UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Now or later: A reinforcement learning model of behavioural delay

Permalink
https://escholarship.org/uc/item/1r09647d

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

Authors
Chebolu, Sahiti
Zhang, Peiyuan
Ma, Wei Ji
et al.

Publication Date
2026

Copyright Information
This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Now or later: A reinforcement learning model of behavioural delay

Sahiti Chebolu (sahiti.chebolu@tuebingen.mpg.de)[1, 2]**, Pei Yuan Zhang**[3]**, Wei Ji Ma**[3, 4] **& Peter Dayan**[1, 2]
[1]Max Planck Institute for Biological Cybernetics, Tuebingen
[2]University of Tuebingen, Tuebingen
[3]Department of Psychology, New York University, New York
[4]Center for Neural Science, New York University, New York

Abstract

(Other) people notoriously delay the initiation and completion of work, sometimes beyond what is optimal. One prominent, and indeed experimentally validated, explanation is based on the fact that rewards delivered in the future are discounted. However, other factors can interact with discounting and affect policies, such as the amount of effort and the probability of successful completion. These have received less empirical attention. Here, we build and fit a new reinforcement learning model to the working trajectories of students over the course of a semester in a real-world task (P. Y. Zhang & Ma, 2024). We show that discounting, effort, and efficacy are all important in explaining students' delays. In addition, the discount factors inferred from task performance correlate significantly with self-reported measures of impulsivity and procrastination, as well as discount rates estimated from a monetary delay discounting task, highlighting that they robustly capture meaningful individual differences in temporal preferences.

Keywords: delay discounting; computational modelling; behavioural delay; procrastination; cognitive effort

Introduction

In daily life, people often have the option of deciding how to distribute their work over time on important tasks such as completing an assignment, finishing a project, or filing taxes. In particular, they choose whether to do them *now or later*. People often delay working on such tasks, and various factors are believed to influence this frequently maladaptive decision.

Delay discounting has attracted special attention (Green et al., 1994; Loewenstein & Elster, 1992). It is the subjective reduction in the value of delayed rewards compared to immediate ones. An implication of discounting is that in (the many) tasks where rewards are only delivered after a deadline, people may be less motivated to engage early for this diminished reward, hence encouraging them to delay in favour of immediate rewards. It is also proposed as an explanation for procrastination (Steel, 2007). A few studies have tested this hypothesis. For instance, the amount of work people complete over time follows a hyperbolic curve, mirroring the shape of the discounting curve (König & Kleinmann, 2005; Steel et al., 2018). Discount rates explicitly estimated from delay discounting tasks, where people have to choose between smaller sooner and larger later monetary rewards, correlate with the extent of delay people display in working for delayed benefits (Reuben et al., 2015; P. Y. Zhang & Ma, 2024).

However, several other factors could interact with discounting, affecting the delay, and it is therefore critical to understand their role. One is the effort cost of working. Discounting implies that one would avoid exerting immediate effort for distant rewards, leading to delays in work (Akerlof, 1991; Le Bouc & Pessiglione, 2022; S. Zhang et al., 2019). Although physical and cognitive effort have been shown to influence how much work humans and animals are willing to perform for a reward and how much time they allocate to a task (Dayan, 2012; Kool & Botvinick, 2018; Müller & Apps, 2019; Niv et al., 2005; Niyogi et al., 2014), the interaction between effort and discounting in directly shaping work allocation over time has not been empirically tested. The subjective probability of successfully completing work (efficacy) also plays a role in deciding whether it is worth doing a task and when it should be done. The greater the probability, the larger the expectation of future reward, and hence the greater the motivation to work earlier (Steel & Weinhardt, 2018). However, if completion before a deadline is mandatory, then the same factor permits a later start, since repetition is less likely to be necessary (Dayan, 2012). Here too, the effect of efficacy on task-based measures of delay has not been studied.

To bridge these gaps, we model behavioural delay in a dataset published by P. Y. Zhang and Ma (2024). It reports when students completed certain requirements for a university course, providing a rich, realistic setting to investigate the effects of the above factors on delay. The original study showed that discount rates (k^{DD}) estimated from a delay discounting task correlated with the extent to which students delayed work. Here, we complement their analysis by fitting a reinforcement learning model consisting of an exponential discount factor (γ), efficacy (η), and effort costs (r_{effort}) to students' trajectories of work to study how these parameters interact to shape students' delays. Firstly, we hypothesise that:

1. the fitted values $\gamma, r_{\text{effort}}, \eta$ are all significantly related to the extent of behavioural delay in the task; and particularly that r_{effort}, η add predictive power above γ in predicting delay

Furthermore, since these factors encompass entire trajectories of work by capturing specific cognitive processes, they provide a richer substrate for relating the real-world task to the scale-based measures and traits of the students (Patzelt et al., 2018). Specifically, we hypothesise that:

2. in addition to predicting behavioural delay, all three parameters also predict students' self-reported tendency to procrastinate, and again, η, r_{effort} improve the correlation

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society* ©2026 the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

3. the fitted γ is significantly related to k^{DD} and scale-based measures of impulsivity and self-control, as it captures people's general tendency to discount future rewards; additionally, it is a more sensitive measure, predicting impulsivity and self-control better than k^{DD}

Methods

Data and code for all analyses are available at https://github.com/SahitiC/delay_cogsci2026

Data

We analysed data from a real-world task published by P. Y. Zhang and Ma (2024). In this task, 194 bachelor students in a psychology course were required to participate in at least 7 hours of experiments over a semester (110 days or about 16 weeks) to pass the course. Each research study took at least 0.5 hours or a multiple thereof, hence defining 0.5 hours as a unit of work. Additional participation beyond 7 hours (14 units) contributed $1/8^{\text{th}}$ of a grade point per unit, up to 4 extra hours. There were more than enough research opportunities, with an average of 15 potential hours of experiments per student.

The published dataset is rich, consisting of each student's working trajectory (i.e., the number of hours of experiments completed each day of the semester), hyperbolic discount rates k^{DD} estimated from a monetary delay discounting paradigm, and questionnaire measures such as the Procrastination Assessment Scale for Students (PASS), Barratt Impulsivity Scale, and Brief Self-Control Scale, among other measurements. In this paper, we excluded students who did no work, those who expressly stated that they dropped out of the requirement, and those who completed more than 22 units, as there was no clear external incentive to do so. This left 160 subjects, all but 3 of whom completed at least 14 units.

Computational model

We model the delays and temporal distribution of work in this task using a Markov decision process (MDP), inspired by Chebolu and Dayan (2024, 2025) and P. Zhang and Ma (2019). The states in the MDP comprise the current time point in weeks ($0 \leq t \leq 15$) and the number of units of work completed so far ($0 \leq s_t \leq 22$). Each week, the individual decides to complete some number of units ($0 \leq a_t \leq 22 - s_t$). Of these, the number of successfully completed units is determined by a Binomial success probability called efficacy (η). The probability of completing $\Delta_{t+1} = s_{t+1} - s_t$ units is given by $P(s_{t+1} \mid s_t, a_t) = \binom{a_t}{\Delta_{t+1}} \eta^{\Delta_{t+1}} (1-\eta)^{a_t - \Delta_{t+1}}$. A number of factors affect efficacy, such as how well individuals can organise their time, uncertainty about the time other commitments will occupy, or lack of desirable tasks in the week. Higher efficacy means more work can be done in a week. There is also an effort cost r_{effort} incurred per unit of work which scales linearly with the amount of work attempted. At the end of the semester, the individual is rewarded with $r_{\text{comp}} = 4$ for every compulsory unit completed, provided a minimum of 14 units is finished. If additional optional work is done, there is a reward of $r_{\text{opt}} = 0.5$ up to 22 units. Individuals weight delayed rewards with an exponential discount factor $\gamma < 1$. We assume the same γ for rewards and efforts. The individual's goal is to maximise the net discounted sum of expected rewards over the semester. We derive the maximising decision policy using dynamic programming. Action selection follows a softmax rule with inverse temperature $\omega = 5$, leading to noisy decisions. We fixed ω because it trades-off with r_{effort} and γ, making the parameters unrecoverable from fits. The final policy specifies the amount of work an individual intends to do each week.

We fit this model to each student's trajectory to estimate the free parameters $(\gamma, \eta, r_{\text{effort}})$. We obtained maximum likelihood estimates (MLE) of the parameters. Since only states s_t are observed but not actions a_t, we derive likelihood $\mathcal{L}$ by marginalising over a_t. Hence $\mathcal{L} = \prod_{t=0}^{T-1} \left[\sum_{a_t} P(s_{t+1}|s_t, a_t, \mathbf{x})\, P(a_t|s_t, t; \mathbf{x}) \right]$ where $\mathbf{x}$ are the parameters. We implemented this procedure in `python` using `scipy.optimize`. We also attempted hierarchical Bayesian model fitting, but there was not enough data per participant for reliable estimates.

Error regression model

Following model fitting, we aim to regress the fitted parameters with the self-reported measures. The estimates of model parameters ($\hat{\mathbf{x}}_i = \{\hat{x}_{ij}\}$ for the parameters j of student i) from single trajectories are inevitably uncertain. We account for this in our regression analyses (for variables y_i such as PASS) by treating model parameters as noisy predictors, rather than using simple linear regression which assumes exact predictors. The resulting model is:

$$y_i \sim \mathcal{N}\left(\beta_0 + \sum_{j=1}^{p} \beta_j x_{ij},\ \sigma_y^2\right), \tag{1}$$

$$\hat{x}_{ij} \sim \mathcal{N}\left(x_{ij},\ \sigma_{ij}^2\right), \quad \text{where} \quad x_{ij} \sim \text{Uniform}(a_j, b_j) \tag{2}$$

where β_0 is the intercept, β_j is the regression coefficient for predictor j, x_{ij} is the latent value of the parameter that we observe noisily as the MLE estimates ($\hat{x}_{ij}$), σ_y^2 is the variance of y_i, σ_{ij}^2 is the variance of $\hat{x}_{ij}$ (estimated from the inverse Hessian of the MLE fit for x_{ij}) and x_{ij} is uniformly distributed with limits $[0, 1]$ for γ, η, and $[-4, 0]$ for r_{effort}. Hence, instead of regressing y_{ij} against a fixed value of x_{ij} to estimate $\boldsymbol{\beta}_i = \{\beta_0, \boldsymbol{\beta}, \sigma_y^2\}_i$, we marginalise over a range of possible values for x_{ij} with the width of the distribution determined by the variance and maximise $P(\mathbf{y} \mid \boldsymbol{\beta}_i) = \prod_{i=1}^{n} \int P(y_i \mid \boldsymbol{\beta}, \mathbf{x}_i) \prod_{j=1}^{p} P\left(\hat{x}_{ij} \mid x_{ij}, \sigma_{ij}^2\right) P\left(x_{ij}\right) \mathrm{d}x_{ij}$. Finally, to test the significance of the regression coefficients, we perform a likelihood ratio test (LRT), comparing this model to a null restricted model, where the coefficient of interest is set to 0. For these regressions, we excluded 7 subjects for which `scipy.optimize` failed to find a local minimum.

Canonical correlation analysis (CCA)

To assess the relationship between questionnaire measures (such as PASS, impulsivity, self-control etc.) on the one hand

and parameter estimates on the other, we performed CCA using `scikit-learn` in `python`. CCA finds linear combinations of variables from each set that maximally correlate (called canonical variables), revealing common dimensions between sets. The loading of a variable in a set on a canonical variable reveals how much it contributes to that dimension.

Results

Model agnostic correlations

In the original study, P. Y. Zhang and Ma (2024) examined associations between students' task delays, monetary delay discount rates k^{DD}, and self-reported procrastination scores (PASS). They used mean unit completion day (MUCD) as the behavioural measure of delay. Here, we re-express their results in mean unit completion week (MUCW). To avoid confounds in the delay measure from varying workloads, they restricted their analysis to students who completed only the minimum requirement of 14 units (N=93). Replicating their findings, there are significant correlations between MUCW & PASS ($r = 0.38$, $p < 0.001$) and MUCW & k^{DD} ($r = 0.30$, $p = 0.006$), but not between k^{DD} & PASS (Table 1).

We repeated this correlational analysis on a broader sample (N=160). For those who finished more than 14 units, we calculated MUCW up to when 14 units are completed. With our larger sample, we found significant correlations between all three quantities, as summarised in Table 1. Hence, there is a positive association between students' extent of delay on the task, their self-reported general tendency to procrastinate, and how much they discount future monetary rewards.

Table 1: Pearson correlations between procrastination score (PASS), behavioural delay (MUCW), and empirical discount rates from a sample of N=93 (used in the original study) and a broader sample of N=160 students.

	N = 93		N = 160	
Correlation	r	p	r	p
MUCW ~ PASS	0.38	0.00016	0.29	0.0002
MUCW ~ k^{DD}	0.30	0.006	0.18	0.029
k^{DD} ~ PASS	0.18	0.096	0.20	0.014

Model simulations

Figure 1A shows the cumulative number of units of work done over the weeks in the semester ('trajectory of work') for a few example students. Notably, there is a diversity of patterns in work distribution.

To demonstrate the range of behaviour our model can capture and the effect of the various parameters, we simulate example trajectories from our generative model while systematically varying the parameters of interest (Figure 1B-D). We also show the mean (in bold) over 1000 simulated trajectories. In Figure 1B, we vary the discount factor γ. We fix efficacy $\eta = 0.9$ and effort $r_{\text{effort}} = -0.3$. When γ is high, future rewards are not as strongly devalued relative to immediate

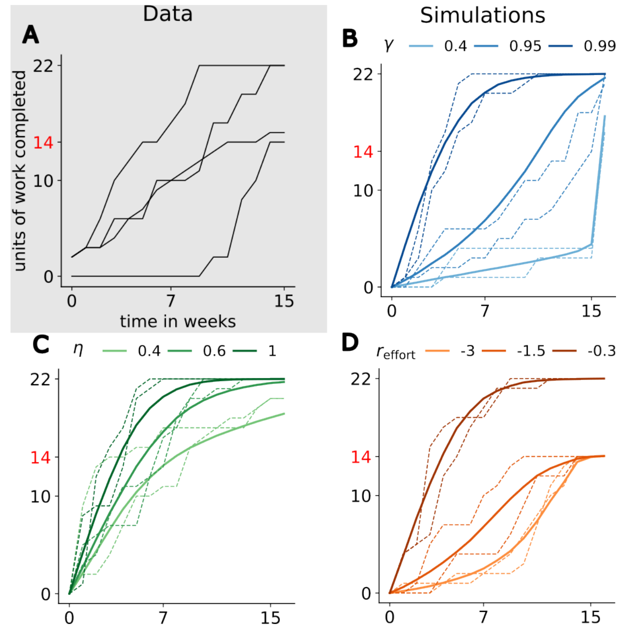

Figure 1: Trajectories of work in the real data (A) and from simulations (B-D). Means (of 1000) simulated trajectories are in bold and sample trajectories are shown as dashed lines in B-D. The number of compulsory units to be done is 14 (marked in red). A. Example trajectories from the data from P. Y. Zhang and Ma (2024). B. As the discount factor (γ) reduces, work is delayed to later in the semester due to a preference for immediate rewards (default $\eta = 0.9; r_{\text{effort}} = -0.3$). C. As efficacy ($\eta$) decreases, more time is taken to complete a unit of work (default $\gamma = 0.99$). D. As the cost of working (r_{effort}) becomes more negative, delay ensues.

rewards. This means it is better to work early to ensure completion and, since $\eta < 1$, to avoid failing. Hence, most work is completed early in the semester (dark blue). As γ decreases, there is a preference to avoid exerting immediate effort for distant rewards from completing the task. This is apparent in the simulated trajectories where more work is done towards the end as γ reduces (in medium and light blue).

As efficacy η decreases, more time is actually required to complete a unit of work. Hence, for a fixed discount and effort ($\gamma = 0.99$, $r_{\text{effort}} = -0.3$), work is completed later for a lower η (Figure 1C). Finally, for fixed $\gamma = 0.99$ & $\eta = 0.9$, delays are larger for more effortful work (Figure 1D). When the effort cost is very high, it is no longer worth the effort to complete optional units (medium and light orange). In sum, the model produces a wide variety of patterns qualitatively similar to those observed in the data. In the next section, we assess how well the model fits real behaviour quantitatively.

Model fits

We fit the model to each student trajectory by maximising the log likelihood of the observed data under the model. From the model fits, we obtain estimates of γ, η and r_{effort} for each

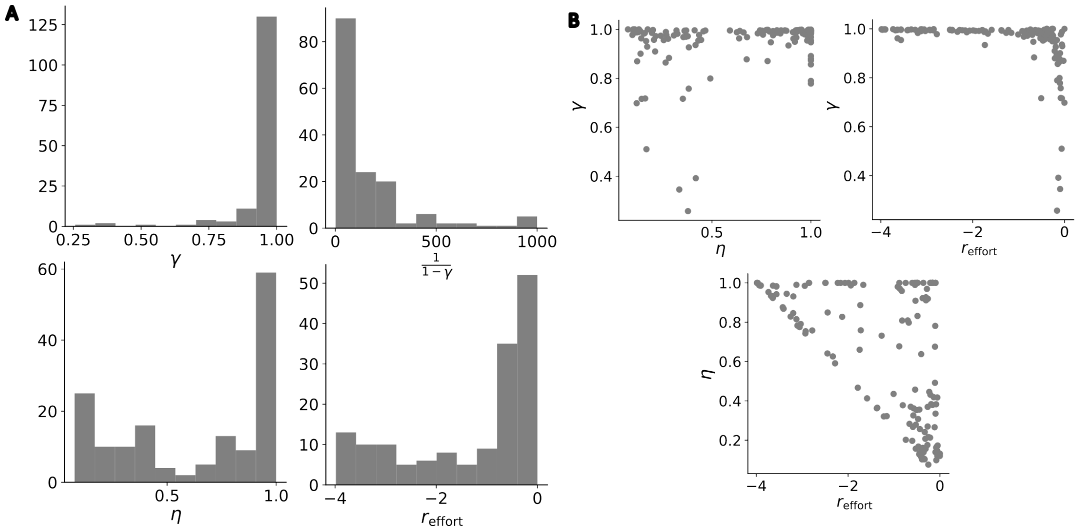

Figure 2: A. Distribution of fitted parameters to the data. B. Pairwise scatter plots showing a lack of correlation between the parameters.

student. The distribution of the fitted parameters in the population is shown in Figure 2A. Discount factors γ are concentrated around 0.8 - 1, with only a few below this range. This apparently narrow distribution still represents a range of temporal preferences as shown by the spread of the effective scale of discounting ($\frac{1}{1-\gamma}$) (Figure 2A). Larger γ implies a longer time horizon. Efficacy η spans the entire range from 0.0 - 1.0 and appears to be bimodal, with one mode close to 0.2 and another close to 1.0. Finally, effort r_{effort} ranges from -4.0 to 0, with a concentration around -1 to 0. The values do not go below this range, as it is no longer worthwhile to work below $r_{\text{effort}} = -4$ (given that $r_{\text{comp}} = 4$). Finally, there is no correlation between the fitted parameters as shown in the pairwise scatter plots (Figure 2B).

We also fit a simpler model with only γ as a free parameter, and with η and r_{effort} fixed to the mean estimates from the full model fits. A model comparison using the Bayesian Information Criterion (BIC) favoured the full three-parameter model (BIC = 5205) over the simpler model (BIC = 6004), indicating that the three-parameter model explains the data better than a discounting-only model.

We ran a recovery analysis on the fitted parameters to ensure that our model fitting procedure can reliably recover the parameters that generated the data. Mirroring the real data, we generated a trajectory for each set of fitted parameters. We then fitted the full delay model to this simulated data, finding a moderately reliable correlation between the recovered parameters and the input parameters. The Pearson correlation for γ is 0.48; for η is 0.63; and for r_{effort} is 0.65.

As a posterior predictive check, we plotted the MUCW computed from the empirical data as a function of each parameter and compared it with the MUCW obtained from the simulated data evaluated over the same parameter values. We calculated the average MUCW from 5 simulated trajectories for each of the fitted parameter sets. As Figure 3A & B show, the simulations closely reproduce the patterns in the original data. Finally, we plotted the total number of units completed and the completion week (for 14 units) for each student over their fitted parameter estimates in the 3D parameter space in Figure 3C & D. These statistics separate well in this space: those with low magnitude of r_{effort}, high η, and high γ complete more units and take the least time to finish. When the magnitude of r_{effort} increases or γ decreases, it is no longer worth doing all 22 units, and work is more delayed.

Statistics

Contribution of model parameters to explaining delay and effort distribution We showed that the three-parameter model is a better fit to the data than a simpler model with only discounting. Further, the fitted three-parameter model also recapitulates key patterns in the data. Next, we test whether all three fitted parameters within the full model are statistically important for explaining real delays in the task. We regressed MUCW with the parameters as predictors.[1] To account for noise in the fitted parameters, we built an error regression model, as described in the Methods. We found that all three fitted parameters are significantly correlated with MUCW (Table 2), with smaller γ, lower η, and larger magnitude of r_{effort}, all predicting more delay. To test whether η and r_{effort} add significant explanatory power over γ, we fitted a restricted model in which the regression coefficients for η and r_{effort} were fixed to zero. A likelihood ratio test revealed a significant difference between the two models, indicating that η and r_{effort} are necessary for explaining some of the variance in working patterns (Table 2). These results confirm our hypothesis that effort costs and success probability are important factors in determining people's delay in the real world.

Regression of procrastination score with fitted parameters Since each of the parameters affects behavioural delay, we also expected all three fitted parameters to have an association with students' general, self-reported tendency to procrastinate on tasks, measured by the PASS scale. However, we found that

[1]Note that the variables in the regression are not normalised, and so coefficients are scaled based on the relative sizes of the variables.

Table 2: Regressions with all three predictors (γ, η, r_{effort}) (a), and comparison to a regression with only γ as a predictor (b)

(a) Regression coefficients

	γ		η		r_{effort}	
Dependent variable	β	p	β	p	β	p
MUCW	-6.60	0.0004	-2.45	0.0009	-0.87	3.83e-5
PASS	-1.48	0.008	-0.11	0.644	0.02	0.787

(b) Likelihood ratio tests

Dependent variable	p
MUCW	0.0001
PASS	0.715

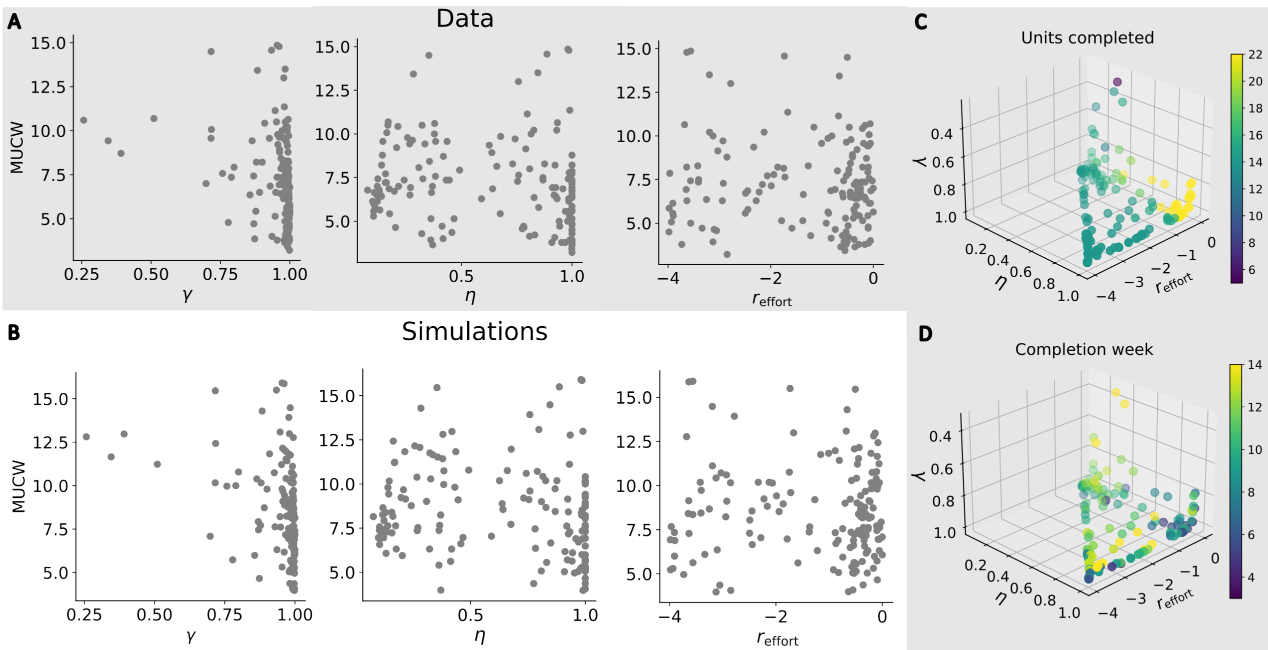

Figure 3: Mean unit completion week (MUCW) plotted against fitted parameters for A. the real data and B. for trajectories simulated using the same parameters. C. Total number of units completed and D. Completion week (for 14 units) for each student over the fitted parameters.

while γ is significantly correlated with procrastination scores, η and r_{effort} are not (Table 2). Furthermore, a likelihood ratio test revealed no significant difference between the full model and a restricted model with only γ (Table 2).

Regressions of other scale measures with fitted parameters Finally, we tested the association between the fitted parameter estimates and conceptually related variables measured explicitly or via self-reports. As before, we performed error regressions. The results are summarised in Table 3.

Firstly, we hypothesised that students' tendencies to discount future rewards would be correlated between the real-world task and the delay discounting task. Indeed, γ and k^{DD} are significantly correlated ($\beta = -0.07$, $p = 0.031$). The negative association is expected, as larger discount rates (and smaller exponential discount factors) imply a greater impatience for rewards.

Next, we hypothesised that those who discount the delayed benefits in the research participation task (low γ) might also be more impulsive (as measured by the Barratt Impulsivity Scale). To enable comparison with k^{DD}, we also included them as a predictor. We found a strongly significant correlation between the two (Table 3, $\beta = -0.88$, $p = 0.001$), that was not present for k^{DD} (in either the original study or with our larger sample size). This is some evidence that the discount factors estimated from model fits might capture temporal preferences more sensitively.

We also hypothesised that those who are more patient for delayed rewards would have a better ability to exert self-control by resisting temptation and maintaining self-discipline. However, we did not find significant correlations between the self-control score (measured by the Brief self control scale) and γ or k^{DD} (Table 3).

Efficacy η captures how much work can be done in a week, which is influenced by time organisation, competing demands, or lack of desirable tasks. Effort costs r_{effort} reflect the cognitive effort required to engage in the task or feelings of aversiveness, such as boredom. The dataset did not include specialised scales that directly assessed these constructs for this specific task and with a clear correspondence to the parameters. However, as an exploratory analysis, we tested the relationships

Regression formula	β	p
$k^{DD} \sim \gamma + 1$	-0.07	0.031
time management $\sim \eta + 1$	-0.33	0.18
task aversiveness $\sim r_{\text{effort}} + 1$	0.02	0.75

	γ		k^{DD}	
Dependent variable	β	p	β	p
impulsivity	-0.88	0.001	-0.08	0.921
self-control	0.95	0.075	-1.71	0.286

Table 3: Regressions between scale measures and fitted parameters

between η and r_{effort} and two predefined subsets of PASS items measuring time management and task aversiveness as reasons for students' general procrastination, respectively, and found no significant correlations in either case.

Finally, we applied canonical correlation analysis (CCA) to identify shared dimensions between the questionnaire measures and the fitted parameters. Specifically, we aimed to test if PASS, impulsivity, and k^{DD} covary with γ, forming a shared dimension. We found one significant canonical correlation ($r = 0.38$, $p = 5.45\text{e-}06$). The canonical loadings of the individual variables from the two sets are in Table 4. PASS, impulsivity, and the fitted γ contribute most strongly to the canonical variate, although other variables also contribute.

Table 4: Canonical loadings of scale measures and parameter fits on the significant pair of canonical variates from CCA.

(a) Questionnaire measures

Variable	Loading
PASS	0.68
k^{DD}	0.44
impulsivity	0.70
self-control	-0.40
time management	-0.49
task aversiveness	-0.002

(b) Fitted parameters

Variable	Loading
γ	-0.94
η	-0.27
r_{effort}	0.45

Discussion

Various factors affect how much people delay working for distant rewards. We introduced a mechanistic model of delay involving discounting of future rewards, a probability of successful completion (efficacy), and effort cost of working. We showed that in addition to discounting, efficacy and effort costs provide significant explanatory power for students' delays in a real-world task. Despite their theoretical importance, the role of these variables in shaping how people distribute work over weeks has not been empirically tested previously. This analysis paves the way for future experiments where these interacting parameters can be explicitly measured and tested.

We also related fitted parameters to self-reported measures in the data. We showed that students' discount rates are correlated across the real-world task and the monetary delay discounting task, suggesting that temporal preferences are shared across domains. This is in line with studies showing that delay discounting has trait-like stability (Odum, 2011). In addition, the fitted γ was significantly associated with students' impulsivity scores, whereas discount rates from the delay discounting task were not, indicating greater sensitivity of the fitted delay model to behavioural variation and that γ captures a common computational mechanism driving students' temporal decisions. Previous studies have shown that impatience for monetary rewards is often correlated with impulsivity (Bjork et al., 2004; Burghoorn et al., 2026; Keidel et al., 2024) (although some have also shown no correlation (Keidel et al., 2021)). Extending this literature, our findings demonstrate that impatience in a different domain, namely working for delayed benefits, is also related to impulsivity.

There are mixed results on the link between lack of self-control and temporal discounting (Duckworth & Kern, 2011). We did not find a relation in this dataset. Although we did find an association with impulsivity, it is important to note that impulsivity and self-control are related but distinct constructs.

An unexpected result was that while all parameters correlated with students' delays, only discount factors were related to procrastination (PASS) scores. One possibility is that delays arising either because tasks take longer to complete (low efficacy) or because effort costs are high, reflect optimal decision-making (and hence are not forms of procrastination). On the other hand, hyperbolic discounting involves a genuine temporal inconsistency where work is delayed more than intended ('intention-action gaps' (Steel, 2007)). Hence, only the latter might be rated as procrastination by students. As we were only interested in the extent of devaluation of future rewards and do not analyse present-bias or intention-action gaps, fitting exponential discount factors sufficed for our purposes. However, had we estimated hyperbolic rates instead, our exponential estimates would likely have been highly correlated with them; hence, we expect a correlation with PASS.

In this paper, we estimated effort and efficacy through model fitting to assess their correlational effects with delay. Future studies can explicitly manipulate how effortful a task is and the time available for completing it to assess the causal relationship. Other factors such as the deadline for completion, the availability and size of other distractions, and task difficulty can also interact with discounting to influence delay, and are hence important targets for future investigation.

Optimisation for model fitting failed for 7 subjects. These students either had sudden jumps in progress or big gaps in work. Our discounting-based model alone might not be able to explain such patterns, and additional mechanisms such as taking breaks due to excess work, possibility of more interesting tasks in the future, etc., might be important to capture these behaviours. Finally, the role of other mechanisms such as steeper discounting of efforts vs. rewards (Akerlof, 1991; Le Bouc & Pessiglione, 2022), and non-linear scaling of effort costs with the amount of work (P. Zhang & Ma, 2019) should also be examined more thoroughly in future empirical studies.

Acknowledgments

We thank Kevin Shen and Surabhi S. Nath for helpful comments and discussions.

References

Akerlof, G. A. (1991). Procrastination and Obedience. *The American Economic Review*, *81*(2), 1–19. https://www.jstor.org/stable/2006817

Bjork, J. M., Hommer, D. W., Grant, S. J., & Danube, C. (2004). Impulsivity in abstinent alcohol-dependent patients: Relation to control subjects and type 1–/type 2–like traits. *Alcohol*, *34*(2), 133–150. https://doi.org/10.1016/j.alcohol.2004.06.012

Burghoorn, F., Roelofs, K., Burk, W. J., Jorgensen, T. D., Scheres, A., & Figner, B. (2026). Intertemporal Impatience Across Mental Health in a Community Sample: A Novel Transdiagnostic Approach. *Journal of Behavioral Decision Making*, *39*(1), e70056. https://doi.org/10.1002/bdm.70056

Chebolu, S., & Dayan, P. (2024). Optimal and sub-optimal temporal decisions can explain procrastination in a real-world task. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *46*.

Chebolu, S., & Dayan, P. (2025). Weighting waiting: A decision-theoretic taxonomy of delay, pacing and procrastination. *PsyArXiv*.

Dayan, P. (2012). Instrumental vigour in punishment and reward. *European Journal of Neuroscience*, *35*(7), 1152–1168. https://doi.org/10.1111/j.1460-9568.2012.08026.x

Duckworth, A. L., & Kern, M. L. (2011). A meta-analysis of the convergent validity of self-control measures. *Journal of Research in Personality*, *45*(3), 259–268. https://doi.org/10.1016/j.jrp.2011.02.004

Green, L., Fry, A. F., & Myerson, J. (1994). Discounting of Delayed Rewards: A Life-Span Comparison. *Psychological Science*, *5*(1), 33–36. http://www.jstor.org/stable/40062338

Keidel, K., Lu, X., Suzuki, S., Murawski, C., & Ettinger, U. (2024). Association of temporal discounting with transdiagnostic symptom dimensions. *npj Mental Health Research*, *3*(1), 13. https://doi.org/10.1038/s44184-024-00060-3

Keidel, K., Rramani, Q., Weber, B., Murawski, C., & Ettinger, U. (2021). Individual Differences in Intertemporal Choice. *Frontiers in Psychology*, *12*. https://doi.org/10.3389/fpsyg.2021.643670

König, C. J., & Kleinmann, M. (2005). Deadline Rush: A Time Management Phenomenon and Its Mathematical Description. *The Journal of Psychology: Interdisciplinary and Applied*, *139*(1), 33–45. https://doi.org/10.3200/JRLP.139.1.33-45

Kool, W., & Botvinick, M. (2018). Mental labour. *Nature Human Behaviour*, *2*(12), 899–908. https://doi.org/10.1038/s41562-018-0401-9

Le Bouc, R., & Pessiglione, M. (2022). A neurocomputational account of procrastination behavior. *Nature Communications*, *13*(1), 5639. https://doi.org/10.1038/s41467-022-33119-w

Loewenstein, G., & Elster, J. (1992, October). *Choice Over Time*. Russell Sage Foundation.

Müller, T., & Apps, M. A. J. (2019). Motivational fatigue: A neurocognitive framework for the impact of effortful exertion on subsequent motivation. *Neuropsychologia*, *123*, 141–151. https://doi.org/10.1016/j.neuropsychologia.2018.04.030

Niv, Y., Daw, N. D., & Dayan, P. (2005). How fast to work: Response vigor, motivation and tonic dopamine. *2005 Annual Conference on Neural Information Processing Systems, NIPS 2005*, 1019–1026.

Niyogi, R. K., Shizgal, P., & Dayan, P. (2014). Some Work and Some Play: Microscopic and Macroscopic Approaches to Labor and Leisure. *PLOS Computational Biology*, *10*(12), e1003894. https://doi.org/10.1371/journal.pcbi.1003894

Odum, A. L. (2011). Delay discounting: Trait variable? *Behavioural Processes*, *87*(1), 1–9. https://doi.org/10.1016/j.beproc.2011.02.007

Patzelt, E. H., Hartley, C. A., & Gershman, S. J. (2018). Computational Phenotyping: Using Models to Understand Individual Differences in Personality, Development, and Mental Illness. *Personality Neuroscience*, *1*, e18. https://doi.org/10.1017/pen.2018.14

Reuben, E., Sapienza, P., & Zingales, L. (2015). Procrastination and impatience. *Journal of Behavioral and Experimental Economics*, *58*, 63–76. https://doi.org/10.1016/j.socec.2015.07.005

Steel, P. (2007). The nature of procrastination: A meta-analytic and theoretical review of quintessential self-regulatory failure. *Psychological Bulletin*, *133*(1), 65–94. https://doi.org/10.1037/0033-2909.133.1.65

Steel, P., Svartdal, F., Thundiyil, T., & Brothen, T. (2018). Examining Procrastination Across Multiple Goal Stages: A Longitudinal Study of Temporal Motivation Theory. *Frontiers in Psychology*, *9*. https://www.frontiersin.org/articles/10.3389/fpsyg.2018.00327

Steel, P., & Weinhardt, J. M. (2018). The building blocks of motivation: Goal phase system. In *The SAGE handbook of industrial, work & organizational psychology: Organizational psychology, Vol. 2, 2nd ed* (pp. 69–96).

Zhang, P. Y., & Ma, W. J. (2024). Temporal discounting predicts procrastination in the real world. *Scientific Reports*, *14*(1), 14642. https://doi.org/10.1038/s41598-024-65110-4

Zhang, P., & Ma, W. J. (2019). Procrastination in Rational Agents. *2019 Conference on Cognitive Computational Neuroscience*. https://doi.org/10.32470/CCN.2019.1351-0

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), e1492. https://doi.org/10.1002/wcs.1492