On the Interaction of Memory and Procrastination: Implications for Reminders, Deadlines, and Empirical Estimation

Keith Marzilli Ericson*

Boston University Questrom School of Business and NBER

February 2016

Abstract

The interaction between present-bias and limited memory can explain why individuals do not act at deadlines and why providing reminders can have large effects. Individuals in my model must choose when and whether to complete a task, but may forget or procrastinate. A calibration exercise shows that assuming perfect memory leads to biased estimates of present-bias because the rate of task completion at the deadline is much lower with imperfect memory. Naïve procrastination explains why individuals do not set up reminders despite large gains to doing so. The model offers guidance for empirical studies of reminders, making a distinction between anticipated and unanticipated reminders: anticipated reminders can induce additional procrastination, lowering both welfare and the probability the task is completed. I then use this framework to show how to optimally set deadlines and time the delivery of reminders to present-biased individuals.

*Email: kericson@bu.edu. I thank Stefano DellaVigna, David Laibson, Ted O'Donoghue, Matthew Rabin, the editor, anonymous reviewers, and seminar participants at the University of California-Berkeley, Harvard University, and ESA-North America for helpful comments. I thank the Williams College Tyng Committee and the National Science Foundation for financial support.

1 Introduction

Present-bias has been used to explain why individuals often do not take a beneficial action, such as enrolling in retirement savings plans: naive procrastinators do not act today because they plan to act tomorrow, but when tomorrow arrives, they delay again. Deadlines seem to offer a simple solution to the procrastination problem, as they remove the possibility of further delay. For instance, adding a deadline to enroll in a retirement savings plan should induce most present-biased individuals to act: the transactions cost of filling out a form are outweighed by the future benefit even with heavy discounting.[1] It is puzzle, then, that individuals often do not act even when there is a deadline. For instance, individuals can typically only switch health insurance plans during an annual open enrollment period, which provides a natural deadline. Yet a large literature shows that individuals don't switch health plans, even though the benefit to switching can be very large– even hundreds or thousands of dollars (see e.g. Handel 2013, Ericson 2014a). Similarly, many students fail to apply for financial aid before important deadlines (King 2004), or fail to submit rebates before they expire (Pechmann and Silk 2013).[2]

At the same time, small reminder interventions have been shown to have an economically significant impact on behavior in a variety of domains, including savings (Karlan, McConnell, Mullainathan, and Zinman, forthcoming), loan repayments (Cadena and Schoar 2011, Karlan, Morten and Zinman, forthcoming), and medication adherence (Nieuwlaat et al. 2014, Bobrow et al. 2016).[3] This presents another puzzle, since individuals can set up their own reminders. Memory technologies are ubiquitous and inexpensive, and include calendars, to-do lists, checklists, and reminder services.

This paper shows, using theory and calibrations, that the interaction of limited memory and procrastination can explain this puzzling behavior. Together, limited memory and present-bias can explain why individuals do not set up reminders despite large gains to doing so, and thus why deadlines may not be effective at ensuring tasks are completed. These interactions also have affect parameter estimates: I show ignoring limited memory will lead to substantial misestimation of present-bias from data on task completion. I offer guidance to empirical work on reminders, showing that the distinction between anticipated and unantici-

[1] For instance, Dellavigna (2009) estimates a gain of at least $1200/year for a typical worker who receives an employer subsidy.

[2] Moreover, deadlines have been shown to have mixed effects on behavior in experiments. Ariely and Wertenbroch (2002) show that deadlines increase performance. However, Burger, Charness, and Lynham (2011) and Bisin and Hyndman (2014) run experiments in which deadlines do not increase task completion rates.

[3] See also appointment show-ups (Guy et al. 2012), gym attendance (Calzolar and Nardotto forthcoming), appointment sign-ups (Altmann and Traxler 2012), rebate claims (Letzler and Tasoff 2014), checklists (Jackson and Schneider 2015), and donations (Damgaard and Gravert 2014).

pated reminders is important: reminders that are anticipated by a present-biased individual can induce her to procrastinate further, making her worse off and less likely to complete the task. Finally, I show how accounting for these interactions can help principals optimally time deadlines and reminders for their agents.

The way individuals attempt to complete tasks that involve costs in the present but benefits in the future depends on whether individuals are time-consistent or present-biased (O'Donoghue and Rabin 1999a, 1999b, 2001). Present-biased individuals overweigh the present relative to the future. Present bias is often modeled with a β-δ quasi-hyperbolic discount function (Laibson 1997).[4] As a result of present bias, individuals procrastinate[5] in doing tasks. While sophisticated individuals who recognize their present-bias will wish to bind their future behavior via commitment devices and deadlines, evidence indicates that individuals are at least partially naive about their present-bias, and fail to predict their future behavior (DellaVigna and Malmendier 2006, Acland and Levy 2015).

When individuals make plans to act in the future, memory is relevant– they may forget about the task. A growing literature shows that individuals not only have limited prospective memory[6]– remembering to take an action– but also may have incorrect beliefs about their memory ability.[7] For instance, Ericson (2011) elicits incentivized forecasts of subjects' subjective probability of remembering to claim a delayed payment, and finds that while their choices imply at least a 75% probability of claim, they only claim the payment about half the time. Similarly, Letzler and Tasoff (2014) find that individuals are overoptimistic about claiming rebates at a delay, claiming about 30% of the time despite believing they would claim about 80% of the time.

I examine individuals who are choosing whether and when to complete a valuable task that entails immediate costs and delayed benefits, extending the present-bias framework of O'Donoghue and Rabin (1999a, 1999b, 2001) to include imperfect memory. I show that the interaction of memory and present-bias is quantitatively important for predicting the time at which individuals complete a task, and as a result, for estimating present-bias parameters from behavior. Data on task completion has often been used to estimate structural models

[4] Both laboratory and field evidence indicate that individuals tend to make relatively impatient choices for decisions that involve immediate costs or benefits, but relatively more patient choices when choosing for the future. Discount rates are high at short horizons, but lower at long horizons, not only for money (Frederick, Loewenstein and O'Donoghue 2002) but also for direct consumption experiences (e.g. McClure et al. 2007 for juice, Read and van Leeuwen 1998 for food).

[5] I use "procrastination" to refer to any additional delay in action induced by present-bias.

[6] Prospective memory is memory for action, and is distinct from retrospective memory (recalling information about the past). See McDaniel and Einstein (2007) for a review.

[7] Note that experience may not remove belief biases if learning is incomplete: see Ali (2011) and Schwartzstein (2014).

of present-bias,[8] but assuming perfect memory leads to biased estimates. I examine an example designed to capture the decision of whether switching health plans, and I show that by mistakenly assuming individuals had perfect memory, a researcher would miss evidence of present-bias. The model with present bias and perfect memory predicts a spike in task completion at the deadline. However, individuals with limited memory and present-bias do not show the spike: they forget to act at the deadline or act earlier so as to avoid forgetting. When estimating parameters from task completion data, an individual with a present-bias parameter $\beta = 0.8$ and 10% probability of forgetting each period is mistakenly estimated to have $\beta = 0.98$ if perfect memory is assumed. This suggests that existing estimates of β from task completion data could be made much more accurate by accounting for limited memory.

I show that present-bias can explain why individuals don't set up inexpensive reminders. The cost of setting up reminders bounds the cost of limited memory for time-consistent individuals with correct beliefs, but that the loss from limited memory can be much larger for present-biased individuals. Present-biased individuals will often procrastinate on setting up a reminder system (e.g. lists, calendars, and smartphones), as doing so entails immediate costs for future benefits. As a result, we should expect to present-biased individuals to be more forgetful than time-consistent individuals.

Present-bias also highlights a new distinction in reminder design: whether reminders should be anticipated or be a surprise. Time-consistent individuals prefer anticipated reminders to surprise ones, because they optimally respond to the flexibility afforded by better memory. However, better memory can enable procrastination, and so anticipated reminders can lower both the probability a present-biased individual completes the task, as well as their welfare. This result provides guidance to empirical studies: when testing the impact of reminders on behavior, it is important to note whether individuals anticipated the reminder intervention when making their plans. A surprise reminder intervention could be quite effective, but once it is scaled up and anticipated, its positive effect could be undone by increased procrastination.

Finally, I examine the optimal timing of reminders and deadlines, which is relevant in many contexts. For instance, firms selling products with recurring purchases (e.g. health insurance) or offering retail promotions must determine when to remind customers about an upcoming decision deadline, as well as when to set the deadline. While previous literature has considered how to set a deadline with present-bias alone[9] and with limited memory

[8] See e.g. Bisin and Hyndman (2014) on student project completion, Acland and Levy (2015) on gym attendance, Paserman (2008) on job search and Fang and Yang (2015) on mammograms.

[9] See Herweg and Muller (2011), who consider deadlines in a different model with continuous effort and multiple tasks, and O'Donoghue and Rabin (1999b) in a principal-agent setting.

alone,[10] I show that accounting for both is quantitatively important. I examine the timing of providing surprise one-shot reminders about an upcoming deadline (e.g. a workplace announcement about the end of an open enrollment period). I show that the optimal time to provide reminders to present-biased individuals can be much later than for time-consistent individuals, because once reminded about a task that had been forgotten, present-biased individuals will procrastinate on completing the task and may forget again.

This paper is connected to the literature on memory and inattention. Holman and Zaidi (2010) analyze the economic consequences of limited memory for time-consistent individuals. They show that in the presence of limited memory, the probability of task completion may decrease in the length of time allocated to it, and that overestimating the probability of remembering can explain the existence of free trials, automatic renewal offers, and rebates. More recently, Taubinsky (2014) analyzes a model of inattentive choice[11], and considers the role of deadlines and reminders, as well as cues and habits. In an experiment, he finds that longer deadlines lower the probability of task completion without reminders, but that reminders are effective and reduce the effect of deadline on task completion (see also Shu and Gneezy (2010)). This paper is distinct from both Holman and Zaidi (2010) and Taubinsky (2014) because it analyzes the interaction of limited memory with present bias. It adds to a growing literature examining the interaction of biases, which can be quite important.[12]

Other work has developed theories of reminders in different contexts. Bernheim and Thomadsen (2005) examine the interaction between imperfect recall memory and anticipatory utility and develop a theory of reminders. However, their reminders are designed to communicate information about the state of the world, rather than the reminders considered here, which bring a task to mind and prompt action. Bordalo, Gennaioli, and Shleifer (2015) examine the interaction of salience and memory, and show that reminders may have an effect by changing the salience of an attribute even if the recipient had remembered that attribute (e.g. a reminder to "save in case of an emergency" might raise the salience of emergencies.) In contrast, the reminders considered in this paper merely bring a task to mind, but otherwise don't communicate information or affect salience.

This paper is organized as follows. Section 2 lays out a task completion model and

[10] See Holman and Zaidi (2010) and Taubinsky (2014), which I discuss in Section 6. Also see Saez-Marti and Sjögren (2008), who show a principal-agent motivation for deadlines for time-consistent individuals who may be distracted.

[11] I refer to "forgetting" and "memory" to refer to not taking a planned action because it does not come to mind. This terminology links to the psychology literature on "prospective memory". However, inattention can also be used to describe the same phenomenon, as in Taubinsky (2014).

[12] For instance, Gottlieb (2014) shows how imperfect recall memory interacts with self-image concerns to produce high levels of risk aversion, and Benabou and Tirole (2004) examine the role of personal rules with present-bias and imperfect recall memory. Hsiaw (2013) and Koch and Nafziger (2014) examine present-biased individuals with reference-dependent preferences, with implications for goal-setting and personal rules.

shows how memory affects the cost threshold at which time consistent and present-biased individuals act. This model serves as the workhorse for the rest of the paper, and presents a simple example to illustrate the ideas. Section 3 shows that allowing for limited memory affects estimates of present-bias parameters from task completion data. Section 4 shows present-biased individuals can face large costs from limited memory, since present-bias affects the decision to invest in reminder systems. I then turn to the design of reminders. Section 5 shows that with present-biased individuals, there is an important distinction to be made between anticipated and unanticipated reminders. Section 6 shows how memory affects when to set deadlines and how present-bias affects when to time unanticipated reminders relative to the deadline. Finally, Section 7 concludes.

2 Model: Completing a Task

2.1 Setup

Consider an individual choosing when to do a particular task that involves costs today (effort) but gives benefits in the future. This is a general set up that captures many economically relevant tasks: for instance, choosing to enroll in a health insurance or retirement plan, paying taxes, or going to the gym. O'Donoghue and Rabin (1999a, 1999b, 2001) use a similar framework to analyze the behavior of present-biased individuals. I adapt this framework to allow for the possibility of imperfect memory.

Each period, the individual decides whether to do a task that can be done only once. There is a deadline: after T periods, the opportunity to do the task disappears. When $T = \infty$, there is no deadline. The individual can only determine her behavior this period, and cannot commit to future actions.[13] If the individual does the task in period t, she pays cost c_t that period and receives benefit y in the next period.[14] The cost of action c_t can vary by period, and is drawn independently each period from a known time-invariant distribution F with associated density function f; I assume F is continuous, differentiable and has positive density throughout the range 0 to δy. (The individual knows her period's c_t value before deciding whether to act.) The cost of action includes immediately borne costs, such as lost time, forgone pleasure, physical or cognitive effort exerted, or lowered consumption; it is not meant to capture financial costs that do not immediately translate into lower consumption.

[13] In some contexts, the use of commitment devices is possible, but here I consider either non-contractible behaviors or an environment with transactions costs high enough to rule out commitment devices. Note that in any case, individuals who are naive about their present-bias would not pay to take up commitment devices.

[14] The value of y may represent a one-time benefit that is received tomorrow, or a flow of benefits received in the future, beginning tomorrow.

Present-bias is captured by quasi-hyperbolic discounting (Laibson 1997, O'Donoghue and Rabin 1999), in which discounted utility is $u_0 + \beta\left(\sum_{t=1}^{\infty} \delta^t u_t\right)$, where u_t is the flow utility in each period t. All future periods beyond the present period are discounted by $\beta \leq 1$, in additional to the standard per period discount factor $\delta \leq 1$. When $\beta = 1$, the individual is time consistent (TC) and discounts the future exponentially, but when $\beta < 1$ the individual has time-inconsistent preferences.

Individuals may not anticipate how present-bias will affect their future behavior. Let the parameter $\hat{\beta}$ capture individual's anticipated level of present-bias in future all periods. The literature suggests that individuals are at least partially naive with $\beta < \hat{\beta}$ (DellaVigna and Malmendier 2006, Acland and Levy 2015). I analyze fully naive present-biased individuals with $\beta < 1$ but $\hat{\beta} = 1$, so that each period in the future they think they will act like a time-consistent individual.[15] I always assume individuals correctly perceive δ.

I thus examine two types of individuals $i \in \{TC, N\}$:

TC: Time Consistent ($\beta = \hat{\beta} = 1$)
N: Naive, Present-biased ($\beta < \hat{\beta} = 1$)

I model limited memory as an exponential decay in the probability the task will be recalled (see Levy and Loftus 1984). Each period, an individual with imperfect memory will forget about the task with some probability. If the individual forgets about the task, she cannot act. If she remembers in one period, the probability she remembers in the next period is ρ. Thus, if the individual remembers about the task today, she will remember it t periods from now with probability ρ^t. In this model, forgetting is an absorbing state: once she forgets about the task, she does not remember again.[16] I also allow for the possibility that individuals have incorrect beliefs about memory ability: the individual believes her probability of remembering to be $\hat{\rho}$. If $\hat{\rho} > \rho$, I say the individual is overconfident.

Because individuals may be present-biased and have incorrect beliefs, the choice of

[15] In a working paper version, I also examined sophisticated individuals who correctly anticipated their future present-bias. While sophisticates' strategies can be complex, memory has similar comparative statics.

[16] This form of limited memory accords with psychological evidence on memory (Mullainathan 2002, Holman and Zaidi 2010): once something is forgotten, it is much less likely to be remembered in the future. Modeling forgetting as a fully absorbing state (as opposed to generalizing the model to allow a positive probability of moving from a state of forgetting to remembering) is a simplification that makes analysis more tractable, akin to how the quasi-hyperbolic discounting model is a simplification of hyperbolic discounting. Moreover, this model of memory is distinct from "slipping the mind", in which an individual forgets this period, but forgetting this period does not affect the probability of remembering in future periods. "Slipping the mind" would also affect the demand for reminders. In the absence of reminders, behavior with "slipping the mind" looks very similar to the case in which the distribution of task costs F entails some probability of drawing a very high cost. However, the welfare implications of, for instance, inducing someone to act if they had drawn a high task cost would differ from inducing them to act if they had forgotten.

welfare criterion must be explicit. I define welfare the way an individual with present bias and correct beliefs would judge their welfare a period before having the opportunity to do the task. That is, welfare is ex ante welfare as judged from the perspective of $\beta = 1$ by an individual who has correct beliefs about their future behavior and memory. Hence, welfare is given by $U_i = \sum_{t=0}^{\infty} \delta^t u_t$. This definition is natural and is useful for evaluating how individuals choose (or would like to choose) task environments; it is often used (Gruber and Koszegi 2002, O'Donoghue and Rabin 2006, Heidhues and Koszegi 2010). However, it differs from other proposed welfare criteria for individuals with present bias (see e.g. Bernheim and Rangel 2009).

2.2 Action Cost Thresholds

An individual's behavior today will depend on what she expects to do in the future, about which she may have correct or incorrect beliefs. To determine how individuals behave, I follow O'Donoghue and Rabin (1999a, 1999b, 2001) and require that individuals follow *perception-perfect strategies.* Perception perfect strategies require that each period's behavior maximizes that period's preferences given beliefs about future strategies, and requires that beliefs be dynamically consistent. With dynamically consistent beliefs, individuals think they will act optimally in the future, given their beliefs about future preferences and strategies. The formal definition is given in the appendix; here I describe the implications of perception-perfect strategies for the types in the model.

Individuals choose their current action to maximize their current value function, given their anticipated future strategies.[17] Here, a strategy takes the form of a cost threshold for action c_t^* : in period t, an individual who remembers acts if and only if $c_t \leq c_t^*$. Thus, the individual who doesn't forget will only act if the cost draw is less than or equal to the cutoff. (I assume that if the individual is indifferent between acting and not, she chooses to act.)[18]

The action threshold c_{it}^* is chosen to maximize an individual's current value function $W_{it}^{\hat{\rho}}(c)$. The threshold will depend on the expected continuation value $EV_{it}^{\hat{\rho}}(c)$, which in turn depends on the distribution of c, the expected future strategies of type $i \in \{TC, N\}$, perceived memory ability $\hat{\rho}$, and period t. Given the recursive structure of the task, we can define the current value functions and perceived continuation value functions. The

[17] Note that the individual only is able to implement the strategy t periods in the future if she remembers, which she anticipates happens with probability $\hat{\rho}^t$; if she forgets, she cannot act. However, once she acts, she gets the benefit of the task next period regardless of whether she would have remembered next period.

[18] For present-biased sophisticates, strategies could be history-dependent (See the discussion in O'Donoghue and Rabin (1999b) for more detail, particularly footnotes 9 and 31.) However, the optimal strategy for a time-consistent individual takes the form of a history-independent action cutoff, as does that of a present-biased naif, since she thinks she will act like a time-consistent individual in the future.

current value function is $W_{it}^{\hat{\rho}}(c) = \begin{cases} \beta\delta y - c & \text{if act} \\ \hat{\rho}\beta\delta EV_{it}^{\hat{\rho}} & \text{if do not act} \end{cases}$, with values of $\beta = 1$ for time consistent individuals and $\beta < 1$ for present-biased individuals. The cutoff strategies are then:

$$\begin{aligned} c^*_{TC,t} &= \delta\left(y - \hat{\rho}EV^{\hat{\rho}}_{TC,t+1}\right) \\ c^*_{N,t} &= \beta c^*_{TC,t} \end{aligned}$$

Present-biased individuals' time inconsistent preferences result in their the current value function $W_{it}^{\hat{\rho}}$ differing from the perceived continuation value function $V_{it}^{\hat{\rho}}$. An individual's continuation value function discounts utils exponentially (since all these periods are in the future), even if he is present-biased. It is a *perceived* continuation value function because it depends on perceived memory $\hat{\rho}$, as well as perceived future strategies. Naifs believe they will follow the time consistent strategy in the future, and so their perceived continuation value is the same as that of the time consistent individual: $V_N^{\hat{\rho}}(c) = V_{TC}^{\hat{\rho}}(c)$.

The perceived continuation value functions $V_i^{\hat{\rho}}$ for each type are then given as follows:

$$\begin{aligned} V^{\hat{\rho}}_{TC,t}(c) &= \max\left\{\delta y - c, \hat{\rho}\delta EV^{\hat{\rho}}_{TC,t+1}\right\} \\ V^{\hat{\rho}}_{N,t}(c) &= V^{\hat{\rho}}_{TC,t}(c) \end{aligned}$$

At the deadline, the continuation value of the task is zero. Only when $\rho = \hat{\rho}$ and $\hat{\beta} = \beta = 1$ are the current value, continuation value, and welfare functions the same.

The strategies (planned and actual action thresholds c_t^*) for a time-consistent individual are that of a standard optimizing individual: they maximize their perceived expected utility $W_{TC}^{\hat{\rho}}$, given their perceived memory $\hat{\rho}$. A time-consistent individual always predicts her behavior correctly, conditional on remembering, but may mispredict her probability of remembering. Naifs maximize their perceived expected utility $W_N^{\hat{\rho}}$, given their perceived memory $\hat{\rho}$, and given that they expect to act like a time-consistent individual in all future periods. Naifs therefore mispredict their behavior, conditional on remembering, and may also mispredict their probability of remembering.

2.3 Memory and the Probability of Action

An individual can only act in period t if the task is still active. I define a task to be "active" whenever the individual did not act in a previous period and the individual has not forgotten. The probability an individual acts in period t is given by:

$$\Pr(\text{act in } t) = F(c_{it}^*)\Pr(\text{task active in } t)$$

The task is always active in period 0 when the individual first has the opportunity to act. Thereafter, the probability the task is active is given by:

$$\Pr(\text{task active in } t) = \rho^t \prod_{j=0}^{t-1} \left(1 - F\left(c_{ij}^*\right)\right)$$

Proposition 1 then shows how memory and beliefs about memory affect behavior.

Proposition 1. *Increasing memory ρ, holding constant beliefs about memory $\hat{\rho}$, raises the probability a task is active in each period t but does not affect the probability of completing an active task. Increasing beliefs about memory $\hat{\rho}$ lowers the minimum cost c_{it}^* at which time consistent and present-biased individuals will act in a given period ($\frac{dc_{it}^*}{d\hat{\rho}} \leq 0$), and thus lowers the probability an individual acts in period t conditional on the task still being active.*

This proposition gives the intuition for how ρ and $\hat{\rho}$ can be separately identified from behavior. Memory itself affects behavior by changing the probability the task is forgotten. Note that beliefs about memory $\hat{\rho}$ only affect behavior via the choice of action cost thresholds. Better perceived memory enables delay, as increasing perceived memory increases the perceived continuation value. The net effect of raising $\hat{\rho}$ on both types is still to decrease c_i^*, and therefore decrease the probability of action in a given period, conditional on the task still being active.

Grundgeiger et al. (2013) shows an example of how the anticipation of future reminders for a task can lower the probability an individual completes an active task in a given period. In a healthcare simulation using nurses, they randomized whether nurses were given visual cues for various important tasks. These visual cues served as reminders, were anticipated by nurses, and so can be interpreted as raising $\hat{\rho}$. Nurses without the anticipated reminders were more likely to finish their current task before dealing with a potential interruption by other staff; nurses who had reminders were more likely to accept the interruption and postpone doing the task.

2.4 An Example: No Benefit of Delay

How memory affects whether present-biased individuals act at a deadline can be illustrated in a simple example. Consider the case in which the cost of completing the task is constant across all periods with $c_t = c$ for all t. There is a deadline at period T after which the opportunity to do the task disappears. I assume c is low enough to make the task worth doing at some point, but high enough to induce procrastination in present-biased individuals with perfect memory. Precisely, I assume the task is "$\beta-$worthwhile" ($\beta\delta y > c$) and that a naif will procrastinate with perfect memory: $c(1-\beta\delta) > (1-\delta)\,\delta\beta y$.

Time-consistent individuals are unaffected by a deadline or memory in this case: so long as ever acting is optimal ($\delta y \geq c$), they act immediately. In contrast, the deadline matters for the present-biased naif. With perfect memory, a present-biased naif will always plan to act tomorrow if she doesn't act today, since she thinks she will be a time-consistent individual in the future. Due to present-bias, she would choose to wait until tomorrow and get utility $\beta\delta(\delta y - c)$ rather than act today and get utility $\beta\delta y - c$. This logic repeats itself until the deadline arrives– at which point, she will act because she can no longer delay. If there is no deadline she will never act.

Adding imperfect memory reduces the probability the naif will act at the deadline in two ways. First, if perceived memory is low enough ($\hat{\rho} \leq \frac{\beta\delta y - c}{\beta\delta(\delta y - c)}$), a naive present-biased individual will also act immediately. A present-biased individual perceives a benefit of $\hat{\rho}\beta\delta(\delta y - c)$ to waiting until tomorrow to act, so as $\hat{\rho}$ declines, the perceived cost of delay is higher. In this way, we see that awareness of limited memory can have a benefit for present-biased individuals. For instance, $\hat{\rho} = 0$ emulates the effect of a commitment device, leading present-biased individuals to act immediately.

Second, for values of $\hat{\rho}$ not low enough to induce immediate action, the naif will procrastinate and only act at the deadline if she remembers. The naif will only remember with probability ρ^T. The probability of forgetting at the deadline can be substantial: even if ρ is close to 1, the deadline may be far away. For instance, with $T = 30$ days and a 2% chance the naif forgets each day ($\rho = 0.98$), the probability of acting at the deadline is only $0.98^{30} = 55\%$.

3 Simulating Behavior and Estimating Parameters

3.1 Method: Estimation via GMM

Armed with a model of how tasks are completed when individuals have present-bias and limited memory, we can then turn to interpreting behavior. In this section, I examine the observed patterns of behavior produced by the interaction of limited memory and present-bias, and show that limited memory can explain patterns of behavior that are difficult to explain by present-bias alone. I focus on estimating parameters from task-completion data, and show that ignoring limited memory gives misleading structural estimates of present-bias (and vice versa).

A number of papers have used data on when tasks are completed to structurally estimate present-bias. For instance, Bisin and Hyndman (2014) structurally estimated values of β in their study on when students complete a task relative to a deadline. In contexts without a deadline, DellaVigna and Malmendier (2006) calibrate a value of β from when individuals cancel gym contracts, and Bernheim, Fradkin, and Popov (2015) estimate values of β that rationalize time to enrollment in retirement savings plans and default taking.[19] Because these models assume perfect memory, a failure to act due to limited memory will affect estimates of present-bias parameters and costs of action. Moreover, individuals that recognize their limited memory may act earlier so as to avoid forgetting, which may create behavior that seems inconsistent with present-bias.

As in Section 2, consider a task with a deadline T. The economist observes only the period the task is completed in (if any). The economist thus observes a vector of task completion probabilities $\mathbf{z} = (z_T, ..., z_0)'$, with z_τ being the probability the task is completed τ periods before the deadline. For simplicity, I assume that δ, y, and the distribution of c are known, but that we cannot observe the realizations of c in any given period.

From this data, we can recover estimates of β, ρ,and $\hat{\rho}$ via Generalized Method of Moments (GMM).[20] Each set of parameters $(\beta, \rho, \hat{\rho})$ produces a vector $\hat{\mathbf{z}}$ of predicted task completion probabilities, which can be calculated analytically or via simulation. We then have a set of moment conditions that can be used to estimate the parameters:

$$E\left[\mathbf{z} - \hat{\mathbf{z}}(\beta, \rho, \hat{\rho})\right] = 0$$

where the probability the task is completed each period z_τ provides the individual moment conditions. Then, we can estimate our parameters via GMM as:

$$\beta, \rho, \hat{\rho} = \arg \min_{\beta, \rho, \hat{\rho} \in [0,1]} \left[\mathbf{z} - \hat{\mathbf{z}}(\beta, \rho, \hat{\rho})\right]' \left[\mathbf{z} - \hat{\mathbf{z}}(\beta, \rho, \hat{\rho})\right]$$

Here, I equally weight each moment. I avoid the complications of sampling variation and assume the observed vector $\mathbf{z}$ comes from an large sample, since based on the model in the previous section, I am able to exactly calculate $\hat{\mathbf{z}}(\beta, \rho, \hat{\rho})$.

This method allows us to determine what would happen if a researcher ignored the role

[19] Relatedly, memory could be added to models in papers estimating present-bias based on when individuals attend the gym (Acland and Levy 2015) or get a mammogram (Fang and Yang 2015). It would also be possible to estimate the interaction of present-bias and limited memory in Damgaard and Gravert (2014)'s study of when individuals give to charity, Tasoff and Letzler (2014)'s examination of rebate claim behavior, and Milkman, Roger, and Bazerman (2009)'s study of when DVDs are returned.

[20] For simplicity, I consider only time-consistent and naive present-biased individuals, so $\hat{\beta} = 1$. The degree of naiveté could also be estimated, either directly from task completion rates or from additional information, such as willingness to use commitment devices.

of memory or present-bias. For instance, researchers ignoring memory (assuming $\rho = \hat{\rho} = 1$) would estimate a β as follows:

$$\beta_{IgnoreMemory} = \arg \min_{\beta \in [0,1]} \left[\mathbf{z} - \hat{\mathbf{z}}(\beta, 1, 1)\right]' \left[\mathbf{z} - \hat{\mathbf{z}}(\beta, 1, 1)\right]$$

Similarly, researchers ignoring present-bias (assuming $\beta = 1$) would estimate:

$$\rho_{IgnorePB}, \hat{\rho}_{IgnorePB} = \arg \min_{\rho, \hat{\rho} \in [0.1]} \left[\mathbf{z} - \hat{\mathbf{z}}(1, \rho, \hat{\rho})\right] \left[\mathbf{z} - \hat{\mathbf{z}}(1, \rho, \hat{\rho})\right]'$$

3.2 Example: Switching Health Plans

Consider an individual changing their health insurance plan from the default choice to their most preferred choice. Let the gain from switching be about twice as large as the transaction costs of making the switch.[21] Then, normalize the value[22] of completing the task to be $y = 1$, and let c be uniformly distributed between 0.4 and 0.6. Let there be an open enrollment deadline in $T = 10$ periods; each period could be thought of as a week. Set $\delta = 0.99$.

Assuming perfect memory with present-bias presents a puzzle. While many people do not in fact switch (Handel 2013, Ericson 2014a), a simple model of moderate present-bias says nearly everyone should take action at or before the deadline– with a large spike in task completion at the deadline. At the deadline, an individual who remembered would act so long as $\beta\delta y > c$; as a result, even if they drew the highest cost of action in the last period, they would act so long as $\beta > 0.61$, which is above typical estimates of β.

Suppose the individuals completing the task actually had $\beta = 0.8$ and $\rho = \hat{\rho} = 0.9$. The solid black line with ■ markers in both panels of Figure 1 plots the resulting probability of action each period: it starts at about 19% in the first period, declines over time to about 1.5% with 1 period left before the deadline, and slightly increases to 4.1% in the very last period, as individuals who still remember the task are induced to act by the deadline. Overall, 71.3% complete the task in the simulated data. I can recover exactly these true parameters when I estimate the model by GMM.

Examine the top panel of Figure 1. The dashed black line with $\times$ markers shows the pattern of behavior that would result from an individual with the correct $\beta = 0.8$ but

[21] For instance, in the U.S. Medicare Part D prescription drug insurance program, the gain for a current enrollee to switching plans has been estimated to be $50-$100 (Ericson 2014); a reasonable estimate for transaction costs of switching is about half that amount. Unfortunately, we don't know much about the timing within the open enrollment period of when people take action to switch plans.

[22] In this model, all that matters is the ratio of benefits y to costs c, since the net benefit of action is $\delta y - c$. Thus, it could be a $10 task with $5 costs or $1000 task with $500 costs. In a more complicated model with errors, however, the frequency of errors may depend on the size of the stakes involved.

$\rho = \hat{\rho} = 1$. It predicts no action until 1 period before the deadline, then a 2% probability of action with 1 period remaining, and then a 98% probability of action in the very final period. Because this simulation ignores the urgency an individual faces as a result of limited memory, it cannot explain why a present-biased individual would act with 10 periods left until the deadline. Moreover, because it neglects the possibility individuals might forget the task, it cannot explain why they do not act at the deadline.

Table 1 shows the estimates that would result if a researcher ignored memory and estimated the level of present-bias (and vice versa). If you assumed perfect memory (fixed ρ and $\hat{\rho}$ at 1), Table 1 shows that the resulting best estimate of β would be 0.98–very close to 1 and substantially above the true $\beta = 0.8$. The researcher would thus miss the existence of present-bias. In the top panel of Figure 1, the gray line with circular markers shows that the predicted pattern of behavior with $\beta = 0.98, \rho = \hat{\rho} = 1$. The predicted pattern, at first glance, seems close to the observed pattern. Each period before the deadline, the predicted probability of action is slightly higher than the observed probability of action (e.g. 20% v. 19% in the first period, declining to 3.5% v. 1.5% in the period before the deadline). It then underpredicts the small spike at the deadline (3.9% versus the observed 4.1%).

There are important differences between the true estimates and the estimates assuming perfect memory. The $\beta = 0.98, \rho = \hat{\rho} = 1$ set of parameters predicts that all individuals will eventually complete the task, despite only 71.3% doing so. Essentially, this set of parameters explains the absence of a spike in task completion at the deadline by assuming away the procrastination problem and having the task be completed earlier. Moreover, this set of parameters predicts that reminders should have no effect on behavior. In contrast, providing a surprise reminder in the last period would have increased task completion to 100% and produced a large spike in completion at the deadline. Similarly, this set of parameters predicts that externally imposed deadlines would not be beneficial, since individuals don't procrastinate much if $\beta = 0.98$. In contrast, with $\beta = 0.8$, deadlines can be quite helpful (see Section 6).

We can conduct a similar exercise for what would happen if the researcher ignored present-bias and set $\beta = 1$. In this case, Table 1 shows that the best estimate of ρ is 0.87, similar to the truth, but also finds $\hat{\rho} = 1.00$: substantial overconfidence in memory. The gray line with circular markers in the bottom panel of Figure 1 shows the pattern of behavior that results from $\beta = 1, \rho = 0.87$, and $\hat{\rho} = 1.00$. It is relatively close to the true line, with an overestimation of the probability of action in the first three periods but a steeper decline. To see why ignoring present bias leads to an estimate of overconfidence in memory, examine the dashed black line with $\times$ markers in the bottom panel, which shows what would result if $\beta = 1$, but with the true $\rho = \hat{\rho} = 0.9$. In this case, an individual would be very likely (73%) to act at the first opportunity, since they recognize they might forget the task. While the reason individuals don't act at the first opportunity is actually due to present-bias, overconfidence in memory can lead to a similar lack of action.

Table 1: Estimated Parameters from Observed Behavior

	β	ρ	$\hat{\rho}$
True Parameters	0.8	0.9	0.9
Ignoring Memory	0.98	Set to 1	Set to 1
Ignoring Present-bias	Set to 1	0.87	1.00

Notes: Estimated via GMM in Stata. Parameters: c is uniformly distributed between 0.4 and 0.6, δ =0.99, y=1.

4 Present-bias Limits Uptake of Memory Aids

One empirical puzzle is that memory is low and reminder interventions seem to matter in many domains, even though the cost of memory aids is relatively small in comparison to the benefits of completing the tasks. Individuals can improve their effective memory through the use of memory technologies, which include planners, personal digital assistants, human assistants, lists, reminder services, etc. Why then does memory still matter? This section examines the choice to invest in one's own memory. It shows that present-biased naifs can fail to take up inexpensive memory aids even if there are large net gains to doing so.

I consider a simple and stark reminder system, which I term "writing it down". As in Section 2, the individual chooses whether to act each period based on their draw of the task cost c_t. However, the individual also has another option each period: if she does not act, she may write the task down. Writing it down entails a cost of w today, which in turn guarantees perfect memory ($\rho = \hat{\rho} = 1$) in all future periods. (I continue to require that individuals of each type play perception-perfect strategies, including the decision of whether to write.)

Proposition 2 shows that while inexpensive memory technologies bound the cost of limited memory for time consistent individuals, the cost of limited memory for a present-biased naif can be as large as the discounted value of the task. I define the cost of limited memory to be the expected utility an individual would receive from getting the task with perfect memory, minus the expected utility an individual would receive from getting the task with imperfect memory.

Proposition 2. *Assume individuals have correct beliefs about memory ($\hat{\rho} = \rho$). The cost w of writing bounds the cost of limited memory for a time consistent individual, but not for a*

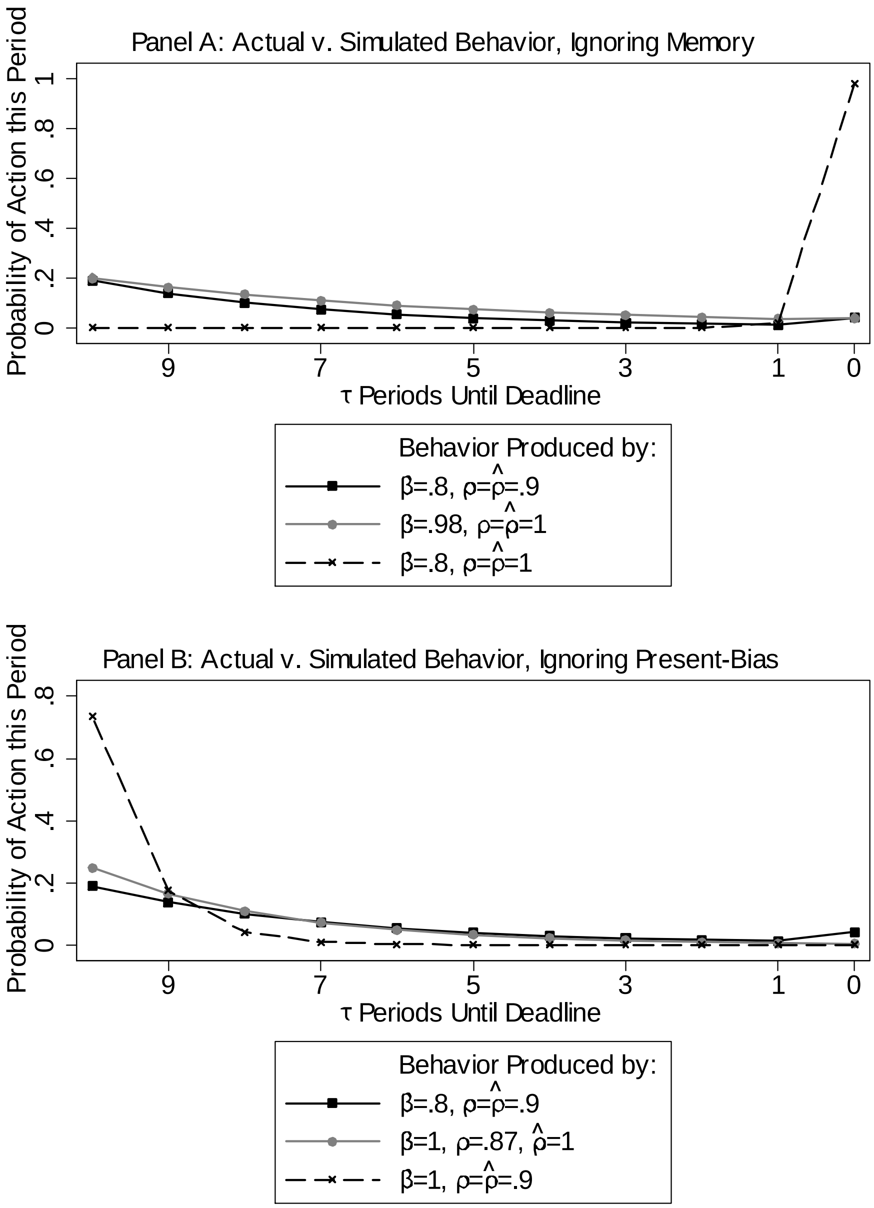

Figure 1: How Ignoring Either Present Bias or Memory Affects Model Estimates. Parameters: c is uniformly distributed between 0.4 and 0.6, $\delta = 0.99$, $y = 1$.

present-biased naif. The cost of limited memory for a present-biased naif can be as large as the net value of the task to a time-consistent individual.

A time consistent individual will write immediately or not all.[23] Since a time-consistent individual with correct beliefs will write it down if it is optimal to do so, it is never beneficial to force them to write, or to subsidize reminder systems. With cheap memory aids, we should expect limited memory to play a relatively unimportant role for time-consistent individuals with correct beliefs.

Present-biased naifs believe they will receive the same benefits as time consistent individuals from memory aids. (Section 5 shows that is not the case.[24]) However, the decision of whether to actually write it down is different with present-bias: the costs of writing are immediate but the benefits come in the future. Present-biased naifs can procrastinate on writing it down. They will never write if the cost of writing is more than the cost of going one period with imperfect memory, because they misperceive their future actions and plan to write tomorrow. Hence, adding the ability to invest in reminders is akin to overlaying one procrastination problem (the choice of whether to write) on top of another (the choice of whether to act).

The large cost of limited memory for a present-biased naif can be shown in the simple example from Section 2.4, in which the cost of completing the task is constant across all periods. Because a time-consistent individual acts immediately in this example, a present-biased naif always thinks she will act tomorrow if she does not act today, so long as she remembers. The naif then has a choice between writing today and remembering tomorrow for sure, versus risking forgetting tomorrow. She will only write today if the cost of writing is low enough: if $w < \beta\delta\left(1-\rho\right)\left(\delta y - c\right)$. If w exceeds that value, she never writes, and each day plans to act tomorrow. As a result, she will act at the deadline T if she remembers, which happens with probability ρ^T. If memory were perfect, she would act at the deadline for sure. Thus, the cost of limited memory in this case is $\left(1-\rho^T\right)\delta^T\left(\delta y - c\right)$. At the deadline gets longer and longer, the cost of limited memory approaches $\delta^T\left(\delta y - c\right)$. Concretely, take the case where $\beta = 0.8$, $\delta = 0.99$, $\rho = 0.98$, and the net value of the task $\delta y - c = \$25$. Then, an individual will not set up a reminder unless the cost of doing so is less than 40 cents ($0.8 \times 0.99 \times 0.02 \times \$25$).

[23] Writing immediately results from the assumption that the cost of writing is constant. This is likely to be roughly true, but some delay in writing could result if the cost of writing fluctuated substantially.

[24] Section 5 shows that increasing memory via anticipated reminders may harm present-biased individuals. As a result, in some cases, naive present-biased individuals may invest in improving their memory, not realizing that doing so will make them worse off. Individuals who are sophisticated or only partially naïve about their present bias may recognize that anticipated reminders can make them worse off, and though reject even free anticipated reminders.

Proposition 2 can explain why individuals do not set up reminders to adhere to medication, a known problem (Nieuwlaat et al. 2014). Adhering to recommended medication entails potentially large health benefits. For instance, Goldie et al. (2003) examined medication adherence for patients with HIV, and found that increasing adherence would lead to gains in terms increased length and quality of life that justified even relatively expensive adherence interventions (e.g. \$500 to \$1000/month). Text messaging reminders, personal organizers, and electronic calendars are very inexpensive, and evidence shows they can be effective in increasing adherence for many diseases, including HIV (Sabin et al. 2015) and high-blood pressure (Bobrow et al. 2016). Why don't patients set up these reminders themselves? Proposition 2 offers an simple explanation: naive procrastination. The perceived cost of waiting a single day to set up a reminder system is low, but the naive individual keeps procrastinating on setting up the reminder.[25]

Proposition 2 also provides an additional explanation individuals are overconfident about the probability they remember to do something. For instance, in Ericson (2011), participants forecasted a 75% chance they would remember to claim a \$20 payment, but only 50% did so. One explanation is that these individuals never planned to use a reminder, and simply overestimated the probability they would remember. Proposition 2 suggests that if these individuals were present-biased naifs, they may have planned to use a reminder but then procrastinated on setting it up and forgot.

5 Unanticipated v. Anticipated Reminders

Providing reminders has been proposed as a policy solution for inaction in many domains. Section 4 suggests that doing so might be valuable, since present-biased individuals may not use available memory aids. However, we lack a theory of how to design effective reminders. In this section, I show that the interaction of memory and procrastination highlights a distinction between reminders that individuals anticipate and reminders that are unanticipated/a surprise. While the theory below shows that this distinction is important, many papers about testing reminders (e.g. for health behaviors, such as pill taking, or for appointment scheduling) do not make it clear to what extent study participants anticipated the reminders.

Memory technologies can modeled in a variety of ways. Here, I consider "reminder systems": technologies that lower the probability of forgetting each period: in the model, they raise ρ. A reminder system can be totally effective (raising ρ to 1) or partially effective.

[25] There is an alternative story that seems implausible: adhering to medication entails side-effects and financial costs that make many patients just about indifferent between taking and not taking medication, despite wide variation in health risk, types of diseases, incomes, and insurance status.

For instance, a to-do list would raise the probability the task is remembered each period, but the to-do list might not be checked every day.

The crucial distinction between anticipated and unanticipated reminder systems is that perceived memory $\hat{\rho}$ and thus action cost thresholds c^*_{it} are unaffected by unanticipated reminders. For instance, a participant in a study who is told they that will receive weekly text message reminders to call to schedule an appointment can rely on future reminders when deciding when to call, while an individual who is not told that they will receive these text messages faces more urgency to call, since they can't be assured that they will remember to act in the future. Precisely:

Definition 1. *Anticipated reminder systems raise perceived memory $\hat{\rho}$ and actual memory ρ by the same amount. Unanticipated reminder systems raise actual memory ρ but hold perceived memory $\hat{\rho}$ constant.*

The following propositions show that anticipated can have very different effects from unanticipated reminder systems. While both types of reminder systems make time-consistent individuals better off, they have different effects on present-biased individuals. The first proposition concerns the observable probability the task is completed, while the second examines welfare.

Proposition 3. *Assume individuals start with correct memory beliefs ($\hat{\rho} = \rho$). For both time-consistent and present-biased individuals, unanticipated reminders always (weakly) raise the probability the task will be completed, while anticipated reminders may raise or lower the probability the task will be completed.*

Proposition 4. *Assume individuals start with correct memory beliefs ($\hat{\rho} = \rho$). Time consistent individuals are always made better off by both types of reminder systems, and prefer anticipated to unanticipated reminders. Present-biased individuals may be harmed by anticipated reminders but will always benefit from unanticipated reminders. Anticipated reminders will are guaranteed to lower the welfare of present-biased individuals if the probability of action in the first period is high enough, and guaranteed to increase welfare if β is close to 1.*

Present-biased individuals face a trade-off that comes from anticipated reminders. Improving memory itself has a welfare gain when there is a potential benefit of delay due to stochastic task costs: the individual is less likely to forget when she waits. Yet when beliefs adjust, better memory also has a cost, because strategies adjust and delay (already inefficiently high for present-biased individuals) is more likely. The harm comes in two ways: first, they may be less likely to complete the task, and second, they may complete the task later (and at higher expected cost). When a present-biased individual is already very likely to act immediately, increasing memory has a strategy response that leads to additional delay, which is a welfare cost, but there is limited direct gain from remembering, since the individual was not delaying often. Conversely, when β is close to 1, they are already near the optimum from a welfare perspective, and so lose less from strategies adjusting; the gain from increasing the probability of remembering to act dominates.

Evidence on how anticipated versus unanticipated reminders affect behavior is limited; I am unaware of any studies directly comparing the two. However, there is some suggestive evidence that anticipated reminders can lower task completion rates differentially for present-biased individuals. For instance, Fernandes and Lynch (2012) examined tasks student participants intended to perform. They compare high v. low "propensity to plan" individuals. It is likely that naive present-biased individuals have a lower propensity to plan (see discussion in Lynch et al. 2010), even though propensity to plan is distinct from time preference. Fernandes and Lynch (2012) had participants set up reminders for themselves, which were thus anticipated by the participants. These reminders led low propensity-to-plan participants both to schedule tasks later and to complete them with lower probability, but increased the probability high propensity-to-plan participants completed the tasks. The propositions can explain this effect: anticipated reminders can harm present-biased individuals but not time-consistent individuals.

Proposition 4 guides our interpretation of how welfare is affected by a change in the probability of task completion due to reminders.[26] For time-consistent individuals (with correct beliefs), reminders that lower or delay the probability of task completion must still be interpreted as making them better off.[27] However, with present-bias, anticipated reminders can enable delay that is harmful to the individual. Thus, in contrast to time consistent individuals, present-biased individuals can indeed be worse off with an anticipated reminder intervention that lowers the probability of making an important health appointment.

5.1 Example: Anticipated Reminders For Switching Health Plans

I show how anticipated reminders affect behavior and welfare using the example from Section 3.2, switching to a preferred health plan. Figure 2 shows how giving anticipated reminders to a present-biased naif with $\beta = 0.8$ affects the probability of completing the task (left panel) and expected utility from the task (right panel). The x-axis varies the values of ρ; beliefs $\hat{\rho}$ are correct and equal ρ. Anticipated reminders of different quality

[26] See also related results comparing attentive versus inattentive individuals in Taubinsky (2014).

[27] How can anticipated reminders can raise the welfare of time-consistent individuals even if they lower the probability of task completion? The intuition is that the reminders enable the individual to wait for a good task draw even at some risk of forgetting, rather than completing the task earlier at high cost.

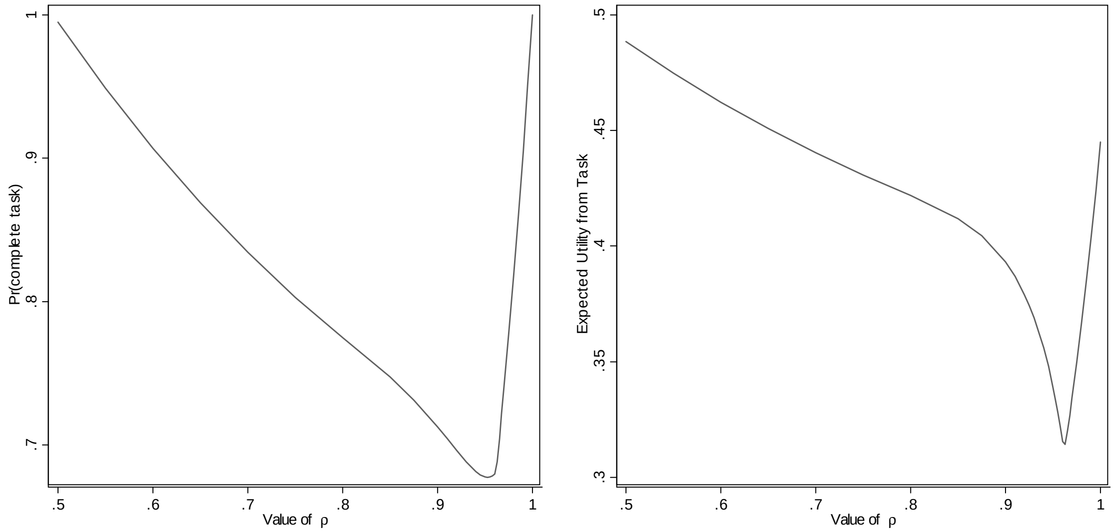

Figure 2: The Effect of Anticipated Reminders on a Present-biased Naif. Parameters: c is uniformly distributed between 0.4 and 0.6, $\delta = 0.99$, $y = 1, \hat{\beta} = 1, \beta = 0.8$.

move the individual from one value of ρ to another. The figure shows that increasing ρ actually decreases expected utility until about $\rho = 0.96$, after which expected utility begins to increase in ρ. (The probability of completing the task moves closely with the expected utility). Thus, imperfect reminders can be quite harmful: a to-do list that is checked only 95% of days might be worse than no to-do list at all. Also note that even perfect reminders leave individuals with lower utility than individuals with relatively poor memory (i.e. an individual is better off with $\rho < 0.7$ than with $\rho = 1$).

6 Deadline Setting and Timing of Reminders

One shot reminders (e.g. announcements in class, email updates, phones) are often used to bring a task back to mind. The effect of these reminders depends on when they are delivered relative to the deadline, and whether they are anticipated in advance. Understanding when to deliver a reminder first requires understanding how deadlines affect behavior, and optimal deadline setting is affected by both memory and present-bias. In this section, I first establish some results about how to set deadlines with limited memory and present bias when a reminder will not be delivered. I then show when to optimally deliver a one-shot reminder, relative to a deadline. Finally, I consider the joint decision of how to optimally set a deadline, given that one can optimally deliver a one-shot reminder.

6.1 Deadlines

While present-bias alone can motivate deadlines (O'Donoghue and Rabin 1999b, Herweg and Muller 2011), time-consistent individuals with perfect memory are never made better off by shorter deadlines (see e.g. Proposition 6 in Taubinsky 2014). Holman and Zaidi (2010) establish that time-consistent individuals with imperfect memory but correct beliefs are also hurt by shorter deadlines, though shorter deadlines can raise the probability the task is completed. However, with incorrect beliefs, shorter deadlines can make time-consistent individuals better off (Holman and Zaidi (2010) and Taubinsky (2014)). To see the intuition, take the limit case in which the individual is certain to forget but believes she is certain to remember. Extending the deadline from today to tomorrow means she may delay the task to wait for a better cost draw, but then will forget.

The results above extend intuitively to the case of limited memory and present bias. Moreover, I show in Proposition 5 that if shortening a deadline increases the probability an individual completes the task, they cannot have perfect memory even if they are present-biased.

Proposition 5. *For individuals with perfect memory, both time consistent and present biased, shorter deadlines lower the probability the task will be eventually completed. With imperfect memory, shorter deadlines can raise the probability of task completion for both time-consistent individuals and present-biased individuals. Shorter deadlines can never raise the welfare of time consistent individuals with correct memory beliefs, but may raise the welfare of present-biased individuals.*

The intuition for how longer deadlines can lower the probability of completing a task but raise welfare is similar to how anticipated reminders can lower task completion and raise welfare. Simply consider an individual whose cost of action each period is either zero or slightly below δy. Prior to the deadline, she only acts with the zero cost draw, but at the deadline she will act at either cost if she remembers, though the net benefit is small if she draws $c \approx \delta y$. A shorter deadline means she is more likely to remember (and act) at the deadline, but would bear a higher expected cost.

Shu and Gneezy (2010) find that longer deadlines for redeeming a gift certificate lowered redemption probability. Similarly, Taubinsky (2014) finds that longer deadlines lower the probability of task completion (when reminders are not provided). This is a puzzle from the perspective of present-bias with perfect memory: Proposition 5 shows that even though lengthening the deadline can harm these individuals, because they delay completing the task, doing so will not lower the probability the task is completed. Yet the results can be explained by imperfect memory: the additional delay from a longer deadline leads to forgetting and a failure to act.

Example: Deadlines for Changing Health Plans.—Proposition 5 implies that the optimal deadline length will need to be calibrated to the environment, including parameters for memory and present-bias. I continue to illustrate the considerations involved with setting a deadline using the example from Section 3.2, switching to a preferred health plan.

I begin by defining notation for the value of a deadline. Let the deadline be given by T, and let τ be the number of periods before the deadline. Define $EU_{i,\tau}$ to be the expected utility (welfare) of getting the task with τ periods before the deadline, with no reminders forthcoming. We have for each type i:

$$EU_{i,\tau} = \int_0^{c^*_{i,T-\tau}} (\delta y - c)\, dF(c) + \left[1 - F\left(c^*_{i,T-\tau}\right)\right] \delta\rho EU_{i,\tau-1}$$

where $c^*_{i,t}$ is the cutoff for acting in period t. (In the equation above, note $t = T - \tau$). Note that the cutoff strategies $c^*_{i,t}$ are chosen based on individuals' current value function, which differs from welfare (unless they are a time-consistent individual with correct beliefs).

In Figure 3, I simulate the utility of different deadline lengths, for different values present-bias, memory, and confidence. (No reminders are delivered in these simulations.) Each panel compares a time-consistent individual to a naive present-biased individual. The left-hand panel examines the case of perfect memory: here, time-consistent individuals gain from longer deadlines, while naifs are made worse off by longer deadlines. Here, with perfect memory, there is an important tradeoff if the level of present-bias is unknown, because the cost to a time-consistent individual from shorter deadlines is similar to the benefit received by naifs from shorter deadlines.

The middle panel shows that with limited memory, longer deadlines are much less valuable for time-consistent individuals than they were in the left panel, as the prospect of forgetting induces them to act quickly. For present-biased naifs, limited memory makes shorter deadlines even more beneficial than they were with perfect memory. Under longer deadlines, the naifs inefficiently delay completing the task, which leads to more forgetting. In this panel, there is little tension between the right deadline for different types of individuals: naifs strongly benefit from short deadlines, while time-consistent individuals are roughly indifferent.

Comparing only the left and middle panels suggests that variation in present-bias is more important than variation in memory for the choice of the deadline, at least in this example. However, the right-most panel shows how large the interaction between overconfidence and present-bias can be. With overconfidence about memory, even time-consistent individuals benefit from shorter deadlines. Moreover, overconfidence has a much bigger effect on present-

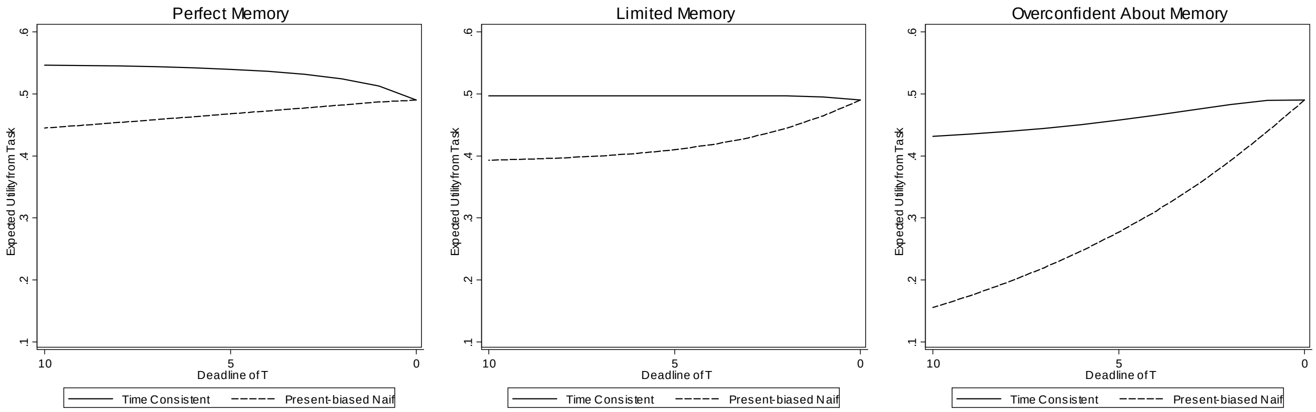

Figure 3: Welfare and Deadlines. Parameters: c is uniformly distributed between 0.4 and 0.6, $\delta = 0.99$, $y = 1$.For time-consistent $\beta = 1$, for naif $\hat{\beta} = 1, \beta = 0.8$. For perfect memory, $\rho = 1$, for limited memory $\rho = \hat{\rho} = 0.9$,for overconfident about memory, $\hat{\rho} = 1, \rho = 0.9$.

biased individuals: with overconfidence, the gain to naifs from shorter deadlines becomes quite large.

6.2 Timing One-Shot Reminders

How should one-shot reminders, such as email announcements, be designed so that they are most effective? Here, I consider the optimal timing for an *unanticipated* reminder relative to a set deadline. Providing unanticipated reminders can only increase welfare for both present-biased and time consistent individuals. (In contrast, recall that for present-biased individuals, anticipated reminders may enable procrastination and make them worse off.)

The optimal time to provide a reminder trades off two effects. First, the probability the reminder will be useful increases as time goes on, since the individual is more likely to have forgotten. Second, the value of the reminder (conditional on it being useful) depends on the number of periods until the deadline. If the reminder comes too close to the deadline, the individual may not have time to act after being reminded (i.e. if they are unlucky and get a high task cost draw); if the reminder comes too early, the individual may delay action (due to procrastination or memory overconfidence) and forget before acting.

Reminders are only useful[28] when the individual has forgotten and not yet acted. The probability a reminder is useful increases as time passes, and can be recursively defined.

[28] I think of the reminder as being "useful" even if the individual does not act after having received the reminder. The usefulness comes from moving the individual from a state of forgetting to a state of remembering.

Note that a reminder is not useful in period 0 when an individual just gets the task. Then, the following equation gives the probability a reminder delivered t periods into the task is useful:

$$\begin{aligned}\Pr\left(useful_t\right) &= \Pr\left(useful_{t-1}\right) + (1-\rho)\left[1 - F\left(c^*_{t-1}\right)\right]\Pr(\text{task active in } t-1) \\ &= (1-\rho)\sum_{k=0}^{t-1}\rho^k \prod_{j=0}^{k}\left(1 - F\left(c^*_{ij}\right)\right)\end{aligned}$$

To understand the logic behind this formula, note that a reminder is always useful if it was useful in the previous period: in these cases, the individual has forgotten, implying she would not be able to act. A reminder becomes newly useful if an individual does not act in a period in which the task was active, and then forgets. This happens $(1-\rho)\left[1 - F\left(c_{t-1}\right)\right]$ of the time that the task is active. Finally, recall that the probability the task was active in period t is the probability the task is remembered (ρ^t) but not acted on previously. Note that the probability a reminder is useful depends directly on t (the number periods after the start of the task that it is delivered), as well as indirectly on the deadline length T, since the deadline affects the action cost thresholds c^*.

When a reminder is useful, it effectively gives the individual the task anew, but with a deadline of the time remaining. Thus, we can apply the results of Section 6.1. Then, for a task with a deadline in period T, the value of a useful reminder given in period t is $EU_{T-t,i}$. The discounted expected value of a reminder t periods into a task with deadline T is then simply given by the following product:

$$ReminderValue_i(t,T) = \delta^t EU_{T-t,i}\Pr\left(useful_t\right)$$

While the probability a reminder is useful is increasing in t, the value of a useful reminder $EU_{T-t,i}$ can increase or decrease with t, since shorter deadlines can sometimes yield higher expected utility when there is procrastination.

Example: Reminders Before Health Plan Enrollment Deadline.—I show how present-bias affects the optimal timing using the example from Section 3.2, switching to a preferred health plan. As before, the cost c_t of doing the task varies from 40% to 60% (uniformly distributed) of the benefit y, and individuals have correct memory beliefs ($\hat{\rho} = \rho = 0.9$). I let the deadline be in period $T = 15$. I distinguish between time-consistent individuals and present-biased naifs with $\beta = 0.8$.

The top panel of Figure 4 shows how the value of one-shot reminders varies by when they are given: it plots $ReminderValue$, separately for time-consistent and present-biased individuals. The middle panel plots the probability the reminder will be useful, $\Pr\left(useful_t\right)$, and the bottom panel plots the value of a useful reminder $EU_{i,\tau}$,the value of getting the task with $\tau = T - t$ periods before the deadline and with no further reminders forthcoming.

For a time-consistent individual, a reminder is most valuable in period 4, but does not vary that much over time. This individual is very likely to act early (their probability of acting in the first period is about 74%), and so if they have not acted by period 4 it is likely they have forgotten, in which case it is valuable to give them a reminder quickly so that they can act without further delay.

In contrast, in this example, the best time to give a reminder to a present-biased naif is in the very last period. The naif has a low probability of acting early on (only 19% in the first period). Thus, even though the reminder is potentially useful because they may have forgotten (Figure 4, middle panel), it will not actually be used because the naif continues to procrastinate. The high value of the reminder in the final period then comes from the fact that the naif is better off without the ability to procrastinate and forget: it is better to get the task with a short deadline than a long deadline (Figure 4, bottom panel). By contrast, the value of the task to a time consistent individual increases as the deadline lengthens, driving the result that their reminder should be early.

6.3 Jointly Choosing Deadlines and Reminders

The previous subsections considered when to set the deadline, given that no reminder was forthcoming, and then when to set a reminder relative to a deadline. The joint decision– how to choose a deadline, given that an optimal one-shot reminder will be delivered– is also of interest.

For each deadline T,we can find t^*_T, is the optimal period in which to deliver the one-shot unanticipated reminder, given a deadline of T. From the previous section, we have:

$$t^*_T = \arg\max_t ReminderValue_i(t, T).$$

Recall that we defined $EU_{i,\tau}$ to be the expected utility of being given the task with τ periods until the deadline, given no reminders forthcoming. Let $EU^{reminder}_{i,T}$ be the expected utility of being given the task with a deadline T and with an optimally timed one-shot unanticipated reminder. Because unanticipated reminders don't affect strategies, this can just be represented as the sum of getting the task without a reminder plus the discounted expected value of a reminder given in period t^*_T :

$$EU^{reminder}_{i,T} = EU_{i,T} + \delta^{t^*_T} ReminderValue_i(t^*, T)$$

Recall that the optimal deadline could be infinite (as is the case for time-consistent individu-

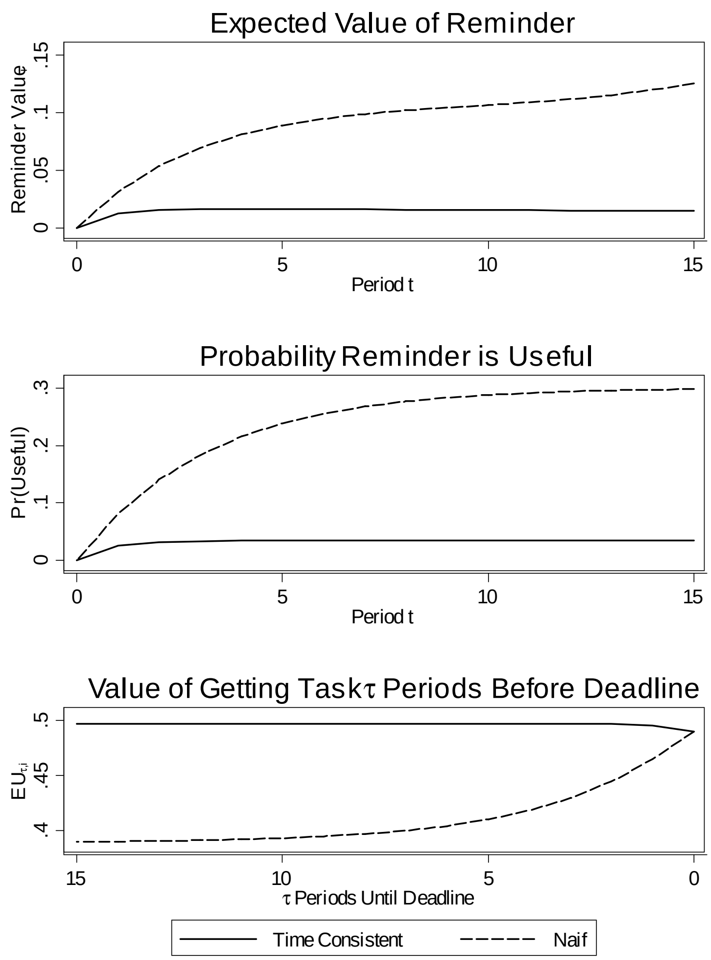

Figure 4: The Timing of One-Shot Reminders, with and without Present-Bias. Assumes reminder is unanticipated, $\hat{\rho} = \rho = 0.9$, $\beta = 0.8, \delta = 0.99$, $y = 1$, c is uniformly distributed between 0.4 and 0.6, and deadline $T = 15$.

als with correct beliefs). If the optimal deadline is finite, it can be found simply by searching over the space of potential deadlines: the optimal deadline $T^* = \arg\max_T EU_{i,T}^{reminder}$.

7 Conclusion

Psychology paints a rich picture of individual decision making. Considering a single phenomenon in isolation may be misleading, which has implications for measuring present-bias from task completion behavior. When procrastinators don't complete a task, it could result from either present-bias or limited memory. Researchers calibrating models of present bias or making inferences about naiveté or sophistication should take into account the effects limited memory may have on individuals' behavior, or estimates will be biased.

In this task completion framework, memory and present bias interact in ways that are interesting, welfare relevant, and that affect reminder and deadline design. Other work could explore how additional biases interact in this framework, such as overoptimism about the distribution of costs (which can manifest itself as the planning fallacy). The interaction of memory and procrastination could be explored in a number of different areas, such as the design of commitment contracts. While this paper focused on a binary decision (individuals either did or did not act), the model could also be extended to an effort choice decision (e.g. as in Herweg and Muller 2011).[29]

These results can also help design better interventions. It can be valuable to tailor reminder interventions based on individuals' level of present-bias. The optimal deadline length varies substantially based on memory. Anticipated reminders are more effective than surprise reminders for time-consistent individuals, but can actually harm present-biased individuals. Moreover, the timing of when to deliver reminders should vary based on present-bias: early reminders help time-consistent individuals complete the task faster, but are wasted on present-bias individuals if they simply procrastinate in response. Beyond just reminders and deadlines, this memory and procrastination framework could also be applied to the design of sales offers (such as rebates) and other incentive contracts. Moreover, just as memory and present-bias affect the use of reminder systems, future work could explore their role in the decision to automate tasks (i.e. having a rent payment automatically deducted versus remembering to send a payment) or choose dynamic renewal defaults (e.g. as in Ericson 2014b), or choose commitment devices (as in Laibson 2015).

[29] I anticipate the results would extend intuitively, since a higher probability of forgetting creates more urgency to act earlier, thereby increasing effort in earlier periods. Thus, while unanticipated reminders won't affect effort choice before the reminder is delivered, while anticipated reminders can lead to lower effort in earlier periods.

REFERENCES

Acland, Daniel and Matthew Levy. 2015. "Naivete, Projection Bias, and Habit Formation in Gym Attendance." *Management Science* 61(1): 146–160.

Ali, S. Nageeb. 2011. "Learning Self-Control." *Quarterly Journal of Economics,* 126(2): 857–93. doi:10.1093/qje/qjr014.

Altmann, Steffen and Christian Traxler. 2012. "Nudges at the Dentist." IZA Discussion Paper 6699.

Ariely, Dan and Klaus Wertenbroch. 2002. "Procrastination, Deadlines, and Performance: Self-Control by Precommitment." *Psychological Science*, 13, 219–224.

Benabou, Roland and Jean Tirole. 2004. "Willpower and Personal Rules", *Journal of Political Economy*, 112 (4): 848-887.

Bernheim, Douglas, Andrey Fradkin, and Igor Popov. 2015. "The Welfare Economics of Default Options in 401(k) Plans." *American Economic Review*, 105(9): 2798-2837.

Bernheim, Douglas and Antonio Rangel. 2009. "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," *Quarterly Journal of Economics* 124:51-104.

Bernheim, Douglas and Raphael Thomadsen. 2005. "Memory and Anticipation." *The Economic Journal*, 115:271-304.

Bisin, Alberto and Kyle. Hyndman. 2014. "Present-Bias, Procrastination and Deadlines in a Field Experiment." NBER Working Paper 19874.

Bobrow, Kirsten et al. 2016. "Mobile Phone Text Messages to Support Treatment Adherence in Adults With High Blood Pressure (StAR): A Single-Blind, Randomized Trial." *Circulation.* doi:10.1161/circulationaha.115.017530

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2015. "Memory, Attention, and Choice." Working Paper.

Burger, N., Gary Charness, and J. Lynham 2011. "Field and Online Experiments on Self-Control," *Journal of Economic Behavior & Organization*, 77, 393–404.

Cadena, Ximena and Antionette Schoar. 2011. "Remembering to Pay? Reminders vs. Financial Incentives for Loan Payments". NBER Working Paper 17020.

Calzolari, Giacomo and Mattia Nardotto. Forthcoming. "Nudging with Reminders? A Field Experiment." *Management Science.*

Carroll, Gabriel, James Choi, David Laibson, Brigitte Madrian, and Andrew Metrick. 2009. "Optimal Defaults and Active Decisions." *Quarterly Journal of Economics* 124,1639-1674.

Damgaard, Mette and Christina Gravert. 2014. "Now or Never! The Effect of Deadlines on Charitable Giving: Evidence from a Natural Field Experiment." Aarhus University Working Paper.

Della Vigna, Stefano, and Ulrike Malmendier. 2006. “Paying Not to Go to the Gym.” *American Economic Review* 96:694-719.

Della Vigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47(2): 315-72.

Ericson, Keith M. 2011. "Forgetting We Forget: Overconfidence and Memory." *Journal of the European Economic Association*, 9(1):43-60.

Ericson, Keith M. 2014a. "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange." *American Economic Journal: Economic Policy*, 6(1): 38-64.

Ericson, Keith M. 2014b. "When Consumers Do Not Make an Active Decision: Dynamic Default Rules and their Equilibrium Effects." *NBER Working Paper* 20127.

Fang, Hanming and Yang Wang. 2015. "Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions." *International Economic Review*, 56(2):565-596.

Fernandes, Daniel and John Lynch. 2012. "Manana: Motivational Effects of Reminders on Accelerating or Delaying Task Completion." Working Paper.

Frederick, Shane, George Loewenstein and Ted O’Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40: 351-401.

Goldie, Sue, et al. 2003. "Projecting the cost-effectiveness of adherence interventions in persons with human immunodeficiency virus infection." *American Journal of Medicine* 115: 632–641.

Gottlieb, Daniel. 2014. "Imperfect Memory and Choice Under Risk." *Games and Economic Behavior* 85:127-158.

Gruber, Jonathan, and Botond Koszegi. 2004. “Tax Incidence when Individuals are Time-Inconsistent: The Case of Cigarette Excise Taxes.” *Journal of Public Economics*, 88(9–10): 1959–87.

Grundgeiger, T. et al. 2013. "Prospective memory in the ICU: the effect of visual cues on task execution in a representative simulation" *Ergonomics* 56(4). doi:10.1080/00140139.2013.765604

Guy, Rebecca. Hocking J, Wand H, Stott S, Ali H, and Kaldor J. 2012. "How effective are short message service reminders at increasing clinic attendance? A meta-analysis and systematic review." *Health Services Research* 47(2): 614-32.

Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review*, 103(7): 2643-82.

Heidhues, Paul and Botond Koszegi. 2010. "Exploiting Naïvete about Self-Control in the Credit Market." *American Economic Review* 100:2279-2303.

Herweg, Fabian and Daniel Muller. 2011. "Performance of Procrastinators: On the Value of Deadlines." *Theory and Decision* 70:329-366.

Holman, Jeffrey. and Farhan Zaidi. 2010. "The Economics of Prospective Memory." University of California Working Paper.

Hsiaw, Alice. 2013. "Goal-setting and Self-control." *Journal of Economic Theory* 148:601-626.

Jackson, C. Kirabo, and Henry S. Schneider. 2015. "Checklists and Worker Behavior: A Field Experiment." *American Economic Journal: Applied Economics*, 7(4): 136-68.

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman. Forthcoming. "Getting to the Top of Mind: How Reminders Increase Saving." *Management Science.*

Karlan, Dean, Melanie Morten, and Jonathan Zinman. Forthcoming. "A Personal Touch: Text Messaging for Loan Repayment." *Behavioral Science and Policy.*

King, J. 2004. "Missed Opportunities: Students Who Do Not Apply for Financial Aid." American Council on Education Issue Brief.

Koch, A. and J. Nafziger. 2014. "Commitment to Self-Rewards." *European Economic Review*, 68: 151-167.

Laibson, David. 1997. “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 112 (2): 443-477.

Laibson, David. "Why Don’t Present-Biased Agents Make Commitments?"*American Economic Review*, 105(5), pp. 267-272.

Letzler, Robert and Joshua Tasoff. 2014. "Everyone Believes in Redemption: Nudges and Overoptimism in Costly Task Completion." *Journal of Economic Behavior & Organization*, 107:107-122.

Levy, R.L., & Loftus, G.R. (1984). Compliance and memory. In J.E. Harris (Ed.) *Everyday Memory, Actions and Absent-Mindedness* (pp. 93-112). New York: Academic Press.

Lynch Jr., J. G., Netemeyer, R. G., Spiller, S. A., & Zammit, A. 2010. "A generalizable scale of propensity to plan: The long and the short of planning for time and for money." *Journal of Consumer Research*, 37(1), 108-128. doi: 10.1086/649907.

McClure, S., Ericson, K., Laibson, D., Loewenstein, G., and Cohen, J. 2007. "Time discounting for primary rewards," *Journal of Neuroscience*, 27:5796-804.

McDaniel, Mark and Gilles Einstein. 2007. *Prospective Memory: An Overview and Synthesis of an Emerging Field.* Sage Publications.

Milkman, Katherine L., Todd Rogers, and Max H. Bazerman. 2009. “Highbrow Films Gather Dust: Time-Inconsistent Preferences and Online DVD Rentals.” *Management Science*, 55 (6): 1047–59.

Mullainathan, Sendhil. 2002. "A Memory-based Model of Bounded Rationality." *Quarterly Journal of Economics*, 117(3): 735-774.

Nieuwlaat et al. 2014. "Interventions for enhancing medication adherence." *The Cochrane Database of Systematic Reviews.* DOI: 10.1002/14651858.CD000011.pub4

O’Donoghue, Ted and Matthew Rabin. 1999a. "Doing it Now or Later," *American Economic Review*, 89(1), pp. 103-124.

O’Donoghue, Ted, and Matthew Rabin. 1999b. “Incentives For Procrastinators.” *Quarterly Journal of Economics* 114: 769-816.

O’Donoghue , Ted and Matthew Rabin. 2001. "Choice and Procrastination," *Quarterly Journal of Economics*, 121-160.

O'Donoghue, Ted, and Matthew Rabin. 2006. "Optimal Sin Taxes." *Journal of Public Economics*, 90(10-11): 1825–49

Paserman, Daniele. 2008. "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation." *Economic Journal*, 117:1418-1452.

Pechmann, Cornelia and Tim Silk. 2013. "Policy and Research Related to Consumer Rebates: A Comprehensive Review." *Journal of Public Policy & Marketing*, 32(2): 255-270.

Read, Daniel and Barbara van Leeuwen. 1998. "Predicting Hunger: The Effects of Appetite and Delay on Choice." *Organizational Behavior and Human Decision Processes* 76(2):189–205.

Reis, Ricardo. 2006. "Inattentive Consumers". *Journal of Monetary Economics*, 53 (8), 1761-1800.

Sabin, Lora. et al. 2015. "Improving Adherence to Antiretroviral Therapy With Triggered Real-time Text Message Reminders: The China Adherence Through Technology Study." *Journal of Acquired Immune Deficiency Syndromes*, 69(5): 551-9. doi: 10.1097/QAI.0000000000000651.

Saez-Marti, Maria, and Anna Sjögren. 2008. "Deadlines and Distractions." *Journal of Economic Theory* 143 (1): 153–76. doi:10.1016/j.jet.2008.02.005.

Schwartzstein, Joshua. 2014. "Selective Attention and Learning." *Journal of the European Economic Association* 12(6): 1423–52. doi:10.1111/jeea.12104.

Shu, Suzanne and Ayelet Gneezy. 2010. "Procrastination of Enjoyable Experiences." *Journal of Marketing Research*, 47(5): 933-944.

Taubinsky, Dmitry. 2014. "From Intentions to Actions: A Model and Experimental Evidence of Inattentive Choice", Harvard University Working Paper.

Supplementary Material for "On the Interaction of Memory and Procrastination"

Keith M Marzilli Ericson

A.I Proofs of Propositions in the Text

Proposition 1. *Increasing memory ρ, holding constant beliefs about memory $\hat{\rho}$, raises the probability a task is active in each period t but does not affect the probability of completing an active task. Increasing beliefs about memory $\hat{\rho}$ lowers the minimum cost c_{it}^* at which time consistent and present-biased individuals will act in a given period ($\frac{dc_{it}^*}{d\hat{\rho}} \leq 0$), and thus lowers the probability an individual acts in period t conditional on the task still being active.*

Proof. First, note that c_{it}^* depends on the perceived continuation value, which does not depend on ρ (only $\hat{\rho}$). The probability a task is active in period t is given by $\rho^t \prod_{j=0}^{t-1} \left(1 - F\left(c_{ij}^*\right)\right)$. It is easy to see that this increases in ρ holding fixed c^*.

The probability an active task is completed is given by $F(c_{it}^*)$, which is unaffected by ρ. To determine how $\hat{\rho}$ affects the probability an individual completes an active task, we must determine $\frac{dc_{it}^*}{d\hat{\rho}}$. Since $c_{i,N}^* = \beta c_{i,TC}^*$, I simply examine $c_{i,TC}^*$. Note that at the deadline T, $c_{TC,T}^* = \delta y$,so $\frac{dc_{TC,T}^*}{d\hat{\rho}} = 0$. In general, in period t, $\frac{dc_{TC,t}^*}{d\hat{\rho}} = -\delta EV_{TC,t+1} - \delta\hat{\rho}\frac{dEV_{TC,t+1}}{d\hat{\rho}}$. We then must examine $\frac{dEV_{TC,t+1}}{d\hat{\rho}}$. Intuitively, this is positive. To see this note that differentiating the formula for continuation value gives:

$$\frac{dEV_{TC,t}}{d\hat{\rho}} = \left[\delta y - c - \delta\hat{\rho}EV_{TC,t+1}\right]\frac{dc_{TC,t}^*}{d\hat{\rho}} f\left(c_{TC,t}^*\right) + \delta\hat{\rho}\left(1 - F\left(c_{TC,t}^*\right)\right)\frac{dEV_{TC,t+1}}{d\hat{\rho}} + \delta\left(1 - F\left(c_{TC,t}^*\right)\right)EV_{TC,t+1}$$

The first term is zero by the envelope theorem, and the last term is positive. Hence, so long as $\frac{dEV_{TC,t+1}}{d\hat{\rho}} \geq 0$,we will have$\frac{dEV_{TC,t}}{d\hat{\rho}} > 0$. Since $\frac{dEV_{TC,T}}{d\hat{\rho}} = 0$, $\frac{dEV_{TC,t}}{d\hat{\rho}} > 0$ and $\frac{dc_{TC,t}^*}{d\hat{\rho}} < 0$ for all $t < T$.

When there is no deadline the problem is an infinite horizon one, and action cost thresholds and continuation values are time invariant. Simplifying the equation above, we have

$$\frac{dEV_{TC}}{d\hat{\rho}} = \delta\hat{\rho}\left(1 - F\left(c_{TC}^*\right)\right)\frac{dEV_{TC}}{d\hat{\rho}} + \delta\left(1 - F\left(c_{TC}^*\right)\right)EV_{TC}$$

or $\frac{dEV_{TC}}{d\hat{\rho}} = \frac{\delta\left(1-F\left(c_{TC}^*\right)\right)}{1-\delta\hat{\rho}\left(1-F\left(c_{TC}^*\right)\right)}EV_{TC} > 0$, and so $\frac{dc_{TC}^*}{d\hat{\rho}} = -\delta EV_{TC} - \delta\hat{\rho}\frac{dEV_{TC}}{d\hat{\rho}} < 0$.

Proposition 2. *Assume individuals have correct beliefs about memory ($\hat{\rho} = \rho$). The cost w of writing bounds the cost of limited memory for a time consistent individual, but not for a* *present-biased naif. The cost of limited memory for a present-biased naif can be as large as the net value of the task to a time-consistent individual.*

Proof. Define $E\tilde{U}_{i,T}$ to be the expected utility (welfare) of getting the task with T periods before the deadline, with the option to write each period. (I use the $\tilde{}$ to distinguish it from the expected utility of getting the task without the option to write).

The loss to an individual due to limited memory is the difference in expected utility from being given the task with deadline T with perfect memory versus imperfect memory: $E\tilde{U}_{iT}^{1} - E\tilde{U}_{iT}^{\rho}$. A time consistent individual who chooses not to act today will write if $-w + \delta E\tilde{U}_{TC,T-1}^{1} > \rho\delta E\tilde{U}_{TC,T-1}^{\rho}$. As a result, a sufficient condition for individuals to write is $-w + E\tilde{U}_{TC,T-1}^{1} > \rho E\tilde{U}_{TC,T-1}^{\rho}$.

The utility of a time-consistent individual who is given the task with imperfect memory is at least as high as the utility from doing the following: as using the action cost threshold for an individual with perfect memory $\tilde{c}^{*}$ in period 0 and then writing it down in period 0 if she does not act. Hence: $E\tilde{U}_{iT}^{\rho} \geq \int_{0}^{\tilde{c}^{*}} (\delta y - c)\, dF(c) + [1 - F(\tilde{c}^{*})] \left[\delta E\tilde{U}_{iT-1}^{1} - w\right]$. Note that $E\tilde{U}_{iT}^{1} = \int_{0}^{\tilde{c}^{*}} (\delta y - c)\, dF(c) + [1 - F(\tilde{c}^{*})] \left[\delta E\tilde{U}_{iT-1}^{1}\right]$, so we have a bound on the loss:

$$E\tilde{U}_{iT}^{1} - E\tilde{U}_{iT}^{\rho} \leq w\,[1 - F(\tilde{c}^{*})]$$

Present-biased naifs only write in period t if $\beta\delta EV_{N,t}^{1} - w > \beta\delta\rho EV_{N,t}^{\rho}\,(optionToWrite)$, where their perceived continuation value if they don't write today accounts for the option to write. An example suffices to show that the cost of limited memory for a present-biased naif can be as large as the net value of the task to a time-consistent individual. Take the case where $c_t = c$ for all t, with a deadline T, and let $\delta = 1$. Let costs be such that the task is "$\beta-$worthwhile" ($\beta y > c$), and memory be high enough that the naif will procrastinate on completing the task:$\rho > \frac{\beta\delta y - c}{\beta\delta(\delta y - c)} = \frac{\beta y - c}{\beta(y-c)}$. We know from Section 2.4 that a present-biased naif will not act in this case until the deadline. With perfect memory, $EU_{N,T}^{1} = (y - c)$, since the individual acts in the last period. With imperfect memory and no option to write, $EU_{N,T}^{\rho} = \rho^{T}(y - c)$, and the loss is $\left(1 - \rho^{T}\right)(y - c)$, which goes to $(y - c)$ as $T \to \infty$. Even with the option to write, the present-biased naif will not write so long as $w > (1 - \rho)\,\beta\,(y - c)$. The naif always thinks she will act tomorrow, and so will only write if $-w + \beta\,(y - c) > \rho\beta\,(y - c)$.

Proposition 3. *Assume individuals start with correct memory beliefs ($\hat{\rho} = \rho$). For both time-consistent and present-biased individuals, unanticipated reminders always (weakly) raise the probability the task will be completed, while anticipated reminders may raise or lower the probability the task will be completed.*

Proof. Unanticipated reminders only affect ρ. By Proposition 1, they thus raise the probability a task is active in each period, but do not change the probability of completing an active task. Since they raise the probability a task is active, and since active tasks are completed with some probability, unanticipated reminders always raise the probability the task will be completed, regardless of whether the individuals are present-biased or time-consistent.

Anticipated reminders lower the probability an active task is completed (since they raise $\hat{\rho}$ and thus lower c^*). However, they raise the probability a task is active in a given period: $\frac{d}{d\rho}\Pr(\text{task active in } t)|_{\hat{\rho}=\rho} > 0$. Recall that:

$$\Pr(\text{task active in } t) = \rho^t \prod_{j=0}^{t-1} \left(1 - F\left(c^*_{ij}\right)\right)$$

By increasing ρ, anticipated reminders raise the probability the task is remembered, which raises the probability the task is active. By increasing $\hat{\rho}$, anticipated reminders lower c^* in all previous periods, thus decreasing the probability the task has been completed previously. For a given period, the change in probability a task is completed is given by the sum of a negative term and a positive term:

$$\frac{\Pr(\text{act in } t)}{d\rho}|_{\hat{\rho}=\rho} = \frac{dc^*_{it}}{d\hat{\rho}} f(c^*_{it}) \Pr(\text{task active in } t) + F(c^*_{it}) \frac{d}{d\rho} \Pr(\text{task active in } t)|_{\hat{\rho}=\rho}$$

For low enough $f\left(c^*_{it}\right)$, the positive effect of ρ on the task being active dominates and $\frac{\Pr(\text{act in } t)}{d\rho}|_{\hat{\rho}=\rho} > 0$. For $F(c^*_{it})$ low enough, the negative effect from $\hat{\rho}$ on c^*_{ii} dominates, and $\frac{\Pr(\text{act in } t)}{d\rho}|_{\hat{\rho}=\rho} < 0$.

To see anticipated reminders have a negative effect on the probability the task is ever completed in any period, consider the two period case with $T = 1$. Then

$$\frac{\Pr(\text{ever act})}{d\rho}|_{\hat{\rho}=\rho} = \frac{\Pr(\text{act in } 0)}{d\rho}|_{\hat{\rho}=\rho} + \frac{\Pr(\text{act in } 1)}{d\rho}|_{\hat{\rho}=\rho}$$

Note that at the deadline, perceived memory is no longer relevant, so $\frac{dc^*_{i1}}{d\hat{\rho}} = 0$. Moreover, the task is active in period 0 with probability 1 and in period 1 with probability $\rho\left(1 - F\left(c^*_{i0}\right)\right)$. Then, we have

$$\frac{\Pr(\text{ever act})}{d\rho}|_{\hat{\rho}=\rho} = \frac{dc^*_{i0}}{d\hat{\rho}} f\left(c^*_{i0}\right)\left[1 - \rho F(c^*_{i1})\right] + F(c^*_{i1})\left(1 - F\left(c^*_{i0}\right)\right)$$

Recall that $\frac{dc^*_{i0}}{d\hat{\rho}} < 0$. Consider starting at the case in which the individual is sure to act in the first period, so $F\left(c_{i0}^{*}\right)=1$. Then, $\frac{\Pr(\text{ever act})}{d\rho}|_{\hat{\rho}=\rho}=\frac{dc_{i0}^{*}}{d\hat{\rho}}f\left(c_{i0}^{*}\right)\left[1-\rho F(c_{i1}^{*})\right]<0$.

A particular example is given in Section 2.4, where $c_t=c$ for all t. Starting at $\hat{\rho}=0$, an anticipated reminder system that raised memory to $\rho>\frac{\beta\delta y-c}{\beta\delta(\delta y-c)}$ but $\rho<1$ actually lowers the probability of action for present-biased individuals: previously, the present-biased agent would immediately act, but now only acts at the deadline with probability ρ^{T}. However, starting at $\rho>\frac{\beta\delta y-c}{\beta\delta(\delta y-c)}$ but $\rho<1$, an anticipated reminder system that raised ρ to 1 increases the probability of action.

Proposition 4. *Assume individuals start with correct memory beliefs ($\hat{\rho}=\rho$). Time consistent individuals are always made better off by both types of reminder systems, and prefer anticipated to unanticipated reminders. Present-biased individuals may be harmed by anticipated reminders but will always benefit from unanticipated reminders. Anticipated reminders will are guaranteed to lower the welfare of present-biased individuals if the probability of action in the first period is high enough, and guaranteed to increase welfare if β is close to 1.*

Proof. For the proof of this proposition, it will be convenient to write EU_{it} as the expected utility for type i in period t, rather than as in the text, where $EU_{i,\tau}$ is the expected utility of getting the task with τ periods before the deadline. To see that unanticipated reminders benefit all types of individuals, recall that unanticipated reminders only change the probability a task is active; they do not affect c_{it}^{*}. Consider an unanticipated reminder that raises ρ. Expected utility at the start of the task can be represented as

$$EU_{i0}=\sum_{t=0}^{T}\Pr(\text{act in }t)\left[\delta y-E\left(c|act\text{ in }t\right)\right]$$

Note that $E\left(c|act\text{ in }t\right)$ is given by $\int_0^{c_{it}^{*}}dF\left(c\right)$,which depends on $\hat{\rho}$ but not on ρ. We can then write $\frac{dEU_{i0}}{d\rho}=\sum_{t=0}^{T}\frac{d\Pr(\text{act in }t)}{d\rho}\left[\delta y-E\left(c|act\text{ in }t\right)\right]$. We know $\frac{d\Pr(\text{act in }t)}{d\rho}>0$ from our previous proposition, and $c_{it}^{*}\leq\delta y$, so $\frac{dEU_{i0}}{d\rho}>0$,and unanticipated reminders benefit both types.

Now consider the effect of anticipated reminders on the welfare of each type. Here, note that $EU_{it}=\int_0^{c_{it}^{*}}\left(\delta y-c\right)dF(c)+\delta\rho\left(1-F\left(c_{it}^{*}\right)\right)EU_{i,t+1}$, and

$$\frac{dEU_{it}}{d\rho}|_{\hat{\rho}=\rho}=\left[\left(\delta y-c_{it}^{*}\right)-\delta\rho EU_{it+1}\right]f\left(c_{it}^{*}\right)\frac{dc_{it}^{*}}{d\hat{\rho}}+\delta\left(1-F\left(c_{it}^{*}\right)\right)EU_{i,t+1}+\delta\rho\left(1-F\left(c_{it}^{*}\right)\right)\frac{dEU_{i,t+1}}{d\rho}|_{\hat{\rho}=\rho}$$

For time-consistent individuals, anticipated reminders have a positive effect on utility. The first term in the above equation equals zero from envelope theorem and definition of c_{TC}^{*}. We know at the deadline $EU_{iT}>0$ and $\frac{dEU_{iT}}{d\rho}|_{\hat{\rho}=\rho}=0$,and so $\frac{dEU_{it}}{d\rho}|_{\hat{\rho}=\rho}>0$ by backward induction. Moreover, we know that time-consistent individuals prefer anticipated reminders to unanticipated reminders, since $\left\{c^*_{TC,1}, c^*_{TC,2}, ...\right\} = \arg\max_{c_t} EU^{\hat{\rho}}_{TC}$, and so will maximize EU^{ρ}_{TC} when $\hat{\rho} = \rho$.

For present-biased individuals, note that $[(\delta y - c^*_{Nt}) - \delta\rho EU_{N,t+1}] > 0$ since $c^*_{Nt} = \beta\left[\delta y - \delta\hat{\rho}EV^{\hat{\rho}}_{TC,t+1}\right]$,and $EU_{N,t+1} < EU_{TC,t+1} = EV^{\hat{\rho}}_{TC,t+1}$. From our previous proposition, $\frac{dc^*_{it}}{d\hat{\rho}} < 0$,and the first term is negative. Thus $\frac{dEU_{it}}{d\rho}|_{\hat{\rho}=\rho} < 0$ as $F(c^*_{it})$ goes to 1 (and the remaining terms go to zero). Similarly, as β goes to 1, the first term goes to zero, and $\frac{dEU_{it}}{d\rho}|_{\hat{\rho}=\rho} > 0$.

For both present-biased and time-consistent individuals, the no deadline (infinite horizon case) has the same logic. Note that $EU_{it} = EU_{it+1}$, so we can rearrange the equation for $\frac{dEU_{it}}{d\rho}|_{\hat{\rho}=\rho}$ to find that

$$\frac{dEU_i}{d\rho}|_{\hat{\rho}=\rho} = \frac{1}{1-\delta\rho\left[1-F\left(c^*_i\right)\right]}\left\{\left[(\delta y - c^*_i) - \delta\rho EU_i\right] f\left(c^*_i\right)\frac{dc^*_i}{d\hat{\rho}} + \delta\left(1-F\left(c^*_i\right)\right)EU_i\right\}$$

and the logic follows as above.

Proposition 5. *For individuals with perfect memory, both time consistent and present biased, shorter deadlines lower the probability the task will be eventually completed. With imperfect memory, shorter deadlines can raise the probability of task completion for both time-consistent individuals and present-biased individuals. Shorter deadlines can never raise the welfare of time consistent calibrated individuals, but may raise the welfare of present-biased individuals or overconfident time-consistent individuals.*

Proof. Compare deadlines of T and $T+1$ periods. Let q_t be the probability the task is completed with t periods remaining, if the task has not been forgotten by that point. With perfect memory, $q_{T+1} = F\left(c^*_{iT+1}\right) + \left(1 - F\left(c^*_{iT+1}\right)\right)q_T$. The individual uses her action cutoff c^*_{iT+1} when there are $T+1$ periods to the deadline. If she does not act, she continues next period just as though she received the task with deadline T. Hence, $q_{T+1} > q_T$, with the inequality strict since $F\left(c^*_{iT+1}\right) > 0$ by the assumption on the distribution F. Note that this logic applies regardless of whether the individual is time consistent or present biased.

With imperfect memory, $q_{T+1} = F\left(c^*_{iT+1}\right) + \rho\left(1 - F\left(c^*_{iT+1}\right)\right)q_T$: the individual can forget the task before she moves to the next period. Hence, it is possible that $q_{T+1} < q_T$. To see this, consider the following distribution on $c : c = 0$ with probability p_0, c is uniformly distributed between 0 and $\beta\delta y$ with probability p_1 so that the distribution has positive density throughout the relevant range, and $c = \beta\delta y - \varepsilon$ with probability $1 - p_0 - p_1$. Consider the case as $p_1$goes to 0 and ε goes to 0. With this distribution, the individual will act for certain when $T = 1$. Considering a deadline of $T + 1$, she will act in the first period with probability p_0 (only when $c = 0$), and in the second period only if she remembers. Hence, $q_{T+1} = p_0 + (1 - p_0)\rho$, which is less than q_T.

Turning to welfare, time-consistent calibrated individuals choose their action cutoffs optimally with $c^*_{i,T+1} = y - \delta\rho EU_{T,i}$ and hence utility with $T+1$ periods until the deadline is at larger than that with T periods: $EU_{T+1,i} > EU_{T,i}$. Because present-biased individuals or overconfident time-consistent individuals do not choose $c^*_{i,T+1}$ optimally, it is possible that $EU_{T+1,i} < EU_{T,i}$. Figure 3 gives a numerical example of such a case for both present-biased individuals and overconfident time-consistent individuals.

A.II Definitions: Perception-Perfect Strategies

A strategy s_t is a plan of action for period t. Define the current value function W in period t, for type $i \in \{TC, N\}$, given memory beliefs $\hat{\rho}$ and anticipated future strategies $\hat{S}^t = \left\{\hat{s}^t_{t+1}, \hat{s}^t_{t+2}, ...\right\}$ as:

$$W^{\hat{\rho}}_{it}|\hat{S}^t = u_t + \beta\left(\sum_{m=1}^{\infty}\hat{\rho}^m\delta^m u_{t+m}\left(\hat{s}^t_{t+m}\right)\right)$$

where $\hat{s}^t_{t+m}$ is the individual's belief in period t about the strategy she will follow in period $t+m$;the resulting utility u_{t+m} is a function of that strategy.

Note that the individual only is able to implement the strategy in period $t+m$ if she remembers, which she anticipates happens with probability $\hat{\rho}^m$; if she forgets, she cannot act and gets $u_{t+m} = 0$. However, once she acts, she gets the benefit of the task next period regardless of whether she would have remembered next period. As a result, I make a slight abuse of notation: If an individual acts m periods in the future, $u_{t+m} = \delta y - c_{t+m}$, and if she acts in the current period t, $u_t = \beta\delta y - c_t$.This way of defining u_{t+m} to include the discounted payoff from next period allows for a simpler representation of the function W : ρ and δ act similarly in the model, except that once an individual acts, the benefit is still discounted by δ but not by ρ.

Definition 2. *Given $\hat{\beta} \leq 1, \hat{\rho} \leq 1$, and δ, a set of beliefs $\left\{\hat{S}^t, \hat{S}^{t+1}, \hat{S}^{t+2}...\right\}$ is dynamically consistent if:*
1) for all $\hat{S}^t, \hat{s}^t_\tau = \arg\max_{a\in A}\hat{W}_{\tau i}\left(\hat{S}^\tau, \hat{\beta}, \hat{\rho}\right)$ for all τ and
2) for all $\hat{S}^t$ and $\hat{S}^{t'}$ with $t < t'$, $\hat{s}^t_\tau =, \hat{s}^{t'}_\tau$ for all $\tau > t'$.
where $\hat{W}_{\tau i} = u_\tau + \hat{\beta}\sum_{m=1}^{\infty}\delta^m\hat{\rho}^m u_{\tau+m}$

Definition 3. *A Perception Perfect Strategy is $S^*\left(\beta, \hat{\beta}, \hat{\rho}\right) = \left(s^*_t\left(\beta, \hat{\beta}, \hat{\rho}\right), s^*_{t+1}\left(\beta, \hat{\beta}, \hat{\rho}\right), ..\right)$ such that there exists dynamically consistent beliefs $\hat{S}\left(\hat{\beta}, \hat{\rho}\right)$ such that $s_t = \arg\max_{a\in A} W_t\left(a, \hat{S}\left(\hat{\beta}, \hat{\rho}\right)\right)$*

$\forall t.$