A Time-Consistent Model for Procrastination

Predicting the Effects of Inertia and Delayed Rewards

Author
B.K. de Koning

Student Number
401219

Study Pogramme
Economics & Business Economics

Institute
Erasmus University Rotterdam

Date
12-06-2016

Supervisor
Dr. J.P.M. Heufer

Abstract

In this bachelor thesis an attempt is made to create a time-consistent model for procrastination using immediately incurred starting costs and a reward that is earned throughout multiple periods, but paid at a set time in the future. The model allows for an accurate and correct prediction of how long a rational agent will procrastinate given his discount factor, the starting costs and the reward per period of effort exerted.

Table of Contents

Introduction

How many times have you decided you were going to start a new project, willing to put in all the effort, only to never start? You know you'll be happier for having done so, and that it is in your best interest to do so, but for some reason you don't.

Not being able to get started with tasks is often described as procrastination and is a very common problem. Starting a task or switching tasks is cognitively straining and makes the new task seem less attractive. At times, this makes people delay starting a task of which they know it is beneficial to start early.

Procrastination is common in both peoples' personal and professional lives. Unpleasant, important tasks, like going to the doctor's office or meeting an important deadline, are delayed, while less important but pleasant things, like watching television or getting coffee for the office, are prioritized for the sole reason that they are easy.

Quasi-hyperbolic discounting is the most commonly used model to explain this phenomenon. However, this model has some drawbacks. First and foremost, it has not been proven to have a ground in reality, as it foregoes starting costs, which is one of the main reasons for procrastination. Aside from that it is time-inconsistent, which means it is unable to correctly predict an agent's behaviour over multiple periods. For that reason, a mathematical, time-consistent, model is built in this thesis. It attempts to explain procrastination in rational human beings, taking starting costs into account. The model will be descriptive and allow for predictions of how long agents will procrastinate.

The paper will be organised as follows. First, the theoretical framework is laid out to make sure definitions are clear and earlier literature is understood. After the theoretical framework, the basic mathematical model will be presented and compared to quasi-hyperbolic discounting. This is followed by the models' real life applications, ideas for future research and lastly the conclusion and discussion.

Theoretical Framework

Inertia and delayed rewards are the main factors influencing procrastination that are discussed in this paper. Together with the model of costs of effort and intertemporal choice models they form the basis for the model proposed later on.

Procrastination

Procrastination is the act of delaying the performance of a task that an agent has to do, even though the agent knows it is not in his future best interest to delay doing the task. There are many reasons for procrastination. Factors that play a role are: the costs associated with starting, how long it takes to receive the reward, the size of the reward, how hard a task is perceived to be, motivation of the agent and many more. (Steel, 2007)

In this paper the focus will lie on the first two factors: task-set inertia and delayed rewards.

Task-set inertia

Being unable to start a new task from a state of rest or from performing a different task is often called task-set inertia. While the new task may not be unpleasant per se, moving from a state of rest to a state of action or from one task to another is often perceived as such. The act of putting the mind or body into action for something not currently being done can be seen as a form of starting costs. Starting costs are incurred once every time an agent switches from a state of rest to a state of action, or from one task to another.

Starting costs have a particularly strong influence when the new task is considered 'hard' or 'cognitively straining', or when the task that is being performed at the moment of switching is very easy. (Allport A., 1994)

Delayed Rewards

Procrastination is often the result of a myopic (present-biased) view as well. If an agent is said to be myopic, it means that to him instant gratification is more important than a (possibly bigger) reward in the long term. Of course, work has to be done before a reward can be reaped, so rewards are often delayed for longer than costs.

Earlier Economic Models of Decision Making

The above proposition is of course related to (optimal) decision-making, an essential part of behavioural economics. There are numerous models that attempt to show how economic agents make decisions based on the information and preferences they have.

Cost of Effort

The simplest model of decision-making is that of costs of effort. It's a utility equation that takes into account both the gains of a performing a certain task and the costs of effort associated with it. If the gains outweigh the costs, the agent will perform the task. If not, the agent will not perform the task.

Take $U(Y,e)$, a continuous differentiable function where Y stands for the material or immaterial gains of performing a certain task, and e stands for the costs of effort. Then $\frac{\partial U}{\partial Y} > 0$ and $\frac{\partial U}{\partial e} < 0$. Utility is rising in Y, the reward, and decreasing in e, the cost of effort. Both Y and e can be functions of other variables, but Y should always be increasing in e so $\frac{\partial Y}{\partial e} > 0$. $U(Y,e)$ has to be either decreasingly increasing in Y: $\frac{\partial^2 U}{\partial Y^2} < 0$ or increasingly decreasing in e: $\frac{\partial^2 U}{\partial e^2} < 0$. (Borjas, 2013)

Intertemporal choice models

Even if $U(Y,e)$ would be positive if the costs would be incurred at the same time as the reward would be received, the agent may still decide against performing a task if the reward is delayed compared to the costs. Intertemporal choice models model an agent that has to make decisions that have an impact on his well-being in the short run and in the long run. Agents often make a choice that does not maximize their total utility, but a choice that maximizes their short term utility, and decreases their long run utility as they discount the latter in the present.

Fisher's intertemporal choice model is a good example. It assumes a two-period model, in which the agent makes a set amount of money every period (Y_t), and is allowed to borrow and save against a constant interest rate (r). It states that utility is a function of the consumption in both period one and period two: $U(C_1, C_2)$. Total consumption is $C = C_1 + C_2$ and depends on the agent's lifetime income. However, $\max C_1 = Y_1 + \frac{Y_2}{1+r}$, whereas $\max C_2 = Y_1 * (1 + r) + Y_2$. This means that saving the entire first period income will lead to the highest total consumption. (Thaler, 1997)

The function of $U(C_1, C_2)$ will dictate how much the agent will spend in the present period, and how much he will spend in the next period. With a non-existent discount rate (assigning equal importance to present and future) the agent will spend all of its money in period 2 (as $\max C_2 > \max C_1$). However, as the discount rate increases, the agent assigns a lower value to future consumption than to present consumption and will thus spend more in the first period.

A high future discount rate shows a preference for immediate gratification. An agent with such a discount rate will save less, and make less healthy choices now for a better future than an agent with a lower future discount rate, ceteris paribus. (Hoch & Loewenstein, 1991)

Time inconsistent preferences

Discount rates do not, however, tell the entire story. Multiple studies have found that people have a hard time making 'correct' decisions when there is an immediate negative side effect or the 'wrong' decision leads to immediate gratification. In one particular study respondents were asked to choose between chocolate and fruit for the day of the study and the week after. 70% of respondents chose chocolate for the day of the study, while only 30% chose chocolate for the week after. (Read & van Leeuwen, 1998)

Another study, which was carried out a year later, showed that the same held true for 'lowbrow' and 'highbrow' movies. When asked to choose a movie to watch for the day of the study, the overwhelming majority chose a lowbrow movie. When asked to choose a movie for four days later only 29% chose a lowbrow movie. (Read, Loewenstein, & Kalyanaraman, 1999)

This phenomenon is often modelled using quasi-hyperbolic discounting. Quasi-hyperbolic discounting adds an extra discount factor to the utility function of an agent to explain present bias, leading to procrastination and immediate gratification. The added variable accounts for the difference in utility between starting today and every single day after today. The prospect of starting today is much worse than the prospect of starting tomorrow, while the prospect of starting tomorrow does not differ much from starting the day after tomorrow.

The utility function of an agent using quasi-hyperbolic discounting is as follows.

$$U(x_0, \dots, x_t) = U(x_0) + \beta * \sum_{t=1}^{T} (\delta^t U(x_t)) \ where\ \delta \in [0,1]\ and\ \beta \in [0,1]$$

(Laibson, 1997).

The mathematical problem with this model is that it is time-inconsistent. An agent following this utility function might decide on performing a certain task in advance, only to reverse his decision and choose another option when the time comes to perform the chosen task. This is called a choice reversal. An agent might think that he will postpone doing a task until tomorrow, but following the same utility function tomorrow, he'll postpone it again. The same goes for every day after that. Because of this, quasi-hyperbolic discounting does not allow for a prediction on when the task will actually be performed.

Aside from that, research has shown that there is very little empirical evidence for quasi-hyperbolic discounting. A constant, fixed, cost for not giving in to procrastination or immediate gratification is a better approximation of present bias than a stronger discounted future utility using β. (Benhabib, Bisin, & Schotter, 2010)

Both the time-inconsistency of quasi-hyperbolic discounting and the fact that it does not represent reality are the reason for the proposition of a different model.

A model of Inertia and Delayed Rewards

In this section an attempt will be made to create a mathematical model that combines inertia and delayed rewards to explain procrastination while remaining time-consistent. Because of time-consistency it will be able to predict exactly when an agent will stop procrastinating.

Basic Mathematical Model

The model assumes that there is an agent that lives in a multi-period world with $c + 1$ periods (including $t = 0$ as a period). In every period t (except for $t = c$) he can choose whether to exert effort (perform a task) or not (so for every t, either e_t = 0 or e_t = 1). For every time he chooses $e_t = 1$, he is compensated with α so that the total reward is $\alpha \sum_{t=0}^{c} e_t$. When choosing to exert effort ($e_t = 1$) the first time, the agent immediately incurs starting costs of β. Choosing to exert effort is costly only the first time so the agent will always exert effort after the first time as long as there is a reward. This is a simplifying assumption discussed more elaborately below. The period $t = c - 1$ is the last period in which the agent can exert effort after which the reward is paid at t = c.

The agent's discount function is δ^{c-t}.

If the agent decides to exert effort for the first time the agent's discounted utility function is D$U(e)$ = $\delta^{c-t} * \alpha \sum_{t=0}^{c} e_t - \beta$.

After having exerted effort for the first time, given the fact that he will exert effort in the following periods as well, D$U(e) = \delta^{c-t} * \alpha \sum_{t=0}^{c} e_t$, which is always positive as long as α > 0. Note that at any point in the model the agent will take into account how many units of effort he will exert in the future if he would decide to exert effort in the current period (t'), and incorporate this into his discount function. This means that if α is positive, $\sum_{t=t'}^{c} e_t$ = $c - t'$.

$DU(e)$ = 0 for $\sum_{t=0}^{c} e_t = 0$. Not doing anything leads to a utility of zero.

Example

The following example shows that an agent may decide to procrastinate, while not being time-inconsistent or irrational.

Let

α = 10

c = 4

δ = 0.5

β = 4

If the agent chooses based on discounted utility, the following result is obtained:

$t = 0: DU(Exert\ effort)\ = 0.5^4\ * 40 - 4 = -1.5 < 0 \rightarrow Will\ not\ exert\ effort \rightarrow DU\ = 0$

$t = 1: DU(Exert\ effort)\ = 0.5^3\ * 30 - 4 = -0.25 \rightarrow Will\ not\ exert\ effort \rightarrow DU\ = 0$

$t = 2: DU(Exert\ effort)\ = 0.5^2\ * 20 - 4 = 1 > 0 \rightarrow Will\ exert\ effort - No\ more\ \beta$

$t = 3: DU(Exert\ effort) = 0.5 * 20 = 10 > 0 \rightarrow Will, of\ course, continue\ to\ exert\ effort$

The agent will start working at t = 2. For a graphical representation of discounted utility before having exerted effort, please see Appendix A.

Note that after incurring β the agent will always exert effort in the following periods, as marginal utility is then always above 0.

Explanation

One might expect the agent to either start working on the project immediately, or never start working on the project at all, as there are only fixed costs. However, in this example the agent only starts working at t = 2.

The reason for this is that while the decrease in reward from not working is -10 for every period, the discount factor increases exponentially as the reward gets closer: the reward is valued more the closer the current period (t) is to the reward period (c). This situation is not necessarily far from reality. In general, salary is paid periodically and increases with every period of effort exerted, while work has to be done every day.

While the agent does procrastinate, he is not ‘irrational’: his decisions are time-consistent. He knows he’ll delay working for some time, and then start. In the earlier periods the agent just values the present comfort more than the future reward, which is a simple preference.

Note that in this example everything is kept constant and marginal costs are omitted to make the effect of task-set inertia as clear as possible. When marginal costs and non-linear rewards are included, the same results as seen in the example can still be found with more plausible parameters. The main requirement is for the rewards to be delayed and starting costs to be incurred immediately.

Necessary Assumptions of the Model

As stated above, not every assumption made in the example is necessary for the model to hold. The assumptions below are the bare essential.

1. There are multiple periods in which the agent can choose to perform or not perform an action that influences a reward that is received at a known point in the future.
    - If there is just one period, it is not possible for the reward to be delayed, and the resulting equation is then a simple cost-benefit analysis. The most important part of the assumption above is that the reward is received at a known point in time regardless of how often the task is performed. If the reward is received for just one period of work, a set amount of time after that period has passed, the agent's decision will remain the same throughout periods. The agent will only change his decision if the 'reward period' comes closer.
2. The future reward is discounted stronger in earlier periods.
    - If the future reward is not discounted stronger in earlier periods, the result will be the same for every single period. Note that the reward need not be discounted exponentially for the model to hold. There just has to be a certain point at which the discount factor is high enough for the discounted utility of the reward to outweigh the starting costs.
3. The reward is increasing in e.
    - If the reward is not increasing in e, there is no reason to perform a task more than once or, in fact, at all. While the reward should always be increasing in e, it need not be linear as shown in the example. An exponentially increasing reward will of course make procrastination less likely, but does not necessarily prohibit it from occurring. Some rewards have diminishing returns to scale. The stronger it diminishes over time, the more likely it is that the agent procrastinates. Diminishing returns to scale are likely to occur when periods get longer and rewards get bigger.

    At a certain income, a little bit of extra money will not make that much of a difference, and the incentive to start working earlier will be smaller.

4. There are costs for performing a certain action that are incurred only once and are incurred immediately
    - Starting costs can, by their definition, be incurred just once. While the same results may be obtained with costs recurring a limited amount of time, inertia is then not what drives the result. This does not mean that exerting effort in later periods has to be costless. There may also be another source of costs, but starting costs should add up to these costs when exerting effort for the first time.

Even if all these assumptions hold, procrastination is not always the result, the delayed reward may be so high that the starting costs are not even too high in the first period. However, if these assumptions do hold, inertia is always at work.

Comparison to Quasi-Hyperbolic Discounting

While quasi-hyperbolic discounting is very similar in that it attempts to explain procrastination as well, the models differ on an essential point: time-consistency.

Quasi-hyperbolic discounting adds another discount factor to an agent's already discounted utility function to explain procrastination. This means that the quasi-hyperbolic discounting model is time-inconsistent, which may lead to choice reversal. As such, an agent will often delay performing a task indefinitely: the agent keeps believing that he will 'start tomorrow' (or at a different particular point in the future).

In the inertia and delayed rewards model, the effect of procrastination stems from the fact that the reward is delayed. The discount factor is time-consistent. Because of the time-consistency, the model of inertia and delayed rewards can accurately predict what the agent will actually do, while quasi-hyperbolic discounting cannot.

Applications and Outlook

The model of inertia and delayed rewards can be applied to multiple situations, as long as one effect is immediate, and one effect is delayed. This section will expand on that with an example of a less straightforward situation. Furthermore, the idea of inertia and delayed rewards can be used to create multiple models, which will also be shown in this section. For further research, it may be very interesting to see whether findings of the model(s) hold in real life.

Applications of the Model

Even under the aforementioned assumptions, the model of inertia and delayed rewards can be applied to numerous situations. In short, the model can be applied to any situation in which there is a task to be performed, that requires some form of initial effort (which results in starting costs) and that task influences a reward at a certain point in the future.

Note that a task is not just defined as the act of doing something, but could also be the act of avoiding an action that has short-term benefits, but long-term negative effects. Missing out on short-term benefits is as a form of opportunity costs, and avoiding a certain punishment can be seen as a reward in and of itself. As such, the model does not need to be altered if there are starting benefits and long-term punishments.

An example of such a situation would be quitting a smoking habit. Whenever the agent decides to quit smoking, he will suffer from withdrawal effects, but these won't last forever. After some time (possibly one period) the costs associated with these effects are gone. However, the sooner he quits smoking the more long-term damage he prevents.

A model for Increasing Costs

The model of inertia and delayed rewards cannot be applied to every situation in which inertia plays a role. For example, some situations have a set deadline that has to be met, no matter the cost. In these scenarios benefits often stay the same, while costs rise. This may also lead to inertia-induced procrastination. While this behaviour can't be explained using the model described in the rest of this paper, it does stem from the same idea, which is why it's included here.

A simple example would be going to the supermarket. Food is necessary for the agent to survive, but he doesn't want to go. However, the longer he waits, the more energy it will cost to get off the couch.

In this case, starting costs should be increasing in t, not e, so costs are $\beta(1 + t)$. It is not possible to exert effort twice, so the reward can be written as α.

The discounted utility function for performing the task can then be written as follows.

$$DU(t) = \delta^{c-t} * \alpha - \beta(1 + t)$$

It is again assumed that not performing the task leads to a utility of zero.

Example

Let

α = 40

c = 4

δ = 0.5

β = 4

$t = 0: DU(Go\ to\ Supermarket) = 0.5^4 * 40 - 4 = -1.5 < 0 \rightarrow Will\ not\ go$

$t = 1: DU(Go\ to\ Supermarket) = 0.5^3 * 40 - 8 = -3 < 0 \rightarrow Will\ not\ go$

$t = 2: DU(Go\ to\ Supermarket) = 0.5^2 * 40 - 12 = -2 < 0 \rightarrow Will\ not\ go$

$t = 3: DU(Go\ to\ Supermarket) = 0.5 * 40 - 16 = 4 < 0 \rightarrow Will\ go$

In this case, the agent decides to go in the last period, as only then the discounted reward is high enough to cover the starting costs. If costs had not been increasing and would have remained constant at 4, the agent would have gone at $t = 1$. Interestingly enough, in the scenario of increasing costs, the agent is the least inclined to go at t = 1. See appendix B for a graphical representation.

If α = 80, another interesting phenomenon can be observed.

$$t = 0: DU(Go\ to\ Supermarket) = 0.5^4 * 80 - 4 = 1 > 0 \rightarrow Wants\ to\ go$$

$$t = 1: DU(Go\ to\ Supermarket) = 0.5^3 * 80 - 8 = 2 \rightarrow Wants\ to\ go$$

$$t = 2: DU(Go\ to\ Supermarket) = 0.5^2 * 80 - 12 = 8 \rightarrow Wants\ to\ go$$

$$t = 3: DU(Go\ to\ Supermarket) = 0.5 * 80 - 16 = 24 \rightarrow Wants\ to\ go$$

If the agent had to choose between either go at $t = 0$ or never go at all, he would choose to go at t = 0. However, because, the agent knows he'll be hungrier later on (discount the future less), he will still decide to go at t = 4. Even when utility is positive, procrastination can be observed.

Is this still procrastination? Probably yes, as going earlier is still less costly and would result in the 'best' hindsight.

Further Research

While the models discussed in this paper are based on arguably reasonable assumptions, as shown in the examples, it has not been verified in reality. A more empirical approach to the effects of task-set inertia may lead to very interesting insights on the topic.

If real life results follow the theoretical model described in this paper, companies and people themselves could use this information to entice their employees or themselves to stop procrastinating. By choosing shorter term rewards (such as weekly paid salaries instead of monthly) or decreasing starting costs, people may be more likely to start earlier and deliver better work.

Conclusion & Discussion

The model of inertia and delayed rewards is a time-consistent explanation for procrastination in the presence of starting costs and delayed rewards. It allows for a correct prediction of an agent's actions when there is a choice between either losing in the short-term and winning in the long-term or winning in the short-term and losing in the long-term.

The model can be applied to any scenario in which there are multiple periods in which an action can be performed, the future is discounted, there is a reward or punishment that is increasing in the amount of times a task is performed and with the first performance of said task the agent incurs starting costs.

Bibliography

Allport A., S. E. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In M. M. Umilta C., *Attention and performance* (pp. 421-452). Cambridge: MIT Press.

Benhabib, J., Bisin, A., & Schotter, A. (2010). Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior*, *69*(2), 205-223.

Borjas, G. J. (2013). Labor Economics. In G. J. Borjas, *Labor Economics.* Harvard University: McGraw-Hill.

Hoch, S., & Loewenstein, G. (1991). Time-Inconsistent Preferences and Consumer Self-Control. 492-507.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 443-477.

Read, D., & van Leeuwen, B. (1998). Predicting Hunger: The Effects of Appetite and Delay on Choice. *Organizational Behavior and Human Decision Processes* , 189-205.

Read, D., Loewenstein, G., & Kalyanaraman, S. (1999). Mixing Virtue and Vice: Combining the Immediacy Effect and the Diversification Heuristic. *Journal of Behavioral Decision Making*, 257-273.

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

Thaler, R. (1997). Irving Fisher: The Modern Behavioral Economist. *The American Economic Review*, 439-441.

Appendices

Appendix A

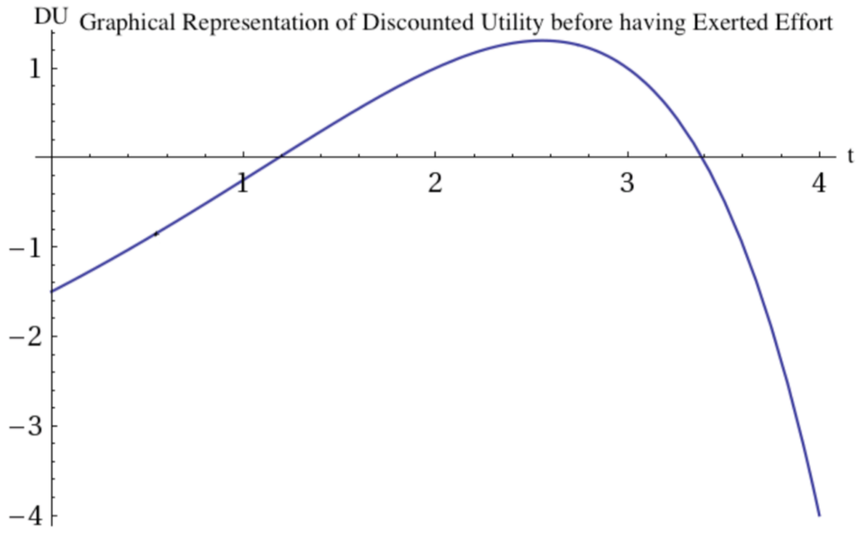

Appendix B

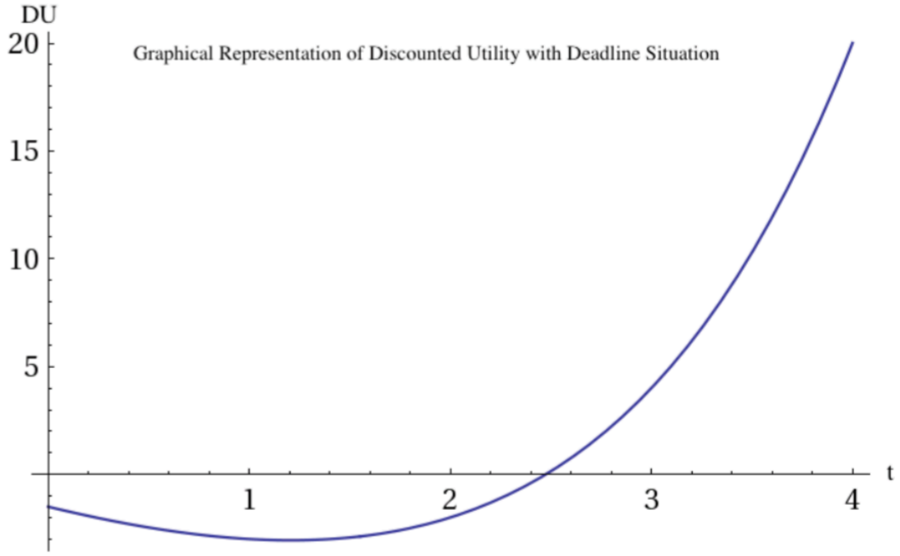
DU
20
15
10
5
Graphical Representation of Discounted Utility with Deadline Situation
t
1
2
3
4