Title
The role of temporal uncertainty in procrastination

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors
Sokratous, Konstantina
Hintze, Arend
Fitch, Anderson
et al.

Publication Date
2026

The role of temporal uncertainty in procrastination

Konstantina Sokratous (ksokratous@missouri.edu)[1], Arend Hintze[2], Anderson K. Fitch[3] & Peter D. Kvam[4]
[1]Department of Psychological Sciences, University of Missouri
[2]Department of Data Analysis, Dalarna University
[3]Department of Psychology, University of Florida
[3]Department of Psychology, The Ohio State University

Abstract

Procrastination, where people engage with an immediately rewarding task when they should work on a task with a delayed payoff, is typically framed as self-regulation failure or lack of motivation. We challenge this view by demonstrating that procrastination emerges naturally from the computational intractability of sequential resource allocation under temporal uncertainty. When time horizons are certain, humans approach mathematically optimal strategies for balancing competing tasks. However, when deadlines are uncertain – requiring integration over exponentially growing outcome spaces – participants systematically over-invest in immediate rewards early, then attempt to compensate too late. The findings reframe procrastination as bounded rationality under computational constraints rather than pure irrationality, suggesting that interventions targeting uncertainty reduction or temporal scaffolding may prove more effective than exhortations to increase motivation or willpower. **Keywords:** cognitive science; procrastination, sequential decision making; temporal uncertainty;

Introduction

Procrastination — often defined as delaying effortful tasks in favor of immediately rewarding activities (Steel, 2007) — is among the most common self-reported behavioral problems, affecting academic achievement, health outcomes, and financial well-being across populations. Despite decades of research, the mechanisms underlying procrastination remain contested. Dominant psychological accounts frame procrastination as self-regulation failure: people know what they should do but fail to do it, whether due to motivational deficits, poor impulse control, or task aversion (Grund & Fries, 2018; Gustavson et al., 2014; Rebetez et al., 2015, 2018). This framework assumes that optimal behavior is obvious and accessible, and therefore procrastination reflects a breakdown in executing known-optimal strategies.

In this paper, we challenge this assumption by re-framing procrastination as an adaptive response to resource allocation problems. When tasks must be completed under uncertain time horizon – a ubiquitous feature of real-world environments – determining the optimal allocation of effort across competing goals is computationally intractable. Rather than representing pure self-regulation failure, procrastination may emerge from the interaction between environmental complexity and cognitive constraints; a signature of bounded rationality rather than irrationality (Hertwig & Pedersen, 2016). This reframing has profound implications. If people cannot reliably compute optimal strategies under uncertainty, then procrastination-like behavior may reflect reasonable heuristics adapted to typical environmental statistics, even when these heuristics produce regret in atypical situations. In terms of Simon's scissors (Gigerenzer et al., 2011; Simon, 1990b), human behavior is after all determined by both the complexity of the environment and our computational capabilities. In this case, the complexity of an environment with temporal uncertainty may simply exceed the computing capabilities of humans (Steel, 2007).

Here, we formalize procrastination as a sequential resource allocation problem where time or effort is modeled as an exhaustible good (Fischer, 1999). Under the assumptions of economic rationality, procrastination can arise from the uncertainty that emanates from two different types of decision-making situations: (1) when an agent exhibits time-consistent preferences but has incomplete information about the way they can allocate their time (i.e. scheduling) and (2) when an agent exhibits time-inconsistent – or hyperbolic – preferences while well-aware and informed of situational constraints (Akerlof, 1991; Ross, 2010). The idea of temporal proximity and that events tend to have less impact as they become available later in time has been central in modeling procrastination, with intertemporal choice models and discounted utility being the most prominent ones (O'Donoghue & Rabin, 1999; O'Donoghue & Rabin, 2001; Steel, 2007).

To test these assumptions, we introduce a novel experimental paradigm where we vary systematically task structures and temporal uncertainty conditions. The strategies that the subjects produced during the task were compared to mathematically derived optimal solutions – or allocations – that maximize expected outcomes. In the following section, we describe in greater detail how the optimal strategies are computed and how this process becomes complicated once temporal or resource uncertainty is introduced. We show that humans typically approach optimal performance when time is certain, but systematically deviate when uncertainty in introduced. Procrastinators do not simply under-perform on resource allocation tasks, as one might expect from a self-regulation explanation; rather, they exhibit a specific temporal signature of over-investing in immediately rewarding activities early in the available time window, then attempting to compensate too late when work tasks become urgent. This temporal misallocation pattern is dissociable from global task performance and emerges most reliably under conditions that mirror real-world procrastination scenarios. Uncertainty

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

combines with task difficulty to determine which task should ideally be prioritized. This highlights that optimizing resource allocation under uncertainty is an exceptionally difficult problem and that both uncertainty and temporal discounting can explain several facets of human procrastination behavior.

Optimal Strategies

We study sequential resource allocation using an investment task where participants allocate time between two activities; Task A provides immediate payoffs with diminishing marginal returns (analogous to "play"), while Task B requires sustained effort before yielding rewards (analogous to "work"). When the total available time t is known with certainty, finding the optimal allocation is straightforward: enumerate all possible time allocations (total $\frac{t}{\Delta_t}$, where Δ_t is the time / effort increment a person considers), compute total payoffs for each, and select the partition that yields the maximum value. For example, $t = 100$ steps (with $\Delta_t = 1$ step) requires evaluating only 101 different allocations. Importantly, under these certain conditions, the *order* of investments is irrelevant – completing play activities first or last yields identical outcomes, provided planned allocation is achieved by the deadline. From this perspective, procrastinating on "work" investments is not inherently good or bad – opting for immediate payoffs first is fine, as long as sufficient time is allocated to complete the delayed-payoff work task.

Temporal uncertainty To understand the difficulty of operating under uncertainty, let us consider the case where an agent might have anywhere from 1-100 time steps – uniformly distributed – to invest in Task A / play and Task B / work. When t is drawn from a distribution (here uniform over 1–100 steps), the order of investments becomes critical; prioritizing one task increases the probability of completing a desired allocation for that task but risks under-investing in the other if time runs out. It is here where procrastination gets interesting; when there is uncertainty about t, the agent must prioritize and think carefully about the *order* of task investments.

Identifying the globally optimal strategy now requires computing expected values across all possible stopping times for all possible sequences – a space of 2^{100} strategies, or more than 10^{30} computations, where only one of these strategies maximizes the expected payoff (e.g, globally optimal strategy). Compared to the condition with temporal certainty, this is a difference of 28 orders of magnitude in terms of computational challenge.

Fortunately, this problem can be simplified by equating the payoffs of trajectories that reach the same point. That is, the sequence A, A, B, B has the same payoff as the sequence A, B, A, B. These two sequences correspond to a "point" in a 2-dimensional space, where one axis is the number of investments into A and the other is the number of investments into B. This allows us to work on a grid like the one shown in Figure 1, taking the value of arriving at each "point" as an input to the algorithm determining an optimal solution.

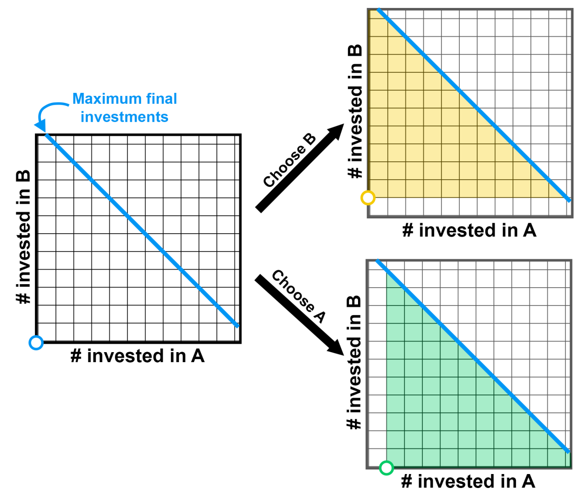

Figure 1: Approach to identifying the optimal next choice at each time step of the investment task. The expected values of two "triangles" of outcomes are compared. To derive the overall optimal solution, this process must be repeated for every investment decision a participant makes.

From here, identifying the optimal solution is a sequential expected-utility-maximization procedure. At each time step in the task, the decision-maker must decide whether to make their next investment into Task A or into Task B. To optimize this decision, they can compute the expected utility of all the outcomes that would be possible if they invested into Task A and compare them to the expected utility of all the outcomes that would be possible if they invested into Task B. This creates two "triangles" as shown by the yellow and green regions of Figure 1. For each shaded region, a decision-maker must calculate the probabilities of stopping at the circumscribed grid points (determined by the distribution of t) and multiply them by the payoff they would receive if they stopped at that point – and then sum across all points in the shaded area to get the overall expected payoff. This allows them to make the payoff-optimizing choice at each step by comparing the expected payoffs between shaded regions and picking the one with a higher value.

Hard problems for humans

The computational complexity of the sequential problem puts it out of reach for optimization by human decision-makers. Yet, bounded rationality theory suggests that intractable problems do not leave organisms helpless; instead they prompt adoption of heuristics or "satisficing" strategies that perform well on average given typical environmental structure, even if they fail to optimize in every instance (Gigerenzer et al., 1999). This "well enough" approach can potentially yield a consistently high level of performance that leaves the decision-maker satisfied with their outcomes (Schwartz et al., 2002).

These strategies may be *ecologically rational*, balancing the structure of the environment and the computational constraints of real human decision-makers to obtain good – if not perfectly optimal – outcomes under constraints (Hertwig et al., 2019; Simon, 1990b).

Methods

We developed a behavioral paradigm in which participants allocated limited time steps between two activities with contrasting payoff structures (2). The payoffs for Task A ("fun") were specified based on a classical diminishing marginal utility function (Bernoulli, 1738), following a concave power function:

$$u_A(t) = 1000 \cdot t^{.4} \tag{1}$$

Conceptually, this represents the way that fun, non-work tasks confer utility; They tend to be immediately rewarding and continue to provide utility the more they are pursued, but with a diminishing marginal utility that means that additional time spent on the activity yields less and less reward (Bernoulli, 1738; Kahneman & Tversky, 1979).

The payoffs for Task B ("work") were varied across conditions, but they all had a consistent functional form, a logistic function, representing the structure of the payoffs for these types of tasks. Specifically, they require some initial investment with no payoff (time spent working) before investments yield significant payoffs. That is, a person must allocate time to the task with little to no initial reward. However, once sufficient time is spent on the task, payoffs grow relatively quickly – similar to a project or task nearing completion. Once a work task is "complete," however, it confers less and less marginal utility until it ultimately stops conferring utility altogether. Therefore, the payoffs for this function were initially low, accelerated as the task neared completion, and then decelerated as it reached completion up to a maximum payoff. This related the amount of time spent on the task t to the payoff $u_B(t)$ formally as:

$$u_B(t) = \frac{1000}{1 + e^{-m(t-\beta)}} \tag{2}$$

At each discrete time step, participants chose whether to invest in Task A or Task B. They received complete information about both payoff functions (displayed graphically in Figure 2) and the exact points they would earn from the next investment in either task. Once their allocated time expired, participants were shown their total points and advanced to the next trial.

Because participants were given a finite amount of time, they had to make their allocation decisions carefully; any time they spent on Task A/fun could not be spent on Task B/work. Thus, the task focused on the *trade-off* between competing goals. At each point in time, therefore, their possible payoffs would be determined both by what they had already done (how much time spent on A vs. B) as well as the next decisions they made about resource allocations.

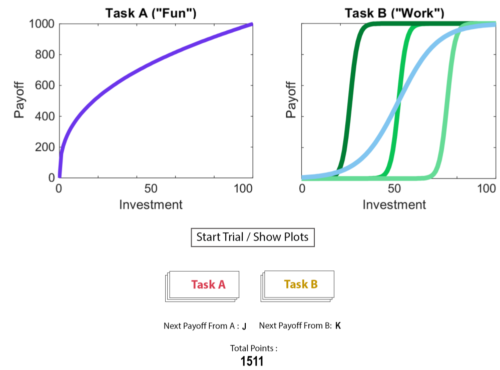

Figure 2: Payoffs for the "fun" (A, top left) and "work" (B, top right) tasks that participants experienced during the task. Since only the relative payoffs between them were important, we fixed the payoffs for Task A and varied the payoffs for Task B, visualized as multiple lines on the right plot.

Conditions

The payoff function for Task A was kept fixed throughout the experiment, while the parameters of payoff function for Task B that were varied across conditions of the experiment (m and β as shown in Equation 2). The task parameter m corresponded to how smoothly the amount of effort or time spent on the work task (Task B) was related to payoffs. The task parameter β corresponded to how long it would take to accomplish half the task – or rather, how many time steps would need to be invested in order to receive 500 points from Task B.

Across the task, we had four different payoff functions for Task B. Varying the difficulty and investment profile of this task by manipulating m and β led to the following payoff structures: (1) standard difficulty ($m = 0.50$ & $\beta = 25$), (2) moderate difficulty ($m = 0.50$ & $\beta = 50$), (3) hard task ($m = 0.50$ & $\beta = 75$), and (4) gradual investment ($m = 0.10$ & $\beta = 50$). The hard task condition was designed to be nearly unachievable within typical time constraints, testing whether participants appropriately abandon unrealistic goals. Overall, this allowed us to explore the effects of the slope and difficulty of the work task without creating too many different conditions.

Time uncertainty In the simplest version of the task with temporal certainty, players had a known, fixed number of time steps to invest. Such a task is rather trivial, as we show in the Optimal Strategies section above. An agent can simply add the payoffs for Task A at j to the payoffs for Task B at $t - j$, where t is the number of steps or amount of time they are allocated in total. We presented participants with three of these certain-deadline conditions, where t was fixed at 30, 70, or 100 steps.

In addition, we explored conditions where t was uncertain. As we demonstrate above, making t uncertain takes the task from a reasonably simple optimization problem, where the order of working on different tasks does not matter, to an extremely complex optimization problem, where the order in which an agent works on Task A and Task B is extremely important. Conceptually, this uncertainty can be uncertainty about the amount of time until the deadline t (deadline uncertainty) or uncertainty about the amount of time one "time step" Δ_t will be (investment uncertainty). Either construal is identical in terms of the formal structure of the task.

In our experiment, we varied this uncertainty by focusing on the deadline t. In the final three time conditions (the first three being certain $t = 30$, $t = 70$, and $t = 100$), the deadline t was drawn by first sampling a normalized variable $n \in [0, 1]$ from a Beta distribution:

$$n \sim Beta(\zeta_1, \zeta_2). \tag{3}$$

and then scaling it to the 1 to 100 timestep range, $t = [100n]$. We used three combinations of ζ_1 and ζ_2: $\zeta_1 = 1$ / $\zeta_2 = 1$ corresponding to a uniform distribution over normalized deadlines; $\zeta_1 = 1$ / $\zeta_2 = 2$ (a linearly decreasing probability on values of t, called the "short" condition), and $\zeta_1 = 2$ / $\zeta_2 = 1$ (a linearly increasing probability in values of t, called the "long" condition). Participants in the experiment were informed exactly what the probabilities of each value of t were and were shown a histogram illustrating how likely it would be that they would get t time steps in each of these conditions. They could, therefore, compute exact probabilities of reaching j time steps into Task A and $t - j$ time steps into Task B – if they wanted.

Put together, the experiment manipulated the two factors orthogonally across 24 conditions; *task payoff structure* (4 levels) and *temporal uncertainty* (6 levels). The 4 payoff conditions ($m = 0.50$ / $\beta = 25$, $m = 0.50$ / $\beta = 50$, $m = 0.50$ / $\beta = 75$, $m = 0.10$ / $\beta = 50$) and 6 termination conditions ($t = 30$, $t = 70$, $t = 100$, t drawn from a uniform, t short / drawn from Beta(1,2), and t long / drawn from Beta(2,1)). This design permitted systematic evaluation of how task difficulty and temporal (un)certainty interact to produce procrastination-like behavior – specifically, conditions where optimal policy prescribes early work investment but participants over – invest in immediate rewards instead. affected their behavior.

Experimental Procedure

Participants were undergraduates recruited from the University of Florida, and they received course credit for completing the experiment. A total of 254 participants (210 women, 42 men, and 2 nonbinary / other / did not respond) completed the experiment; an additional 7 began the experiment but were removed for failing to complete it in its entirety.

The experiment was hosted online using JavaScript embedded within Qualtrics. At the beginning of the experiment, participants completed informed consent and were introduced to the task with two practice trials.

After completing practice trials to familiarize themselves with the task, participants moved on to complete trials of the experiment. Trials were divided into blocks by termination condition so that the 30, 70, 100, short, long, and uniform termination conditions were all presented in separate blocks of the experiment. At the beginning of each block, participants were given information about the termination condition – either that they had a certain number of time steps to invest (30, 70, or 100) or that the number of time steps would be drawn according to a particular distribution (short, long, uniform). For the latter time-uncertainty conditions, they were shown a bar graph that showed the exact percentage of stopping after each number t of time steps. During debriefing, participants were asked if they understood these instructions; none reported any issues.

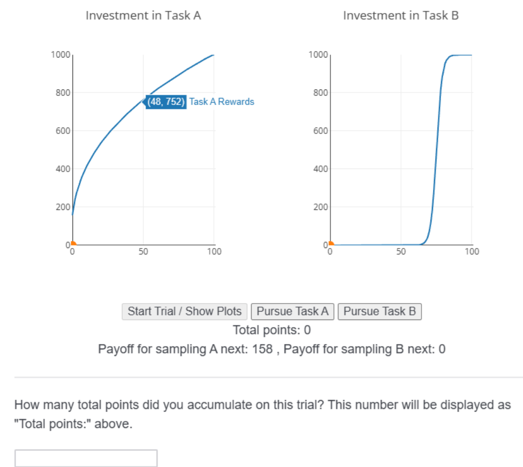

Figure 3: Visualization of the behavioral task as it was presented to subjects. They were permitted to Pursue Task A or Task B at each time point until the termination time, at which point the buttons were disabled and they were permitted to continue. Participants were then asked to report the total number of points they received on that trial as an attention check, after which they proceeded to the next trial (set of payoffs / time condition) until completing the experiment.

Once they received block instructions, participants completed 4 trials in each condition (1 trial of each payoff combination). At any point during a trial, participants could mouse over the graphs representing the payoffs for Task A and Task B, displaying the number of time steps (x value) and the payoff for investing that number of time steps into the corresponding task (y value). Participants were also told explicitly how many points they would receive for investing in either task next. This gave them complete information about the payoffs on each and every trial. We recorded the exact order of investments into Task A and Task B, including the number, proportion, and timing of when they clicked on the buttons to invest in each one. We also collected an attention check at the end of each trial – shown at the bottom of Figure 3 – that asked participants to report the number of points they accumulated during the trial. Trials where participants entered the wrong number or did not enter a number were removed from analysis.

Statistical Analysis

For data analysis on the experimental data, we looked primarily at the sequence of investments that participants made into Task A and Task B. This was compared against the mathematically derived optimal solutions.

Behavior on each trial generated a sequence of binary choices between Task A (fun) and Task B (work). In the uncertain time conditions, each sequence of decisions varied in length. To analyze behavior on the task, we first plot a heatmap of the sequence of choices that each participant made in each condition as a 2-D trajectory relative to the optimal solution. This gave us a qualitative idea of the general behavioral tendencies generated by each condition and the sample as a whole and is shown in Figure 4.

We additionally compute a "Proportion under the curve" [PUC] metric, which describes what proportion of trajectories correspond to procrastination behavior. For the certain termination conditions, this simply corresponds to the proportion of participant trajectories that were to the right of the optimal point – i.e., the trajectories that could never reach the optimal payoff because they had invested too much into Task A. For the uncertain conditions, the PUC was calculated as the proportion of trajectories that were to the right of / under the optimal trajectory, indicating they had invested too much into Task A up to that point in time relative to optimality. Participants in the uncertain trials could potentially be "procrastinating" for some proportion of a trial by over-investing in Task A, but then catch up and move above the curve later in a trial by (over)investing into Task B. The PUC metric therefore provides a good measure of the prevalence of procrastination behavior as it directly quantifies the proportion of time participants spent over-investing into immediate payoffs (Task A) relative to delayed ones (Task B).

Results

The behavioral trajectories across the 24 combinations of 4 payoff function conditions (rows) and 6 termination conditions (columns) are plotted as a heatmap in Figure 4.

Overall, procrastination appeared anytime there was sufficient time to make progress on Task B (i.e., the optimal solution wasn't to simply invest everything into Task A), and was especially common in the uncertain conditions when Task B would be difficult to complete in the allotted time and needed to be finished first. This corresponds to the classic characterization of procrastination behavior. When a "work" task should be completed immediately because it may be challenging to finish in the allotted time, participants do not begin work on it immediately but instead spend time procrastinating by investing time into an alternative ("play") task instead. However, this behavior typically manifests when there is some uncertainty about how long the tasks will take – or how much time there is in total.

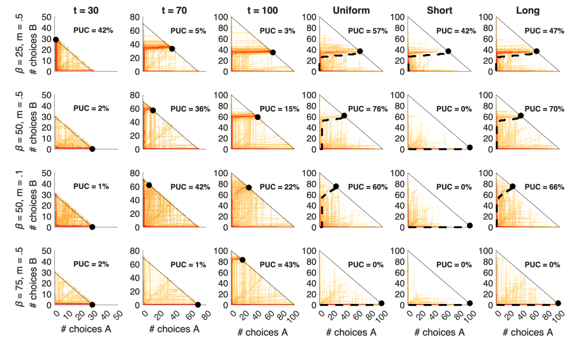

Figure 4: Heatmap of trajectories of human behavior for each condition depicting how often different balances of Task A (x) and Task B (y) were observed across all participants in the experiment. Darker / redder areas correspond to more common balances of behavior. Black dots correspond to the optimal strategy, and dotted black lines in the plots on the right side correspond to the optimal sequences of task investments.

Aside from the relative prevalence of procrastination behavior, there are a few other interesting aspects to Figure 4. Participants tended to be highly effective and close to optimal in most of the certain time conditions, with a few exceptions (e.g., the gradual-investment work task $m = 0.10$). It was mainly when uncertain time horizons were introduced that procrastination behavior regularly emerged – indicating that uncertainty is a key component driving these behaviors. Finally, most participants were able to figure out when a task was unachievable – which the hard work task condition ($\beta = 75$) tended to be. Human performance may therefore be driven in large part by recognizing when tasks are unachievable and (appropriately) avoiding them when this is the case. The perception of a task as being achievable or unachievable in the allotted time should therefore be a key component driving procrastination behavior.

Discussion

This paper set out to explore procrastination as a complex resource allocation problem, investigating human behavior in the context of different cumulative payoffs and uncertain temporal deadlines. Our results highlight the unavoidable computational difficulty in decision-making tasks that require balancing immediate and delayed rewards, especially when these are combined with complex environments. The computational limitations humans experience can be overcome by evolutionary pressures, reinforcement learning, and explicit instruction to create strategies that work well, even if they result in procrastination relative to the optimal solutions (Winter, 2000).

A key observation is that human participants performed best - i.e., less procrastination, earning the most rewards - in conditions where time was certain, consistently approaching the optimal solutions in these cases. When temporal uncertainty is not present, the computational challenge is reduced and humans are capable of (approximately) rational resource allocation. This aligns with previous research indicating that individuals can maximize rewards when the environment is predictable and task demands are clear (Balcı et al., 2011; Barendregt et al., 2022; Drugowitsch et al., 2016; Kilpatrick et al., 2018; Tanaka et al., 2006). However, when faced with uncertainty about the time available to complete tasks or when the difficulty of tasks increases, then humans consistently exhibit procrastination behavior. Precisely, procrastination manifested primarily when the following two conditions were met; sufficient time existed to make progress on both tasks (making immediate Task A investment suboptimal), and uncertainty about the time horizon created ambiguity about task achievability. This pattern was especially pronounced when Task B (work task) was difficult to complete in the allotted time and required early initiation. This corresponds precisely to the classic characterization of procrastination: when a "work" task should be completed immediately because it will be hard/take long to finish, participants instead invested time in an alternative "play" task. Critically, this behavior emerged most reliably under temporal uncertainty, indicating that uncertainty is not merely a contributing factor but a key component driving procrastination behavior. The stark performance difference between certain and uncertain conditions suggests that procrastination is not primarily a failure of self-regulation or motivation, but rather a response to the computational complexity introduced by environmental uncertainty.

Another important finding was that most participants accurately recognized when tasks were unachievable, as evidenced by appropriate avoidance of the hard work task condition. Human performance may therefore be driven in large part by implicit assessments of task achievability within the perceived time constraints. The perception of a task as achievable or unachievable in the allotted time emerged as a key component driving procrastination behavior. This finding extends theories of bounded rationality by suggesting that humans employ heuristics not only to simplify complex computations, but also to rapidly categorize tasks along an achievability dimension. When cognitive resources required to compute an optimal solution are unavailable, or when uncertainty about future rewards is high, individuals may favor immediate payoffs as a way to reduce problem complexity. This aligns with satisficing strategies proposed by researchers for managing the overwhelming complexity of real-world decision-making tasks (Simon, 1990a).

The emergence of procrastination across different task conditions points to a structural explanation rooted in computational intractability. While optimal strategies are mathematically achievable under certain conditions, humans consistently struggled to fully optimize behavior when faced with uncertainty about deadlines or task difficulty. The prevalence of procrastination may thus reflect an adaptive response to environments where exact optimization is computationally prohibitive (Lam et al., 2013). When faced with intractable problems involving multiple competing goals under uncertainty, defaulting to immediate rewards may serve as a reasonable heuristic that reduces cognitive load, even if it produces regret in hindsight.

This perspective reframes procrastination not as pure irrationality, but as boundedly rational behavior emerging from the interaction between cognitive constraints and environmental complexity (Hertwig et al., 2019). Such an interpretation is consistent with ecological rationality frameworks suggesting that apparently suboptimal behaviors may be well-adapted to the statistical structure of natural environments, even if they fail in artificial laboratory tasks designed to expose specific decision biases (Gigerenzer & Gaissmaier, 2011; Simon, 1990a).

Several limitations must be acknowledged. Our experimental paradigm simplifies real-world procrastination by focusing on a single task environment with well-defined payoffs and explicit time constraints. Real-world procrastination often involves ill-defined goals, ambiguous deadlines, and emotional factors such as task aversiveness or anxiety that our paradigm did not capture. Future work should explore how emotional and motivational factors interact with the computational mechanisms identified here, examine whether training or feedback can help individuals develop better allocation strategies under uncertainty, and investigate the neural substrates supporting both global diligence and temporal misallocation components of procrastination.

References

Akerlof, G. A. (1991). Procrastination and obedience. *The American Economic Review*, *81*, 1–19.

Balcı, F., Freestone, D. M., Simen, P., deSouza, L., Cohen, J. D., & Holmes, P. (2011). Optimal temporal risk assessment. *Frontiers in Integrative Neuroscience*, *5*.

Barendregt, N. W., Gold, J. I., Josić, K., & Kilpatrick, Z. P. (2022). Normative decision rules in changing environments. *eLife*, *11*.

Bernoulli, D. (1738). Exposition of a new theory on the measurement of risk. *Econometrica: Journal of the Econometric Society*, 23–36. https://doi.org/10.2307/1909829

Drugowitsch, J., Wyart, V., Devauchelle, A.-D., & Koechlin, E. (2016). Computational precision of mental inference as critical source of human choice suboptimality. *Neuron*, *92*, 1398–1411.

Fischer, C. (1999). *Read this paper even later: Procrastination with time-inconsistent preferences* (tech. rep.). Resources for the Future.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451–482. https://doi.org/10.1146/annurev-psych-120709-145346

Gigerenzer, G., Hertwig, R., & Pachur, T. (2011). *Heuristics: The foundations of adaptive behavior*. Oxford University Press.

Gigerenzer, G., Todd, P. M., & A. B. C. Research Group. (1999). *Simple heuristics that make us smart*. Oxford University Press USA.

Grund, A., & Fries, S. (2018). Understanding procrastination: A motivational approach. *Personality and Individual Differences*, *121*, 120–130.

Gustavson, D. E., Miyake, A., Hewitt, J. K., & Friedman, N. P. (2014). Genetic relations among procrastination, impulsivity, and goal-management ability. *Psychological Science*, *25*, 1178–1188.

Hertwig, R., & Pedersen, A. P. (2016). Finding foundations for bounded and adaptive rationality. *Minds and Machines*, *26*, 1–8.

Hertwig, R., Pleskac, T. J., & Pachur, T. (2019). *Taming uncertainty*. MIT Press.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–292.

Kilpatrick, Z. P., Holmes, W. R., Eissa, T. L., & Josić, K. (2018). Optimal models of decision-making in dynamic environments. *Current Opinion in Neurobiology*, *58*, 54–60.

Lam, S.-W., Ng, T. S., Sim, M., & Song, J.-H. (2013). Multiple objectives satisficing under uncertainty. *Operations Research*, *61*(1), 214–227.

O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. *The American Economic Review*, *89*, 103–124.

O'Donoghue, T., & Rabin, M. (2001). Choice and Procrastination*. *The Quarterly Journal of Economics*, *116*(1), 121–160. https://doi.org/10.1162/003355301556365

Rebetez, M. M. L., Rochat, L., Barsics, C., & der Linden, M. V. (2018). Procrastination as a self-regulation failure: The role of impulsivity and intrusive thoughts. *Psychological Reports*, *121*, 26–41.

Rebetez, M. M. L., Rochat, L., & der Linden, M. V. (2015). Cognitive, emotional, and motivational factors related to procrastination: A cluster analytic approach. *Personality and Individual Differences*, *76*, 1–6.

Ross, D. (2010). Economic models of procrastination. *The thief of time*, 28–50.

Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., & Lehman, D. R. (2002). Maximizing versus satisficing: Happiness is a matter of choice. *Journal of personality and social psychology*, *83*(5), 1178.

Simon, H. A. (1990a). Bounded rationality. *Utility and probability*, 15–18.

Simon, H. A. (1990b). Invariants of human behavior. *Annual Review of Psychology*, *41*, 1–19. https://doi.org/10.1146/annurev.ps.41.020190.000245

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

Tanaka, S. C., Samejima, K., Okada, G., Ueda, K., Okamoto, Y., Yamawaki, S., & Doya, K. (2006). Brain mechanism of reward prediction under predictable and unpredictable environmental dynamics. *Neural networks : the official journal of the International Neural Network Society*, *19 8*, 1233–41.

Winter, S. G. (2000). The satisficing principle in capability learning. *Strategic management journal*, *21*(10-11), 981–996.