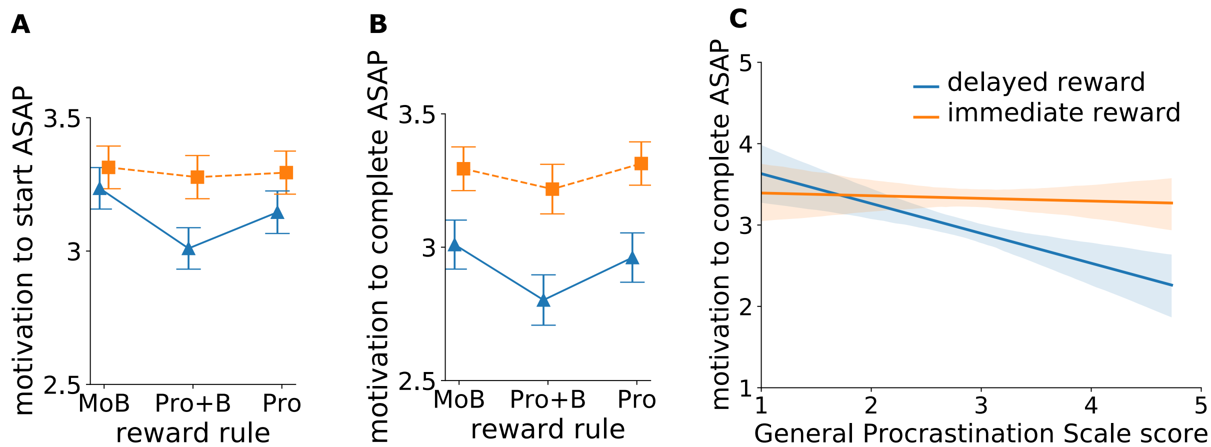

Figure 4.8: The effects of reward timing and reward rule on levels of motivation to start the task as soon as possible and to complete the task as soon as possible. (A) Offering immediate rewards increased the level of motivation to start as soon as possible ($F(1, 605) = 6.41$, $p = 0.012$). (b) Offering immediate rewards increased the level of motivation to complete the task as soon as possible ($F(1, 605) = 23.14$, $p < 0.00001$). (c) Offering immediate rewards increased the level of motivation to complete the task as soon as possible in those people who generally procrastinated more ($F(2, 449) = 6.56$, $p = 0.011$).

($F(1, 449) = 6.87$, $p = 0.009$). Participants with a higher general tendency to procrastinate reported a lower level of motivation to start working as soon as possible. There was no main effect of reward rule ($F(2, 449) = 2.01$, $p = 0.14$) or any interactions (between reward timing (RT) and reward rule (RR): $F(2, 449) = 0.22$, $p = 0.80$; between RT and General Procrastination Scale score (GPS): $F(1, 449) = 1.67$, $p = 0.20$; between RR and GPS: $F(2, 449) = 0.18$, $p = 0.83$; between RT, RR and GPS: $F(2, 449) = 0.24$, $p = 0.79$).

Motivation to Complete the task as soon as possible. There was a main effect of reward timing ($F(1, 605) = 23.14$, $p < 0.00001$) (Fig. 4.8B). Participants reported a higher level of motivation to start the task as soon as possible given immediate reward treatment ($M = 3.30$, $SD = 0.81$) than that given delayed reward treatment ($M = 3.13$, $SD = 0.80$). There was no main effect of reward rule ($F(2, 605) = 1.53$, $p = 0.22$) or interaction ($F(2, 605) = 0.27$, $p = 0.76$).

After we include the General Procrastination Score, there was a main effect of reward timing ($F(1, 449) = 25.73$, $p < 0.000001$), a main effect of the General Procrastination Scale Score ($F(1, 449) = 8.73$, $p = 0.003$), and the interaction between reward timing and General Procrastination Scale Score ($F(2, 449) = 6.56$, $p = 0.011$) (Fig. 4.8C). This interaction suggests that offering immediate rewards increased the level of motivation to complete the task as soon as possible in those participants who generally procrastinated more. Simple slope analysis shows that when the General Procrastination Scale Score is above 2.55, immediate reward significantly increases the participant's level of motivation to complete the task as soon as possible than delayed reward.

There was no main effect of reward rule ($F(2, 449) = 1.38$, $p = 0.25$), or any other interactions (RT*RR: $F(2, 449) = 0.29$, $p = 0.75$; RR*GPS: $F(2, 449) = 0.66$, $p = 0.52$; RT*RR*GPS: $F(2, 449) = 0.46$, $p = 0.63$). The non-significant three-way interaction suggested that the interaction between reward timing and the General Procrastination Scale score is consistent across all three reward rules.

In summary, offering immediate rewards increased participant's motivation to start the task as soon as possible and increased the motivation to complete the task as soon as possible for those who generally procrastinate more.

4.3.4 Offering immediate rewards reduces procrastination

For data analysis, we first did regression for reward timing and reward rule and their interaction without including the General Procrastination Scale Score, and then we included the General Procrastination Scale Score and all the possible interactions associated with the General Procrastination Scale Score. Again, the reason we include the General Procrastination Scale Score is to see if the effect of reward manipulations differs on people with different levels of procrastination tendencies in general.

For task starting day and task completion, we first did linear regression, and then we did generalized linear regression to confirm the qualitative result from the linear regression since they are discrete outcomes and their distributions violate normal distribution (Fig. 4.9).

For task starting day and task completion day, we also tested whether the effect of reward manipulations on procrastination is mediated by motivation.

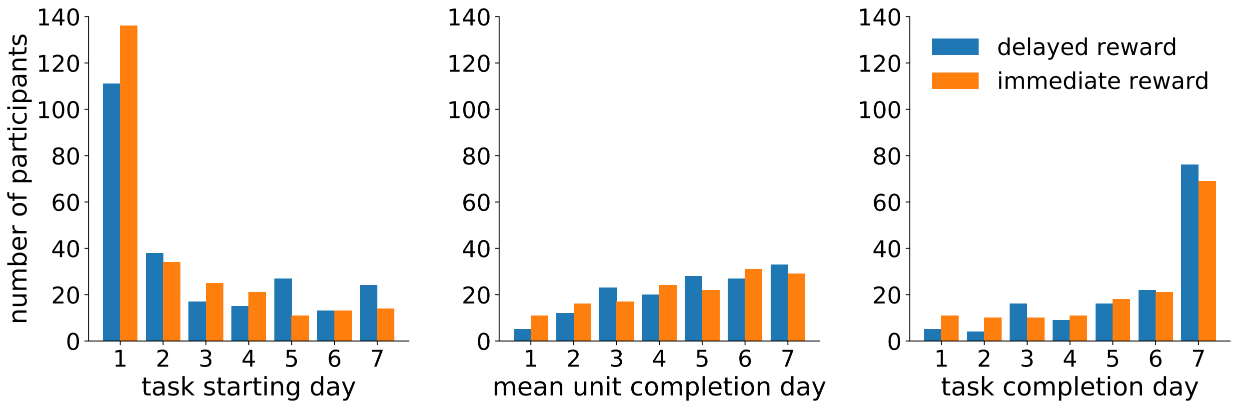

Figure 4.9: Histogram of indices of procrastination. (a) task starting day (b) mean unit completion day (c) task completion day

Task starting day. There was a main effect of reward timing ($F(1, 493) = 5.68$, $p = 0.018$), no main effect of reward rule ($F(2, 493) = 0.796$, $p = 0.45$) or interaction ($F(2, 493) = 0.048$, $p = 0.95$) (Fig. 4.10). We also conducted quasi-Poisson regression, and the result was consistent with the linear regression: main effect of reward timing ($\chi^2(1) = 5.72$, $p = 0.017$), no main effect of reward rule ($\chi^2(2) = 1.59$, $p = 0.45$) or interaction ($\chi^2(2) = 0.16$, $p = 0.92$), which suggests that offering an immediate reward helps start the task earlier, regardless of the reward.

After we include General Procrastination Score, we found a marginal main effect of reward timing ($F(1, 441) = 3.06$, $p = 0.081$); this seemingly inconsistent result with the main effect found in the above regression excluding General Procrastination Score is due to the decreased sample size (from 499 to 453) because 46 participants did not work on the post-task survey and did not have General Procrastination Score. This main effect of rewarding timing on task starting day is stronger in those participants who did not have a General Procrastination Score than those who had a General Procrastination Score. Consistent with above regression, there was a main effect of General Procrastination Score ($F(1, 441) = 12.66$, $p = 0.0004$), no main effect of reward rule ($F(2, 441) = 0.92$, $p = 0.40$) or interaction (RT*RR: $F(2, 449) = 0.12$, $p = 0.89$; RT*GPS: $F(1, 449) = 2.19$, $p = 0.14$; RR*GPS: $F(2, 449) = 1.92$, $p = 0.15$; RT*RR*GPS: $F(2, 449) = 0.80$, $p = 0.45$).

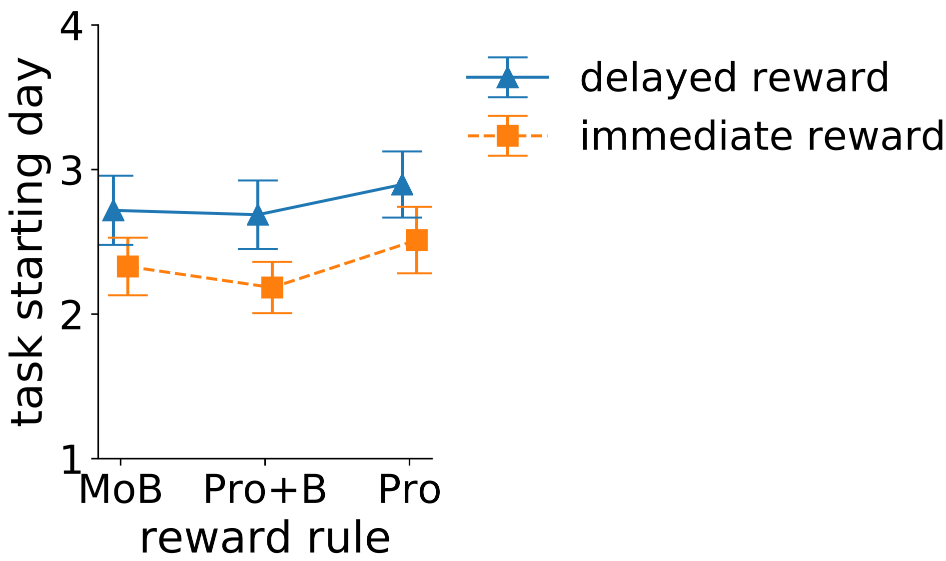

Figure 4.10: The effect of reward timing and reward rule on task starting day: offering immediate rewards helps people started the task earlier ($F(1, 493) = 5.68$, $p = 0.018$).

Since we found that there is a main effect of reward timing on task starting day, we further asked if this effect is mediated by participants' motivation to start the task as soon as possible. We conducted a simple mediation effect analysis using *mediation* package (R Core Team et al., 2013). The reward timing significantly predicted the motivation ($a = 0.23$, $SE = 0.07$, $t = 3.35$, $p < 0.001$). Controlling for reward timing, the mediator motivation significantly predicted the task starting day ($b = -0.33$, $SE = 0.12$, $t = -2.84$, $p = 0.0047$). The reward timing also predicted the task starting day ($c' = -0.36$, $SE = 0.18$, $t = -2.00$, $p = 0.046$). The effect of reward timing on the task starting day is partially mediated by motivation ($ab = -0.076$, 95% CI=$[-0.16, -0.01]$, $p = 0.016$).

Mean unit completion day. Excluding General Procrastination Scale Score, there was no effect of reward timing ($F(1, 292) = 1.24, p = 0.27$) (Fig. 4.11A), reward rule ($F(2, 292) = 0.17, p = 0.84$) or interaction ($F(2, 292) = 0.038, p = 0.96$).

After we include General Procrastination Score, there was a main effect of the General Procrastination Scale Score ($F(1, 281) = 22.9, p < 0.001$) and interaction between the General Procrastination Scale Score and reward timing ($F(1, 281) = 8.38, p = 0.004$) (Fig. 4.11B). This in-

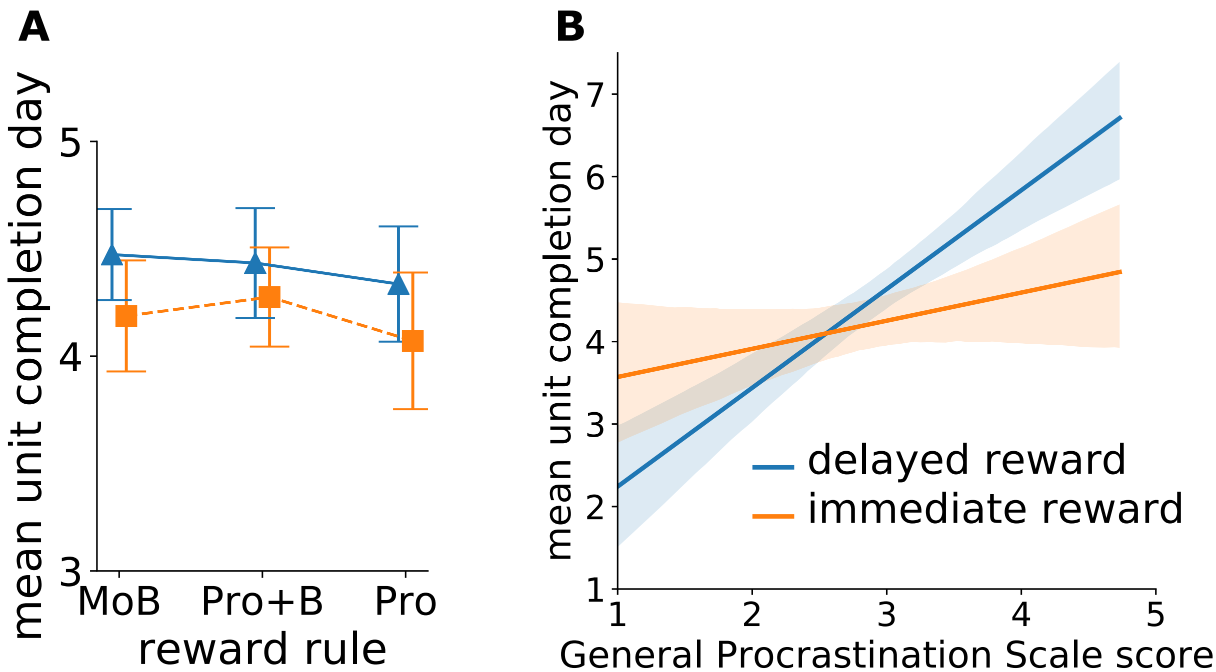

Figure 4.11: The effects of reward timing and reward rule on mean unit completion day. (A) There is no main effect of reward timing on mean unit completion day ($F(1, 292) = 1.24, p = 0.27$) (b) There is an interaction between reward timing and the General Procrastination Scale score ($F(1, 281) = 8.38, p = 0.004$) indicating that offering immediate rewards help those people who generally procrastinate more to complete units of work earlier.

teraction suggested that immediate reward helps to complete individual units of work earlier in people who have a greater tendency to procrastinate in general, regardless of the reward rule (no three-way interaction RT*RR*GPS: $F(2, 281) = 0.19, p = 0.83$). Simple slope analysis indicates that when the General Procrastination Scale Score is above 3.44, immediate reward rather than delayed reward significantly helps participants to complete individual units of work earlier. There was no main effect of reward rule ($F(2, 281) = 0.07, p = 0.93$) or other interactions (RT*RR: $F(2, 281) = 0.02, p = 0.98$; RR*GPS: $F(2, 281) = 0.70, p = 0.50$).

Since we found the interaction between reward timing and the General Procrastination Scale Score, we further asked if this interaction is mediated by participants' motivation to complete the task as soon as possible. We first tested, among those participants who completed the task, if the interaction predicted their motivation to complete the task as soon as possible and found a

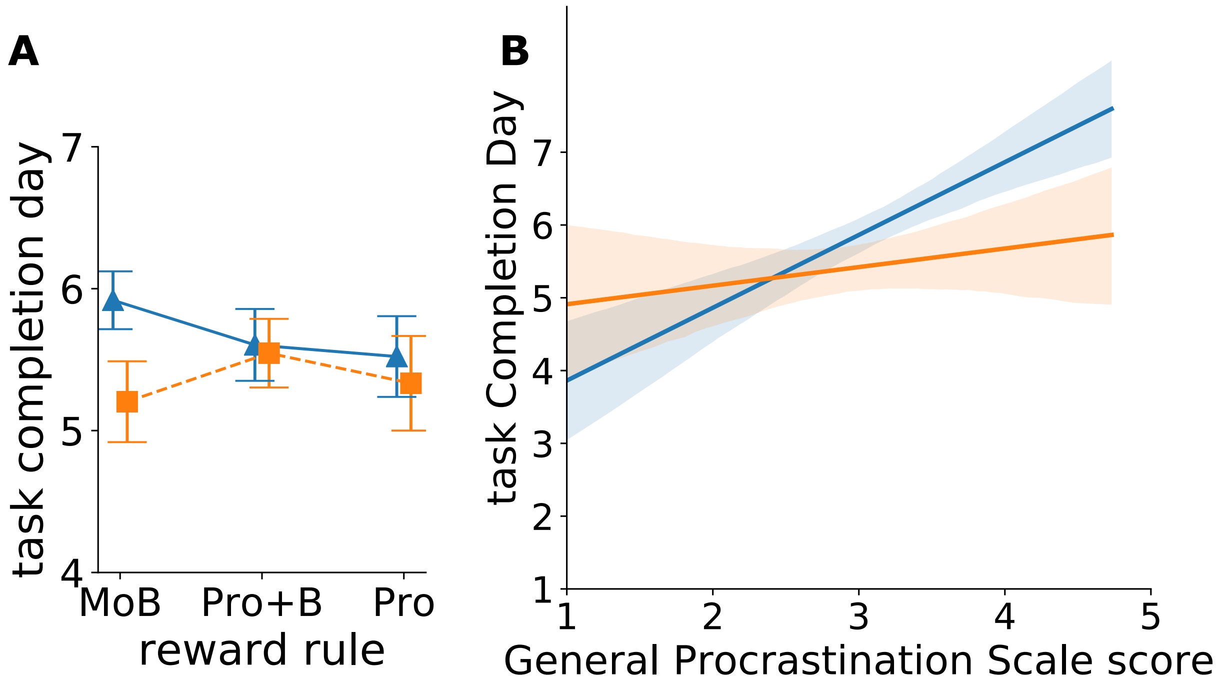

Figure 4.12: The effects of reward timing and reward rule on task completion day. (A) There is no main effect of reward timing on mean unit completion day ($F(1, 292) = 2.30, p = 0.13$) (B) There is an interaction between reward timing and the General Procrastination Scale score ($F(1, 281) = 4.44, p = 0.036$) indicating that offering immediate rewards help those people who generally procrastinate more to complete the task earlier.

negative result ($F(1, 298) = 2.64, p = 0.11$). If there is no effect of interaction on motivation, then the interaction is not mediated by motivation.

Task completion day. The results for Task Completion Day are qualitatively the same as that for Mean Unit Completion Day. Excluding General Procrastination Scale Score, there was no effect of reward timing ($F(1, 292) = 2.30, p = 0.13$) (Fig. 4.12A), reward rule ($F(2, 292) = 0.18, p = 0.84$) or interaction ($F(2, 292) = 0.90, p = 0.41$). We also conducted quasi-Poisson regression where we reshaped the outcome as $7 -$ Task Completion Day to capture the shape of the quasi-Poisson distribution, and the result was consistent with the linear regression: no main effect of reward timing ($\chi^2(1) = 2.34, p = 0.13$), no main effect of reward rule ($\chi^2(2) = 0.36, p = 0.84$) or interaction ($\chi^2(2) = 1.93, p = 0.38$).

After we include General Procrastination Score in the linear regression, there was a main effect of General Procrastination Scale Score ($F(1, 281) = 13.65, p < 0.001$) and interaction between General Procrastination Scale Score and reward timing ($F(1, 281) = 4.44, p = 0.036$) (Fig. 4.12B). This interaction suggested that immediate reward helps to complete the task earlier in people who have a greater tendency to procrastinate in general, regardless of the reward rule (no three-way interaction RT*RR*GPS: $F(2, 281) = 0.46, p = 0.63$). Simple slope analysis indicates that when the General Procrastination Scale Score is above 4.08, immediate reward rather than delayed reward significantly helps participants to complete the task earlier. There was no main effect of reward rule ($F(2, 281) = 0.09, p = 0.92$) or other interactions (RT*RR: $F(2, 281) = 0.64, p = 0.53$; RR*GPS: $F(2, 281) = 0.19, p = 0.83$).

We confirmed the above results in quasi-Poisson regression: main effect of General Procrastination Scale Score $\chi^2(1) = 13.28, p < 0.001$, and interaction between General Procrastination Scale Score and reward timing $\chi^2(1) = 6.94, p = 0.008$. Simple slope analysis indicates that when the General Procrastination Scale Score is above 4.45, immediate reward than delayed reward significantly helps participants to complete the task earlier.

In summary, offering immediate rewards helped people start the task earlier and helped those who generally procrastinate more complete the task and units of work earlier, both of which held true regardless of reward rules.

4.4 Discussion

We created a novel experimental paradigm, BORE, to study if offering immediate rewards upon completion rather than delayed rewards helps reduce procrastination and if the benefit is consistent across three reward rules that represent real-world scenarios. We found that offering immediate rewards helped people start the task earlier and helped those who generally procrastinate more complete the task and units of work earlier, both of which held true regardless of reward rules.

Regarding the effect of the reward rule on procrastination, we found that more people completed the task and persisted in working on the task in Make-or-Break condition compared to the Proportion condition, a trend that was also observed in the Proportion Plus Bonus condition. However, we did not find the effect of the reward rule on procrastination, as the interaction between reward timing and the reward rule did not yield significant results. This lack of significance might be due to an insufficient sample size, which typically necessitates a larger number to discern a substantial interaction effect. Future research, with the benefit of a larger sample size, may elucidate whether the reward rule indeed holds a considerable effect on procrastination.

Rewarding immediately upon task completion is one strategy to make future rewards more tangible for people. Zhang and Ma, 2019 (or this thesis Chapter 2: A normative account of temporal dynamics of procrastination) have suggested a theory encompassing various potentially effective interventions of immediate rewards for ongoing progress besides rewarding immediately upon task completion. Future research should delve deeper into exploring the impacts of these alternative immediate reward strategies on procrastination. For instance, investigating the effects of immediate rewards given after achieving specific milestones or even upon completion of each unit of work progress could be instrumental. These nuanced approaches may offer more insights into effectively mitigating procrastination.

BORE, the novel experimental paradigm we created, will pave the way to test these alternative immediate reward strategies on procrastination. To implement rewarding immediately after each milestone in BORE, for example, we can send monetary reward immediately after people completed every third of the task (i.e., 30 paragraphs). To implement rewarding immediately after completing each unit of work, we will deliver monetary reward at the completion of each paragraph. In practice, inspired by gamification in self-regulation problems in daily life (Milkman et al., 2014) and real-world examples of pseudo-reward (e.g. points) in online learning platforms, we can use visual aid of a virtual bank instructing the participants that the money they earn is saved in a virtual bank plus the possibility of cashing out anytime by having a cash-out button displayed in the header of the web page.

To conclude, we provided the first empirical evidence supporting the long-standing theory that offering immediate rewards can help reduce procrastination, and the novel experimental paradigm we designed will pave a way for testing the effects of other potential strategies on procrastination suggested in the literature.

5 The cognitive process underlying procrastination

There is so much **richness** in the time course of progress!

We see very systematic shapes in the time courses:

People working steadily, ramping up a bit, or ramping up a lot.

We need a computational model that generates a time course of progress

and accounts for this qualitative variation!

We have a candidate –a normative model– from chapter 2.

Does it fit the data well?

If not, what would be a better model to uncover the underlying cognitive process?

5.1 Introduction

What is the cognitive process underlying procrastination? Cognitive neuroscientists have extensively studied the cognitive and neural mechanisms of mental effort (Kool and Botvinick, 2018; Shenhav et al., 2017) in the field of cognitive control. The commonly used experimental paradigm asked people to choose repeatedly between performing a high-demand task for a larger amount of dollars and performing a low-demand task for a smaller amount. They found that exerting mental effort is costly. However, it is unclear how those studies can inform us about how people make mental efforts in their daily lives for long-term projects such as writing articles or building software that often extends over long periods of time, ranging from days to months. Making sustained efforts during those intervals is essential for people to achieve their goals and for a society to function efficiently.

Neuroscientists have used reinforcement learning to study how animals allocate time between work and leisure over time (Niyogi, Shizgal, et al., 2014; Niyogi, Breton, et al., 2014). Animals, however, tend to maximize instantaneous reward rate, in contrast to humans, who allocate effort over time towards a temporally distant goal.

There is also a growing number of efforts to account for and explain goal pursuit using computational process models (Lieder et al., 2019; Prystawski et al., 2022; Singhi et al., 2023). Still, these theoretical works focus on difficult decision-making problems rather than the simple temporal problems we face every day, like working to finish a project.

In reviewing the literature that aims to unveil the cognitive processes behind procrastination, a discernible gap is evident. On the one hand, several existing theories attempt to predict the time course of progress (Fischer, 1999; Fischer, 2001) besides our work (Zhang and Ma, 2019). However, these theories have not been successfully corroborated by empirical data, casting uncertainty on their ability to elucidate the cognitive mechanisms driving procrastination.

On the other hand, there is a wealth of data illustrating the intricacies and richness of the

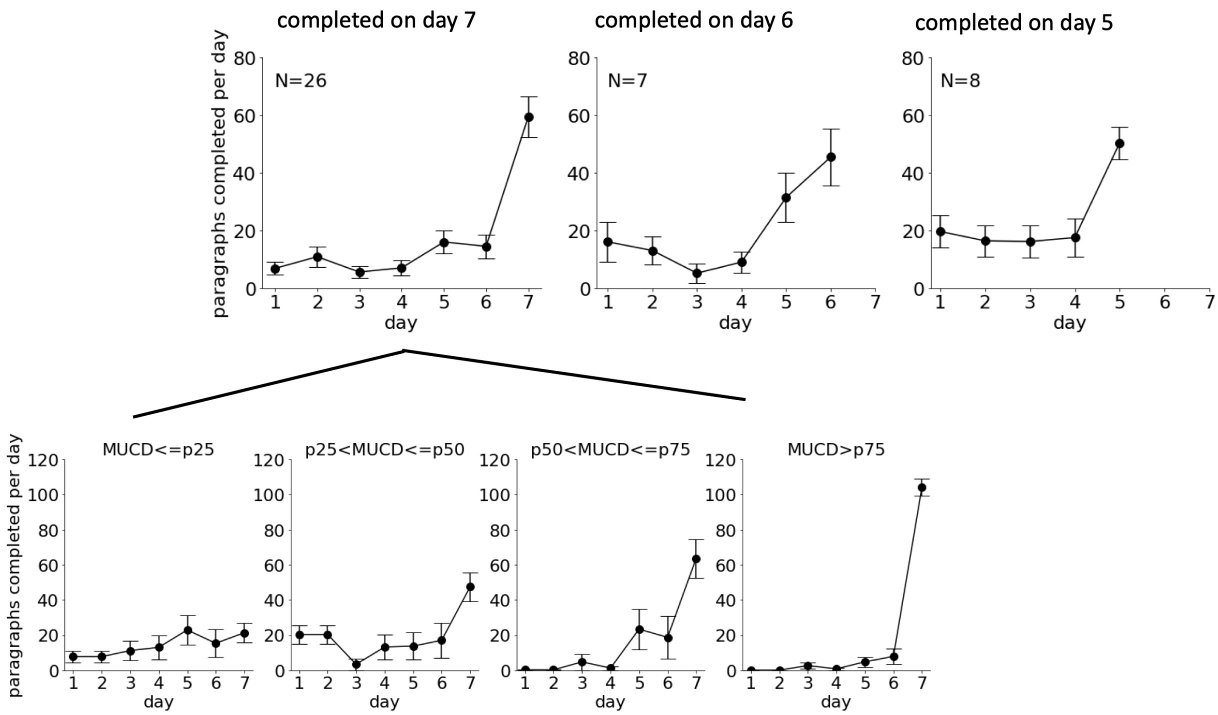

Figure 5.1: Systematic shapes in the time course of progress in Make-or-Break and delayed reward condition. Upper panel: (from the left to the right) the averaged time courses of progress across all the participants who completed the task on day 7, on day 6, and on day 5. Lower panel: the averaged time courses of progress across all the participants who completed the task on day 7 and it is sorted according to the four quantiles of the mean unit completion day from the left to the right.

time course of progress (Schouwenburg and Groenewoud, 2001; Steel et al., 2018; Konradt et al., 2021; Moon and Illingworth, 2005; Vangsness and Young, 2020) besides our work (Zhang and Ma, 2023a). However, there is a lack of models capable of encapsulating and accurately accounting for these characteristic time courses of progress. Noteworthy is an attempt that employed a cognitive model that adequately fit the task completion day (Raphaël and Mathias, 2022). However, this singular focus on a specific time point fails to capture the broader, dynamic time course of progress.

In this study, we aim to bridge this gap. We proposed two models and fit them into the data that we collected in the BORE task in Chapter 4. Before we jump right into the modeling part, we want to highlight the systematic shapes in the time courses of progress we collected in the BORE task. If we split the time courses of progress into different task completion days and average over the participants, we see these very characteristic time courses (Fig. 5.1 upper panel). People tended to work more on the last day. If we look closely at those who completed on the last day and sorted them according to the four quantiles of the mean unit completion day, we see people working steadily, ramping up a bit, or ramping up a lot (Fig. 5.1 lower panel). We need a computational process model that generates a time course of progress and accounts for this qualitative variation. The process model that fits the data well will uncover the underlying cognitive process of procrastination.

The first model we proposed is the rational model, which we have discussed intensively in Chapter 2. The second is the roll-out model, which we will elaborate on further below.

5.2 Models

What both models share is a sequential decision-making framework and a computational goal. They differ in how the value of making additional progress is computed. First, I will outline the shared sequential decision-making framework and the computational goals, then discuss the differences.

We assume that the time course of work progress results from a sequential decision-making process. Each day, people make a decision between working and having fun. If they choose to work, they make progress, and more work leads to more progress. This progress increases the likelihood of earning a reward at the end. However, working also involves the cost of mental effort, which is higher for more aversive tasks. If they choose to have fun instead, they make no progress and also pay no cost.

The computational goal each day is to maximize the value of additional progress while minimizing the associated costs. Moreover, as observed in Chapters 3 and 4, the value of making additional progress is temporarily discounted.

5.2.1 Rational model

In the rational model, the value of making additional progress is derived from the Bellman equation (Bellman, 1957). To fit the model to the time courses of progress, we add softmax noise to the Bellman equation

$$p_t(a|s) = \frac{e^{\beta Q_t(s,a)}}{\sum_{i=0}^{120} e^{\beta Q_t(s,i)}}, \tag{5.1}$$

where $p_t(a|s)$ is the probability of working on a number of paragraphs a on day t given the participant has already completed s paragraphs until day $t-1$, and β is the inverse temperature which captures the stochasticity of the choice data. In total, the rational model has four parameters: discount rate δ, maximum cost $C_{\max}$, exponent of the cost function k, and noise β. We used maximum-likelihood estimation to estimate the model parameters.

5.2.2 Roll-out model

Different from the rational model, the roll-out model computes the value of making additional progress by simulating future time courses. This approach is very similar to the roll-outs in the Monte Carlo tree search. To compute the value of nodes in the game tree, the algorithm performs multiple playouts, also known as roll-outs. In each roll-out, the game is played out to the very end by selecting moves at random. The final game result of each playout is then used to weigh the nodes in the game tree, ensuring that better nodes are more likely to be chosen in future playouts. This concept inspired the name of our cognitive model.

We assume that people will have an idea of how much work they will do on average in the next few days. It is because they may have a sense of how much time they are going to have for this task and/or how busy they are going to be in the next few days. The amount of work they will do on average in the next few days is modeled as a Poisson distribution with a mean number of events represented as n_{par}. From the roll-outs, people can predict their final work status, estimating the

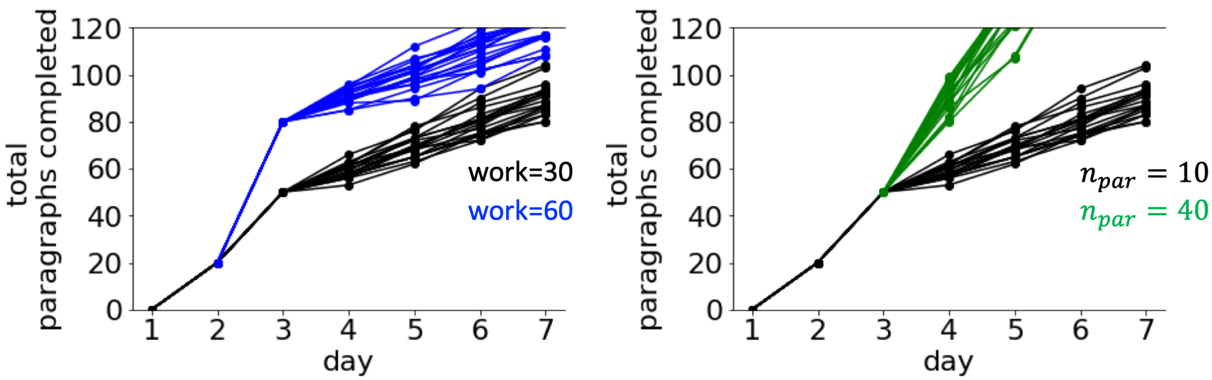

Figure 5.2: Roll-outs illustration. Left panel: simulated cumulative time courses of progress from day 3 to day 7, given 20 paragraphs completed by day 2. The blue curve represents 60 additional paragraphs completed on day 3, and the black curve represents 30. Right panel: simulated cumulative time courses of progress from day 3 to day 7 based on anticipated number of paragraphs completed per day from day 3 to day 7 (green: 40 paragraphs per day, black: 10 paragraphs per day.

likelihood of completing the task (and on which day) and the final proportion completed if the task is incomplete in the end. Based on these predictions and the specific experimental conditions, they can then calculate the expected reward.

In general, across all the experimental conditions, the decision on how much work to do today can affect the expected rewards. Choosing to do a small amount of work decreases the likelihood of receiving a reward, resulting in a lower expected reward (Fig. 5.2 left panel). Conversely, choosing to do a large amount of work increases the likelihood of earning the final reward, yielding a higher expected reward (Fig. 5.2 left panel). Furthermore, one's anticipation of their average workload in the upcoming days can also impact their decision. If a person anticipates a significant amount of work in the future, they can opt to do less work today (Fig. 5.2 right panel).

We formulate the value of making additional progress as the net utility $U_t(s, a) = ER_t - \text{cost}$, where $U_t(s, a)$ is the net utility of working a number of paragraphs on day t, given the participant has already completed s paragraphs by day $t-1$. The cost is given by cost $= C_{\max} a^{\lambda}$, which is the same as that in the rational model. Depending on the reward timing and the reward rule, ER_t is computed differently in each of the six experimental conditions as described below.

In the Make-or-Break and delayed reward condition,

$$ER_t = \sum_{t_{end}=t}^{T} p_{\text{complete},t_{end}} \cdot \frac{R_{\max}}{1 + k(8 - t)}, \tag{5.2}$$

where ER_t is the expected reward at day t, $p_{\text{complete},t_{end}}$ represents the percentage of completing the task at day t_{end}, calculated through simulating future time courses. Here, $R_{\max}$ is normalized to 1, representing the maximum reward (equivalent to \$36 in the reading task). The parameter k indicates the discount rate, with 0 signifying no discounting, and larger values of k indicating stronger temporal discounting. In the delayed reward condition, the reward is given one day after the deadline (the 8th day), regardless of the task completion day, leading to a consistent discount factor of $\frac{1}{1+k(8-t)}$,independent of task completion day t_{end}.

In the Make-or-Break and immediate reward condition,

$$ER_t = \sum_{t_{end}=t}^{T} p_{\text{complete},t_{end}} \cdot \frac{R_{\max}}{1 + k(t_{end} - t)}. \tag{5.3}$$

What differs in the immediate reward condition, compared to the delayed reward condition, is the timing of the reward. In the immediate reward condition, the reward is given immediately after task completion, making the discounting dependent on the task completion day, represented as $\frac{1}{1+k(t_{\text{end}}-t)}$.

In the Proportional Plus Bonus and delayed reward condition,

$$ER_t = \sum_{t_{end}=t}^{T} p_{\text{complete},t_{end}} \cdot \frac{R_{\max}}{1 + k(8 - t)} + \sum_{\text{finalprop}=0}^{119/120} p_{\text{incomplete},\text{finalprop}} \cdot \frac{\text{finalprop} \cdot R_{\max}/2}{1 + k(8 - t)}, \tag{5.4}$$

where $p_{\text{incomplete},\text{finalprop}}$ represents the probability of not completing the task, calculated through simulating future time courses. finalprop denotes the final proportion completed if the task is incomplete. What differs in the Proportional Plus Bonus condition, compared to Make-or-Break condition where participants receive no reward for incomplete tasks, is that participants receive a reward for incomplete tasks, but it is limited to half of the maximum reward. This reward for incomplete tasks contributes to the second term in the right-hand side of the formula. In addition, rewards for incomplete task is given one day after the deadline (on the 8th day), making the discounting factor consistently $\frac{1}{1+k(8-t)}$ for both immediate and delayed reward conditions.

In the Proportional Plus Bonus and immediate reward condition,

$$ER_t = \sum_{t_{end}=t}^{T} p_{\text{complete},t_{end}} \cdot \frac{R_{\max}}{1+k(t_{end}-t)} + \sum_{\text{finalprop}=0}^{119/120} p_{\text{incomplete,finalprop}} \cdot \frac{\text{finalprop} \cdot R_{\max}/2}{1+k(8-t)}. \quad (5.5)$$

In the Proportional and delayed reward condition,

$$ER_t = \sum_{t_{end}=t}^{T} p_{\text{complete},t_{end}} \cdot \frac{R_{\max}}{1+k(8-t)} + \sum_{\text{finalprop}=0}^{119/120} p_{\text{incomplete,finalprop}} \cdot \frac{\text{finalprop} \cdot R_{\max}}{1+k(8-t)}. \quad (5.6)$$

What differs in Proportional condition, compared to the Proportional Plus Bonus condition, is that the reward given for incomplete tasks is not limited to half of the maximum reward. Instead, it is proportional to the final proportion completed. This is why $R_{max}/2$ in Proportional Plus Bonus condition is replaced with R_{max} in Proportional condition.

In the Proportional and immediate reward condition,

$$ER_t = \sum_{t_{end}=t}^{T} p_{\text{complete},t_{end}} \cdot \frac{R_{\max}}{1+k(t_{end}-t)} + \sum_{\text{finalprop}=0}^{119/120} p_{\text{incomplete,finalprop}} \cdot \frac{\text{finalprop} \cdot R_{\max}}{1+k(8-t)}. \quad (5.7)$$

Similar to the rational model, we add softmax noise to the net utility to fit the data of the time course of progress.

$$p_t(a|s) = \frac{e^{U_t(s,a)}}{\sum_{i=0}^{120} e^{U_t(s,i)}}, \quad (5.8)$$

where $p_t(a|s)$ is the probability of work a number of paragraphs on day t given that s paragraphs are completed by day $t-1$, and β is the inverse temperature which captures the stochasticity of the choice data. In total, the roll-out model has five parameters: mean number of paragraphs in the future n_{par}, discount rate k, maximum cost $C_{\max}$, exponent of cost function γ, and noise β. We used maximum-likelihood estimation to estimate the model parameters.

5.3 Results

We first present the simulation results from the roll-out model, focusing on the variation of the mean number of events of the Poisson distribution, n_{par}, since this is a new parameter introduced in the roll-out model, in contrast to other parameters that are shared with the rational model. The simulations indicate that when a person anticipates a lower amount of work on average in the upcoming days, they tend to work more during the earlier days and reduce their workload in the later days (Fig. 5.3 lighter green curves), showing a decreasing trend in additional progress over time. In contrast, if a person anticipates a significant amount of work on average in the subsequent days, there appears to be a ramping up in the time course of progress (Fig. 5.3 darker green curves).

Note that these simulations differ from those of the rational model. The rational model only displays steady working or ramping up, never showing a decreasing trend in the amount of work over time. In contrast, the roll-out model can predict a more diverse range of the temporal patterns of the work progress, including the decreasing trend in the time courses.

Next, we present the results of model fitting. Both models were fitted to each individual's time course of progress. To display these results, we put all the experimental conditions together and grouped the time courses of progress based on whether the task was incomplete or complete. Within the "complete" group, we further grouped the time courses based on different task completion days. A well-fitting model should exhibit a good fit across all groups.

The rational model provided a poor fit to all the groups (Fig. 5.4), particularly failed to predict

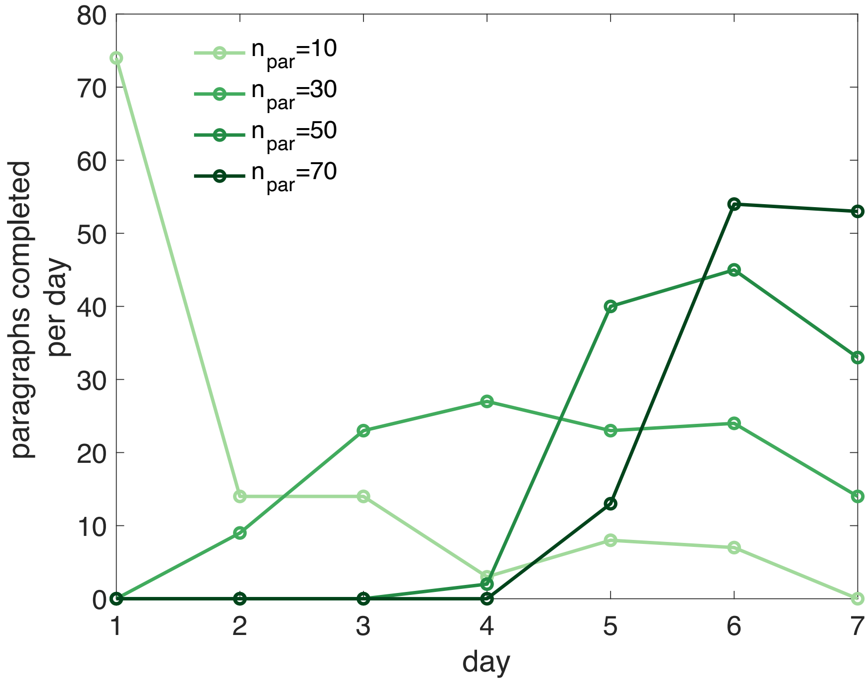

Figure 5.3: Roll-out model simulation results from varying the mean number of events of the Poisson distribution n_{par}.

the ramping up observed when tasks were completed earlier than the last day. In contrast, the roll-out model displayed a good fit with the data (Fig. 5.5). The former model comparison showed that the roll-out model outperformed the rational model by a sum of ΔAIC of 1461.5 with bootstrapped 95% CI: [1207.4, 1713.7] and a sum of ΔBIC of 1583.7 with bootstrapped 95% CI: [1713.7, 1836.0] (Fig. 5.6). It's worth noting that the BIC is more suitable for our data set than the AIC because BIC takes into account the number of observations, unlike the AIC. The number of observations varies, being seven for incomplete tasks, and is equal to the days taken to complete the task for completed tasks.

To see if the roll-out model fits well enough to the various characteristic shapes of the time courses, we grouped the time courses into four quantiles of the mean unit completion day within each group. The roll-out model captured well the systematic shapes of the time courses (Fig. 5.7 and Fig. 5.8).

Why did the rational model provide a poor fit, whereas the roll-out model provide a good fit?

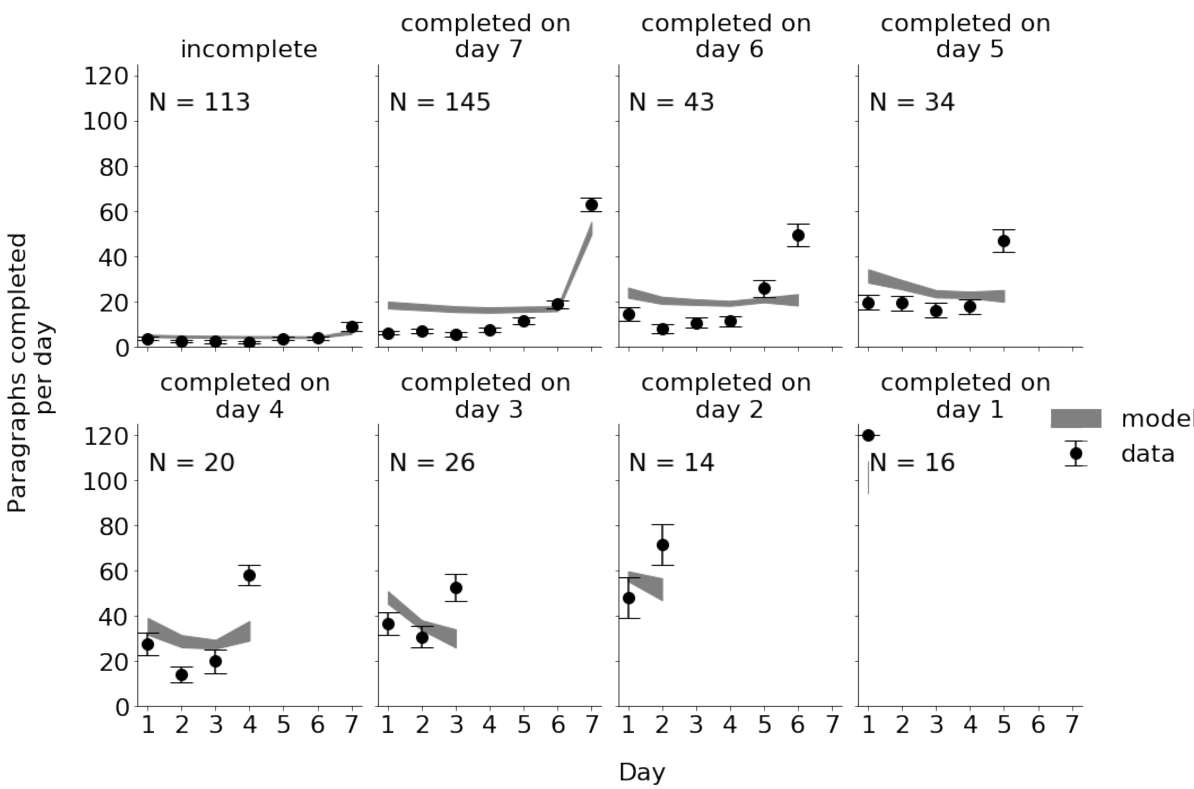

Figure 5.4: Rational model fitting results grouped into task incomplete, task completed on day 7, day 6, day 5, day 4, day 3, day 2 and day 1.

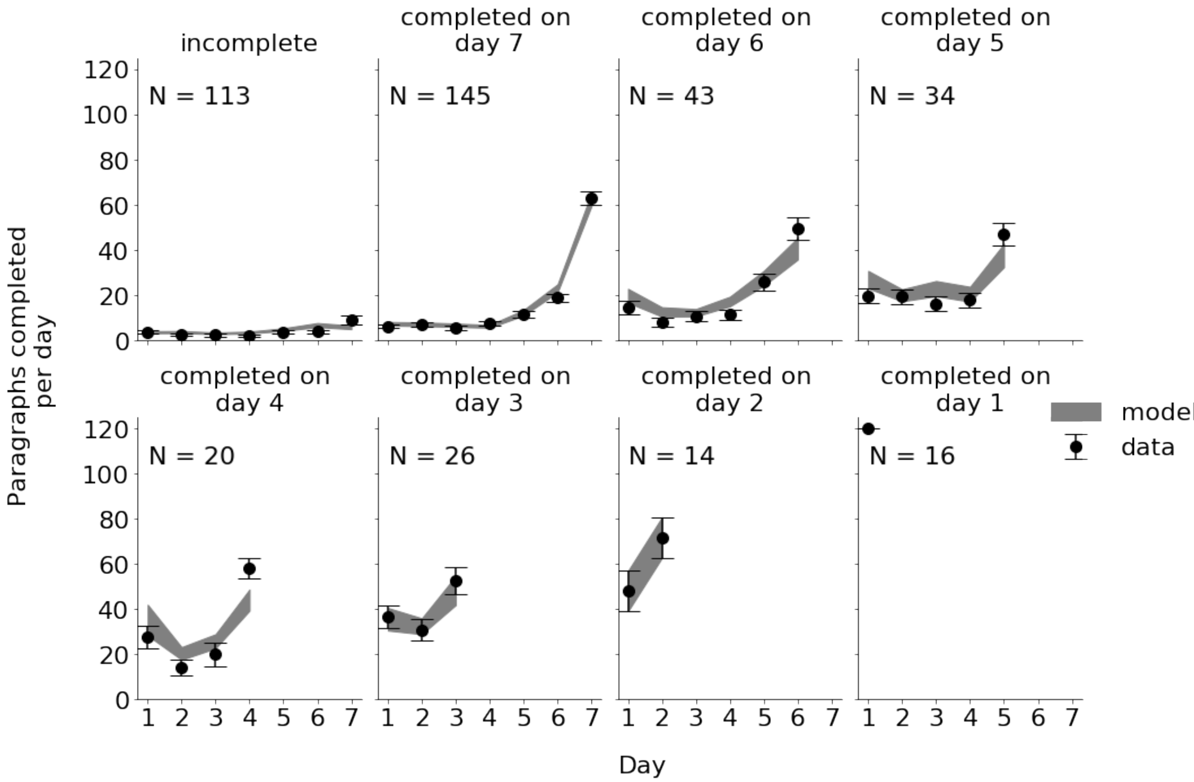

Figure 5.5: Roll-out model fitting results grouped into task incomplete, task completed on day 7, day 6, day 5, day 4, day 3, day 2 and day 1.

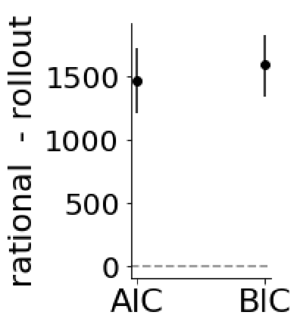

Figure 5.6: Model comparison between the rational model and the roll-out model.

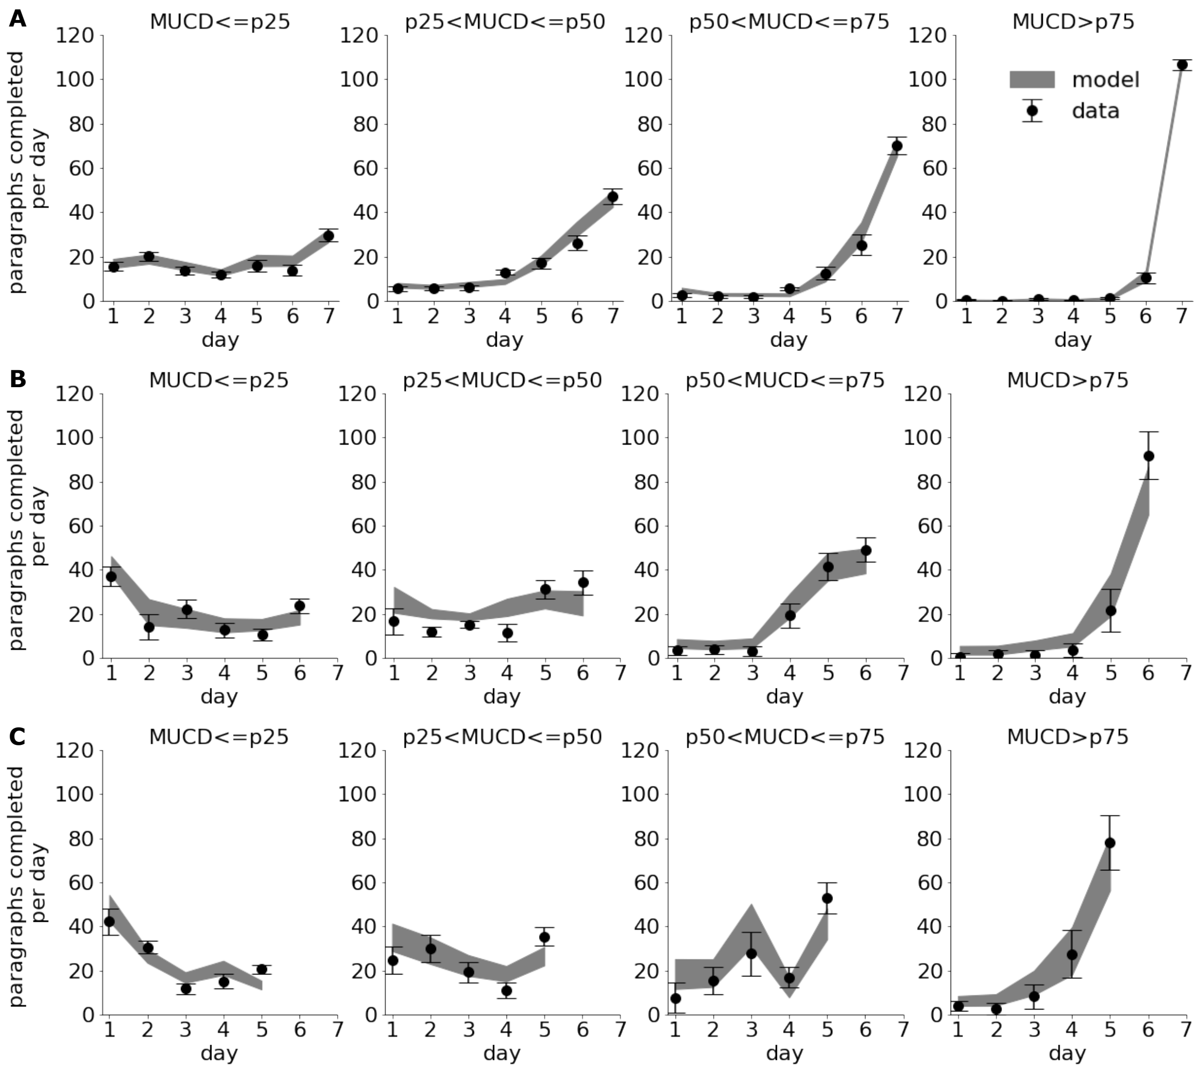

Figure 5.7: Roll-out model fitting results for the characteristic shapes of the time courses of progress that are grouped into four quantiles of the mean unit completion day. (A) The group where the task was completed on day 7. (B) The group where the task was completed on day 6. (C) The group where the task was completed on day 5.

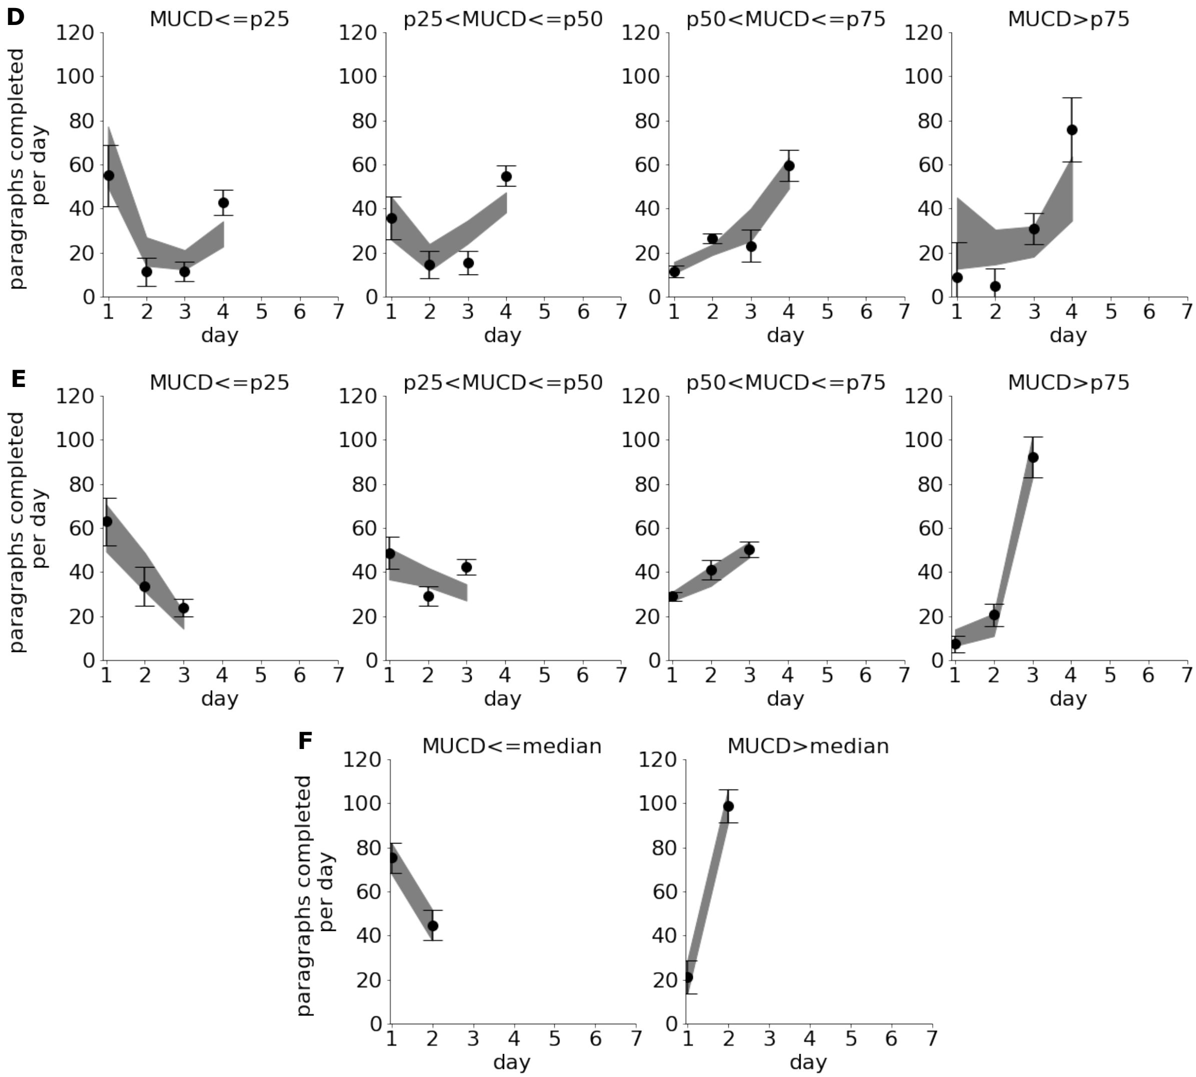

Figure 5.8: Roll-out model fitting results for the characteristic shapes of the time courses of progress that are grouped into four quantiles of the mean unit completion day. (D) The group where the task was completed on day 4. (E) The group where the task was completed on day 3. (F) The group where the task was completed on day 2. Since the number of participants is low, instead of grouping the participants into four quantiles of mean unit completion day, we grouped into two halves based on the median of the mean unit completion day.

We could have some insights from the simulation results of each model.

The rational model failed to predict the ramping up observed in groups where the task was completed earlier than the last day. The rational model, using the Bellman equation (Bellman, 1957), aims to maximize expected rewards while minimizing cumulative costs over seven days. To reduce the cumulative cost over seven days, the rational model predicts that people should take advantage of the given total days to reduce the cumulative cost. Even the weakest temporal discounters should work for seven days instead of completing the task before the last day. Therefore, the rational model fails to predict any time courses for early task completion, treating the ramping up observed on certain days, such as the sixth day, as noise rather than a meaningful pattern.

On the other hand, the roll-out model fits well the early task completion groups. The simulations suggest that by varying the parameter n_{par}, the roll-out model can predict a trend of decreasing progress over time, allowing it to capture the essence of early completion alongside the noise parameter. This nuanced approach enables the roll-out model to capture the observed ramping up in early task completion.

In summary, we found some evidence against people behaving rationally. We had some evidence for people simulating their future work progress, and they have an idea of how much work they will do on average in the next few days.

5.4 Discussion

Where do our computational models of the time course of progress stand in terms of contributions to the field of cognitive science?

We reviewed a broad range of literature related to our topic, including long-term goal pursuit, cognitive labor and leisure. For instance, let's consider the literature on cognitive labor. There is a clear gap in modeling work that accurately reflects real-world behavior and understands its cognitive processes. On the one hand, cognitive neuroscientists have extensively studied the cognitive mechanisms of mental effort in the field of cognitive control. However, their experimental paradigm was very far from the real world. The commonly used experimental paradigm asked people to come to the lab and repeatedly choose among hundreds of trials between performing a high-demand task for a greater reward and a low-demand task for a lesser reward. They found that exerting mental effort is costly. However, it is unclear how those studies can inform us about how people make mental efforts in their daily lives for long-term projects such as writing articles or building software that often extends over long periods of time, ranging from days to months. On the other hand, some studies examine how people allocate time between cognitive labor and leisure. They ask people to report the mean times spent in work or leisure without examining the detailed temporal dynamics of work. The model was also on a macroscopic timescale that limits understanding of the underlying processes.

Our computational model showed a successful attempt to bridge this gap. It helps us understand the cognitive process of naturalistic behavior in a longitudinal study.

The roll-out model fits better than the rational model. However, there are deviations between the model fitting and the data in the roll-out model. For example, the roll-out model failed to capture the ramping up in incomplete tasks in Make-or-Break conditions (Fig. 5.9). The data shows that people worked more on the last day despite the Make-or-Break reward rule, where no rewards are given for incomplete tasks. However, the roll-out model predicts that people should not work much if their progress is far away from completing the task by the last day.

A plausible explanation for this deviation lies in the decision-making timing within the roll-out model, which is structured on a daily basis. In contrast, real-world decision-making might occur at various times throughout the day. Given that the task needs approximately three hours to complete, people may evaluate their possibility of completing the task up to three hours before midnight of the last day. One potential adjustment to the roll-out model involves sequential decisions every three hours on the last day, as opposed to maintaining a daily basis across all

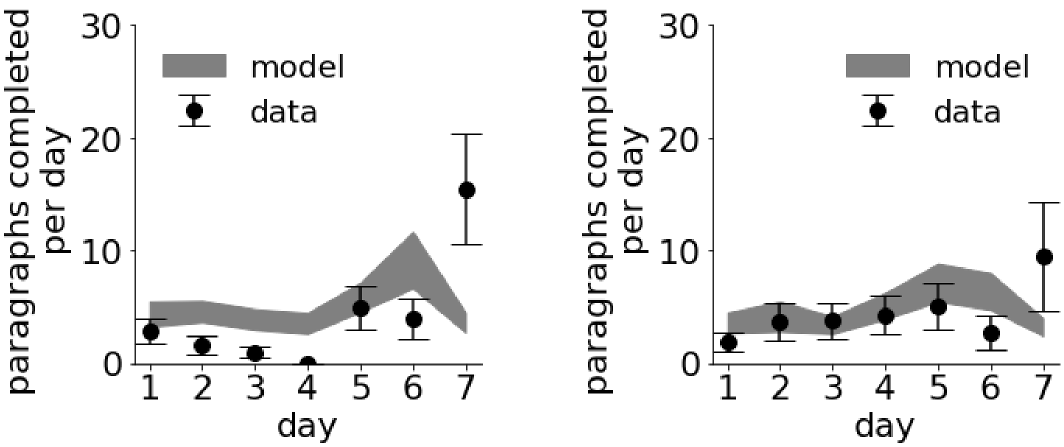

Figure 5.9: Deviations between the model fitting and the data in the roll-out model. Left panel: time courses of progress in the group of tasks incomplete in the Make-or-Break and delayed reward condition. Right panel: time courses of progress in the group of tasks completed in Make-or-Break and immediate reward condition.

days.

Nevertheless, we found some evidence against people behaving rationally. We had some evidence for people simulating their future work progress, and they have an idea of how much work they will do on average in the next few days.

Future work should look into the necessity of the cognitive components of the roll-out model by conducting model lesion studies. Additionally, while we evaluated these two models, there may exist alternative models worth considering, such as non-utility models. These alternative models could incorporate components like planning and plan adjustment based on discrepancies between planned and actual progress, as indicated by participants' self-reported reflections on how they worked on the task in post-surveys.

6 Conclusion

6.1 Summary of dissertation

We studied the dynamic nature of procrastination. We characterized procrastination through the time course of work progress and uncovered its underlying cognitive mechanisms.

In Chapter 2, we proposed a normative account of the time course of progress. We assumed that the time course of progress arises from a sequential decision-making process. On each day, people decide whether to work now (and, if so, how much) or later. If they decide to work now, they pay the cost of investing mental effort immediately but also make progress, and more work leads to more progress. If they decide not to work, they make no progress and also pay no effort cost. The optimal amount of work on each day is derived from the Bellman equation (Bellman, 1957), which assumes that a person's goal on each day is to maximize the discounted value gained by making progress while minimizing the immediate effort cost. This normative model predicted three patterns of procrastination: a delay in the beginning and then ramping up, working at the last minute, and not working at all, and several correlates of procrastination, including perfectionism, the shape of cost function, temporal discounting, and the total given time. Last, this theory reproduced the effect of interventions on reducing procrastination or improving performance, replicating what was observed in the empirical literature.

In Chapter 3, we examined if discount rates were associated with behavioral levels of procrastination in a real-world task. We used a long-term real-world task to measure procrastination;

critically, we measured the entire time course of work progress instead of only completion time, allowing us to compute a fine-grained metric of procrastination. We found a positive correlation between individuals' degree of future reward discounting and their level of procrastination, suggesting that temporal discounting is a cognitive mechanism underlying procrastination.

In Chapter 4, we tested if offering immediate rewards helps reduce procrastination regardless of reward rules. We created a novel experimental paradigm named BORE (Boring Online Reading Experiment) that mimics real-world procrastination while still allowing for manipulation. We utilized a between-subject design, crossing two levels of reward timing (either delayed or immediate upon task completion) and three levels of reward rules (Make-or-Break, Proportional Plus Bonus, and Proportional). We found that Make-or-Break conditions led to a higher completion rate and more persistence in working on the task. Offering immediate rewards motivated people to start the task earlier and complete the task earlier, regardless of reward rules. Moreover, behaviorally, offering immediate rewards helped people start the task earlier and helped those who generally procrastinate more complete the task and units of work earlier, both of which held true regardless of reward rules.

In Chapter 5, we uncovered the cognitive process underlying procrastination. We proposed two models and fit them to the data that we collected in Chapter 4. The first model was the rational model, which we have discussed intensively in Chapter 2. The second was the roll-out model, inspired by the roll-outs in the Monte Carlo tree search. We found that the rational model provided a poor fit to the data, while the roll-out model fit the data quite well. Therefore, we found some evidence against people behaving rationally. We had some evidence for people simulating their future work progress, and they had an idea of how much work they will do on average in the next few days.

Taken together, we understand better the dynamic nature of procrastination and its underlying cognitive mechanisms. In addition, we offer implications for reducing procrastination. This thesis shows a successful attempt at applying cognitive science to the real world.

6.2 Future work

The BORE paradigm sets a solid foundation for testing various interventions on procrastination and examining the underlying cognitive processes of each intervention. For future work, we propose four additional experiments to assess the effects of different cognitive components on procrastination:

1. Testing the highest immediate reward level by delivering rewards immediately upon the completion of each task paragraph.

2. Setting subgoals with or without immediate rewards, setting subgoals with or without time limits, setting subgoals with varying patterns of inter-subgoal intervals, and whether or not they are self-imposed.

Furthermore, in all experiments, we will perform the same analyses mentioned in Chapters 4 and 5 regarding the measures and the computational models.

6.2.1 Experiment 1: Giving immediate reward upon unit completion

The association between temporal discounting and procrastination we observed in Chapter 3 suggested that if we brought a future reward temporally closer, then that would help reduce procrastination. One way to bring immediate reward is to deliver upon task completion, which we have studied in Chapter 4. Under this immediate rewards intervention, the progress the participants make before task completion is not immediately rewarded. Our theory proposed in Chapter 2 suggested alternative immediate reward interventions that reward the progress people make before task completion. Moreover, gamification in self-regulation problems in daily life (Milkman et al., 2014) and real-world examples of pseudo-reward (e.g., points) in online learning platforms showed that rewarding progress increased people's persistence in activities. Altogether, these

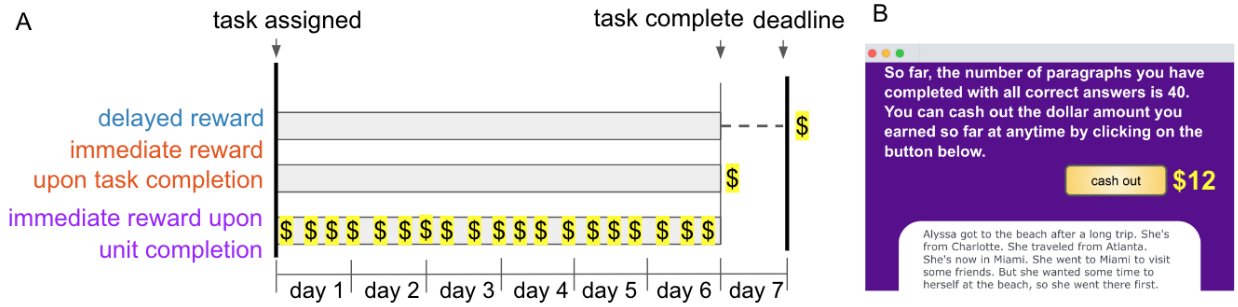

Figure 6.1: Experimental design of reward timing with varying levels of immediacy. (A) Control condition is the delayed reward and the first experimental condition is offering immediate rewards upon task completion and the second experimental condition is offering immediate reward upon each work unit completion. (B) Example display for immediate reward upon unit completion, showing a virtual bank and a cash-out button.

suggested the potential of rewarding progress in reducing procrastination. Future work should examine the effects of rewarding progress on procrastination.

In **Experiment 1**, we propose an experimental paradigm where we will add one manipulation that increases the immediacy of reward to an extreme by immediately delivering rewards at the completion of each paragraph (Fig. 6.1A). Together with the delayed reward condition and immediate reward upon task completion, we will have three levels of immediacy in reward delivery. Also, we will introduce two reward rules: Proportional Plus Bonus and Proportional. We will leave out the Make-or-break condition, in which progress is not rewarded. This leads to a 3 by 2 design. We will study how it affects the temporal dynamics of work, especially when compared with immediate reward upon task completion, and if it will help subjects start the work earlier, complete the task and units of work earlier, and get more work done.

To implement the immediate reward at the completion of each paragraph, we will instruct the participants that the money they earn is saved in a virtual bank. They will be able to cash out at any time. The cash-out button will always be visible in the header of the webpage (Fig. 6.1B).

6.2.2 Setting subgoals

Setting subgoals is useful to reduce procrastination in daily life, for example, by setting a daily test in a self-paced course (Tuckman, 1998; Wesp, 1986) or by splitting the writing of an essay into three sections (Loebenstein, 1996). For future work, we first ask what ingredients of subgoals help people complete the task earlier. The ingredients we consider (Exp. 2) are 1) giving a reward immediately after the completion of a subgoal; 2) setting a time limit to complete a subgoal, which has previously proven effective in goal pursuit (Ariely and Wertenbroch, 2002).

Second, we ask which pattern of inter-subgoal intervals is most effective: equal, increasing, or decreasing intervals (Exp. 3). Lastly, we will ask if people benefit from setting subgoals by themselves (Exp. 4). In all three experiments, we will use Proportional Plus Bonus as the reward rule because it leads to a higher completion rate and meets the necessary requirement of rewarding along the way.

6.2.2.1 Experiment 2: Varying deadline and immediate reward associated with subgoals

We will have five conditions: one control condition without any subgoals and four experimental conditions with assigned subgoals (Fig. 6.2A). For those four experimental conditions, we will equally divide the task into four subgoals. We then manipulate two ingredients of each assigned subgoal: with or without an immediate reward following the completion of a subgoal, and with or without a deadline on each subgoal.

We will define four subgoals for subjects, each consisting of 30 paragraphs. Before they start the task, they will be presented with the first subgoal using an illustration of a mountain climber (Fig. 6.2B). To help them focus on the current subgoal, the number of paragraphs they have completed and the number of paragraphs left to achieve the first subgoal will be presented on the header of the webpage. After they achieve the first subgoal, they will be presented with the

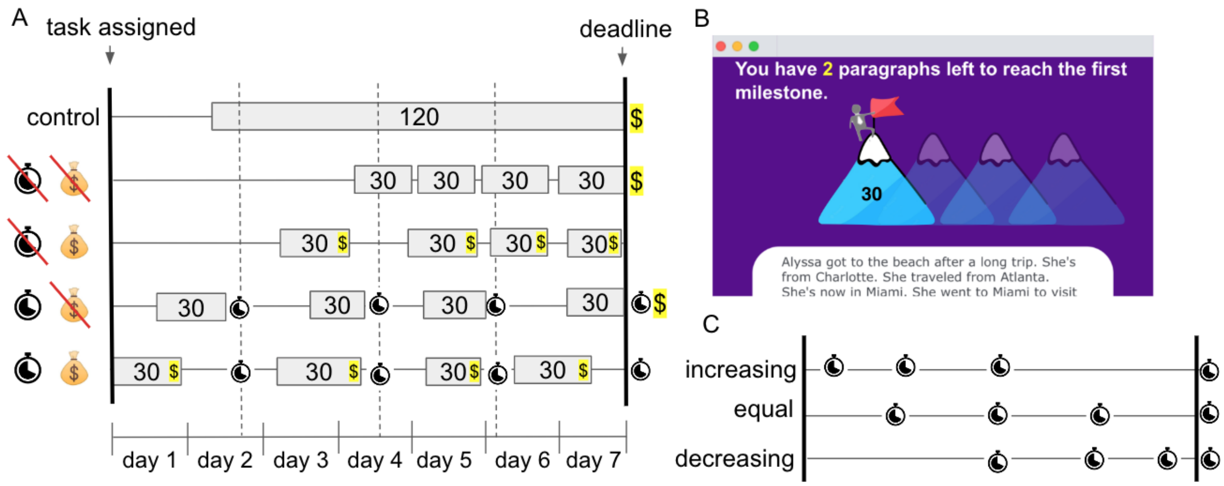

Figure 6.2: Experimental design of varying deadlines and immediate rewards associated with subgoals. (A) The experiment consists of five conditions. In the four treatment conditions, a total of four subgoals, each with 30 paragraphs, are untimed or timed and come with or without immediate reward. Each grey box shows when a subject starts working towards a subgoal and when a subject receives a subgoal, and the width of the box represents the duration spent on each subgoal. (B) Example display of the instruction: a climber facing four mountains. (c) Treatment conditions in the experiment differ in the nature of the inter-subgoal intervals.

second subgoal, etc.

The deadline for each subgoal is proportional to the deadline of the task; it takes up to 42 hours (168 hours divided by four subgoals) to achieve each subgoal. If the deadline passes, but the subgoal is not achieved, the payoff for the amount of work within that subgoal will be proportional to the work amount but halved (the Proportional Plus Bonus reward rule).

By comparing each experimental condition with the control condition, we will test which ingredients are necessary for subgoals to help subjects start the task earlier, complete the task and units of work earlier, and get more work done. We will also test whether combining immediate reward with a deadline is more effective than either in helping participants reduce procrastination and get more work done.

6.2.2.2 Experiment 3: Varying subgoal schedule

We will test which pattern of inter-subgoal intervals is most effective in helping people reduce procrastination. We will have four conditions: one control condition without subgoals and three conditions with different subgoal schedules: equal, increasing, and decreasing intervals (Fig. 6.2C). We will apply the most effective condition that we will find in Exp. 2 to determine whether the reward will be immediate and whether a deadline will be set. We will test whether participants assigned to increasing intervals complete the task earlier than those in an equal-interval condition and whether participants in the descending-interval condition complete the task later than those in the equal-interval condition. We will also compare the least effective subgoal schedule with the control condition to test if the least effective subgoal schedule is even worse than not having subgoals at all.

6.2.2.3 Experiment 4: Self-imposed subgoals

In this experiment, we will let subjects set their own subgoals with the most effective subgoal ingredients determined from Exp. 3. Subjects will be randomly assigned to one of two conditions. In Condition 1, participants set their subgoal at the beginning of the task, not allowing them to change it along the way (fixed subgoals in advance). In Condition 2, participants set their subgoal along the way (adjustable subgoals along the way). Specifically, the first time they access the task, they will be asked to set the first subgoal, and after they achieve their first subgoal, they will be asked to set the second subgoal, etc.

Besides these two experimental manipulations of self-imposed subgoals, we will also include the most effective experimenter-set subgoals from Exp. 3 and a control condition without subgoals, which leads to a total of four conditions. We will first test whether self-imposed subgoals, both fixed in advance and adjustable along the way, are more effective than experimenter-set subgoals. Second, we will test if the least effective self-imposed subgoal condition is even worse than the control condition without a subgoal. Third, we will test whether a self-imposed subgoal along the way is more effective than one that is fixed in advance. Fourth, we will examine how close the self-imposed subgoal, both fixed in advance and adjustable along the way, is to the most effective subgoal schedule we will find in Exp. 3. Last, we will examine how individual correlates of the cognitive process affect how people set their subgoal schedule in self-imposed subgoal conditions.

6.3 Relation to broader literature

Researchers name the subject—time course of progress—differently. The subject of our study is the time course of progress. One aspect I find quite interesting is that the existing literature refers to this subject by different names. Steel et al. refer to it as 'pacing style' (Steel et al., 2018), and they cite Roe, who calls for more attention to this subject under the term 'temporal footprint of work' (Roe, 2014). Moon and Illingworth refer to it as temporal changes in procrastination (Moon and Illingworth, 2005). Vangsness and Young refer to it as the (dynamic) process of task completion (Vangsness and Young, 2020). Konradt et al. refer to it as the allocation of effort over time (Konradt et al., 2021). The variety of terms used to describe this subject indicates that it is a relatively unexplored area of research, leading to inconsistency in how researchers refer to the subject. This also adds to the difficulty of conducting a literature review, as we might miss some papers that study the same topic.

Goal gradient effect. Clark Hull, in 1932, proposed the Goal Gradient Effect, which states that as people get closer to a reward, they speed up their behavior to reach their goal faster (Hull, 1932). We also observe this goal gradient effect in the BORE. We calculated the additional days participants took to complete another 20 paragraphs, from the first paragraph to the last 120th paragraph. If they completed another 20 paragraphs on the same day, the additional days were 0. We found that, on average, it took participants fewer days to complete another 20 paragraphs closer to completing the task. In specific, it took them another 0.37 days to complete the first 20 paragraphs, another 0.34 days to complete the second 20 paragraphs, another 0.22 days for the third 20 paragraphs, another 0.19 days for the fourth 20 paragraphs, another 0.15 days for the fifth 20 paragraphs, and another 0.06 days for the last 20 paragraphs. If we varied the interval of the number of paragraphs (e.g., another ten paragraphs), we got the same result.

Stuck in the middle effect. Regarding the shape of the time course of progress, researchers have previously studied the time course of progress with three data points: the amount of work in the beginning, in the middle, and in the end (Bonezzi et al., 2011; Koo and Fishbach, 2014). They claim a stuck-in-the-middle effect where less work is found compared to the beginning and the end. We observed this stuck-in-the-middle effect in a few participants who completed the task on the fourth day (Fig. 5.8D). However, most participants ramped up in their time courses of progress, and more importantly, there was a great diversity in the shapes of the time courses of progress observed in the BORE.

Deliberate delay. Chu and Choi defined a new construct called active procrastination (Chu and Choi, 2005), which is the deliberate deferral of tasks to the last minute, resulting in positive outcomes despite the delay. This deliberate delay is different from passive procrastination, which is often fueled by fear of failure and anxiety. Later on, Chowdhury and Pychyl categorized this deliberate procrastination into two granular concepts (Chowdhury and Pychyl, 2018). One is called arousal delays, which occur when a person decides they are more motivated to do something at the last minute. This could be due to the fact that people are more focused or faster at working. The other is called purposeful delay, which indicates that people are framing the question and thinking about an issue or creative work before getting down to the act of writing or producing something.

Future work should test whether people have arousal delays using our experimental paradigm, BORE. Specifically, we can ask participants about their intended day of completing the task before giving it and see if some participants' intended day of completing it is close to the last day.

We will ask these participants the reason for their later intended completion day. We will see whether they are more focused (less distracted behavior on webpages), and we will ask them to report their motivation level to see if they are more motivated to work at a later time close to the deadline.

Macroscopic approaches to cognitive labor and leisure. Kool and Botvinick have studied people's decisions about how much time they spend on cognitive labor versus leisure (Kool and Botvinick, 2014). This partial allocation has characterized behavior on a macroscopic timescale, reporting and studying the mean times spent at work or leisure. Yet, this study focuses on a coarse split between work and leisure instead of examining the detailed temporal dynamics of work. Our thesis work studied the microscopic decision process between cognitive labor and leisure, which can be viewed as an extension of their work in a temporal manner.

On the pursuit of multiple goals. Researchers in the fields of applied psychology and organizational psychology have done a lot of work investigating how people allocate time while facing multiple tasks with different deadlines and factors that influence the dynamic pursuit of multiple goals over time (Schmidt and DeShon, 2007; Schmidt et al., 2009; Schmidt and Dolis, 2009; Neal et al., 2017; Ballard, Farrell, et al., 2018; Ballard, Vancouver, et al., 2018; Ballard, Yeo, Neal, et al., 2016; Ballard, Yeo, Loft, et al., 2016). However, the tasks used in these studies are not longitudinal and do not closely resemble real-life scenarios. We propose a variant of the BORE paradigm to address this gap. This variation involves multiple reading tasks, each with its own deadline and varying levels of difficulty, to more accurately simulate the real-life pursuit of multiple goals.

Bibliography

Ackerman, D. S., & Gross, B. L. (2005). My instructor made me do it: Task characteristics of procrastination. *Journal of Marketing Education*, *27*(1), 5–13.

Akerlof, G. A. (1991). Procrastination and obedience. *Am. Econ. Rev.*, *81*(2), 1–19.

Alfi, V., Parisi, G., & Pietronero, L. (2007). Conference registration: How people react to a deadline. *Nat. Phys.*, *3*(11), 746–746.

Altmann, E. M., & Gray, W. D. (2002). Forgetting to remember: The functional relationship of decay and interference. *Psychol. Sci.*, *13*(1), 27–33.

Analytis, P. P., Wu, C. M., & Gelastopoulos, A. (2019). Make-or-break: Chasing risky goals or settling for safe rewards? *Cogn. Sci.*, *43*(7), e12743.

Ariely, D., & Wertenbroch, K. (2002). Procrastination, deadlines, and performance: Self-Control by precommitment.

Ballard, T., Farrell, S., & Neal, A. (2018). Quantifying the psychological value of goal achievement. *Psychon. Bull. Rev.*, *25*(3), 1184–1192.

Ballard, T., Vancouver, J. B., & Neal, A. (2018). On the pursuit of multiple goals with different deadlines. *J. Appl. Psychol.*, *103*(11), 1242–1264.

Ballard, T., Yeo, G., Loft, S., Vancouver, J. B., & Neal, A. (2016). An integrative formal model of motivation and decision making: The MGPM*. *J. Appl. Psychol.*, *101*(9), 1240–1265.

Ballard, T., Yeo, G., Neal, A., & Farrell, S. (2016). Departures from optimality when pursuing multiple approach or avoidance goals. *J. Appl. Psychol.*, *101*(7), 1056–1066.

Batistuzzo, M. C., Sheshachala, K., Alschuler, D. M., Hezel, D. M., Lewis-Fernández, R., de Joode, N. T., Vriend, C., Lempert, K. M., Narayan, M., Marincowitz, C., Lochner, C., Stein, D. J., Narayanaswamy, J. C., van den Heuvel, O. A., Simpson, H. B., & Wall, M. (2022). Cross-national harmonization of neurocognitive assessment across five sites in a global study. *Neuropsychology*.

Beck, B., Koons, S., & Milgrim, D. (2000). Correlates and consequences of behavioral procrastination: The effects of academic procrastination, Self-Consciousness, Self-Esteem and Self-Handicapping. *J. Soc. Behav. Pers.*, *15*(5), 3.

Bellman, R. (1957). A markovian decision process. *Journal of Mathematics and Mechanics*, *6*(5), 679–684.

Beswick, G., Rothblum, E. D., & Mann, L. (1988). Psychological antecedents of student procrastination. *Aust. Psychol.*, *23*(2), 207–217.

Blais, A.-R., & Weber, E. U. (2006). *A Domain-Specific Risk-Taking (DOSPERT) scale for adult populations.*

Blunt, A. K., & Pychyl, T. A. (2000). Task aversiveness and procrastination: A multi-dimensional approach to task aversiveness across stages of personal projects. *Pers. Individ. Dif.*, *28*(1), 153–167.

Bonezzi, A., Brendl, C. M., & De Angelis, M. (2011). Stuck in the middle: The psychophysics of goal pursuit. *Psychol. Sci.*, *22*(5), 607–612.

Buchanan, J., Summerville, A., Lehmann, J., & Reb, J. (2016). The regret elements scale: Distinguishing the affective and cognitive components of regret. *Judgm. Decis. Mak.*, *11*(3), 275–286.

Buehler, R., & Griffin, D. (2003). Planning, personality, and prediction: The role of future focus in optimistic time predictions.

Bulley, A., Lempert, K. M., Conwell, C., Irish, M., & Schacter, D. L. (2022). Intertemporal choice reflects value comparison rather than self-control: Insights from confidence judgements. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, *377*(1866), 20210338.

Burger, N., Charness, G., & Lynham, J. (2011). Field and online experiments on self-control. *J. Econ. Behav. Organ.*, *77*(3), 393–404.

Burka, J., & Yuen, L. M. (2007). *Procrastination: Why you do it, what to do about it now*. Hachette UK.

C., C. B. J., Wendelien, v. E., Rutte, C. G., & Roe, R. A. (2007). A review of the time management literature. *Personnel Review*, *36*(2), 255–276.

Cerezo, R., Esteban, M., Sánchez-Santillán, M., & Núñez, J. C. (2017). Procrastinating behavior in Computer-Based learning environments to predict performance: A case study in moodle. *Front. Psychol.*, *8*, 1403.

Chowdhury, S. F., & Pychyl, T. A. (2018). A critique of the construct validity of active procrastination. *Pers. Individ. Dif.*, *120*, 7–12.

Chu, A. H. C., & Choi, J. N. (2005). Rethinking procrastination: Positive effects of “active” procrastination behavior on attitudes and performance. *J. Soc. Psychol.*, *145*(3), 245–264.

Constantin, K., English, M. M., & Mazmanian, D. (2017). Anxiety, depression, and procrastination among students: Rumination plays a larger mediating role than worry. *J. Ration. Emot. Cogn. Behav. Ther.*

De Paola, M., & Scoppa, V. (2015). Procrastination, academic success and the effectiveness of a remedial program. *J. Econ. Behav. Organ.*, *115*, 217–236.

Dewitte, S., & Schouwenburg, H. C. (2002). Procrastination, temptations, and incentives: The struggle between the present and the future in procrastinators and the punctual. *Eur. J. Pers.*, *16*(6), 469–489.

Diver, P., & Martinez, I. (2015). MOOCs as a massive research laboratory: Opportunities and challenges. *Distance Education*, *36*(1), 5–25.

Eckert, M., Ebert, D. D., Lehr, D., Sieland, B., & Berking, M. (2016). Overcome procrastination: Enhancing emotion regulation skills reduce procrastination. *Learn. Individ. Differ.*, *52*, 10–18.

Elvers, G. C., Polzella, D. J., & Graetz, K. (2003). Procrastination in online courses: Performance and attitudinal differences. *Teach. Psychol.*, *30*(2), 159–162.

Enns, M. W., Cox, B. J., Sareen, J., & Freeman, P. (2001). Adaptive and maladaptive perfectionism in medical students: A longitudinal investigation. *Med. Educ.*, *35*(11), 1034–1042.

Farnham, B. R. (2021). *Roosevelt and the munich crisis: A study of political Decision-Making*. Princeton University Press.

Fedus, W., Gelada, C., Bengio, Y., Bellemare, M. G., & Larochelle, H. (2019). Hyperbolic discounting and learning over multiple horizons.

Fee, R. L., & Tangney, J. P. (2000). Procrastination: A means of avoiding shame or guilt? *J. Soc. Behav. Pers.*, *15*(5; SPI), 167–184.

Ferrari, J. R. (1993). Procrastination and impulsiveness: Two sides of a coin? In W. G. McCown (Ed.), *The impulsive client: Theory, research, and treatment , (pp* (pp. 265–276). American Psychological Association, ix.

Ferrari, J. R. (2001). Procrastination as self-regulation failure of performance: Effects of cognitive load, self-awareness, and time limits on 'working best under pressure'. *Eur. J. Pers.*, *15*(5), 391–406.

Ferrari, J. R., & Scher, S. J. (2000). Toward an understanding of academic and nonacademic tasks procrastinated by students: The use of daily logs. *Psychology in the Schools*, *37*(4), 359–366.

Ferrari, J. R., & Tice, D. M. (2000). Procrastination as a Self-Handicap for men and women: A Task-Avoidance strategy in a laboratory setting. *J. Res. Pers.*, *34*(1), 73–83.

Fiore, N. (2007). *The now habit: A strategic program for overcoming procrastination and enjoying guilt-free play*. Penguin.

Fischer, C. (1999). Mread this paper even later: Procrastination with time% inconsistent preferences, nresources for the future.

Fischer, C. (2001). Read this paper later: Procrastination with time-consistent preferences. *J. Econ. Behav. Organ.*, *46*(3), 249–269.

Flandrin, P. (2010). AN EMPIRICAL MODEL FOR ELECTRONIC SUBMISSIONS TO CONFERENCES. *Advs. Complex Syst.*, *13*(03), 439–449.

Flett, G. L., Blankstein, K. R., Hewitt, P. L., & Koledin, S. (1992). COMPONENTS OF PERFECTIONISM AND PROCRASTINATION IN COLLEGE STUDENTS. *Social Behavior and Personality: an international journal*, *20*(2), 85–94.

Flett, G. L., Hewitt, P. L., & Martin, T. R. (1995). Dimensions of perfectionism and procrastination. In J. R. Ferrari, J. L. Johnson, & W. G. McCown (Eds.), *Procrastination and task avoidance: Theory, research, and treatment* (pp. 113–136). Springer US.

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

Froese, A. D., Nisly, S. J., & May, R. M. (1984). The effects of task interest and difficulty on procrastination. *Trans. Kans. Acad. Sci.*, *87*(3-4), 119–128.

Frost, R. O., & Marten, P. A. (1990). Perfectionism and evaluative threat. *Cognit. Ther. Res.*, *14*(6), 559–572.

Glimcher, P. W., & Fehr, E. (2013). *Neuroeconomics: Decision making and the brain*. Academic Press.

Gopher, D., Armony, L., & Greenshpan, Y. (2000). Switching tasks and attention policies. *J. Exp. Psychol. Gen.*, *129*(3), 308–339.

Green, L. (1982). Minority students' self-control of procrastination. *J. Couns. Psychol.*, *29*(6), 636–644.

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

Haghbin, M., McCaffrey, A., & Pychyl, T. A. (2012). The complexity of the relation between fear of failure and procrastination. *J. Ration. Emot. Cogn. Behav. Ther.*

Haycock, L. A., McCarthy, P., & Skay, C. L. (1998). Procrastination in college students: The role of self-efficacy and anxiety. *J. Couns. Dev.*, *76*(3), 317–324.

Hensley, L. C. (2014). Reconsidering active procrastination: Relations to motivation and achievement in college anatomy. *Learn. Individ. Differ.*, *36*, 157–164.

Holland, T. (2001). The perils of procrastination. *Far East. Econ. Rev.*

Hotaling, J. M., & Busemeyer, J. R. (2012). DFT-D: A cognitive-dynamical model of dynamic decision making. *Synthese*, *189*(S1), 67–80.

Howell, A. J., & Watson, D. C. (2007). Procrastination: Associations with achievement goal orientation and learning strategies. *Pers. Individ. Dif.*, *43*(1), 167–178.

Howell, A. J., Watson, D. C., Powell, R. A., & Buro, K. (2006). Academic procrastination: The pattern and correlates of behavioural postponement. *Pers. Individ. Dif.*, *40*(8), 1519–1530.

Hull, C. L. (1932). The goal-gradient hypothesis and maze learning. *Psychological review*, *39*(1), 25.

Huys, Q. J. M., Lally, N., Faulkner, P., Eshel, N., Seifritz, E., Gershman, S. J., Dayan, P., & Roiser, J. P. (2015). Interplay of approximate planning strategies. *Proc. Natl. Acad. Sci. U. S. A.*, *112*(10), 3098–3103.

Janssen, T., & Carton, J. S. (1999). The effects of locus of control and task difficulty on procrastination. *J. Genet. Psychol.*, *160*(4), 436–442.

Ji Won You. (2015). Examining the effect of academic procrastination on achievement using LMS data in e-learning. *J. Educ. Techno. Soc.*, *18*(3), 64–74.

Johnson, P. E., Jr, Perrin, C. J., Salo, A., Deschaine, E., & Johnson, B. (2016). Use of an explicit rule decreases procrastination in university students. *J. Appl. Behav. Anal.*, *49*(2), 346–358.

Kegley, C. W. (1989). The bush administration and the future of american foreign policy: Pragmatism, or procrastination? *Pres. Stud. Q.*, *19*(4), 717–731.

Keinan, R., & Bereby-Meyer, Y. (2012). “leaving it to chance”—passive risk taking in everyday life. *Judgm. Decis. Mak.*, *7*(6), 705–715.

Kerdijk, W., Cohen-Schotanus, J., Florentine Mulder, B., Muntinghe, F. L. H., & Tio, R. A. (2015). Cumulative versus end-of-course assessment: Effects on self-study time and test performance.

Kim, K. R., & Seo, E. H. (2015). The relationship between procrastination and academic performance: A meta-analysis. *Pers. Individ. Dif.*, *82*, 26–33.

Kim, S., Fernandez, S., & Terrier, L. (2017). Procrastination, personality traits, and academic performance: When active and passive procrastination tell a different story. *Pers. Individ. Dif.*, *108*, 154–157.

Kirby, K. N., Petry, N. M., & Bickel, W. K. (1999). Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *J. Exp. Psychol. Gen.*, *128*(1), 78–87.

Klassen, R. M., Krawchuk, L. L., & Rajani, S. (2008). Academic procrastination of undergraduates: Low self-efficacy to self-regulate predicts higher levels of procrastination. *Contemp. Educ. Psychol.*, *33*(4), 915–931.

Konradt, U., Ellwart, T., & Gevers, J. (2021). Wasting effort or wasting time? a longitudinal study of pacing styles as a predictor of academic performance. *Learn. Individ. Differ.*, *88*, 102003.

Koo, M., & Fishbach, A. (2014). Dynamics of self-regulation: How (un)accomplished goal actions affect motivation. *Motivation Science*, *1*(S), 73–90.

Kool, W., & Botvinick, M. (2014). A labor/leisure tradeoff in cognitive control. *J. Exp. Psychol. Gen.*, *143*(1), 131–141.

Kool, W., & Botvinick, M. (2018). Mental labour. *Nat Hum Behav*, *2*(12), 899–908.

Krause, K., & Freund, A. M. (2014a). How to beat procrastination. *Eur. Psychol.*, *19*(2), 132–144.

Krause, K., & Freund, A. M. (2014b). Delay or procrastination – a comparison of self-report and behavioral measures of procrastination and their impact on affective well-being. *Pers. Individ. Dif.*, *63*, 75–80.

Lamwers, L. L., & Jazwinski, C. H. (1989). A comparison of three strategies to reduce student procrastination in PSI.

Lay, C. H. (1986). At last, my research article on procrastination. *J. Res. Pers.*, *20*(4), 474–495.

Lay, C. H. (1987). Publication delay: An analysis of journal days, reviewer days, and author days to revision. *Can. J. Behav. Sci.*, *19*(3), 324–331.

Lee, N. C., Krabbendam, L., Dekker, S., Boschloo, A., de Groot, R. H. M., & Jolles, J. (2012). Academic motivation mediates the influence of temporal discounting on academic achievement during adolescence. *Trends in Neuroscience and Education*, *1*(1), 43–48.

Lempert, K. M., Mechanic-Hamilton, D. J., Xie, L., Wisse, L. E. M., de Flores, R., Wang, J., Das, S. R., Yushkevich, P. A., Wolk, D. A., & Kable, J. W. (2020). Neural and behavioral correlates of episodic memory are associated with temporal discounting in older adults. *Neuropsychologia*, *146*, 107549.

Liborius, P., Bellhäuser, H., & Schmitz, B. (2019). What makes a good study day? an intraindividual study on university students' time investment by means of time-series analyses. *Learning and Instruction*, *60*, 310–321.

Lieder, F., Chen, O. X., Krueger, P. M., & Griffiths, T. L. (2019). Cognitive prostheses for goal achievement. *Nat Hum Behav*, *3*(10), 1096–1106.

Lim, J. M. (2016). Predicting successful completion using student delay indicators in undergraduate self-paced online courses. *Distance Education*, *37*(3), 317–332.

Lopez-Guzman, S., Konova, A. B., Louie, K., & Glimcher, P. W. (2018). Risk preferences impose a hidden distortion on measures of choice impulsivity. *PLoS One*, *13*(1), e0191357.

Lord, R. G., Diefendorff, J. M., Schmidt, A. M., & Hall, R. J. (2010). Self-regulation at work. *Annu. Rev. Psychol.*, *61*, 543–568.

Lubbers, M. J., Van Der Werf, M. P. C., Kuyper, H., & Hendriks, A. A. J. (2010). Does homework behavior mediate the relation between personality and academic performance? *Learn. Individ. Differ.*, *20*(3), 203–208.

Ludwig, P., & Schicker, A. (2018). *The end of procrastination: How to stop postponing and live a fulfilled life*. St. Martin's Essentials.

Malatincová, T. (2015). The mystery of "should": Procrastination, delay, and reactance in academic settings. *Pers. Individ. Dif.*, *72*, 52–58.

Marshall, K., Forbes, A., Kearns, H., & Gardiner, M. (2008). When a high distinction isn't good enough: A review of perfectionism and self-handicapping. *Australian Educational Researcher*, *35*(3), 21–36.

Martinez, S.-K., Meier, S., & Sprenger, C. (2017). Procrastination in the field: Evidence from tax filing, 78.

McCrea, S. M., Liberman, N., Trope, Y., & Sherman, S. J. (2008). Construal level and procrastination. *Psychol. Sci.*, *19*(12), 1308–1314.

McElroy, B. W., & Lubich, B. H. (2013). Predictors of course outcomes: Early indicators of delay in online classrooms. *Distance Education*, *34*(1), 84–96.

Milgram, N., Marshevsky, S., & Sadeh, C. (1995). Correlates of academic procrastination: Discomfort, task aversiveness, and task capability. *J. Psychol.*, *129*(2), 145–155.

Milgram, N. A., Dangour, W., & Ravi, A. (1992). Situational and personal determinants of academic procrastination. *J. Gen. Psychol.*, *119*(2), 123–133.

Milkman, K. L., Minson, J. A., & Volpp, K. G. M. (2014). Holding the hunger games hostage at the gym: An evaluation of temptation bundling. *Manage. Sci.*, *60*(2), 283–299.

Moon, S. M., & Illingworth, A. J. (2005). Exploring the dynamic nature of procrastination: A latent growth curve analysis of academic procrastination. *Pers. Individ. Dif.*, *38*(2), 297–309.

Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychol. Rev.*, *86*(3), 214–255.

Neal, A., Ballard, T., & Vancouver, J. B. (2017). Dynamic Self-Regulation and Multiple-Goal pursuit. *Annu. Rev. Organ. Psychol. Organ. Behav.*, *4*(1), 401–423.

Niermann, H. C. M., & Scheres, A. (2014). The relation between procrastination and symptoms of attention-deficit hyperactivity disorder (ADHD) in undergraduate students. *Int. J. Methods Psychiatr. Res.*, *23*(4), 411–421.

Niyogi, R. K., Breton, Y.-A., Solomon, R. B., Conover, K., Shizgal, P., & Dayan, P. (2014). Optimal indolence: A normative microscopic approach to work and leisure. *J. R. Soc. Interface*, *11*(91), 20130969.

Niyogi, R. K., Shizgal, P., & Dayan, P. (2014). Some work and some play: Microscopic and macroscopic approaches to labor and leisure. *PLoS Comput. Biol.*, *10*(12), e1003894.

O'Brien, W. K. (2000). *Applying the transtheoretical model to academic procrastination* (J. P. Carbonari, Ed.; Doctoral dissertation). University of Houston. Ann Arbor, United States, search.proquest.com.

O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. *Am. Econ. Rev.*, *89*(1), 103–124.

O'Donoghue, T., & Rabin, M. (2001). Choice and procrastination. *Q. J. Econ.*, *116*(1), 121–160.

Onji, K. (2013). Estimating the effects of procrastination on performance: A small sample study. *J. Socio Econ.*, *44*, 85–90.

Owens, S. G., Bowman, C. G., & Dill, C. A. (2008). Overcoming procrastination: The effect of implementation intentions1. *J. Appl. Soc. Psychol.*, *38*(2), 366–384.

Park, S. W., & Sperling, R. A. (2012). Academic procrastinators and their self-regulation. *Psychology*, *03*(01), 12–23.

Parthasarathi, T., McConnell, M. H., Luery, J., & Kable, J. W. (2017). The vivid present: Visualization abilities are associated with steep discounting of future rewards. *Front. Psychol.*, *8*, 289.

Pehlivanova, M., Wolf, D. H., Sotiras, A., Kaczkurkin, A. N., Moore, T. M., Ciric, R., Cook, P. A., Garcia de La Garza, A., Rosen, A. F. G., Ruparel, K., Sharma, A., Shinohara, R. T., Roalf, D. R., Gur, R. C., Davatzikos, C., Gur, R. E., Kable, J. W., & Satterthwaite, T. D. (2018). Diminished cortical thickness is associated with impulsive choice in adolescence. *J. Neurosci.*, *38*(10), 2471–2481.

Pittman, T. S., Tykocinski, O. E., Sandman-Keinan, R., & Matthews, P. A. (2008). When bonuses backfire: An inaction inertia analysis of procrastination induced by a missed opportunity. *J. Behav. Decis. Mak.*, *21*(2), 139–150.

Prystawski, B., Mohnert, F., Tošić, M., & Lieder, F. (2022). Resource-rational models of human goal pursuit. *Top. Cogn. Sci.*, *14*(3), 528–549.

Putnik, Z., Ivanović, M., Budimac, Z., & Bothe, K. (2013). Analysis of students' behaviour based on participation and results achieved in wiki-based team assignments. *Proceedings of the 6th Balkan Conference in Informatics*, 179–186.

Pychyl, T. A., & Flett, G. L. (2012). Procrastination and self-regulatory failure: An introduction to the special issue. *J. Ration. Emot. Cogn. Behav. Ther.*, *30*(4), 203–212.

R Core Team, R., et al. (2013). R: A language and environment for statistical computing.

Rakes, G. C., & Dunn, K. E. (2010). The impact of online graduate students' motivation and self-regulation on academic procrastination [Accessed: 2023-10-31].

Raphaël, L. B., & Mathias, P. (2022). A neuro-computational account of procrastination behavior. *Nat. Commun.*, *13*(1), 5639.

Reuben, E., Sapienza, P., & Zingales, L. (2015). Procrastination and impatience. *Journal of Behavioral and Experimental Economics*, *58*, 63–76.

Richardson, M., Burges, C. J. C., & Renshaw, E. (2013). MCTest: A challenge dataset for the open-domain machine comprehension of text [Accessed: 2023-10-29].

Roberts, M. S., Fulton, M., & Semb, G. (1988). Self-Pacing in a personalized psychology course: Letting students set the deadlines. *Teach. Psychol.*, *15*(2), 89–92.

Roe, R. A. (2014). Time, performance and motivation. *Time and work: How time impacts individuals*, *1*, 63–110.

Rothblum, E. D., Solomon, L. J., & Murakami, J. (1986). Affective, cognitive, and behavioral differences between high and low procrastinators. *J. Couns. Psychol.*, *33*(4), 387–394.

Russell, D. (1982). The causal dimension scale: A measure of how individuals perceive causes. *J. Pers. Soc. Psychol.*, *42*(6), 1137–1145.

Saleem, M., & Rafique, R. (2012). Procrastination and Self-Esteem among university students. *Pakistan Journal of Social & Clinical Psychology*, *9*(3).

Scent, C. L., & Boes, S. R. (2014). Acceptance and commitment training: A brief intervention to reduce procrastination among college students. *J. College Stud. Psychother.*, *28*(2), 144–156.

Schmidt, A. M., & DeShon, R. P. (2007). What to do? the effects of discrepancies, incentives, and time on dynamic goal prioritization. *J. Appl. Psychol.*, *92*(4), 928–941.

Schmidt, A. M., & Dolis, C. M. (2009). Something's got to give: The effects of dual-goal difficulty, goal progress, and expectancies on resource allocation. *J. Appl. Psychol.*, *94*(3), 678–691.

Schmidt, A. M., Dolis, C. M., & Tolli, A. P. (2009). A matter of time: Individual differences, contextual dynamics, and goal progress effects on multiple-goal self-regulation. *J. Appl. Psychol.*, *94*(3), 692–709.

Schouwenburg, H. C. (1992). Procrastinators and fear of failure: An exploration of reasons for procrastination. *Eur. J. Pers.*, *6*(3), 225–236.

Schouwenburg, H. C., & Groenewoud, J. (2001). Study motivation under social temptation; effects of trait procrastination. *Pers. Individ. Dif.*, *30*(2), 229–240.

Schouwenburg, H. C., & Lay, C. H. (1995). Trait procrastination and the big-five factors of personality. *Pers. Individ. Dif.*, *18*(4), 481–490.

Schunk, D. H., & Swartz, C. W. (1993). Goals and progress feedback: Effects on self-efficacy and writing achievement. *Contemp. Educ. Psychol.*, *18*(3), 337–354.

Senecal, N., Wang, T., Thompson, E., & Kable, J. W. (2012). Normative arguments from experts and peers reduce delay discounting. *Judgm. Decis. Mak.*, *7*(5), 568–589.

Senécal, C., Koestner, R., & Vallerand, R. J. (1995). Self-Regulation and academic procrastination. *J. Soc. Psychol.*, *135*(5), 607–619.

Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a rational and mechanistic account of mental effort. *Annu. Rev. Neurosci.*, *40*, 99–124.

Singhi, N., Mohnert, F., Prystawski, B., & Lieder, F. (2023). Toward a normative theory of (self-)management by goal-setting.

Sirois, F. M. (2007). "i'll look after my health, later": A replication and extension of the procrastination–health model with community-dwelling adults. *Pers. Individ. Dif.*, *43*(1), 15–26.

Sirois, F. M. (2014). Procrastination and stress: Exploring the role of self-compassion. *Self Identity*, *13*(2), 128–145.

Sirois, F. M., Melia-Gordon, M. L., & Pychyl, T. A. (2003). "i'll look after my health, later": An investigation of procrastination and health. *Pers. Individ. Dif.*, *35*(5), 1167–1184.

Solomon, L. J., & Rothblum, E. D. (1984). Academic procrastination: Frequency and cognitive-behavioral correlates. *J. Couns. Psychol.*

Steel, P. (2007). The nature of procrastination: A meta-analytic and theoretical review of quintessential self-regulatory failure. *Psychol. Bull.*, *133*(1), 65–94.

Steel, P., Brothen, T., & Wambach, C. (2001). Procrastination and personality, performance, and mood. *Pers. Individ. Dif.*, *30*(1), 95–106.

Steel, P., & Klingsieck, K. B. (2016). Academic procrastination: Psychological antecedents revisited. *Aust. Psychol.*, *51*(1), 36–46.

Steel, P., & König, C. J. (2006). Integrating theories of motivation. *AMRO*, *31*(4), 889–913.

Steel, P., Svartdal, F., Thundiyil, T., & Brothen, T. (2018). Examining procrastination across multiple goal stages: A longitudinal study of temporal motivation theory. *Front. Psychol.*, *9*, 327.

Stöber, J. (2001). Worry, procrastination, and perfectionism: Differentiating amount of worry, pathological worry, anxiety, and depression. *Cognit. Ther. Res.*, *25*(1), 49–60.

Strunk, K. K., & Steele, M. R. (2011). Relative contributions of self-efficacy, self-regulation, and self-handicapping in predicting student procrastination. *Psychol. Rep.*, *109*(3), 983–989.

Tice, D. M., & Baumeister, R. F. (1997). Longitudinal study of procrastination, performance, stress, and health: The costs and benefits of dawdling. *Psychol. Sci.*, *8*(6), 454–458.

Urban, T. (2016). Inside the mind of a master procrastinator. *Retrieved May*.

van Eerde, W., & Klingsieck, K. B. (2018). Overcoming procrastination? a meta-analysis of intervention studies. *Educational Research Review*, *25*, 73–85.

Vangsness, L., & Young, M. E. (2020). Turtle, task ninja, or time waster? who cares? traditional Task-Completion strategies are overrated. *Psychol. Sci.*, *31*(3), 306–315.

Wesley, J. C. (1994). Effects of ability, high school achievement, and procrastinatory behavior on college performance. *Educ. Psychol. Meas.*, *54*(2), 404–408.

Wesp, R. (1986). Reducing procrastination through required course involvement.

Wieber, F., & Gollwitzer, P. (2010). Overcoming procrastination through planning.

Wohl, M. J. A., Pychyl, T. A., & Bennett, S. H. (2010). I forgive myself, now I can study: How self-forgiveness for procrastinating can reduce future procrastination. *Pers. Individ. Dif.*, *48*(7), 803–808.

Woolley, K., & Fishbach, A. (2016). For the fun of it: Harnessing immediate rewards to increase persistence in Long-Term goals. *J. Consum. Res.*, *42*(6), 952–966.

Woolley, K., & Fishbach, A. (2017). Immediate rewards predict adherence to Long-Term goals. *Pers. Soc. Psychol. Bull.*, *43*(2), 151–162.

Wrosch, C., Scheier, M. F., Miller, G. E., Schulz, R., & Carver, C. S. (2003). Adaptive self-regulation of unattainable goals: Goal disengagement, goal reengagement, and subjective well-being. *Pers. Soc. Psychol. Bull.*, *29*(12), 1494–1508.

Wylie, G., & Allport, A. (2000). Task switching and the measurement of "switch costs". *Psychol. Res.*, *63*(3-4), 212–233.

Xin, Y., Xu, P., Aleman, A., Luo, Y., & Feng, T. (2020). Intrinsic prefrontal organization underlies associations between achievement motivation and delay discounting. *Brain Struct. Funct.*, *225*(2), 511–518.

Xu, P., González-Vallejo, C., & Xiong, Z. H. (2016). State anxiety reduces procrastinating behavior. *Motiv. Emot.*, *40*(4), 625–637.

Yang, X., Liu, R.-D., Ding, Y., Hong, W., & Jiang, S. (2021). The relations between academic procrastination and self-esteem in adolescents: A longitudinal study. *Curr. Psychol.*

Yu, L. Q., Lee, S., Katchmar, N., Satterthwaite, T. D., Kable, J. W., & Wolf, D. H. (2017). Steeper discounting of delayed rewards in schizophrenia but not first-degree relatives. *Psychiatry Res.*, *252*, 303–309.

Zhang, P., & Ma, W. J. (2019). Procrastination in rational agents. *2019 Conference on Cognitive Computational Neuroscience.*

Zhang, P., & Ma, W. J. (2023a). The dynamic nature of procrastination. *2023 Conference on Cognitive Computational Neuroscience.*

Zhang, P., & Ma, W. J. (2023b). *Temporal discounting predicts procrastination in a real-world task.*

Zhang, S., & Yu, A. J. (2013). Forgetful bayes and myopic planning: Human learning and decision-making in a bandit setting. *Adv. Neural Inf. Process. Syst.*, *26.*

Zhang, S., & Feng, T. (2020). Modeling procrastination: Asymmetric decisions to act between the present and the future. *J. Exp. Psychol. Gen.*, *149*(2), 311–322.