

Modeling procrastination as rational metareasoning about task effort

Shobhit Jagga (shobhitj@iitk.ac.in)

Dept of Computer Science and Engineering, IIT Kanpur, India

Narayanan Srinivasan (nsrini@iitk.ac.in)

Dept of Cognitive Science, IIT Kanpur, India

Nisheeth Srivastava (nsrivast@iitk.ac.in)

Dept of Computer Science and Engineering, IIT Kanpur, India

Abstract

Current theories of procrastination argue that people put things off into the future with the expectation that they will be better able to do them later. In this paper, we rationalize such expectations within the framework of evidence accumulation models of the choice process. Specifically, we show that it is rational for observers to adopt lower decision thresholds for choices with weak evidence for any alternative, and that observers learning to estimate optimal decision thresholds for tasks that involve decisions will find it reasonable to put the tasks off until the threshold has been sufficiently lowered by time-varying urgency. We designed a computational model and an experiment to differentiate our theory from more general expectancy based temporal motivation accounts. Both simulation and experimental results support our proposal, indicating a large role for choice difficulty in people's self-assessed estimates for how likely they are to procrastinate any given task.

Keywords: procrastination; race models; metacognition; rational analysis; computational modeling

Introduction

Procrastinating, putting important tasks off into the future usually with sub-optimal outcomes, is extremely common (Steel, 2007). Research in workplace settings reveals that up to a quarter of the average US worker's work hours are spent procrastinating (Nguyen, Steel, & Ferrari, 2013). Procrastination in academic settings is so widespread that a particular form of it - putting things off until deadlines and then working urgently to meet them - is separately identified as 'student syndrome' (Schouwenburg, 1995).

Steel has proposed a temporal motivation based meta-theory of procrastination, accommodating several interesting empirical phenomena related to procrastination (Steel, 2007). On this account, peoples' motivation to do a task is directly proportional to their expectation of receiving a reward from it, and inversely proportional to the delay with which they expect to receive this reward (Steel, Svartdal, Thundiyil, & Brothen, 2018). While simple, this theory captures important aspects of procrastination, e.g. a negative skew in start times of tasks, greater procrastination for tasks with delayed rewards, etc.

Current computational models of procrastination predominantly adopt the assumptions of the temporal motivation account of procrastination. For instance, Akerlof's model considers procrastination to arise from the extra cost of doing a task now rather than later (Akerlof, 1991). In his model, at each time step, the person formulates a plan to minimise the

difference between cost and reward associated with the task. An additional factor multiplicatively amplifies the cost for the current time step. The formulated plan emits the optimum start time for the task. Because of the salience-weighted extra cost associated with the current time step, the model suggests starting tasks later than the present moment.

Similarly, O'Donoghue & Rabin's model assumes that people have an additional cost associated with doing something now versus later, but have varying degrees of belief about how much this cost, which they interpret as problems with self-control, is likely to be at any time in the future (O'Donoghue & Rabin, 2001). Observers naive about their future self-control problems are likely to decide to put things off at any given point in time, following the same mechanism as in the Akerlof model.

Modeling procrastination as a direct consequence of intertemporal reward discounting fails to capture important aspects of the phenomenon. Crucially, the temporal motivation explanation does not offer any explanation for the negative effect of procrastination on task performance and downstream affective correlates, such as anxiety and guilt (Ferrari & Emmons, 1995; Ferrari & Díaz-Morales, 2007). Put simply, temporal motivation accounts predict later start times, but not worse performance, and guilt and anxiety about worse performance, on procrastinated tasks. Recently, Sirois and Pychyl (2013) have proposed a theoretical unification of the temporal and affective views on procrastination by explaining procrastination as a failure in predicting the efficacy of one's future self. Following an argument similar to the mathematical assumptions made in O'Donoghue and Rabin (2001), they argue that procrastination essentially arises from an inability to appreciate that one's future self will be as beset by anxiety and worry about the task as one's present self.

But even this theoretical extension of the temporal explanation is inadequate to explain the *selective* nature of procrastination. Across activities as diverse as filing tax returns, applying for research grants, or submitting conference papers (!), the pattern of submissions shows a great rush to beat deadlines, while activities like checking email, playing video games and pursuing hobbies do not, even within the same person. As Ferrari and Díaz-Morales (2007) document in their meticulous surveys, a large fraction of procrastinators, whom they call *arousal* procrastinators, procrastinate deliberately because they feel that they work more efficiently under

deadline pressure. Such a positive view of procrastination is difficult to reconcile with existing theoretical and computational accounts of the phenomenon.

In this paper, we present a theory and a computational model of procrastination that jointly predicts later start times and worse performance on procrastinated tasks followed by an experiment to test our predictions. Our theory rationalizes procrastination as rational meta-reasoning about the expected effort required for completing tasks, successfully operationalizing an arousal view of procrastination, and better characterizing the difference between tasks that are more prone and less prone to procrastination.

The procrastination model

We present a computational model of procrastination, using evidence accumulation race models as our basic building block. The central insight that drives our model is that, when evidence strength for decision sub-tasks that make up the complete task is weak, the speed-accuracy trade-off to determine where the evidence threshold should be set for sub-tasks becomes heavily weighted in favor of speed. Simply put, spending more time collecting evidence is sub-optimal when the evidence is not there.

This insight centrally drives our hierarchical model of task start times, wherein sub-task evidence thresholds are selected by optimizing a composite objective function of task error and duration, and task error and duration in turn are obtained by performing the sub-tasks using a particular value of the sub-task evidence thresholds. We describe our model in greater detail below.

Modeling unforced choices

We model the completion of the target task in our account as requiring the completion of a series of sub-tasks, and model the completion of each sub-task as requiring the processing of information before deciding on the right way of doing the sub-task. This processing is modeled as a two alternative unforced choice task, solved by the model observer using bounded integration choice process models (Usher & McClelland, 2001; Bogacz & Gurney, 2007).

Unlike in classical forced choice paradigms, the accumulator models in our account emit no choice if evidence does not breach the evidence threshold within a fixed number of time steps. In the event that no choice is emitted, the corresponding sub-task is taken to not have been finished, and the model observer has to retry it.

For all the simulation results we describe, we consider the target task to be a sequence of 10 sub-tasks. Each sub-task is modelled as a binary decision problem, solved using a race model that accumulates evidence on every discrete time step of the integration process. The amount of evidence collected per time step is a function of a drift rate parameter that is a proxy for evidence strength, and a diffusion parameter that is held constant across all our experiments.

Each sub-task is further parameterized by an evidence threshold and a time threshold. The time threshold controls

the upper time limit to which the model agent accumulates evidence for a particular sub-task trial. If the evidence threshold is breached before the time threshold is reached, the sub-task is considered completed with one of the two outcomes selected. Otherwise, it has to be redone, and the additional time affects the total task duration, which is simply the sum of all sub-task durations. We measure task error as the fraction of sub-tasks on which the model’s response fails to match the response predicted by the sign of the drift rate parameter.

Urgency saves Buridan’s ass

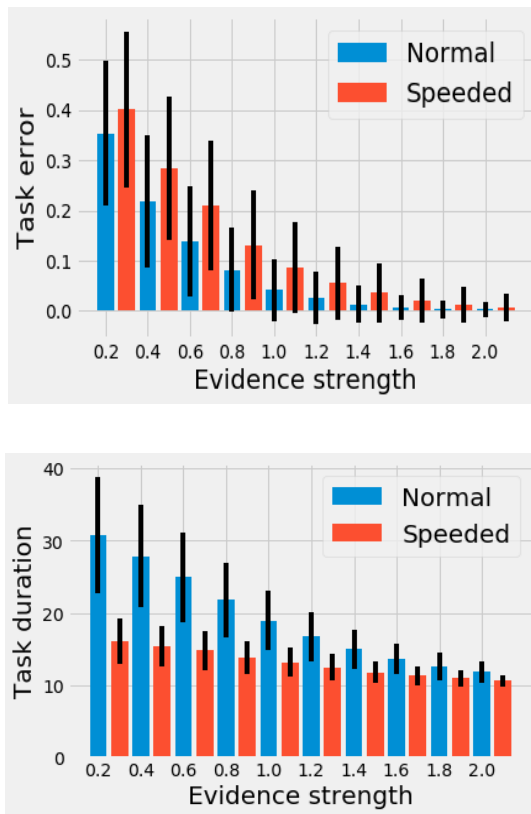


Figure 1: Illustrating an asymmetry in speed-accuracy trade-offs as a function of evidence strength in accumulator models.

The simulation results shown in Figure 1 were obtained by running 1000 trials of the target task for each of 10 discrete evidence strength levels. A key observation here is the large asymmetry in the speed-accuracy trade-offs for such a task under conditions of weak evidence for the individual sub-tasks.

When the evidence for the sub-tasks is weak, very little additional accuracy is gained by using a high evidence threshold as opposed to a lower one (compare the large increase in time for the left-most bars from the speeded to the normal condition in Figure 1(bottom) with the small drop in task error resulting as a consequence in Figure 1(top)). For instance, at an evidence strength of 0.2, doubling the time spent on the task improves the error rate from 0.4 to 0.35. In contrast, at

an evidence strength of 1.0, increasing the time spent on the task by 25% more than halves the error rate.

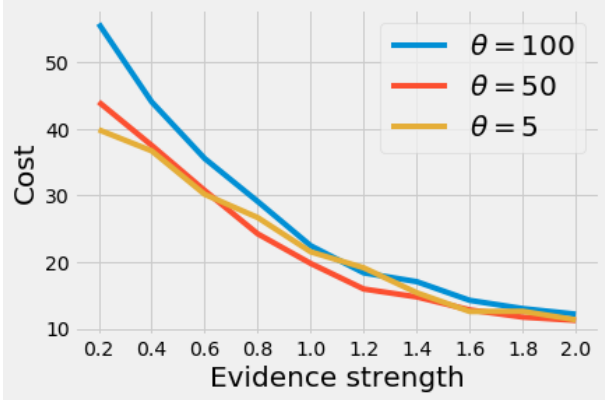


Figure 2: Differential costs for normal and speeded thresholds

Thus, for the same simulated task, an observer model motivated by a composite cost $:= \alpha$ task duration + β task error may find it rational to use lower evidence thresholds under weak evidence conditions (see Fig 2). Assuming that evidence thresholds are lowered by urgency, and that urgency increases via proximity to deadlines, we retrieve a model of start time inference, as we describe below.

Meta-reasoning about evidence thresholds

The simulations above suggest how observers might acquire an understanding of how to set decision thresholds given experience with the strength of evidence they are likely to have in the task. For any task, they’d first, over repeated exposures to the task under varying conditions of urgency and thus varying evidence thresholds, identify a point on the speed-accuracy curve, parameterized by τ, ϵ in our account, they’d be satisfied with, and then map this to the threshold θ likely to yield this performance, as judged by prior experience on other tasks.

In our model, we implement this calculation as a linear optimization problem of estimating the threshold value as

$$\operatorname{argmin}_{\theta} \alpha\tau(\theta) + \beta\epsilon(\theta), \quad (1)$$

where α, β are scale parameters.

In its present form, this model estimates the evidence threshold $\hat{\theta}$ for sub-tasks that observers are expected to perform. Tasks on which observers intend to spend a lot of time will have high estimates of $\hat{\theta}$ and vice versa. To obtain start time predictions from this model, we further assume that the urgency of the task increases with time, causing the evidence threshold expected to be used by the observer to behave as

$$\theta(t) = \theta_0 \exp(-\theta_0^{-1}t), \quad (2)$$

where θ_0 is a free parameter (we use $\theta_0 = 100$ for our simulations).

Given this specification of the threshold $\theta(t)$ at every time point of the trial, we obtain the probability of starting the task at time t as

$$p(t) = \frac{1}{1 + \exp(\lambda(\theta(t) - \hat{\theta}))}, \quad (3)$$

where λ is a model parameter. We use $\lambda = 0.1$ for our simulations. This model predicts procrastination for settings where $\hat{\theta}$ is low, and non-procrastination when $\hat{\theta}$ is high.

For the simulated results presented in this paper, we assume perfect foreknowledge of $\hat{\theta}$ on the part of the observer. That is, our results don’t engage with how observers learn to procrastinate.

Simulation Results

The key commonality across existing computational models of procrastination is the fact that procrastination emerges from additional cost being attributed to doing something *now* as opposed to later (Akerlof, 1991; O’Donoghue & Rabin, 2001). But what could the nature of such a cost be? Our model offers a possible answer.

Our model endogenously and simultaneously predicts a negative skew in task start times and an inverse correlation between task performance and task start times as a consequence of the posited relationship between evidence strength and evidence thresholds.

Figure 3 demonstrates these predictions quantitatively. The left panel plots task start time probabilities for two different values of the observer’s evidence threshold estimate. As we have outlined above, these estimates are expected to be lower for tasks in which observers have to make decisions with weak evidence. Thus, for such tasks, the task start time probability behaves as shown by the red line in Figure 3(left).

The middle panel of Figure 3 plots a start time distribution obtained by running 10000 simulations of the task, using start time probabilities emitted by the model for both high (blue) and low (red) evidence threshold conditions. The histograms clearly demonstrate a propensity for our model observer to procrastinate tasks under weak evidence conditions.

The right panel of Figure 3 plots the average error fraction for the task as a function of when the observer chose to begin the task within the available time window, for a given level of evidence strength. Clearly, starting the task later leads to worse performance on the task.

An Experimental Test

The key novelty of our model, in contrast with prior modeling efforts in this area (Akerlof, 1991; O’Donoghue & Rabin, 2001; Steel, 2007), lies in the relationship we posit between decision difficulty and procrastination. Such a relationship is not unexpected on theoretical grounds, given that *arousal* has been documented as a clear cause of procrastination by previous psychological research (Ferrari & Emmons, 1995; Ferrari & Díaz-Morales, 2007). However, while experimental studies have documented a relationship between task difficulty

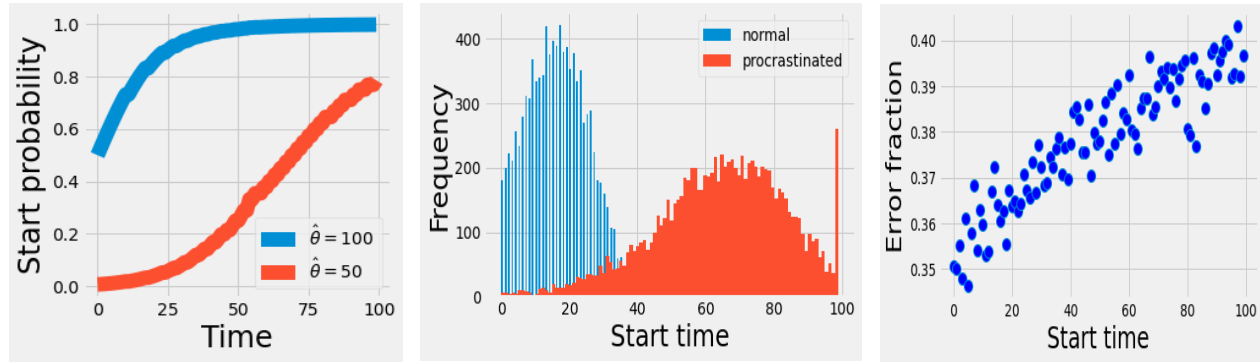


Figure 3: Our model predicts that (left) task start times will be low at the beginning and increase exponentially when observers’ estimates of the optimal evidence thresholds for a particular task are low, (middle) such settings will lead to delays in starting tasks as well as failures in starting tasks altogether, and that (right) later task start times will be correlated with worse performance on the task.

and procrastination more generally (Steel et al., 2018), the specific role of choice difficulty has not been experimentally examined. Since choice difficulty plays a central role in our explanation for the mechanism of procrastination, we tested its effect on procrastination with an experiment.

Design

We designed an experiment with the basic aim of differentiating generic task difficulty from choice difficulty. To achieve this differentiation, we developed a grid-search game, inspired by the spatial foraging paradigm of Hills and Hertwig (2010), wherein participants flip over grid cells and see what reward each cell contains. The purpose of the task is to find a maximally rewarding cell, hidden somewhere in the grid. Participants are only allowed to flip over grid cells contiguous to grid cells that have already been flipped.

As in our simulation, successfully completing the task requires a sequence of directional decisions about which grids to flip over. As illustrated in Figure 4A, rewards can be distributed on the grid in ways that force observers to make this sequence of decisions with either strong (top, easy grid) or weak (bottom, hard grid) directional evidence. Orthogonal to this manipulation of choice difficulty, it is possible to make the overall task easier or harder by forcing participants to search for the maximally rewarding cell in a smaller or larger grid. Crucially, varying the size of the grid does not influence the strength of evidence available to the observer while selecting which direction to go while flipping cells.

We conducted our experiment as a 2×2 within-subject design, varying choice difficulty in two levels of easy and hard, and grid size in two levels of grids of side 15 and 20 respectively. We generated all experiment grids procedurally. To generate easy grids, we first selected a random grid cell to place the maximal reward in, then arranged the remaining rewards as a function of Manhattan distance from this cell, as shown in Figure 4A (top).

Thirty-six university students participated in the experi-

ment for course credit. The study protocol and methods were reviewed and approved by an IRB.

An experiment trial consisted of one grid search task, which ends either when the maximal reward is found, or the time which was specified in the beginning runs out. After every trial, the participant is asked to report their propensity to procrastinate on the task they had just performed on a 5-point Likert scale. Each participant completed two trials per condition of the experiment, thus completing eight trials in all, and providing eight Likert scale procrastination propensity responses. The order of presentation of the four conditions was counterbalanced across participants.

Results

We expected a strong effect of choice difficulty on procrastination ratings, and this expectation was justified by our results (see Figure 4B). Table 1 presents the results of a Bayesian repeated measures ANOVA, treating procrastination ratings as interval-scaled for the purpose¹. The best model that accounts for the procrastination ratings presumes independent contributions to these ratings from choice difficulty and grid size, along theoretically expected directions. Notably, the analysis demonstrates a decisive effect of choice difficulty on the ratings ($BF_{inclusion} = 25,363$), but only a weak effect of grid size-dependent task difficulty ($BF_{inclusion} = 4.6$) and no evidence for an interaction between these two sources of difference ($BF_{inclusion} = 0.284$). These results are consistent with self-expectations of procrastination being largely determined by the difficulty of decisions within the task, as assumed in our theory.

Additionally, we counted the number of times someone had to stop and think about what to do in any given grid search trial (quantified as a 2SD deviation from the mean RT of that trial), and found that participants who experienced more

¹Conclusions from ANOVA on Likert scales are highly robust to high degrees of non-linearity between the scale and the latent construct (Norman, 2010).

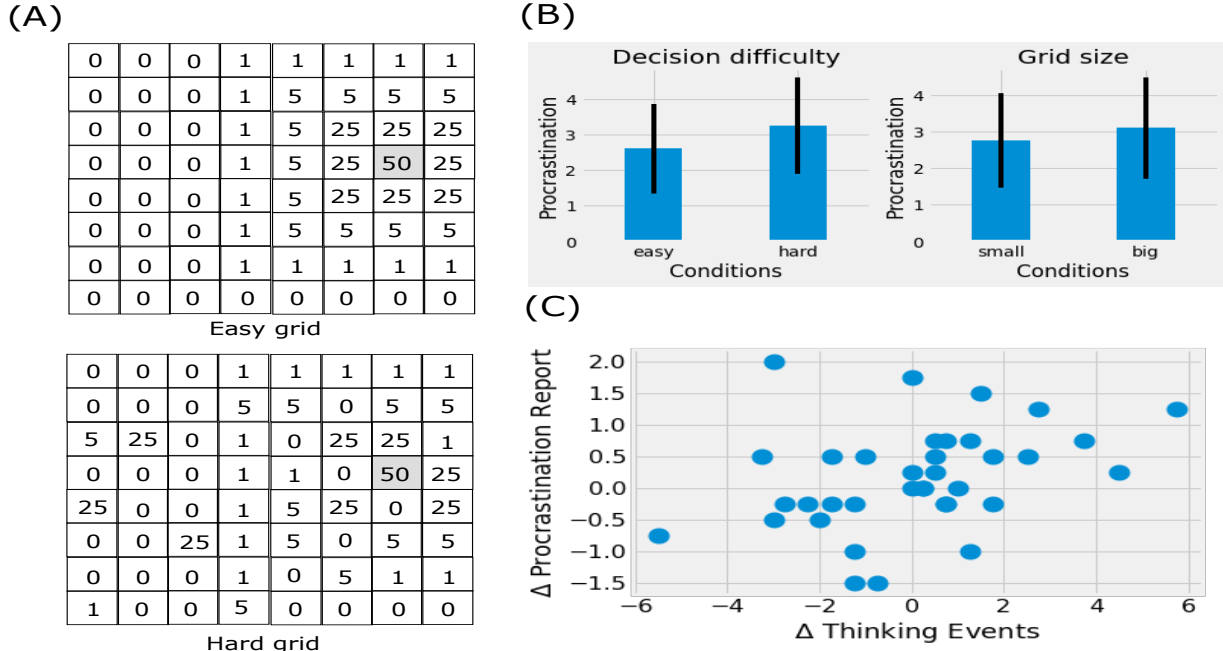


Figure 4: Design and results of an experiment differentiating choice and task difficulty. (A) Searching a grid for a hidden maximal reward (in gray) can be easy or hard, depending on how sub-maximal rewards are distributed (B) average procrastination ratings elicited after grid search in different conditions and (C) plotting the extent to which a participant’s procrastination ratings changed from one elicitation to the next for each condition as a function of the difference in the number of times they had to stop and think about where to move next during grid search across the two elicitations for the same condition.

such *thinking* events in their second experience with a particular experiment condition than the first, gave higher procrastination ratings to the second trial than the first, and vice versa. Figure 4C illustrates this by plotting change in rating from first to second against change in number of *thinking* events, averaged across all four conditions, for all participants. The positive correlation observable in the graph ($\rho = 0.36, p = 0.038$), supports the view that participants’ experience of being unable to decide significantly affects their judgment of how much they would be likely to put it off if encountered again.

Table 1: Results of Bayesian repeated measures ANOVA on procrastination likelihood responses for our experiment. BF₁₀ compared to the best model.

Models	BF ₁₀
choice + grid	1.000
choice + grid + choice * grid	0.284
choice	0.215
grid	3.649e-5
Null model (incl. subject)	1.141e-5

Discussion

We present a rational explanation for procrastination, and simulate it in a computational model. The explanation is that people rationally believe that it is inefficient to begin a task that contains choices for which they possess insufficient evidence. For such choices, we show that the speed-accuracy trade-off asymmetrically leans towards speed, in the sense that the marginal value of spending more time in making such decisions is low because low evidence will lead to high error rates anyway. Thus, we propose that procrastination is primarily caused by people believing that they will be more efficient in doing something later; this belief is justified for tasks that are characterized by difficulty in making choices.

Our characterization of the source of procrastination remains consistent with earlier proposals, such as Steel’s temporal motivation theory (Steel, 2007), while adding significant nuance. Earlier explanations have set great store by the argument that activities that are more likely to be procrastinated are ones wherein rewards arise after considerable delay, while distractor activities offer immediate reward (Steel et al., 2018). While such a framing seems intuitively reasonable - someone expected to submit a work report might find their social media feed more engaging - it is inadequate (Sirois & Pychyl, 2013). Procrastinators are generally beset with guilt and anxiety about putting things off (Ferrari & Díaz-Morales, 2007), something that shouldn’t happen if all they are doing is picking more rewarding activities to perform at all points

in time. Furthermore, distractor activities are frequently not selected for being rewarding, e.g. people playing *idle* games, where the whole point of the game is to see how far they can play it without being bored of it (Cutting, Gundry, & Cairns, 2019), or experienced as being rewarding, even if they are advertised as such (Dhir, Yossatorn, Kaur, & Chen, 2018).

Our explanation for procrastination, delivered via a computational model, argues for it being caused fundamentally by a mismatch between what people think they know how to do, and what they are expected to do. This argument is consistent with existing temporal motivation accounts of procrastination, the most recent variants of which ascribe procrastination to the belief that one's future self may be better suited to perform some tasks (Sirois & Pychyl, 2013; Steel et al., 2018). Our work characterizes specific circumstances under which such future expectations are rational - when the person believes that they will be better able to make decisions under heightened uncertainty and hence with compressed evidence thresholds.

For perceptual decisions on short time-scales, it has previously been shown that observers are able to change decision thresholds to optimize a joint function of error and response time (Bogacz, Hu, Holmes, & Cohen, 2010). It has also been shown that, for difficult tasks, observers take too long to decide than is optimal to maximize their reward rate (Starns & Ratcliff, 2012). Our proposal points to the existence of such a trade-off on much longer time-scales, with procrastination emerging as a consequence of similar sub-optimality.

Our empirical demonstration of the validity of this explanation for procrastination is naturally limited by our use of a specific experimental context and the elicitation of stated rather than revealed preferences. However, we note that it is also consistent with interesting sociological observations, viz. the fact that procrastination is endemic among workers meant to operate at the boundaries of their knowledge such as students and researchers (Schouwenburg, 1995), greater procrastination among white-collar workers than blue-collar workers (Hammer & Ferrari, 2002), and the observation that the societal prevalence of procrastination appears to be increasing with the complexity of worker roles (Milgram, 1992). Testing this proposal in more naturalistic settings, and using revealed rather than stated preferences, constitutes a strong direction for future work.

References

- Akerlof, G. A. (1991). Procrastination and obedience. *The American Economic Review*, *81*(2), 1–19. Retrieved from <http://www.jstor.org/stable/2006817>
- Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation*, *19*(2), 442–477.
- Bogacz, R., Hu, P. T., Holmes, P. J., & Cohen, J. D. (2010). Do humans produce the speed–accuracy trade-off that maximizes reward rate? *The Quarterly Journal of Experimental Psychology*, *63*(5), 863–891.
- Cutting, J., Gundry, D., & Cairns, P. (2019). Busy doing nothing? what do players do in idle games? *International Journal of Human-Computer Studies*, *122*, 133–144.
- Dhir, A., Yossatorn, Y., Kaur, P., & Chen, S. (2018). On-line social media fatigue and psychological wellbeing—a study of compulsive use, fear of missing out, fatigue, anxiety and depression. *International Journal of Information Management*, *40*, 141–152.
- Ferrari, J. R., & Díaz-Morales, J. F. (2007). Procrastination: Different time orientations reflect different motives. *Journal of Research in Personality*, *41*(3), 707–714.
- Ferrari, J. R., & Emmons, R. A. (1995). Methods of procrastination and their relation to self-control and self-reinforcement: An exploratory study. *Journal of Social Behavior and Personality*, *10*(1), 135.
- Hammer, C. A., & Ferrari, J. R. (2002). Differential incidence of procrastination between blue and white-collar workers. *Current psychology*, *21*(4), 333–338.
- Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological science*, *21*(12), 1787–1792.
- Milgram, N. (1992). El retraso: Una enfermedad de los tiempos modernos. *Boletín de psicología*(35), 83–102.
- Nguyen, B., Steel, P., & Ferrari, J. R. (2013). Procrastination's impact in the workplace and the workplace's impact on procrastination. *International Journal of Selection and Assessment*, *21*(4), 388–399.
- Norman, G. (2010). Likert scales, levels of measurement and the “laws” of statistics. *Advances in health sciences education*, *15*(5), 625–632.
- O'Donoghue, T., & Rabin, M. (2001, 02). Choice and procrastination. *The Quarterly Journal of Economics*, *116*, 121–160. doi: 10.1162/003355301556365
- Schouwenburg, H. C. (1995). Academic procrastination. In *Procrastination and task avoidance* (pp. 71–96). Springer.
- Sirois, F., & Pychyl, T. (2013). Procrastination and the priority of short-term mood regulation: Consequences for future self. *Social and personality psychology compass*, *7*(2), 115–127.
- Starns, J. J., & Ratcliff, R. (2012). Age-related differences in diffusion model boundary optimality with both trial-limited and time-limited tasks. *Psychonomic bulletin & review*, *19*(1), 139–145.
- Steel, P. (2007). The nature of procrastination: A meta-analytic and theoretical review of quintessential self-regulatory failure. *Psychological bulletin*, *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*, 327.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological review*, *108*(3), 550.