

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Optimal and sub-optimal temporal decisions can explain procrastination in a real-world task

Permalink

<https://escholarship.org/uc/item/2mq517js>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Chebolu, Sahiti

Dayan, Peter

Publication Date

2024

Peer reviewed

Optimal and sub-optimal temporal decisions can explain procrastination in a real-world task

Sahiti Chebolu (sahiti.chebolu@tuebingen.mpg.de)
MPI for Biological Cybernetics

Peter Dayan (dayan@tuebingen.mpg.de)
MPI for Biological Cybernetics, University of Tuebingen

Abstract

Procrastination is a universal phenomenon, with a significant proportion of the population reporting interference and even harm from such delays. Why do people put off tasks despite what are apparently their best intentions, and why do they deliberately defer in the face of prospective failure? Past research shows that procrastination is a heterogeneous construct with possibly diverse causes. To grapple with the complexity of the topic, we construct a taxonomy of different types of procrastination and potential sources for each type. We simulate completion patterns from three broad model types: exponential or inconsistent temporal discounting, and waiting for interesting tasks; and provide some preliminary evidence, through comparisons with real-world data, of the plausibility of multiple types of, and pathways for, procrastination.

Keywords: procrastination; computational modeling; temporal decisions; temporal discounting; naturalistic data

Introduction

Many of us have finished assignments in the last minute, repeatedly put off going to the gym, sat on tax returns for weeks and worse. Procrastination is widespread, affecting some 80% of students and 20% of adults (Steel, 2007). Many suffer effects on their health (Sirois, 2007) and finances (O'Donoghue & Rabin, 1998), and most procrastinators wish to reduce it (O'Brien, 2002). But what exactly is procrastination and what are the underlying mechanisms?

Procrastination has long been recognised as a heterogeneous construct, with multiple psychological, personality correlates and possibly many explanations (Steel & Klingsieck, 2016). This understanding has informed the formulation of myriad questionnaires assessing multiple dimensions of interest comprising procrastination: delay of work, unnecessary or unreasonable delays, action-intention gaps, irrationality (i.e., failing to maximise utility), suffering consequences like missing deadlines or stress due to rushing, etc (Lay, 1986; McCown, Johnson, & Petzel, 1989; Steel, 2010).

Concurrently, types of procrastination have been proposed based on factors or clusters extracted from personality traits, procrastination questionnaires and self-reports. Thus, separate factors corresponding to traits of anxiety and neuroticism, impulsivity and disorganisation, low-energy and depression have been associated with different types of high procrastinators (Lay, 1986; McCown et al., 1989). Analysis of reasons for procrastination reveals factors like fear of failure and task aversiveness or avoidance (Grunschel, Patrzek, & Fries, 2013; Solomon & Rothblum, 1984). Schouwenburg

(2004) suggest that impulsivity forms the common base for all procrastination types, while neuroticism and extraversion affect how each type is manifested.

However, while this approach is useful in identifying important variables, it is unclear if these factors indeed correspond to distinct types and the axis on which their differences lie (do they map onto distinct mechanisms, consequences or other criteria?) Furthermore, the factors remain agnostic to specifics of the task where delays occur, focusing on general procrastination measures and traits. In reality, people procrastinate across various task structures: immediate or delayed rewards (Rozenal & Carlbring, 2014; Shu & Gneezy, 2010), with and without deadlines, or presence or absence of uncertainty about aspects of a task (Fischer, 1999). Finally, despite recognizing procrastination's complexity, it is often reduced to this singular definition: 'voluntary delay of an intended course of action despite expecting to be worse off' (Steel, 2007). We suggest this leaves out other types of problematic delays such as planning a course of action that leaves insufficient time in the first place (without any defections) or losing time in the process of making a decision.

In this paper, we propose an alternative approach to address procrastination diversity, drawing from reinforcement learning and decision-making principles. We categorize procrastination types based on commitment, adherence (or lack thereof) to a delay, and how delays contribute to problems (through missed deadlines, irrational choices, or flawed valuation processes). This includes the popular definition mentioned before but extends beyond it. Furthermore, it allows classification of mechanisms based on why a procrastination decision was made within the task's context. In the second part of the paper, we provide evidence for existence of various of these types using real-world data, demonstrating multiple plausible explanations for students' procrastination aligned with each type in our taxonomy.

Taxonomy

Definition

We operationalise procrastination in three main flavours:

1. Sticking to an intended delay, leaving insufficient time to complete a task. These could be delays intended to optimise utility, or miscalculations in the attempt to do so.
2. Delaying to a later time in spite of intending to act earlier.

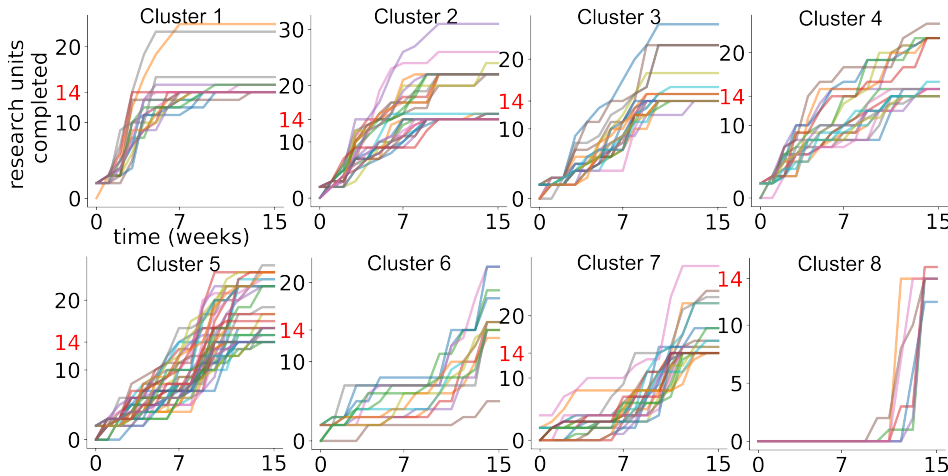


Figure 1: Eight clusters capture a variety of styles of work allocation by students in a real-world task (P. Zhang & Ma, 2023). Plots show the cumulative number of hours of work completed by each week of the semester. The threshold of 14 units (7 hours) is marked in red.

3. Delays from not committing to a time of action and hence missing a deadline, or losing time in the process.

Delays can be at the start, during, or at the end of tasks.

Mechanisms

For each definition, we can ask what the reasons for the respective decisions might be, naturally giving rise to a classification of mechanisms: that is, why commit to a delay that leaves insufficient time, why defect on an intended time of action or why not commit to one? The reasons for each type of decision would, of course, depend on the task structure.

In the deadline task we consider, we propose three (or many) possible mechanisms for students’ procrastination: prioritizing immediate rewards over delayed ones due to discounting, repeated delays from steeper discounting of efforts over rewards and delays from waiting for more interesting tasks with uncertain timing. Each corresponds to a type from our taxonomy. We then compare model simulations to patterns in task data to see which of these are plausible explanations (and hence if multiple types exist).

Methods

Task We analysed and modelled real-world data from P. Zhang and Ma (2023). In this study, 194 bachelor students in a Psychology course had to participate in at least 7 hours of experiments over a 16-week semester (110 days) to receive course credit. The task had a deadline at the end of the semester, with each research study requiring a minimum time of 0.5 hours or a multiple thereof, hence defining a unit of work. Additional participation above 7 hours contributed $1/8^{th}$ of a grade point per unit, up to 4 extra hours of work. Abundant research opportunities (average of 15 hours per student) implied they did not compete (P. Zhang & Ma, 2023).

Data P. Zhang and Ma (2023) only included 93 students who completed exactly 7-7.5 hours, to avoid confounds in their measure of procrastination. They found a correlation between this measure and discount factors of the students, albeit without categorising the patterns of completion.

We only excluded participants who dropped out, leaving us with data from 173 students. We first conducted a model-agnostic analysis of how students distributed their efforts through the semester. We normalised the trajectories by the total number of credits completed, since students completed different numbers. Employing k-means clustering with a Euclidean distance metric, we identified 8 clusters using the elbow method, capturing different styles and shapes (Figure 1). Consistent with prior research on pacing styles (Gevers, Rutte, & Van Eerde, 2006; Konradt, Ellwart, & Gevers, 2021), we identified 5 broad styles of patterns: early completion (cluster 1, 2, 3), steady completion (cluster 4, 5), action at the beginning and end (cluster 6), bulk of work at an intermediate point (cluster 7) and completion towards deadline (cluster 8). Some clusters share styles but differ in the momentary rates of working (trajectory slopes). In the following modeling work, we use simulations to explore mechanisms contributing to each pattern through qualitative comparisons.

Models To model temporal allocation of work in this task, we used Markov Decision Processes (MDPs). MDPs provide a normative framework to model sequential decisions where summed, discounted long-term net-utility should be maximised given task dynamics (Dayan & Daw, 2008; Sutton & Barto, 2018). In our case, the solution is an optimal policy specifying the utility-maximising allocation of work through the semester. Optimal policies are derived under assumptions such as exact knowledge of the decision problem, time-consistent discounting etc., but we also relax them in specific cases to derive sub-optimal policies that might mimic real behavior in the task.

We formalised the task as follows. The agent has 16 weeks ($0 \leq t \leq 15$) to complete at least 14 units and up to 22 units of work. Each week, the agent is in a state ($0 \leq s_t \leq 22$) indicating the number of units completed so far. The agent can decide to complete some number of units ($0 \leq a_t \leq 22 - s_t$) each week, with a Binomial success probability or efficacy (η) governing the actual number of units completed. This is a probabilistic interpretation of effectiveness or average rate of unit completion, where higher η means more work

can be completed per week on average. Stochasticity could arise from uncertainty about available time for work due to other unknown tasks or their durations. Some students may be more effective in organizing their time, resulting in higher efficacy. The Binomial transition probabilities are given by:

$$P(s' | s, a) = \binom{a}{s' - s} \eta^{s' - s} (1 - \eta)^{a - s' + s}$$

Every unit of work incurs an immediate effort cost r_{effort} , while remaining time not used for work ($22 - a_t$ units) is used to ‘shirk’ with reward r_{shirk} . This includes alternative tasks such as other university work, relaxing, chores, etc. Effort may seem more tedious or fatiguing with greater amount of work per week, akin to vigour costs (Niv, Daw, Joel, & Dayan, 2007). P. Zhang (2024) operationalised this by making effort costs convex, rising more quickly with more units of work per week: $r_{\text{effort}}(a) = r_{\text{effort}} a^k$, where a is the no. of units of work, k is convexity of the effort function, and r_{effort} is effort for a unit of work ($a = 1$).

Finally, the reward associated with finishing each unit r_{unit} , is only delivered once 14 units have been completed. Without explicit rewards, there is ambiguity about reward timing: some students may perceive rewards upon hitting the 14-unit requirement, while others might perceive it only with a formal confirmation of grades at the end of the semester. Indeed, in verbal responses, some students indicated that they felt the consequences of finishing lay too far in the future to motivate initial effort. We simulate both types of schedule separately.

In all cases, we find the optimal value function by recursively maximising the (discounted) sum of current and future rewards or the returns ($\mathbb{E} [\sum_{t=0}^T \gamma^t R(s, a, s')]$) as given by the Bellman optimality equation (Sutton & Barto, 2018). Action selection follows a softmax rule, leading to noisy decisions and ultimately, noisy trajectories of completion with the inverse temperature parameter, β controlling the extent of the noise (and is one source of inconsistency with the optimum).

Results

Unless specified otherwise, $r_{\text{unit}} = 4.0$ and $r_{\text{extra}} = r_{\text{unit}}/4$. Rewards from shirking and efforts from working were $r_{\text{shirk}} = 0.1$ and $r_{\text{effort}} = -0.3$, delivered immediately at the time of action. Softmax parameter $\beta = 7.0$. Typically, effort convexity $k = 1$ specifying a linear relationship between cost and work. We chose combinations of rewards and efforts such that the relative utility still makes it worthwhile to work, while avoiding effort from becoming too cheap. This ensures a need for balancing rewards and efforts, optimising allocation of work rather than continuously working until task completion. We chose a relatively high value of β so that trajectories are not too noisy and reflect the underlying policy, but can still account for inherent noisiness in people’s choices. We vary other parameters to show their effects on patterns of work allocation. For most results, we plot the average trajectory (over 1000 runs) in bold along with standard deviation as a shaded region and a few sample trajectories in dashed lines. The threshold of 14 units is shown in red on the y-axis.

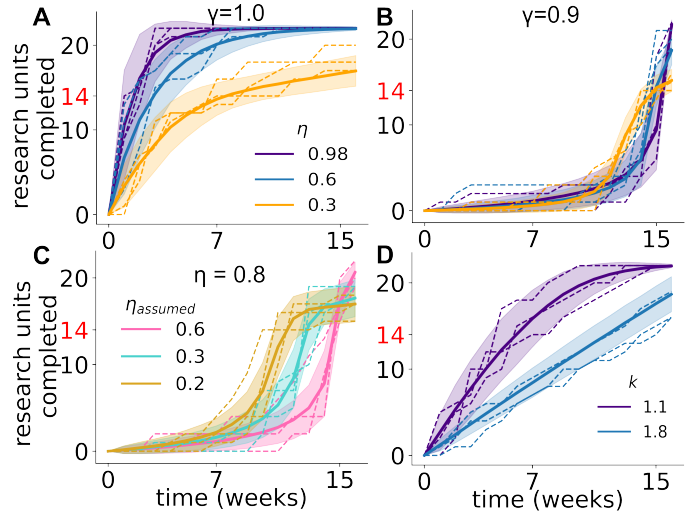


Figure 2: Mean and sample trajectories of work when rewards come at the deadline: A. With no discounting ($\gamma = 1$), it is best to finish as early as the efficacy (η) allows. B. $\gamma < 1$ makes it better to delay to the end instead due to preference for immediate rewards from shirking. C. Underestimating efficacy means one would finish earlier than one would nominally wish. D. Vigour costs incentivise spreading out work through the semester.

Exponential discounting with delayed rewards

In the first model, there is a common exponential discount factor γ for positive and negative rewards, and we assume the rewards from completing a minimum of 14 units (or more) come at the end of the semester.

Temporal discounting induces a temporal preference for working later Say there is no discounting ($\gamma = 1$). This means a reward is as valuable later in time as it is immediately. This implies it is always better to work than shirk as long as the efforts are worth the final rewards (to overcome the effect of efficacy $\eta < 1$). Hence, it is optimal to finish working before shirking. The decisions we show are noisy however, due to the softmax rule. Figure 2A shows that lower the efficacy, longer it takes to finish the task in practice. This lack of discounting along with varying levels of efficacy is one explanation for early completion patterns in clusters 1-4.

On the other hand, if delayed rewards are discounted, with $\gamma < 1$, there is a temporal preference for shirking and obtaining rewards immediately over working and paying effort costs to secure a distant reward of the credits. Therefore it becomes optimal to put off work until the end of the semester when $\gamma = 0.9$, as seen in Figure 2B.

Efficacy affects the extent of delay When $\gamma = 0.9$, efficacy controls how late a subject can afford to delay working, while allowing a reasonable chance of finishing as shown in Figure 2B. Therefore, higher the efficacy, the longer work can be delayed. Some students did mention they delayed requirements

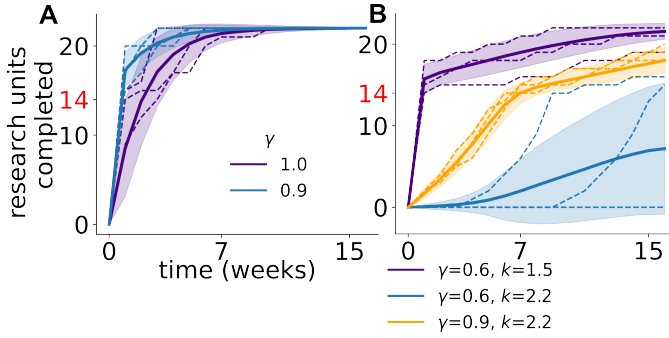


Figure 3: Patterns of working with immediate reward at threshold. A. There are no delays any longer even with $\gamma < 1$. B. However, with relatively high convexity in effort functions, it can become optimal to work later in the semester.

because they foresaw a quick completion. At very low efficacies (e.g., $\eta = 0.3$), working beyond the 7-hour requirement is no longer worthwhile for the low-reward extra units. These trajectories with discounting can replicate late completion as found in cluster 8 and some trajectories in cluster 6 which have shallower rise (maybe due to lower efficacy).

A gap between real and assumed efficacy leads to overestimation of delay So far, we assumed subjects calculate their best course of action based on perfect knowledge of their abilities. However, what if they underestimate their efficacy, that is, overestimate the average time it will take them to finish? With $\gamma = 0.9$, high actual efficacy $\eta_{\text{real}} = 0.8$ and if subjects have a low self-efficacy ($\eta_{\text{assumed}} < 0.8$), they plan to work earlier than if they had correctly estimated their efficacy. The greater the gap between the real and assumed efficacy, the earlier the subjects finish before the deadline, contributing to patterns like those in cluster 7 and even 6, where the bulk of work is postponed to an intermediate point (Figure 2C).

Convex effort costs could explain steady completion Finally, what if the relationship of effort costs to amount of work is more convex and not linear (P. Zhang, 2024)? This makes it prohibitive to do a lot of work at once, inducing a preference for spreading the work out and doing a little bit every week. With $\gamma = 1$ and $\eta = 0.8$, the greater the convexity, the steadier is the time course of working as shown in Figure 2D, replicating flatter patterns as in clusters 3-5.

Immediate rewards at threshold

Together, combinations of varying discount factors, efficacies and convexities along with delayed rewards could replicate most patterns found in this task. However, some students might perceive rewards as arriving immediately upon completing the 7 hour-requirement, rather than at the end. For this set of simulations, all parameter settings remain the same, except the timing of rewards for completion. We assume that rewards come immediately after the threshold of 14 units (and then with extra units as any are completed). We set $\eta = 0.8$.

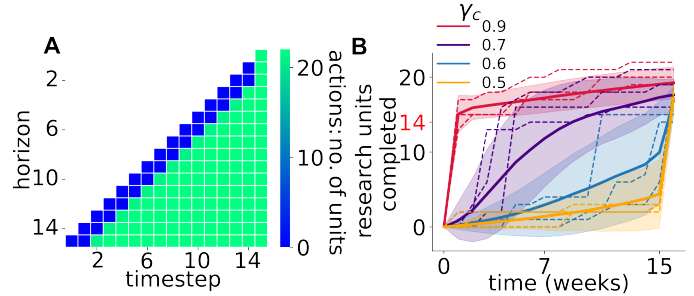


Figure 4: Delays from steeper discounting of efforts than rewards. A. The policy is time-inconsistent – the plan for when to start working changes at every time step. B. This leads to repeated delaying of work. Here, the yellow curve corresponds to the deterministic policy shown in A.

Immediate rewards at threshold eliminate delays in working due to discounting As before, in the absence of discounting ($\gamma = 1$), present rewards are as valuable as future rewards and hence there is a drive to finish work early. Now, even with discounting ($\gamma < 1$), there is no delay in working since the rewards are not delayed until the end of the semester. If anything, there is an opposite tendency to work and obtain rewards as soon as possible due to preference for immediate rewards, expediting work even more than the no discounting case as shown in Figure 3A. The reward schedule itself, irrespective of personal discount factors could be another explanation for early completion patterns in clusters 1 and 2.

Vigour costs can lead to delays in working As mentioned, when the convexity $k > 1$, it is better to spread work out over several weeks and so there is a delay until the requirement of 14 units is completed. This leads to an interesting effect where it is better to delay working in the beginning due to discounting of these temporarily delayed rewards. The effect can be seen in the progress lines that are curved upwards when discount factors are steep ($\gamma = 0.6$) and effort functions are relatively more convex ($k = 2.2$) as seen in Figure 3B. This could explain curved up patterns in clusters 6 and 7. For shallower discount factors or convexity, this effect disappears and now, higher the convexity, the flatter the trajectory, similarly to steady completion patterns in clusters 3-5.

Differential discounting of efforts and rewards

While discounting and convex costs can explain some delays and patterns, they fall short in explaining why someone might delay most work until the end of the semester as in the patterns in cluster 8, even with immediate rewards. We propose an alternate route to procrastination from time-inconsistent decisions stemming from non-exponential discounting, corresponding to the the second type in our taxonomy. This has previously been used to explain procrastination (Fischer, 1999; O'Donoghue & Rabin, 2001). By considering different exponential discount factors for rewards and costs, we show how this leads to defections and delays to the end of the

semester with immediate rewards (Le Bouc & Pessiglione, 2022). Defections disappear with larger rewards per unit, so this mechanism works only when rewards from task completion are relatively smaller. Thus we set $r_{\text{unit}} = 1.0$ and $\eta = 0.8$, with other parameter values unchanged, and rewards are received immediately upon reaching a threshold.

Optimisation with different discount factors Previously, with a single discount factor, the goal was to maximise $\mathbb{E} [\sum_{t=0}^T \gamma^t R(s, a, s')]$. However, now the objective changes to the following: $\mathbb{E} [\sum_{t=0}^T \gamma_r^t R(s, a, s') + \gamma_c^t C(s, a, s')]$ since there are different discount factors (γ_r and γ_c) for rewards and costs. This separation eliminates the structure that allowed solving the optimisation problem recursively before. We can still find the best sequence of actions to take from each point in time, but the optimal sequence of actions might be different in the future timesteps – leading to time-inconsistencies.

Repeated delays from temporal inconsistencies lead to procrastination when efforts are discounted more steeply than rewards Say $\gamma_r = 0.9$ and $\gamma_c = 0.5$, i.e., steeper discounting for future efforts than rewards. This leads to time-inconsistent policies. Figure 4A shows deterministic policies (without softmax noise) for remaining timesteps at each horizon for $s = 0$, when no work has been completed yet. At $t = 0$ or horizon = 15, the best policy is to begin at $t = 2$, but by $t = 2$ or horizon = 13, it becomes rational to delay until $t = 4$ and so on. Hence, repeated delays dictate that one only starts at the end of the horizon. In some sense, there is a constant underestimation of how much effort one’s future self will feel itself exerting to do the task. Corresponding trajectories for this parameter setting are shown in yellow in Figure 4B.

With smaller discrepancies between the two discount factors, defections are less pronounced, with delays not necessarily to the very end as shown in the violet and blue trajectories in Figure 4B. In comparison, when $\gamma_r = \gamma_c$, it is optimal to finish as soon as possible (shown in red), given the immediate rewards. The range of trajectory types captured here can account for early completion (clusters 1-4), intermediate delays (cluster 6), and deadline completion (clusters 7, 8).

Waiting for interesting tasks

In the final model, we explore a route to delaying corresponding to the third type in our taxonomy. Here, one declines to commit to a time of action at the very start. This doesn’t stem from discounting of future rewards. Some students indicated they waited and signed up for studies they found interesting. We introduce a probabilistic process to model the uncertain appearance of (relatively rare) interesting tasks. For simplicity, say there is a low-reward state where rewards for completing work comes at threshold as before, and a high-reward state, where in addition, immediate ‘interest’ rewards are available for completing interesting work. Alternatively, high-reward states can be seen as high motivational states that might be rare for some students. Transitions between the two states are governed by probabilities ($P(H|L)$ and $P(L|H)$), in-

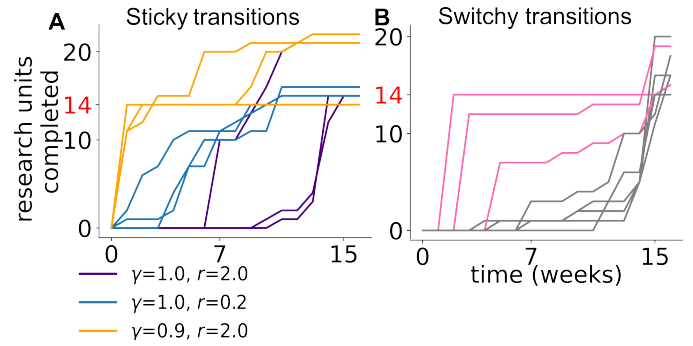


Figure 5: Delays from waiting for larger rewards from interesting tasks. A. Interesting tasks that hold a higher reward are rare. When $\gamma = 1$, it is worth waiting for them (in purple) but with $\gamma < 1$, it is better to work immediately (in yellow). B. When interesting tasks also disappear quickly (switchy transitions), there can be gaps between successive bouts of working as in the pink trajectories

dependent of the progress and actions of the agent. The agent can work in either state to progress with binomial probability as before. Probability of the high reward state coming up is low ($P(H|L) = 0.05$). We set $\eta = 0.5$.

Waiting for interesting rewards leads to delays in the absence of discounting

With $\gamma = 1$, no interesting tasks initially (starting in low reward state), and sufficiently interesting rewards ($r_{\text{interest}} = 2$), it is optimal to wait for these rare tasks, thereby delaying work until they arise. Higher the efficacy, the longer one can afford to wait. This leads to patterns where the agent works a lot at an intermediate point when higher rewards come up or waits for higher rewards without success, then works towards the end, replicating patterns in clusters 7 and 8 (Figure 5A). With low $r_{\text{interest}} = 0.2$, it is not worth waiting anymore, eliminating delays, and most work occurs at the beginning, resembling clusters 1-3.

Introducing delay discounting (here, $\gamma = 0.9$) actually eliminates such delays. This is because it looks more attractive to seize the smaller rewards currently available instead of waiting for higher rewards that have a low chance of appearing in the future. This may be an explanation for why some procrastinators actually have low discounting (P. Zhang & Ma, 2023), which facilitates waiting in contrast to earlier explanations where delays are associated with higher discounting.

Short-lived interesting tasks lead to logit-shaped trajectories

In previous simulations, we set $P(L|H) = 0.05$, meaning high reward opportunities lasted a long time once they come up. However, interesting tasks might quickly deplete, setting $P(L|H) = 0.95$. We set $\gamma = 1.0$, $r = 2.0$. Now, not all units are completed during the brief period of high rewards. Consequently, the agent must wait until they appear again, leading to a logit-shaped pattern with gaps between bouts of significant progress, like pink trajectories in Figure 5B. These could account for bimodal completion patterns in cluster 6.

Discussion

Procrastination is a heterogeneous construct associated with a variety of factors and correlates. In this work, we categorized procrastination types based on delay structures and proposed a classification of underlying mechanisms. We simulated completion patterns using three model types: exponential temporal discounting, non-exponential discounting and waiting for interesting tasks, and found that they qualitatively reproduce real-world procrastination patterns. This suggests plausibility of multiple procrastination types and pathways.

In contrast to earlier attempts of categorisation which relied on finding factors from questionnaire data, we constructed a taxonomy based on the nature of commitment or adherence to a delay and the manner in which it leads to a problem. In addition to the conventional meaning of contradicting one's own intentions, our broader definition of procrastination includes adhering to decisions to postpone and delays from not committing to a timing of action, even when intending to act. These decisions can be irrational due to miscalculations or gaps in knowledge, but can also be rational, and still lead to issues like missing deadlines, for instance due to discounting of delayed rewards or waiting for more interesting tasks. Indeed, a few students in the task failed to meet the deadline. While the rest could finish, our taxonomy can still be applied to the variety of ways in which people allocate work in time.

Prior mechanistic explanations primarily focused on discounting effects, including subjective reduction in value with delay (Fischer, 2001; Steel, 2007; S. Zhang, Liu, & Feng, 2019) and time-inconsistent decisions from non-exponential discounts (Fischer, 1999; Steel, 2007). P. Zhang (2024) utilized a controlled reading comprehension experiment to model work allocation by incorporating discounting and assuming that participants maintain a dynamic estimate of future work completion. We used a similar formalism but model a different dataset from P. Zhang and Ma (2023). We demonstrated that various mechanisms, not limited to discounting, can explain completion patterns. Going beyond past work, we explored the role of defections and the hitherto overlooked element of uncertainty about timing of interesting tasks in mediating delays. Future research should consider other explanations like forgetting, or defections from unexpected changes like appearance of other work or learning that it is unexpectedly effortful. There can also be other reasons to work earlier than later including the mental burden of unfinished tasks, a preference for leisure after completing tasks, or anticipation of reduced free time later in the semester.

The relationship between decision-theoretic mechanisms we explored here and psychological theories, motivational and cognitive factors associated with procrastination is intriguing. For instance, impulsivity, often linked to procrastination may arise from discounting future rewards, driving distractibility. Oppositely, discipline and perseverance on tasks (which procrastinators lack) may stem from internal costs of unfinished tasks, or devices to counteract preference reversals. Organisational skills might relate to long-

term planning, effective and efficient task prioritisation and time management, some of which is captured by the efficacy parameter in our models. Motivationally, low energy could relate to generally low task rewards which drive procrastination in all proposed mechanisms. Intrinsic motivation might be rewards associated with learning something new. Here, this might be from tasks students find interesting or from feelings of accomplishment or satisfaction in contributing to ongoing research. Extrinsic motivation comes from external rewards like final course grade or future academic or career goals (Rebetez, Rochat, & Van der Linden, 2015).

We qualitatively compared simulation results with empirical completion patterns. While explicit model fitting to data is a crucial avenue for future work, our replication of patterns are useful for highlighting the important components for each pattern type and could already show that multiple explanations are plausible for them. This is an essential initial step in understanding the diverse processes involved in procrastination. However, there are also some caveats to model fitting, as we had to make simplifying assumptions due to a lack of information on factors like distribution of interesting vs boring tasks, availability of tasks in time and the stochastic process underlying completion, which may impact the informativeness of exact fits. Secondly, we acknowledge that convexity and efficacy have somewhat overlapping effects in controlling the average rate of completion, although the former is from vigour costs and latter from probabilistic failure.

Since multiple causes could possibly lead to each of the patterns, how can we disambiguate between the models? Targeted experimental manipulations in the future could provide clarity. To address reward schedule ambiguity, rewards can be delivered explicitly at specific times. For example, immediate rewards would eliminate delays due to discounting and would shift any observed delays to be attributed to alternative mechanisms, such as defections from inconsistent discounts or waiting for better options. To further distinguish these, we could conduct an experiment with uniformly boring tasks, ruling out the issue of waiting for interesting tasks. P. Zhang (2024)'s controlled experiments have both these features and would be interesting targets for future work. Finally, the differential discount factors relies on relatively low reward sizes. Varying reward sizes to check if delays persist at higher rewards, could indicate the involvement of other mechanisms.

As noted earlier, people procrastinate under a variety of task settings, even when there is no real semester-like deadline. Examples include making an appointment with a doctor or going to gym. It would be interesting to examine these. Additionally, literature typically focuses on easy tasks, but procrastination is also linked to feelings of inadequacy, fear of failure, and anxiety (Lay, Knish, & Zanatta, 1992; Pychyl & Flett, 2012), which are more likely to occur in challenging tasks like preparing for a public talk or a tough exam. Therefore, future work should explore work allocation in other types of tasks, which might allow discovery of other mechanisms for procrastination and work allocation in general.

Acknowledgments

SC and PD are funded by the Max Planck Society. SC is also supported by the International Max Planck Research School MMFD (IMPRS-MMFD). PD is additionally funded by the Humboldt foundation. We would also like to thank the anonymous reviewers for their helpful comments on an earlier version of the manuscript.

References

- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience*, 8, 429–453.
- Fischer, C. (1999). *Read This Paper Even Later: Procrastination with Time-Inconsistent Preferences* (Tech. Rep.). Resources for the Future.
- Fischer, C. (2001). Read this paper later: procrastination with time-consistent preferences. *Journal of Economic Behavior & Organization*, 46(3), 249–269.
- Gevers, J. M., Rutte, C. G., & Van Eerde, W. (2006). Meeting Deadlines in Work Groups: Implicit and Explicit Mechanisms. *Applied Psychology*, 55(1), 52–72.
- Grunschel, C., Patrzek, J., & Fries, S. (2013). Exploring reasons and consequences of academic procrastination: An interview study. *European Journal of Psychology of Education*, 28(3).
- Konradt, U., Ellwart, T., & Gevers, J. (2021). Wasting effort or wasting time? A longitudinal study of pacing styles as a predictor of academic performance. *Learning and Individual Differences*, 88, 102003.
- Lay, C. H. (1986). At last, my research article on procrastination. *Journal of Research in Personality*, 20, 474–495.
- Lay, C. H., Knish, S., & Zanatta, R. (1992). Self-handicappers and procrastinators: A comparison of their practice behavior prior to an evaluation. *Journal of Research in Personality*, 26(3), 242–257.
- Le Bouc, R., & Pessiglione, M. (2022). A neurocomputational account of procrastination behavior. *Nature Communications*, 13(1), 5639.
- McCown, W., Johnson, J., & Petzel, T. (1989). Procrastination, a principal components analysis. *Personality and Individual Differences*, 10(2), 197–202.
- Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology*, 191(3), 507–520.
- O'Brien, W. K. (2002). *Applying the transtheoretical model to academic procrastination* (Vol. 62). Unpublished doctoral dissertation, ProQuest Information & Learning, US.
- O'Donoghue, T., & Rabin, M. (2001). Choice and Procrastination. *The Quarterly Journal of Economics*, 116(1), 121–160. (Publisher: Oxford University Press)
- O'Donoghue, T., & Rabin, M. (1998). Procrastination in preparing for retirement.
- Pychyl, T. A., & Flett, G. L. (2012). Procrastination and Self-Regulatory Failure: An Introduction to the Special Issue. *Journal of Rational-Emotive & Cognitive-Behavior Therapy*, 30(4), 203–212.
- Rebetez, M. M. L., Rochat, L., & Van der Linden, M. (2015). Cognitive, emotional, and motivational factors related to procrastination: A cluster analytic approach. *Personality and Individual Differences*, 76, 1–6.
- Rozental, A., & Carlbring, P. (2014). Understanding and Treating Procrastination: A Review of a Common Self-Regulatory Failure. *Psychology*, 5(13), 1488–1502.
- Schouwenburg, H. C. (2004). *Procrastination in Academic Settings: General Introduction*. Washington, DC, US: American Psychological Association.
- Shu, S. B., & Gneezy, A. (2010). Procrastination of Enjoyable Experiences. *Journal of Marketing Research*, 47(5), 933–944.
- Sirois, F. M. (2007). "I'll look after my health, later": A replication and extension of the procrastination-health model with community-dwelling adults. *Personality and Individual Differences*, 43(1), 15–26.
- Solomon, L. J., & Rothblum, E. D. (1984). Academic Procrastination: Frequency and Cognitive-Behavioral Correlates. *Journal of Counseling Psychology*.
- Steel, P. (2007). The nature of procrastination: a meta-analytic and theoretical review of quintessential self-regulatory failure. *Psychological Bulletin*, 133(1), 65–94.
- Steel, P. (2010). Arousal, avoidant and decisional procrastinators: Do they exist? *Personality and Individual Differences*, 48(8), 926–934.
- Steel, P., & Klingsieck, K. B. (2016). Academic Procrastination: Psychological Antecedents Revisited. *Australian Psychologist*, 51(1), 36–46.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning, second edition: An Introduction*. Cambridge, MA: MIT Press.
- Zhang, P. (2024). *The dynamic nature of procrastination*. Doctoral dissertation, Department of Neural Science New York University, New York.
- Zhang, P., & Ma, W. J. (2023). Temporal discounting predicts procrastination in a real-world task. *PsyArXiv*.
- Zhang, S., Liu, P., & Feng, T. (2019). To do it now or later: The cognitive mechanisms and neural substrates underlying procrastination. *WIREs Cognitive Science*, 10(4).