Title
Top of Mind in Task-Based Environments and Choice Under Risk

Author
Zaidi, Farhan

Publication Date
2010

Peer reviewed|Thesis/dissertation

Top of Mind in Task-Based Environments and Choice Under Risk

by

Farhan Zaidi

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Matthew Rabin, Chair
Professor Stefano DellaVigna
Professor Miguel Villas-Boas

Fall 2010

Abstract

Top of Mind in Task-Based Environments and Choice Under Risk

by

Farhan Zaidi

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Matthew Rabin, Chair

I examine psychological biases that influence "top-of-mind" status in task-based environments and choice under risk, and the implications of those biases for predicting and explaining consumer behavior in the field. I develop (in joint work with Jeff Holman) a model of prospective memory, defined as the capacity to recall tasks to be carried out in the future. Motivated by the economics and psychology literature on overconfidence, we introduce memory overconfidence into the model, and show that it leads to 1) inefficiently low rates of task completion, and 2) the prediction that the probability of task completion may vary inversely with the length of time allocated to completing the task. We discuss two instances where consumers face tasks that broadly fit our model – submitting rebates and cancelling negative-option subscriptions.

I then present empirical evidence of this latter instance through a field experiment with a membership-based website. I find that implementing "camouflage" pricing – that is, charging amounts such as \$20.13 or \$19.83 rather than "standard" amounts such as \$20 or \$19.95 – significantly lowers cancellation rates. Assuming that credit card statements can serve as a reminder to cancel unused subscriptions, and that the strength of this reminder is greater with standard amounts that tend to "stand out" on a bill and draw more attention than camouflage amounts, the lower cancellation rates in the camouflage-pricing groups are consistent with the prediction of our prospective memory model.

I also consider how what is top of mind may affect consumers' choices under risk. I exploit a "natural experiment" from the trading card industry to find evidence that existing models of choice under risk may understate the extent to which salient outcomes affect the valuation of risky prospects. I track auction sales of a particular card that experienced a sudden jump in value, as well as auction sales of the sealed, unopened card boxes into which this card was randomly inserted at known odds. I estimate the price jumps for both the single card and the unopened boxes, and back out the "decision weight" that would reconcile these price movements. The estimated decision weight is well beyond the range of previous experimental estimates. I discuss how consumers' attention to particularly salient or vivid outcomes (i.e., the extent to which they are top of mind) may provide an explanation for the result.

Acknowledgements and Dedication

I would like to thank:

- My committee members, Stefano DellaVigna, Matthew Rabin, and Miguel Villas-Boas, for their feedback, suggestions, guidance, and encouragement;
- Jeff Holman for his partnership in co-authoring the first chapter of this dissertation;
- David Card, Keith Ericson, Botond Koszegi, David Laibson, and Justin Sydnor for many helpful comments;

and finally, I would like to thank and dedicate this dissertation to my parents, Sadiq and Anjum, who made their children their life's work.

Table of Contents

Overview

The Psychology and Economics field encompasses a growing body of theory and evidence on how human cognition and subsequent behavior can deviate from the standard economic model.[1] In this dissertation, I examine psychological biases that influence what is top of mind for consumers in task-based environments and choice under risk. In each case, there is evidence from the psychology literature that what draws our attention and occupies the top of mind is not always consistent what the standard model dictates, and this inconsistency has important implications for predicting and explaining behavior in the field. In what follows, I present the psychological basis for these deviations from the standard model, provide empirical results consistent with these deviations, and, where applicable, discuss potential modeling alternatives to the standard decision-making framework.

In Chapters 1 and 2, I expand on the role of top-of-mind in task-based environments. I develop (in joint work with Jeff Holman) a model of prospective memory, where agents can only perform certain tasks if they are recalled, or top of mind. Imperfect memory and overconfident memory beliefs can lead to delay or even failure to execute welfare-improving tasks, as the importance of being top of mind may be underestimated by agents (by virtue of their overconfidence that the task will resurface to top of mind later). In the model, an agent faces some task with stochastic cost c_t, fixed benefit b, and T periods until some exogenously imposed deadline. The agent can only execute the task at time t if the task is recalled in that period. The agent sets a threshold cost each period based on her expectations of whether she will recall and carry out the task in future periods. If the task is recalled at time t, and the draw from the cost distribution is below this threshold, the task is executed. Memory overconfidence is incorporated into the model, defined as overestimating the likelihood of recall in future periods. There are three key results from the model:

1) The hazard rate – the probability of executing the task in a given period, conditional on not having executed it already – can be decreasing over time with imperfect prospective memory, but must be increasing otherwise;
2) Overconfident beliefs about prospective memory lower the probability of, and expected utility from, task completion;
3) For sufficiently poor prospective memory, and sufficiently overconfident memory beliefs, lengthening a deadline may make agents *less* likely to complete tasks and make them worse off ex post.

We discuss two extensions of the model – memory aids and time discounting. We model memory aids as the option to purchase recall with certainty in some subset of periods before the deadline. The incorporation of memory aids reveals another potential detriment of memory overconfidence – for an increasing task benefit, the welfare loss from overconfident memory beliefs can be made arbitrarily large, while for agents with correct memory beliefs, the welfare loss is bounded by the cost of a complete memory aid. Present-biased preferences are built into the model by assuming the costs of the task are immediate while the benefits are deferred, and

[1] See Angner and Loewenstein (forthcoming) for a survey of the field's historical origins and philosophical underpinnings, and DellaVigna (2009) for a recent review of empirical work.

that the agent discounts all future utility flow by a factor β. Naïve time-inconsistency has positive interaction effects with memory overconfidence – procrastination further encourages overconfident agents to defer tasks, making their overestimate of the continuation value of deferral more costly. Importantly, however, result 3) above – the possible negative correlation between task completion rates and length of deadline – cannot be generated from present-biased preferences alone.

Chapter 1 concludes with a discussion of economic scenarios where fallible prospective memory and memory overconfidence may be exploited by firms to the detriment of consumers – in particular, rebates and trial offers. We present anecdotal evidence of the high degree of ex post consumer regret and frustration over these tactics, and discuss how this evidence runs counter to standard model explanations of their usage by firms (such as price discrimination and learning).

Chapter 2 is an empirical study that demonstrates how firms can exploit consumer forgetting in a task-based environment. I conduct a field test with a membership-based website that finds that implementing "camouflage" pricing – that is, charging amounts such as \$20.13 or \$19.83 rather than "standard" amounts such as \$20 or \$19.95 – lowers the odds of cancellation by roughly 10% over the first 4 months of membership. Assuming that consumers' credit card statements serve as a reminder of the subscription and of the option to cancel, this result suggests that camouflage amounts – by virtue of blending in with other charges such as groceries, gas, or taxed clothing purchases – serve to obscure this reminder. Furthermore, the cancellation pattern over the course of the billing cycle – in particular, a spike in cancellations over the first few days – provides suggestive evidence of subscriber forgetfulness and of the reminder effect from credit card statements.

In Chapter 3, I discuss how what is top of mind may affect choice under risk. I exploit a "natural experiment" from the trading card industry to show that for a lottery-type good with a particularly salient "grand prize," the jackpot possibility appears to be overly weighted by consumers in a dynamic pricing environment – perhaps because the grand prize scenario draws so much of consumers' attention when they evaluate the good. I track auction sales of a particular card that experienced a sudden jump in value, as well as auction sales of the sealed, unopened card boxes into which this card was randomly inserted at known odds. In abstract terms, this experiment tests how the market price for a lottery ticket (i.e., an unopened box) with fixed, known odds of winning adjusts when the size of the "grand prize" (the particular card) increases. I estimate the price jumps for both the single card and the unopened boxes, and back out the "decision weight" that would reconcile these price movements. I find that, while the probability of obtaining the particular card was 1 in 60 (or ~1.7%), the estimated decision weight is roughly 0.27. The full range of documented experimental estimates of this probability weight – between 0.03 and 0.10 – can be ruled out with 99% confidence. One interpretation of the result is that outcomes that are salient and draw top-of-mind attention in risky scenarios (such as the grand prize when evaluating a lottery-type good, or the worst-case scenario when evaluating disaster insurance) may be valued in a manner not fully captured by expected utility theory or alternate models of choice under risk. The vividness of certain outcomes – effectively, extent to which they are top of mind when choices under risk are made – may explain the disparity between past experimental calibrations and this field result.

In the chapters that follow, I focus on the marketing implications of my research – that is, how consumer attention can impact market outcomes, and how firms can exploit their understanding of consumers' top-of-mind thinking to their advantage. However, this research may have important policy applications as well. Just as firms attempt to manipulate consumer attention to obscure subscription renewal decisions or to increase demand for risky goods,[2] policymakers can design programs to direct attention towards potentially welfare-improving choices. Undersaving is one public policy issue where getting to top of mind is being examined as a potential remedy. Karlan, McConnell, Mullainathan, and Zinman (2010) find that reminders increase saving, and that reminders that reference a specific future expenditure (perhaps making the benefit of saving particularly salient) are especially effective. Kearney, Tufano, Guryan and Hearst (2010) survey the evidence on how prize-linked savings accounts (which pool interest payments and raffle them off periodically as cash prizes) have increased saving rates, suggesting that the prospect of winning a lottery may create a greater incentive to save than simple interest payouts.[3] Another policy issue where getting to consumers' top of mind may prove beneficial is energy efficiency. In-home, real-time "smart" electric meters make the costs of energy choices particularly salient at the point of consumption. Ongoing studies are testing how these devices, by getting energy expenses to consumers' top of mind, may impact behavior in everything from turning off unused lights to replacing old, inefficient appliances (Kahn 2009). Continued research along these lines could yield new policy instruments that address consumers' inattention and subsequently elicit welfare-improving choices.

[2] That is, by making positive outcomes particularly salient – such as how exercise equipment companies feature extreme success stories in their infomercials.

[3] Perhaps because potential savers misperceive or overweight the odds of winning a cash prize, or perhaps because cash prizes are simply more effective at garnering potential savers' attention.

Chapter 1: The Economics of Prospective Memory

1.1 Introduction

Human memory is far from perfect. Beginning with Ebbinghaus's groundbreaking work on the "forgetting curve" in 1885, the psychology literature contains extensive evidence supporting this claim. Yet, in economics, the critical role that fallible memory plays in coloring individual decision-making has been largely neglected until recently. This neglect may lead to the misattribution of certain observed anomalies to other types of biases.

One such anomaly is the failure of individuals to carry out projects with small, immediate costs and large, deferred benefits. Several empirical studies have documented this finding and explained it by time-inconsistent preferences.[4] For instance, DellaVigna and Malmendier (2006) suggest that the low cancelation rates observed with automatically renewed health club memberships may be an example of status quo bias generated by hyperbolic time-discounting. Madrian and Shea (2001) and Choi, Laibson, Madrian, and Metrick (2006) suggest that this same status quo bias keeps employees from updating the often-suboptimal default enrollment options in 401(k) plans.

While we agree that time-inconsistency plays a role in generating this inefficient behavior, we believe that this explanation is incomplete, particularly when considering projects with deadlines. While present-biased preferences may lead to delay in executing beneficial tasks such as canceling an unused health club membership or modifying a benefits plan, such preferences often cannot explain the failure to carry out such tasks altogether. For example, consider an individual who has 30 days to mail in a $50 consumer rebate. It is calibrationally feasible to derive from present-biased preferences that, for the first 29 days, she will defer completing and mailing out the rebate form, preferring to do it "tomorrow" over "today." However, on the 30th day, her choice is between "today" and "never," and justifying task omission in such cases often requires assuming an extremely low β (i.e., extreme present-biased preferences).

Assuming stochastic costs and naïvete with respect to present-biased preferences, as in O'Donoghue and Rabin (1999b), can lead to inefficiently low task execution rates, even for feasible values of β. However, these assumptions are inconsistent with the empirical finding that task execution rates may vary inversely with the length of deadline. Shafir and Tversky (1992) offered students $5 to return a long questionnaire by a given date. The students were randomly assigned to one of three deadline groups -- the first group was given 5 days, the second group 3 weeks, and the third group no definite deadline. The respective rates of return for the three groups were 60%, 42%, and 25%. Procrastination can explain the drop in the return rate between the groups with a deadline and the group with no deadline – when no deadline is given, students never reach the "now" or "never" decision, and naïve hyperbolics will defer completing the survey each period, always planning on doing it next period. However, with a finite deadline, a model with stochastic costs and naïve hyperbolic time-discounting would predict a positive

---

[4]In the quasi-hyperbolic form, time-inconsistent agents have a discount rate of $\beta\delta$ between the current period and the next period, and δ between all future periods. Thus, β captures the special salience of the present.

relationship between the likelihood of task execution and the length of deadline. Having more time gives agents more chances to draw a low-cost realization, and thus the likelihood of executing the task would increase.

Perhaps most importantly, the time-inconsistency explanation simply feels incomplete when we relate it to our personal experiences in failing to execute basic, time-sensitive tasks. Consider the last time you failed to mail in a rebate, cancel a free trial offer, or return a borrowed movie or book on time. What happened? A common response is simply, "I forgot." In this paper, we explore this type of fallible memory as an alternative (or more accurately, complement) to time-inconsistent preferences in explaining inefficient behavior with respect to completing tasks with deadlines.

We build on the work of Mullainathan (2002), which developed a retrospective memory-based model of bounded rationality, by extending the analysis to prospective memory. The distinction, as explained by psychologist P.E. Morris (1992), is that "prospective memory is memory for intentions, for actions that we wish to carry out in the future, while retrospective memory is recall of information from the past." There is continued debate among psychologists on whether a distinction should be made between these two "types" of memory, given that the underlying mechanisms that determine the success or failure of remembering a past memory or a future task may be the same. Nonetheless, from the perspective of economics, we view these as distinct types of memory recall, with different implications for individual decision-making behavior.

Section 1.2 surveys the existing psychology literature on fallible memory, and prospective recall in particular. The experimental evidence on prospective memory is quite limited. The most common studies – "postcard studies" – instruct subjects to mail in postcards at (or before) some specified future date, and monitor how return rates vary with time frame, cues, and monetary incentives. Unfortunately, many of these studies have limited sample sizes and other methodological concerns, and are uninformative on how prospective memory beliefs correlate with actual recall performance. More recent work in the economics and marketing literature (Silk 2004 and Ericson forthcoming) addressed some of these concerns.

In Section 1.3, we develop a baseline model of prospective memory (PM), and apply it to a choice problem where the agent faces some task with stochastic cost c_t, fixed benefit b, and T periods until some exogenously imposed deadline. The agent can only execute the task at time t if the task is recalled in that period. We adopt a memory function with the rehearsal property described in Mullainathan (2002) – that is, the probability of recall is lower if the event (or in this case, the task) was forgotten in the recent past. The agent sets a threshold cost each period based on her expectations of whether she will recall and carry out the task in future periods. If the task is recalled at time t, and the draw from the cost distribution is below this threshold, the task is executed. We discuss the time pattern of hazard rates for agents with perfect and imperfect memories, and demonstrate an empirically testable prediction for fallible memory – hazard rates must be increasing for agents with perfect memory, but may be decreasing for agents with imperfect memory.

In our baseline model, we assume that agents, despite their fallible memory, have correct beliefs over their PM process. In Section 4, we introduce the possibility that agents have systematically incorrect PM beliefs, as suggested by experimental evidence in Silk (2004) and Ericson (forthcoming) on subjects' incentive-compatible memory beliefs and subsequent performance. In particular, agents may be overconfident in their prospective memory. We define prospective memory overconfidence as overestimating the base likelihood of recall in future periods, or underestimating the effect of forgetting the task in any period (or series of periods) on the subsequent probability of recall. We view PM overconfidence as a form of projection bias (Loewenstein, O'Donoghue, and Rabin 2003). That is, individuals "project" their current memory state onto all future periods, and develop their expectations for future recall and behavior accordingly. PM overconfidence can also be viewed as analogous to other forms of "information projection" such as hindsight bias and curse of knowledge.

In Section 1.5, we discuss the effects of PM overconfidence on agent behavior in our model. First, PM overconfidence reduces welfare. When an agent is overconfident with respect to prospective memory, the perceived continuation value of deferring tasks is higher than the true continuation value – the agent will inefficiently defer tasks, overoptimistically relying on future recall and execution. Second, in our key result, PM overconfidence increases the likelihood that extending the deadline will be to the detriment of the agent – that is, that both the ex ante expected utility and probability of task execution decrease as T increases. While the optimal strategy requires agents to decrease threshold costs when the deadline is extended, overconfident agents, in overestimating the benefit of the extended deadline, reduce threshold costs by more than they should. Under certain conditions (in particular, if agents are sufficiently forgetful and overconfident), the cost of this over-selectivity in early periods outweighs the benefit of having more periods to remember the task, and agents are made worse off. Thus, our model provides an explanation for the Shafir and Tversky student survey result.

In Section 1.6 we discuss two extensions of the model – memory aids and time discounting. We incorporate memory aids by assuming that, for some cost, agents can ensure recall in some subset of the periods before the deadline. Depending on the values of the memory parameters, agents will either choose to backload reminders (to allow for lower threshold costs in early periods) or spread out reminders (to mitigate the adverse effect of successive forgetting). We also show that, for an increasing benefit from task completion, the welfare loss from fallible memory for overconfident agents can be made arbitrarily large, while for agents with correct PM beliefs, the welfare loss is bounded by the cost of a complete memory aid.

We next incorporate time-discounting into the model, and demonstrate that the Shafir and Tversky result cannot be explained by hyperbolic discounting alone. At the same time, we show that naïve time-inconsistency has positive interaction effects with PM overconfidence – procrastination further encourages agents to defer tasks, making their overestimate of the continuation value of deferral more costly. Indeed, we believe that individuals' failure to execute basic, time-sensitive tasks in the real world is not due solely to imperfect memory or present-biased preferences, but rather some combination of the two.

Rebates and free trial offers are two "real world" mechanisms by which firms appear to be exploiting naive forgetting by consumers. In Section 1.7, we provide an overview of these prospective memory failures in the marketplace. We present anecdotal evidence on the high degree of ex post consumer regret and frustration over rebates and trial offers, and discuss how this evidence runs counter to the rational-model explanations for these marketing tactics. We believe these areas are ripe for potential field studies identifying and calibrating the effects of prospective memory failures and overconfidence (Section 1.8). We conclude with some thoughts on potential theoretical extensions of our model.

1.2 The Psychology Literature on Retrospective and Prospective Memory

1.2.1 Forgetting and The Retention Function

Much of the psychology literature on memory focused on retrospective recall. Ebbinghaus (1885) is widely recognized as the first experimental study to attempt to measure the rate of forgetting, and is credited for providing the first empirical evidence that memory retention is nonlinear in time.[5]

Since then, the memory retention function has been estimated in experiments of various contexts. Studies have varied the time frame (from minutes to years) and the content of information to be remembered (e.g., words, faces, foreign languages, skills), and have generally found evidence consistent with Ebbinghaus's pioneering work. Levy and Loftus (1984) remark:

"Many researchers, beginning with Ebbinghaus, have assumed forgetting to be exponential over time, t, between learning and test, that is:

$$P_t = e^{-kt}$$

Such a function would follow from the reasonable assumption that information in memory, like the content of many other physical systems, is lost at a rate that is proportional to the amount remaining in the system."

Rubin and Wenzel (1996), in a meta-study of 210 published data sets on retrospective recall, find that the exponential form $p_t = be^{-mt}$ is one of the best 2-parameter fits of the retention function.[6]

1.2.2 Prospective Memory: Components and Types

Meacham and Singer (1977) are credited with coining the term "prospective memory." Cohen (1993) describes the components of prospective memory as "remembering what the planned action is, remembering to perform it, and remembering when and where to do it." It has been

---

[5]Ebbinghaus taught himself a list of 13 nonsense syllables ("nonsense" to rule out recall by association), and tested himself on the list at various time intervals. His metric for retention was a "saving rate," – that is, the ratio of the number of repetitions needed to relearn the entire list and recite it twice from memory to the number of repetitions that were needed to learn and recite the list initially.

[6]Rubin and Wenzel find three other functional forms that fit well: logarithmic ($p_t = b - m\log(t)$), power ($p_t = bt^{-m}$), and hyperbolic ($p_t = \frac{m}{t+b}$)

noted that PM tasks actually contain both a prospective and a retrospective component. Baddeley (1997) defines the distinction between these components as "when" something should be remembered, versus "what." For instance, if the task is to pass along a message to someone, the retrospective component is remembering the message, while the prospective component is remembering that we have a message for that person when we see them.

While PM tasks contain a retrospective component, they have a number of characteristics that distinguish them from pure retrospective recall. For one, PM tasks generally have low information content – Harris (1984) notes that "the information to be recalled may be trivially easy, but remembering to recall at all may be the difficulty." Also, while retrospective recall is cued or prompted by events in the present – e.g., being asked for someone's phone number from memory, or partaking in an activity that requires the recall of some previously acquired skill – cues play much less of a role in prospective recall. While PM tasks are occasionally cued (e.g., driving by the grocery store and remembering that you have to buy groceries), these cues are generally random and thus affect the pattern of recall in a less predictable manner.[7]

1.2.3 Experimental Studies of Prospective Memory

Prior to Silk (2004) and Ericson (forthcoming), most prospective memory experimental studies focused on the impact of exogenous factors on PM success rates (such as deadline length, subject age, and reminders provided by the experimenter), rather than testing for relevant PM belief and performance parameters in an economic decision-making framework. The most common were postcard studies, where subjects are asked to return postcards on certain dates or within some time interval (usually without incentives). Wilkins (1976) asked 34 subjects to return one card each, from 2 to 36 days later, and found no effect of length of interval on performance. Meacham and Leiman (1982), in a similar study, found that later cards were less likely to be posted than earlier cards, and that providing subjects with cues (such as colored tags on their key rings) improved return rates. Orne (1970) and Meacham and Singer (1977) found a similar effect of cues on return rates, and also found that monetary incentives improved return rates.[8] Since these early postcard studies, there has been a growing psychology literature testing subjects' prospective memory in "semi-naturalistic" (as opposed to laboratory) settings – see Table 7.2 in McDaniel and Einstein (2007) for a summary of these experiments and the measured prospective memory success rates.

Silk (2004) advanced this body of evidence by running a series of experiments that "examine consumers' purchase and post-purchase behavior with a real rebate offer." Silk lists a number of findings from the experiments, including, 1) increasing the rebate reward increased takeup rates, but had "a weaker effect"on redemption rates; 2) increasing the length of the redemption period increased redemption confidence and takeup rates, but reduced redemption rates; and 3) there is

[7] Cues may affect the likelihood of recall, if not the pattern. That is, all other things equal, PM tasks with more frequent cues will be recalled more often. Additionally, one type of cue that can affect the pattern of recall is the use of memory aids (e.g., planners, Post-It notes, etc.), which we discuss in Section 1.6.

[8] Other studies (such as Shepard and Mosley 1976 and Gates and Colburn 1976) have focused on prospective memory with respect to appointment-keeping. Levy and Claravall (1977), in a study of medical patients needing regular check-ups at varying intervals, found that reminders increased compliance most for patients with the longest intervals between appointments.

a surprisingly positive correlation between randomized redemption effort levels and redemption rates (that is, redemption rates were higher in conditions where the rebate form was made more cumbersome and lengthy).

Another key area of research in the psychology literature has been the relationship between retrospective and prospective memory. In one of the best-known studies (Wilkins and Baddeley 1978), medical subjects were tested on free recall of unrelated words, and then monitored in their pill-taking. Surprisingly, they find that subjects who did worse in the retrospective memory test (word recall) remembered to take their pills at a higher rate.[9] Other studies have found a positive correlation between performance in retrospective memory tests and PM tasks.[10]

1.3 Baseline Model

1.3.1 Definitions

We consider an individual facing some task with cost c and benefit b, and T periods until an exogenously imposed deadline. In each period, there is some uncertainty about whether the individual will recall the task. An individual will execute the task in a period if 1) it is recalled; and 2) the expected utility from carrying out the task in that period is greater than the perceived expected utility from deferring the task and relying on future recall and execution.

The probability of recall is not fixed across periods – rather, it depends on the history of previous recall. Our model incorporates a simple form of the rehearsal property, where the probability of recall in some period t (p_t) depends on the number of successive recall failures in the preceding periods.[11] Define N_t as the number of successive, immediately preceding periods in which the task was not recalled (i.e., if the task was recalled at $t-1$, then $N_t = 0$; if the task was forgotten at $t-1$ but recalled at $t-2$, then $N_t = 1$; and so on). Then:

$$p_t = \theta^{N_t} p,$$

where $p, \theta \leq 1$. The term p captures the base rate of remembering – that is, the likelihood of remembering the task in the first period, and in any period preceded by a period of recall. The term θ captures the factor by which memory decays (i.e., the factor by which the likelihood of recall is reduced) with each additional period of forgetting. Note that this form leads to exponential decay in expectation of future recall, consistent with the empirical evidence on

[9] The authors suggest this result is driven by the greater use of memory aids by the poor retrospective memory test performers, though they do not track the use of memory aids by the subjects.
[10] See Cockburn and Smith (1988) and Wilson, Cockburn, Baddeley and Hiorns (1989).
[11] While the rehearsal property is more commonly associated with retrospective recall, it has a clear application to PM tasks. Having a task "top-of-mind" increases the chances that it will stay "top-of-mind" and be completed before the deadline. Ellis and Nimmo-Smith (1993), in a study of a pulse PM task, found that "preservation" – the remembering that one has to carry out an action hours or even days prior to the planned time – improves the likelihood of correct recall.

retention discussed in Section 1.2. Our modeling of the rehearsal property is a slightly more nuanced version of that adopted in Mullainathan (2002), which essentially assumes:

$$p_t = \begin{cases} p & \text{if } m_{t-1} = 1 \\ \theta p & \text{if } m_{t-1} = 0, \end{cases}$$

where m_t is a random variable equal to 1 if the task is recalled in period t and 0 otherwise (i.e., $p_t = E(m_t)$). While Mullainathan assumes that that the recall probability is only contingent on recall in the immediately preceding period, we assume that each successive recall failure will lower the subsequent probability of recall by a factor θ.

We use a search model structure, similar to O'Donoghue and Rabin (1999b), and assume the cost of task completion is stochastically determined each period. Let c be a random variable with CDF $F(.)$, with support $[\underline{c}, \overline{c}]$ and $\underline{c} > 0$. Let c_t denote the draw from this distribution at time t – that is, the cost of task execution at t. We assume the benefit is fixed at b, regardless of when the task is completed. In this section, we do not incorporate time-discounting into the agent's preferences – whenever the agent completes the task, she receives utility $b - c_t$. An agent's strategy is a sequence $s = (s_1, s_2, \cdots, s_T)$ where each s_t denotes the cost threshold for task execution. The task is executed in t if the task is recalled in t and $c_t \leq s_t$. We denote first-best strategy for given memory parameters (p, θ) by $s^*(p,\theta) = (s_1^*(p,\theta), s_2^*(p,\theta), \cdots, s_T^*(p,\theta)))$. This concept of a first-best strategy is referred to in O'Donoghue and Rabin (1999b) as a “perception-perfect strategy,” and defined as “a strategy that in all periods (even those after the activity is performed) a person chooses the optimal action given her current preferences and her perceptions of her future behavior.”

Furthermore, define the following:

1. m_t as a random variable equal to 1 if the task is recalled in period t and 0 otherwise, with $m = (m_1, \cdots, m_T)$. We assume $m_1 = 1$ (the task is recalled with probability 1 in the first period, which can be interpreted as the period in which the agent becomes aware of the task). Note that since m_t is binary the set of all possible recall outcomes is just 2^T;
2. $\mu_t(m, p, \theta)$ as the probability from the time t perspective of memory sequence m conditional on $m_t = 1$;
3. $\pi_{t,t'}(s, p, \theta \mid m)$ as the probability from the time t perspective that the agent will not complete the task before $t' > t$ given the memory sequence m if the agent chooses to defer the task in t and follows strategy s thereafter;
4. $V_t(s, p, \theta)$ as the utility from the time t perspective of the task, given that it has not been completed before time t, and that the agent follows strategy s.[12] V_t is a random variable

[12] We consider the effects of present-biased preferences on agent behavior in Section 1.6.

that depends on current and future values of c and m, and we will write its time t expectation in a Bellman equation below.

Note that

$$\pi_{t,t'}(s,p,\theta \mid m) = \begin{cases} 1 & \text{if } t' = t+1 \\ \prod_{i=t+1}^{t'-1} 1 - m_i F(s_i) & \text{if } t' > t+1. \end{cases}$$

That is, the probability of arriving in period $t+1$ without having completed the task conditional on deferring in period t is 1, and for every following period is the joint probability of drawing a cost above the threshold for each period that the task is recalled (i.e. where $m_t = 1$).

We can write the following Bellman equation for $E_t(V_t)$ by recognizing that the value of a task from the time t perspective is the maximum of executing in t and deferring to $t+1$, which sets up a recursive system.

$$V_t = \max\{m_t(b-c_t), V_{t+1}\}, \text{and}$$
$$E_t(V_t(s,p,\theta)) = \max[m_t(b-c_t), E_t(V_{t+1}(s,p,\theta))].$$

The continuation value from deferring in period t is the second term in the max expression. It is equal to the probability of task execution multiplied by the expected utility conditional on task execution in each proceeding period from $t+1,...,T$. Two alternative expressions for V_t involve V^1 and V^0 terms, one using same period V s and the other next period V s:

$$V_t(s,p,\theta) = \max[m_t(b-c_t), pV_{t+1}^1 + (1-p)V_{t+1}^0)]$$
$$= m_t V_t^1 + (1-m_t)V_t^0$$

Finally, let $Q(s,p,\theta)$ denote the ex ante probability of task execution over all T periods, such that

$$Q(s,p,\theta) = \sum_{t'=1}^{T} \sum_{m \in 2^T} \mu_t(m,p,\theta)\pi_{0,t'}(s,p,\theta \mid m) m_t F(s_t).$$

Define $\hat{p}$ and $\hat{\theta}$ as the agent's beliefs about the memory process. The agent will solve for her strategy s based on these beliefs by backwards induction. Note that in periods where $m_t = 0$ and the agent forgets the task, the value of s_t is not relevant. So we define the optimal strategy only in periods for which $m_t = 1$:

$$s_t^*(\hat{p},\hat{\theta}) = \begin{cases} b & \text{if } t = T, m_t = 1 \\ b - E_t(V_{t+1}(s,\hat{p},\hat{\theta}) \mid m_t = 1) & \text{if } 1 \le t < T, m_t = 1 \end{cases} \quad (1.1)$$

1.3.2 Hazard Rates

We now consider how imperfect memory affects the time-pattern of hazard rates, where the hazard rate at t is defined as the probability of executing a task at period t' without having executed the task. To do so, we first distinguish between agents with perfect memory, and agents with imperfect memory, but accurate beliefs about memory.

We define a "Perfect Rememberer" (PR) as an agent with $p = \hat{p} = \theta = \hat{\theta} = 1$ – that is, an agent with perfect memory and correct beliefs about memory.

We define a "Sophisticated Forgetter" (SF) as an agent with imperfect memory, that is with $p < 1$ and $\theta < 1$, but correct beliefs about memory, that is $p = \hat{p}$ and $\theta = \hat{\theta}$.

Define h_t to be the hazard rate of task completion at time t, and note that

$$h_t(s,p,\theta) = Prob(m_t = 1 \; and \; c_t \le s_i) = E(m_t)F(s_t)$$

That is, the hazard rate equals the probability of remembering the task at t conditional on forgetting the task or drawing a cost over the threshold in each previous period, multiplied by the probability that the current period cost is below the current period threshold.

Proposition 1 *For a PR, hazard rates strictly increase in time.*

This follows directly from the fact that s_t^*, the PR's threshold cost, is increasing in time, which can be seen in the iterative solution for s^* in equation (1.1). Since $E(m_t) = 1$ for all t for PRs, and $F(s_t^*)$ is increasing in t, it follows that $h_t(s^*,1,1)$ is increasing in t.

Proposition 2 *For any SF, with a sufficiently large T hazard rates will strictly decrease in time over some interval.*

The intuition for Proposition 2 is that as T grows large, the hazard in the last period converges to zero, because the fraction of the surviving population that remember in the last period goes to zero. However, the hazard in the first period is bounded below, even for arbitrarily large T.

1.3.3 Three-Period Example

To illustrate the solution employed by rational agents, and the resulting time pattern of hazard rates, we consider an example where $T = 3, b = 1, c \sim U[0,2]$.

First consider the PR's backwards-induction reasoning. In the last period of the model, the PR will set $s_3^* = 1$ and will thus execute the task in period 3, conditional on having arrived at period 3 without already completing the task, with probability 0.5.

Now consider the period 2 problem. The continuation value $V_2(s^*)$ from deferring equals $F(s_3^*)(b - E(c_3 \mid c_3 \le s_3^*)) = 0.5(1-0.5) = 0.25$. Thus the PR will set $s_2^* = 1 - 0.25 = 0.75$. Finally, in period 1 the PR will set $s_1^* = b - F(s_2^*)(b - E(c_2 \mid c_2 \le s_2^*)) - (1 - F(s_3^*)F(s_3^*)$ $(b - E(c_3 \mid c_3 \le b))) = 1 - (0.375(1-0.375)) - (0.625)(0.5)(1-0.5) = 0.61$. This implies an ex ante expected utility $V_0(s^*) = 0.48$ and the probability of task completion $Q(s^*) = 0.78$. The solution for the PR is summarized in Table 1.1 below.

Table 1.1: PR Behavior for *b* = 1, c_t ~ *U*[0,2], *T* = 3

Period	Optimal Threshold cost (s_t^*)	Hazard rate (h_t)	Ex ante probability of executing at *t*	Ex ante EU from Executing at *t*
1	0.61	0.31	0.31	0.21
2	0.75	0.38	0.26	0.16
3	1.00	0.50	0.21	0.11
Total			0.78	0.48

Now consider a SF. Assume the memory process is characterized by $p = 0.5, \theta = 0.2$. In period 3, conditional on recall the SF acts just like the PR: as long as the realization of c_3 is less than or equal to b, she will execute that task. In our notation she uses $s_3^*(0.5,0.2) = s_3^*(1,1) = b = 1$.

However, in period 2 the SF realizes that her continuation value is lower due to imperfect memory and uses $s_2^*(0.5,0.2) > s_2^*(1,1)$ accordingly. In particular $V_2(s^*(p,\theta), p,\theta) =$ $pV_2(s^*(1,1),1,1) = (0.5)(0.25) = 0.125$, and so $s_2^* = 0.875$.

In period 1, the SF uses $s_1 = b - V_1(s)$, where

$$\begin{aligned} V_1(s) &= \sum_{t'=2}^{3} \sum_{m \in 2^3} \mu_1(m, p, \theta) \pi_{1,t'}(s, p, \theta \mid m) m_t F(s_1)(b - E(c_{t'} \mid c_{t'} \le s_{t'})) \\ &= pF(s_2)(b - E(c \mid c \le s_2)) + [(1-p)\theta p + p^2(1 - F(s_2))] F(s_3)(b - E(c \mid c \le s_3)) \\ &= (0.5)(0.4375)(1 - 0.4375) + (0.05 + 0.14)(0.5)(1 - 0.5) \\ &= 0.17. \end{aligned}$$

So, $s_1^* = 0.83$. this implies an ex ante expected utility $V_0(s) = 0.19$ and a probability of task completion $Q(s) = 0.33$. Table 1.2 summarizes SF behavior.

Note that in this particular example, hazard rates are decreasing in time for the SF and increasing for the PR.

Table 1.2: SF Behavior for p = 0.5, θ = 0.2, b = 1, c_t ~ U[0,2], T = 3

Period	Optimal Threshold cost (s_t^*)	Hazard rate (h_t)	Ex ante probability of executing at t	Ex ante EU from Executing at t
1	0.83	0.21	0.21	0.12
2	0.87	0.11	0.08	0.05
3	1.00	0.06	0.04	0.02
Total			0.33	0.19

Figure 1.1 shows the sensitivity of the ex ante probability of task completion Q to p and θ. Figure 1.2 shows the sensitivity of V_0 to p and θ.

Figure 1.1: Sensitivity of Q to p and θ (for T = 3, b = 1, c ~ U[0,2])

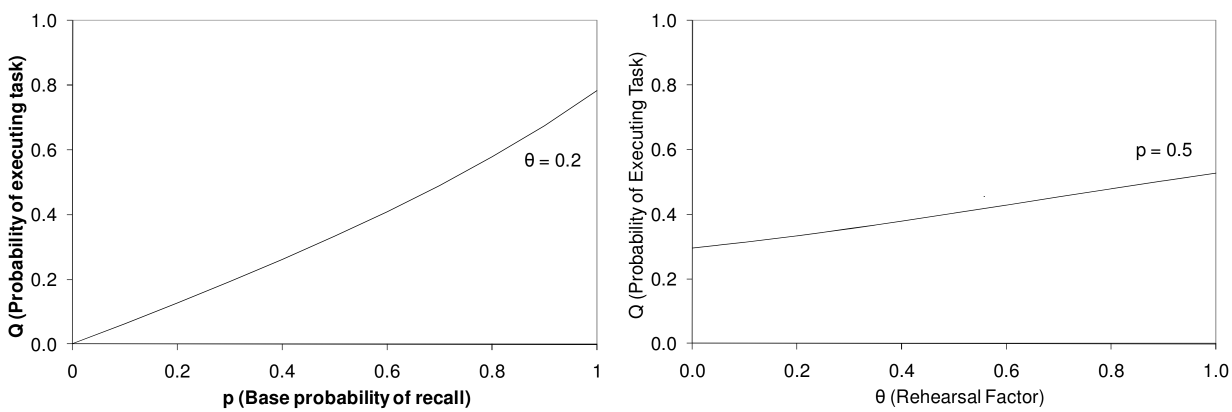

Figure 1.2: Sensitivity of V_0 to p and θ (for T = 3, b = 1, c ~ U[0,2])

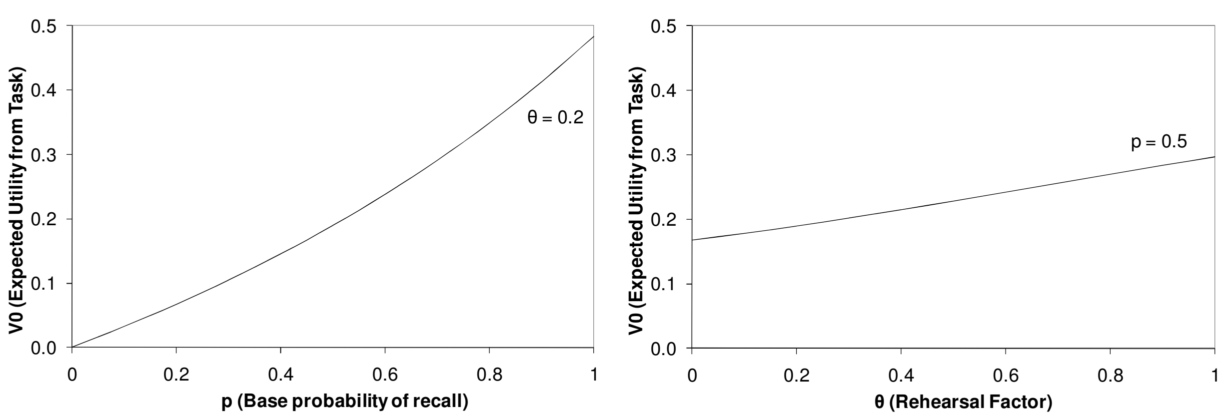

1.4 Overconfidence in Prospective Memory

Thus far we have assumed that agents, despite their fallible memory, have correct beliefs over their prospective memory processes, and account for this perfectly in calculating threshold costs each period. We now discuss augmenting the model by allowing for systematic errors in prospective memory beliefs. Specifically, we consider prospective memory overconfidence – in the context of the model, $\hat{p} > p$ and/or $\hat{\theta} > \theta$.

1.4.1 Overconfidence Literature

Experimental and survey studies have found that the average individual exhibits overconfidence over a wide range of skills, abilities, and personal traits. Most people believe they are more intelligent (Wylie 1979), more sociable and better leaders (Cook and Frank 1996), more productive (Parker, Taylor, Barnett and Martens 1959), and better drivers than the average person (Svenson 1981). To the extent that remembering things is a skill or ability like those above, assuming overconfidence in memory is a natural extension of these findings.

Most of the evidence on memory overconfidence, until recently, has been in retrospective contexts – that is, individuals are overconfident in the accuracy of their recollections.[13] For instance, eyewitnesses tend to be overconfident in the accuracy of lineup identifications of criminal suspects, and the details of their trial testimony (see Loftus 1979 and Brewer and A.Risworth 2002). However, two recent studies – Silk (2004) and Ericson (forthcoming) – find evidence of prospective memory overconfidence in experimental designs that broadly fit the task model presented in Section 1.3.

Silk contains an experimental condition where subjects are offered a rebate and are asked for their "redemption confidence" (i.e., a subjective probability of redemption), as well as their expectation of the overall redemption rate in the population (i.e., a base rate estimate). Those who choose to accept the rebate offer are, on average, 93.5% confident they will submit it, while the ultimate redemption rate is only 60%. Their population (or base rate) estimates are much more in line with the true redemption rates, suggesting that prospective memory is yet another trait where individuals systematically evaluate themselves as being above average. On the other hand, those who reject the rebate offer have subjective probability rates that are slightly lower than both their base rate estimates and the eventual redemption rate, suggesting heterogeneity in prospective memory overconfidence.

Ericson uses an incentive-compatible design to elicit subjects' beliefs that they will remember to claim a future contingent payment, then compares these beliefs with actual claim behavior. Subjects are asked for their preference between a conditional payment of \$20 (contingent on the subject sending an email to the experimenter in a 5-day window 3-4 months in the future) and an automatic payment \$x, where x varies from \$5 to \$20 in \$0.75 increments (each subject gives their preference between the contingent payment and the automatic payment for each potential value of x). Subjects are then randomly assigned to one of three conditions – 1) receiving a \$20

[13] Mullainathan (2002) makes a related assumption that individuals have limited memories, but assume their recollections to be perfectly accurate.

automatic payment; 2) receiving a \$20 contingent payment; or 3) the choice they made between the \$20 contingent payment and a \$x automatic payment for a randomly selected value of x. The threshold values of \$x for which subjects switch their preference from the contingent payment to the automatic payment are used to infer subjects' beliefs about their likelihood of claiming the contingent payment. The claim probability implied by the preferences of the subjects in the contigent payment condition was 0.76, while the actual claim rate was 0.53 – a significant difference at the 99.9% level and evidence of prospective memory overconfidence in the subject pool.

1.4.2 Overestimating p vs Overestimating θ

Overestimating p can be interpreted as being overconfident in the base level of future recall probability. In the extreme case, when $\hat{p} = 1$, the agent believes that she will recall the task with probability 1 in every future period – that is, that she believes she is a perfect rememberer. Cohen (1993) notes that "people's beliefs about their own [prospective] memories are based on their experience and success and failure in everyday life." In this light, overestimating p can be viewed as a "recency bias" – an agent who correctly remembers the task in period t may upwardly revise PM beliefs above the true value of p for periods $t+1,...,T$.

Overestimating θ can be interpreted as underestimating the rate at which a specific memory weakens once forgotten. In the extreme case, when $\hat{\theta} = 1$, the agent believes the recall probability in all future periods is identical (and equal to $\hat{p}$). Overestimating θ can be interpreted as a form of projection bias, as defined by Loewenstein, O'Donoghue, and Rabin (2003). While these authors use the term to describe the tendency to project current preferences on future selves, in this case agents project the current memory state onto future selves. When an agent is in a state of recall, as she must be whenever assessing whether to perform the task or not, the likelihood of recall in the next period is p. If an agent believes $\hat{\theta} > \theta$, she underestimates the effect that forgetting will have on eroding her current memory state – that is, she estimates the probability of recall for every possible N_t to be closer to the next period probability of recall than it actually is. As with agents exhibiting projection bias, agents overestimating θ underappreciate how a change in circumstances will affect some parameter.

Despite these different interpretations, the effects of overestimating either parameter in our model are similar – in both cases, agents overestimate the continuation value of deferral and set inefficiently low threshold costs – and thus we do not emphasize the distinction in what follows.

1.4.3 Prospective Memory Overconfidence as Information Projection

Both instances of prospective memory overconfidence – that is, overestimating p or θ – are closely related to other examples of "information projection" such as hindsight bias[14] and curse of knowledge[15]. Hindsight bias is the tendency to feel like we knew the outcome of some

[14] See Fischoff (1975).
[15] See Camerer, Loewenstein and Weber (1989).

probabilistic event (e.g., a football game, political election, or business investment) all along. It is generated by the inability to fathom ever being without the information that we have after the fact – i.e., the inability to imagine having been in a different information state. PM overconfidence is similarly caused by the difficulty of imagining being without information we have in the present, only now projecting that information state on the future rather than the past. "Curse of Knowledge" refers to the inability to ignore information in making economic decisions, even when internalizing that information is harmful. Thaler (2000) writes of the curse of knowledge, "once we know something, we can't ever imagine thinking otherwise. This makes it hard for us to realize that what we know may be less than obvious to others who are less informed." In the case of prospective memory overconfidence, the (potentially) less informed "others" are our future selves. We find it difficult to imagine not knowing (or remembering) what we know (or are aware of) in the present, and thus are overoptimistic in our capacity to later recall information that is currently top of mind.

1.5 Effects of Prospective Memory Overconfidence on Agent Behavior

We can immediately establish several results about PM overconfidence and ex ante task completion and welfare:

Proposition 3 *For any agent with imperfect memory and overconfident beliefs, the ex ante likelihood of task completion decreases in the level of overconfidence.*

Proposition 4 *For any agent with imperfect memory and overconfident beliefs, the ex ante expected utility of a task is decreasing in the level of overconfidence. Here we assume the ex ante utility is computed knowing the true memory parameters, and independently of the memory beliefs.*

The intuition for these results is that agents with PM overconfidence overestimate the continuation value of task deferral, and thus set cost thresholds below the optimal level every period before T. In our notation, $s_t(\hat{p},\hat{\theta}) < s_t(p,\theta) \forall t < T$. The lower probability of task completion follows immediately from following a strategy with lower threshold costs (holding the true memory parameters constant).

To see that utilities are lower, note that these lower threshold costs imply that for every period before T, an overconfident agent defers the task for a range of cost realizations when the first-best strategy would dictate executing the task. Thus, it must be that $V_0(s_t(\hat{p},\hat{\theta}), p,\theta) \leq V_0(s(p, \theta), p,\theta)$. More generally, cost thresholds are decreasing in the level of PM overconfidence, as we show in the appendix.

In comparing the behavior of agents with PM biases to rational agents (PRs and SFs), we find it useful to define an agent exhibiting extreme bias with respect to PM beliefs. We define a ``Naive Forgetter'' as an agent with $p < 1, \theta < 1$ but $\hat{p} = 1$. NFs incorrectly anticipate perfect recall in all future periods. Note that for NFs, $s = s^*$ – that is, they employ the same set of threshold costs as PRs, only this strategy is inefficient given their true memory parameters.

1.5.1 Some Results Involving Length of Deadline

We now investigate the effect of lengthening the deadline on task completion for the PR, SF, and NF. We update the notation by adding time subscripts to $s, s_t, h, V_t, \text{and} Q$ to represent the number of periods in the particular problem.

Proposition 5 *For a PR, the probability of task completion and the expected utility of the task are strictly increasing in the time allotted. In our subscripted notation, we claim that* $V_{0,T+\delta}(s^*_{T+\delta}) > V_{0,T}(s^*_T)$ *and* $Q_{T+\delta}(s^*_{T+\delta}) > Q_T(s^*_T)$ *for all* T, δ.

Proposition 6 *For a SF, the expected utility of the task is strictly increasing in the time allotted for any* $p > 0$.

Proposition 7 *For a NF, given any two deadline lengths* T *and* $T+\delta$,

1. *There exist sufficiently poor memory parameters such that the probability of task completion is lower for the longer deadline. That is, there exists some* p', θ' *such that for* $0 < p < p'$ *and* $\theta < \theta'$, $Q_{T+\delta}(s^*_{T+\delta}, p, \theta) < Q_T(s^*_T, p, \theta)$

2. *There exist sufficiently poor memory parameters such that the ex ante expected utility from task completion is lower for the longer length of deadline. That is, there exists some* p', θ' *such that for* $0 < p < p'$ *and* $\theta < \theta'$, $V_{0,T+\delta}(s^*_{T+\delta}, p, \theta) < V_{0,T}(s^*_T, p, \theta)$.

For the intuition behind this result, consider the case where *p* is small and $\theta = 0$. In this case, the overall probability of task execution is approximately equal to the probability of task execution in the first period, since the probability of task execution at $t = 1$ is proportional to *p*, and for all other periods, is proportional to some higher order power of *p*. In particular, the ex ante probability of completing the task in period k is proportional to p^k. But the probability of task execution in period 1, $pF(s^*_{1,T})$, is decreasing in the length of deadline T, since $s^*_{1,t+\delta} < s^*_{1,t}$ implies $F(s^*_{1,t+\delta}) < F(s^*_{1,t})$. Thus, for *p* small and $\theta = 0$, $Q_{T+\delta}(s^*_{T+\delta}, p, \theta) \approx pF(s^*_{1,T+\delta}) < pF(s^*_{1,T}) \approx Q_T(s^*_T, p, \theta)$. Since $Q_{T+\delta}(s^*_{T+\delta}, p, \theta)$ and $Q_T(s^*_T, p, \theta)$ are continuous functions of p and θ, $Q_{T+\delta}(s^*_{T+\delta}, p, \theta) - Q_T(s^*_T, p, \theta)$ is continuous in *p* and θ and thus the result holds for some range of sufficiently p, θ.

The analogous result for ex ante expected utility can be demonstrated in a similar manner. The ex ante expected utility from executing in period 1 – $pF(s^*_{1,T})E(c \mid c < s^*_{1,T})$ – approximates overall expected utility for small *p* and $\theta = 0$, since it is proportional to *p* and expected utility contributions for all future periods are proportional to some higher order power of *p*. As long as $s^*_{1,t} < b$ (which must always hold), the ex ante expected utility from executing in period 1 is decreasing in *t*, since $s^*_{1,t}$ is decreasing in *t*. Again, by the continuity of $V_{0,T+\delta}$ and $V_{0,T}$ in p, θ, their difference is continuous in p, θ, thus expected utility for a NF is decreasing in the length of deadline for some range of sufficiently small p, θ.

1.5.2 Two- to Three- Period Example

Table 1.3 shows the change in the probability of task completion and expected utility for SFs as the number of periods is increased from 2 to 3 (for $p = 0.5$, $\theta = 0.2$, $b = 1$, $c \sim U[0,2]$). Consistent with Proposition 6, expected utility increases for all combinations of p, θ.

Table 1.3: Change in Probability of Task Completion and Expected Utility for SF, 3 Periods vs 2 Periods, $b = 1$, $c_t \sim U[0,2]$

	Change in Probability of Task Completion					Change in Expected Utility				
p/θ	0.0	0.2	0.5	0.8	1.0	0.0	0.2	0.5	0.8	1.0
0.2	0.000	0.005	0.023	0.051	0.074	0.001	0.004	0.014	0.029	0.041
0.5	0.005	0.019	0.051	0.090	0.114	0.009	0.019	0.039	0.062	0.076
0.8	0.037	0.051	0.077	0.101	0.110	0.044	0.053	0.069	0.084	0.090
1.0	0.096	0.096	0.096	0.096	0.096	0.093	0.093	0.093	0.093	0.093

Table 1.4 shows the change in the probability of task completion and expected utility for NFs as the number of periods is increased from 2 to 3. Consistent with Proposition 7, the probability of task completion and expected utility decrease for low values of p, θ.

Table 1.4: Change in Probability of Task Completion and Expected Utility for NF, 3 Periods vs 2 Periods, $b = 1$, $c_t \sim U[0,2]$

	Change in Probability of Task Completion					Change in Expected Utility				
p/θ	0.0	0.2	0.5	0.8	1.0	0.0	0.2	0.5	0.8	1.0
0.2	(0.014)	(0.011)	0.004	0.030	0.051	(0.003)	0.000	0.010	0.025	0.037
0.5	(0.021)	(0.006)	0.028	0.068	0.093	0.004	0.014	0.034	0.058	0.073
0.8	0.022	0.037	0.066	0.092	0.102	0.042	0.051	0.068	0.083	0.090
1.0	0.096	0.096	0.096	0.096	0.096	0.093	0.093	0.093	0.093	0.093

1.6 Extensions of the Model

1.6.1 Memory Aids and Reminders

Sophisticated agents accurately recognize the welfare loss from their fallible memory. As such, they may have a demand for memory aids which, at some cost, ensure (or increase the likelihood of) remembering in future periods. This can be thought of as analogous to sophisticated hyperbolic discounters' demand for commitment devices.

We define a memory aid $a = (a_1, a_2, \ , a_T)$, where $a_t \in \{0,1\}$ for all $t \in \{1...T\}$. If $a_t = 1$ then the agent remembers the task in period t with probability 1. Define a^1 as the complete memory aid such that $a_t^1 = 1$ for all $t \in \{1...T\}$, and a^0 as the "no aid" choice such that $a_t^0 = 0$ for all $t \in \{1,...,T\}$. Let A denote the set of all 2^T possible memory aids, and $\kappa(a)$ denote the utility-price of memory aid a, with $\kappa(a^0) = 0$.
We assume that in period 0 the agent is presented with the full menu of memory aids A and prices $\kappa(a)$ for each $a \in A$. The agent can choose to implement any memory aid at cost $\kappa(a)$, and is then completely committed to that memory aid for the full length of the problem. A memory aid cannot be adjusted at any time after $t = 0$.

We now define the following memory-aid-augmented versions of V_t and s:

1. $U_t(s, p, \theta, a)$ as the continuation value at t from executing strategy s given memory parameters p, s, and memory aid a; and
2. $\psi(p, \theta, a)$ as the first-best strategy given p, θ, a.

Note that $\psi(p, \theta, a^0) = s(p, \theta)$ and $U_t(\psi(p, \theta, a^0), p, \theta, a^0) = V_t(s(p, \theta), p, \theta)$.

Given a complete menu of memory aids at time 0, define a^* as the memory aid chosen by the SF. Since the SF correctly perceives her lack of memory, we can define this memory aid as one satisfying the following property:

$$a^* = \arg\max_{a \in A} \left[U_0(\psi(p, \theta, a), p, \theta, a) - \kappa(a) \right]$$

The following Proposition examines the welfare loss to a SF of her forgetting.

Proposition 8 *With a complete menu of memory aids, a SF's welfare loss relative to that of an otherwise identical PR is bounded by $\kappa(a^1)$ as $b \to \infty$.*

When one views $\kappa(a^1)$ as the cost of upgrading to perfect memory, this proposition is obvious.

Memory aids provide three benefits for a forgetful agent. First, they increase the ex ante expected utility contribution for every period with a reminder, since they increase the probability of recall in those periods to 1. But they provide two other subtler, yet important benefits. Memory aids allow agents to "gamble" with lower threshold costs in any period before a reminder. Consider the case where *p* is small. Without memory aids, threshold costs will be close to *b* in every period, since the SF recognizes the small chance of recall. However, with any reminder, the SF knows that she will remember the task for sure in some period, and thus can more aggressively set a threshold cost away from *b* without worrying about wasting an opportunity to execute the task. The third benefit of a memory aid is that it resets the memory process to the state of recall (i.e., a reminder at t implies $N_{t+1} = 0$), and thus increases the chances of recall in all future periods. This benefit could be thought of as keeping a low θ from "kicking in" – that is, mitigating the adverse effect of repeated forgetting on the probability of recall.

What does the optimal memory aid look like, in terms of the number and sequencing of reminders? It depends on the relative size of these three effects. The first two effects suggest that if the price of reminders across time is constant, later reminders will be preferred to earlier reminders. Setting the probability of recall in any period to 1 is more beneficial in later periods, since it buys more recall. Without memory aids, the ex ante probability of recall is declining in time, and so a later reminder increases that period's recall by a larger amount than an earlier reminder would. Furthermore, later reminders provide agents with more periods to gamble with low threshold costs.

However, the third effect, that is, mitigating the effect of θ on the likelihood of recall, points to another reminder structure. In particular, this effect would induce an agent to implement an intermittent reminder, where the agent spreads reminders out as much as possible, reducing the welfare loss caused by successive forgetting. For instance, if an agent can choose any *T/2* reminders, choosing reminders every other period will eliminate θ from the problem entirely. Thus, for low values of θ, this third effect can dominate the first two effects, and the agent will prefer to spread reminders out rather than concentrate them close to the deadline.

1.6.2 Three-Period Example with Memory Aids

To illustrate the SF's solution with a complete menu of memory aids, we return to our 3-period example, where $b = 1$, and $c : U[0,2]$. Let a^{ijk} denote the memory aid with $a_1 = i$, $a_2 = j$, and $a_3 = k$. For example, a^{011} denotes a memory aid with reminders in the 2nd and 3rd periods.

Table 1.5 below summarizes the SF's willingness-to-pay for the 7 possible memory aids, for 4 sets of memory parameters: $(p,\theta) \in \{(0.1,0.1),(0.1,0.9),(0.9,0.1),(0.9,0.9)\}$.

Table 1.5 illustrates how the relative sizes of p and θ can dictate whether the agent will prefer to spread out or backload reminders. Consider the relative WTPs for a single reminder. When p is low and θ is high the impact of θ on the probability of recall each period is minimized[16] and the agent strictly prefers later deadlines (see column 2). However, in all other scenarios, when either θ is low, or θ is high but p is also high, the agent prefers the reminder in the middle period. Note that, with memory aid a^{010}, θ is eliminated from the problem: the probability of recall in any period depends only on p, since the number of successive forgettings entering any period is always 0. While the expected utilities with the two other single-reminder aids are increasing in θ, the expected utility with a^{010} is constant with respect to θ. Thus, for a given p, lower values of θ improve the value of a^{010} relative to a^{100} and a^{001}.

[16]This is true because θ only appears in the expressions of expected utility when interacted with p.

Table 1.5: SF's WTP for Memory Aids, b = 1, c_t ~ U[0,2], T = 3

	$p = 0.1$		$p = 0.9$	
	$\theta = 0.1$	$\theta = 0.9$	$\theta = 0.1$	$\theta = 0.9$
1 reminder				
a^{100}	0.2354	0.2086	0.0413	0.0122
a^{010}	0.2472	0.2110	0.0555	0.0143
a^{001}	0.2372	0.2111	0.0451	0.0140
2 reminders				
a^{110}	0.3695	0.3333	0.0652	0.0240
a^{101}	0.3704	0.3342	0.0662	0.0250
a^{011}	0.3709	0.3347	0.0667	0.0255
3 reminders				
a^{111}	0.4544	0.4182	0.0759	0.0347

In contrast to SFs, NFs will never purchase memory aids. The reason is analogous to the explanation for why naive hyperbolic discounters do not use commitment devices. Since NFs expect to remember the task with probability 1 in all future periods, any memory aid with cost greater than 0 will be perceived as a strictly welfare-reducing choice. This leads directly to the proposition below:

Proposition 9 *Even with a complete menu of memory aids, a NF's welfare loss relative to that of an otherwise identical PR grows arbitrarily large as* $b \to \infty$.

Thus, incorporating memory aids into the model amplifies the cost of PM overconfidence. Without memory aids, overconfidence is welfare-reducing because agents set excessively low threshold costs. With memory aids, overconfidence is doubly costly – threshold costs are still too low, plus the agent abstains from purchasing welfare-improving memory aids.

Incorporating memory aids into the model provides an explanation for the experimental finding that subjects' retrospective and prospective memory performance may be negatively correlated (Wilkins and Baddeley 1978). Consider a group of subjects who are SFs, but with different true memory parameters. If memory aids are sufficiently costly, they will only be employed by agents with poor memories. Thus, we may observe that agents who have poorer memories (as tested through retrospective recall) may perform better on prospective memory tasks, but only because they are more likely to employ memory aids.

1.6.3 Present-Biased Preferences and Imperfect Memory

Earlier, we asserted that the Shafir and Tversky (1992) result, that the probability of task completion could be decreasing in the length of deadline, cannot be derived from hyperbolic discounting alone. We will explicitly show this by incorporating time-discounting into our prospective memory model, and proving that any hyperbolic discounter (naive or sophisticated) with perfect memory will be made strictly better off by lengthening the deadline.

We choose a basic quasi-hyperbolic form of discounting, where the discount factor β between the present and all future periods is less than one, while the discount factor between all future periods $\delta = 1$. We set $\delta = 1$ for analytical tractability, and maintain that choosing a δ close to, but less than 1 will not meaningfully alter our results.
Furthermore, we assume, as in O'Donoghue and Rabin (1999b), that upon task execution, costs are incurred immediately, while benefits are delayed. Thus, the utility from executing a task in period t is now $\beta b - c_t$ (as opposed to $b - c_t$ in our baseline model).

Hyperbolic discounters determine their threshold costs each period in a similar manner to agents in our baseline model, with some important distinctions. First, note that with hyperbolic discounting in any period t for given memory parameters p, θ, the value of a particular task which we denote $V_t^H(s, p, \theta, \beta)$ is less than $V_t(s, p, \theta)$. In particular,

$$V_t^H(s, p, \theta, \beta) = \beta V_t(s, p, \theta).$$

Furthermore, for partially or fully naïve hyperbolic agents,[17] anticipated threshold costs are greater than true threshold costs. Intuitively, naïve hyperbolic agents believe that they will not discount the deferred benefit as greatly in the future, and thus predict threshold costs will be higher than they are. Define $\hat{s}_t^H$ as the anticipated threshold cost in period t, from the perspective of any period prior to t. Then:

$$\hat{s}_t^H = \begin{cases} \hat{\beta} b & \text{if } t = T \\ \hat{\beta} b - E_t(V_{t+1}^H(\hat{s}^H, \hat{p}, \hat{\theta}, \hat{\beta})) & \text{if } 1 \le t < T, \end{cases}$$

while

$$s_t^H = \begin{cases} \beta b & \text{if } t = T \\ \beta b - E_t(V_{t+1}^H(\hat{s}^H, \hat{p}, \hat{\theta}, \hat{\beta})) & \text{if } 1 \le t < T, \end{cases}$$

Note that a naïve hyperbolic agent with $\hat{\beta} = 1$ anticipates the same threshold costs as a SF in our model. That is, for $\hat{\beta} = 1$,

$$\hat{s}_t^H = \begin{cases} b & \text{if } t = T \\ b - E_t(V_{t+1}^H(\hat{s}^H, \hat{p}, \hat{\theta}, 1)) & \text{if } 1 \le t < T, \end{cases}$$

[17] Defined in O'Donoghue and Rabin (1999a) as agents who anticipate the value of β in future periods – $\hat{\beta}$ – to be greater than its current value, β.

which is exactly the set of threshold costs for a SF with PM beliefs $\hat{p}, \hat{\theta}$ (see equation 1.1). Given this, the naïve hyperbolic's actual set of thresholds is:

$$s_t^H = \begin{cases} \beta b & \text{if } t = T \\ \beta[b - E_t(V_{t+1}(\hat{s}^H, \hat{p}, \hat{\theta}, 1)) & \text{if } 1 \le t < T, \end{cases}$$

or simply $\beta s_t(\hat{p}, \hat{\theta})$, which is the fraction β of the optimal cost threshold at t if the "true" memory parameters were $\hat{p}, \hat{\theta}$.

Note, for a sophisticated hyperbolic agent with $\hat{\beta} = \beta$, $\hat{s}_t^H = s_t^H$ for all $t \in \{1, ...T\}$. For a given β, the sophisticated hyperbolic agent's anticipated and executed cost thresholds are below the naïve hyperbolic agent's anticipated thresholds, and above the naïve hyperbolic agent's actual thresholds.

We now assess the welfare effect of lengthening the deadline on perfect-remembering hyperbolic discounters:

Proposition 10 *For any hyperbolic discounting agent with* $p = \hat{p} = \theta = \hat{\theta} = 1$*, the probability of task completion and expected utility from task completion are strictly increasing in the time allotted.*

The logic behind this proposition is the same as for Proposition 5 and the proof in the appendix is similar.

While our length of deadline result cannot be generated solely from hyperbolic time-preferences, we do not mean to imply that hyperbolic time preferences do not importantly influence agent behavior in our model. Rather, we believe there are positive interaction effects between naïvete with respect to fallible memory and procrastination. In particular, for agents with imperfect memory and hyperbolic time preferences, the cost of fallible memory is increasing in the tendency to procrastinate (as captured by β), and vice versa. Table 1.6 shows the ex ante probability of task execution and the non-discounted expected utility over a range of p and β for naive-forgetting, naive-hyperbolic agents, for our 3-period model with $b = 1$, $c_t \sim U[0,2]$, and $\theta = 0.2$:

Table 1.6: Probability of Task Completion and Ex Ante (Non-Discounted) Expected Utility for NF Naive-Hyperbolics, $b = 1$, $c_t \sim U[0,2]$, $T = 3$, and $\theta = 0.2$:

	Probability of Task Completion				Expected Utility			
p/β	0.2	0.5	0.8	1.0	0.2	0.5	0.8	1.0
0.2	0.019	0.047	0.073	0.089	0.018	0.039	0.053	0.060
0.5	0.067	0.156	0.233	0.278	0.062	0.127	0.166	0.180
0.8	0.143	0.324	0.469	0.546	0.132	0.262	0.328	0.345
1.0	0.215	0.478	0.676	0.777	0.198	0.384	0.468	0.483

Note, for example, that for $\beta = 0.8$, the probability of completion goes from 0.469 to 0.073 as p goes from 0.8 to 0.2, a reduction of 84%. For $\beta = 0.2$, the probability goes from 0.143 to 0.019, a reduction of 87%. Similarly, fix $p = 0.8$ – the percentage reduction in the probability of completion for $\beta = 0.2$ vs $\beta = 0.8$ is 70%, while for $p = 0.2$ the reduction is 74%. Similar interactions hold for expected utility. These results suggest that the mixed partials with respect to p and β of both $\log Q$ and $\log V_0$ are positive. In words, the percentage reduction in the ex ante probability of task execution and non-discounted expected utility as p is lowered are decreasing in β, and decreasing in p as β is lowered.

Note however, that the interaction of fallible memory and procrastination is dependent on the agent's relative sophistication with respect to her memory parameters and time preferences. As discussed in Ericson (2010), for a sophisticated forgetter, fallible memory may effectively act as a commitment device against procrastination, elucidating that the "now or later" choice is actually "now or never."

1.7 Prospective Memory Failures in the Marketplace

1.7.1 Mail-In Rebates

Many consumer goods companies (particularly electronics manufacturers) offer mail-in rebates, ranging in value from less than \$1 to several hundred dollars. To claim the rebate, purchasers generally must fill out an information card and mail it (with proof of purchase) to a provided address within some time frame after the purchase. This time frame generally varies from 4 weeks to 3 months. Experts estimate that the rebate market has grown from under \$1 billion to \$4-\$10 billion in the past decade, and that the number of rebates offered are in the hundreds of millions.[18] However, typically only 5% of outstanding rebates are ever claimed by consumers, and while claim rates rise with the value of the rebate, redemption rates are typically less than 40% even for "big-ticket" electronics items with rebate values greater than \$20.[19]

While the traditional explanation for rebates is that they are a form of price discrimination by retailers, an increasing body of evidence suggests that a significant amount of rebate non-redemption (or "breakage") comes from consumers who initially planned on redemption – a dynamic not captured by the price discrimination model. In addition to the experimental studies cited earlier (Silk 2004 and Ericson forthcoming), Silk and Janiszewski (2003), in a survey of 35 marketing managers at firms that use rebates, find that those managers estimate that nearly two-thirds of breakage is attributable to customers that had intended to redeem the rebate, but for some reason or another, failed to do so. One possible reason for this breakage is that firms are exploiting naïve forgetting by consumers. Naïve forgetters fully internalize rebate discounts in making purchase decisions, failing to account for the prospective memory lapses that often result in non-submission. Firms may have recognized that through this naïve forgetting, rebates increase the effective downward price elasticity of demand, and consistent with profit-maximization, have made greater use of them.

[18] See Grow and Chatwal (2005).
[19] See Spencer (2002).

There are other explanations for rebates beyond price discrimination and naive forgetting, both on the consumer and firm sides. Consumers may be naïve not about their forgetfulness, but rather their present-biased preferences (O'Donoghue and Rabin 1999b). That is, consumers may purchase goods expecting to complete the rebate process, not anticipating their tendency to procrastinate on such onerous tasks in the future. This explanation may hold true for lesser value rebates, but is calibrationally inconsistent with low claim rates for rebates worth $100 or more. Chen, Moorthy, and Zhang (2005) develop a theory of rebates for perfectly rational consumers, where rebates facilitate "utility arbitrage" through state-dependent pricing. However, such a rational explanation is inconsistent with anecdotal evidence of consumers' ex post regret and frustration over rebate programs. The New York Times reported that rebate complaints to the Better Business Bureau increased nearly three-fold from 2001 to 2005 (from 964 to 2,715), and US News and World Report more recently documented that the complaint volume cleared 4,500 in 2007.[20] Furthermore, due to pressure from consumers and advocacy groups, legislation has been proposed in California and Texas and passed in New York to regulate rebate practices. The New York law addresses rebate form availability, establishes a 2-week minimum redemption period, requires rebate rewards to be paid within 60 days of submissions, and restricts retailers from misleading advertising of post-rebate prices.[21]

On the firm side, rebates have been justified on grounds other than price discrimination, such as collecting consumer data. While these explanations may have some validity, rebates would not likely be as commonly offered if claim rates were closer to 100%. Larissa Hall, the vice president of marketing at Buy.com, stated that "the reason we can offer rebates is because not everybody will redeem them" (Chuang 2003). Analysts tracking consumer goods companies have claimed that "rebates are a good business plan only when consumers fail to claim them,"[22] and fulfillment companies often market their services to manufacturers by touting the low redemption rates for promotions they have administered.

One prediction of our model is that firms should lengthen rebate deadlines, or even restrict rebate submissions until after some grace period, to induce more forgetting by naive consumers. While we have no systematic evidence to this effect, we have observed instances where new rebate practices do seem geared towards further capitalizing on forgetting by consumers. For instance, a number of fulfillment companies are now offering manufacturers and retailers "deferred" rebate programs, where consumers have to file initial paperwork soon after the purchase, but then must file an additional claim as much as a full year later. These "deferred" rebate programs generally offer substantially greater discounts than regular rebates, but these discounts are more than offset by the lower claim rates over the lengthened rebate process. Additionally, several on-line retailers have offered free-after-rebate deals – essentially a 100% rebate – while initially pricing goods above the suggested retail price. Indeed, a number of Internet retailers made such promotions the core of their business strategy.[23] As with deferred rebate plans, free-after-rebate deals can be explained as an attempt by manufacturers and retailers to further capitalize on naive forgetting by consumers. Edwards (2009), in commenting on this increasingly evident strategy

[20] See Tugend (2006) and Palmer (2008).
[21] See Edwards (2009).
[22] See McLaughlin (2002).
[23] See Chen et al. (2005) for a more detailed discussion of this sales tactic.

by firms, notes that one potential issue with the New York rebate law is that "it does not set any maximum time period....[and] some scholarship has suggested that longer deadlines can lead to fewer rebate redemptions due to increased chances for consumer procrastination and forgetfulness."

1.7.2 Free Trial Offers

A growing practice among firms with subscription-based products is to offer consumers "risk-free" trials. Consumers are not billed and may cancel at any time during the trial. However, if the trial offer is not canceled before its expiration, it is automatically renewed at the full subscription price. An agent with imperfect memory may learn during the trial that she would not be willing to renew the subscription, but simply forget to cancel before the deadline, and incur a loss as a result.[24] Overconfidence in prospective memory could further exacerbate the welfare loss from unwanted renewals. We collected data on a free trials run by a subscription-based website that appears consistent with the key prediction of our model incorporating memory overconfidence – that task incompletion caused by naive forgetting may be increasing in the time allotted before the deadline. The website randomly e-mailed prospective subscribers one of two trial offers – one for 3 days, and one for 7 days. Trial subscriptions were automatically renewed on a monthly basis at the end of the trial. Post-trial retention was 28% for the 3-day trial group, and 41% for the 7-day trial group. The difference in retention rates is significant at 99% confidence.[25]

The proliferation of credit cards over the past 20 years has greatly facilitated the recurrent billing process and opened up virtually the entire US consumer population to be targeted by negative option marketers.[26] Publishers and other companies selling subscription-based products and services have not only unilaterally marketed free trials, but partnered with third parties in marketing free trials as a "bonus" item. For instance, most major US airlines allow customers to redeem miles for short, 3-6 month magazine subscriptions that are automatically renewed at the regular subscription rate. Retailers (particularly online) often offer these subscriptions as a "free bonus" with any purchase.[27] A relatively recent phenomenon is "incentive" marketing, where intermediary companies offer consumers gifts for signing up for multiple trial offers at once. These companies are paid \$40-\$60 per free trial signup, leaving them a comfortable profit

[24] In Chapter 2, I describe a field experiment with a subscription-based website which finds that varying the strength of the reminder that subscribers receive on their credit card bill has a significant effect on cancelation rates – evidence that forgetting does play a role in consumer behavior with respect to subscription-based (or more broadly, "negative-option") goods.

[25] The free trials were offered to consumers on the website's mailing list, most of whom were past subscribers. To the extent that past subscribers already have a sense for the quality of the website, the learning/information explanation for the existence of free trial offers (discussed later in this section) is less likely to explain the disparity in retention rates.

[26] Before credit cards were widespread, customers generally paid for subscriptions by check or money order. According to a publishing executive we interviewed on this subject, payment rates for free trial customers who were billed for automatic extensions through the mail were extremely low, reducing the attractiveness of this strategy.

[27] One example of this is a promotion run by the Sam Goody music store chain, where customers making a purchase were offered a three-month free trial subscription to Entertainment Weekly magazine. Customers filed a class-action suit against the retailer in August 2003, claiming they were not informed that their billing information would be passed on to the publisher and that their credit cards would automatically be charged for a full subscription at the end of the trial period.

margin when rewarding 5-10 free trial signups with a $100-$200 gift.[28] One such company, Gratis Internet, grossed over $20M in 2005, and reported giving away over 20,000 iPods to consumers as rewards for their signups.[29]

Anecdotal evidence suggests that firms are increasing the length of their free trials, possibly to increase naive forgetting by consumers. Many weekly magazines (such as Sports Illustrated) have increased the standard length of their free trials from 4 issues to 3 months. In the past 5 years, AOL has increased the length of their dial-up service free trial offer from 30 to 50 days. In addition, a new class of health-related products – weight-loss supplements, smoking-cessation aids, even teeth whiteners – has emerged via direct marketing channels such as the internet and infomercial programming. Their manufacturers offer extended free or reduced-price trials – of lengths running in months rather than days or weeks – ostensibly to give their new customers as much time as possible to test and experience the benefits of their products, but effectively as a way to more fully capitalize on consumer forgetting. This extended-trial tactic has garnered such artificially high takeup rates that it has elicited a high volume of consumer complaints to the Better Business Bureau, Federal Trade Commission, and credit card companies. "Acai Supplements and Other 'Free' Trial Offers" ranked #1 on the BBB's List of the Top 10 Scams and Rip-Offs in 2009.[30] The BBB reported that one company selling acai berry supplements received more complaints in 2009 than the entire airline industry combined.[31] Visa recently removed 100 companies that were marketing "deceptive free trial offers" from its authorized vendor list, after a rash of customer complaints and chargeback requests.[32] The extent of ex post consumer regret and frustration over these offers is not only suggestive that consumers are incurring a loss of welfare from this marketing strategy, but also runs counter to the rational explanation for free trials -- that they simply offer consumers a free or low-cost opportunity to learn about their product before making a greater investment. It is hard to imagine consumers generating this volume of complaints just from learning about their preferences.

While the welfare loss from any one free trial transaction may be small, the aggregate loss generated by this exploitation of naive forgetting may be quite significant. Magazine industry revenues topped $20 billion in 2009,[33] and many other products and services use negative option models, including telecommunications providers, music and book clubs, and information services (such as credit report agencies). Visa, in a survey of 1,000 cardholders in September

---

[28] See Whitaker (2004).

[29] See Inc.com company profile at http://www.inc.com/inc5000/2007/company-profile.html?id=2005018

[30] See full BBB Bulletin at http://www.bbb.org/us/article/bbb-lists-top-10-scams-and-rip-offs-of-2009-14436

[31] See Willis (2009).

[32] See full notice to Visa customers at http://usa.visa.com/personal/security/learn-the-facts/deceptive-marketing.html?ep=v_sym_negativeoption. Chargebacks are the main drawback of this marketing tactic, according to a publishing executive that we interviewed on the subject of automatic renewals. Chargebacks occur when customers call their credit card issuer directly to dispute a charge. The card issuer then charges the merchant back for the amount, unless the merchant files a response. For most merchants (including magazine publishers), the administrative cost of filing responses is too great to contest each chargeback, and thus chargebacks usually stand. If a merchant's chargebacks exceed some threshold (usually around 2% of all transactions) in a billing period, the merchant is assessed a large penalty (on the order of $100,000).

[33] According to year-end numbers published by the Publishers Information Bureau at http://www.magazine.org/advertising/revenue/by_mag_title_ytd/pib-4q-2009.aspx.

2009, found that nearly 30% had been "duped" by some form of negative option marketing. With over 575 million credit cards and 500 million debit cards issued to 175 million cardholders in the US by year-end 2009,[34] an extrapolation of that 30% to the broader credit-card-carrying consumer population would suggest that the welfare loss from consumers' prospective memory failures is pervasive and considerable.

1.8 Conclusion

As discussed in Sections 1.2 and 1.4, empirical evidence on prospective memory beliefs and performance is somewhat limited. The experimental results in Silk (2004) and Ericson (forthcoming) are strongly suggestive that naïve forgetting impacts consumer behavior in the field. Rebates and free trial offers are two ways that firms appear to exploit naïve forgetting to consumers' detriment, and more field work in these areas would be valuable in calibrating the overall welfare loss from these deviations from the standard model. Ideally, field tests could be arranged in a "revealed-preference" design, giving subjects the choice between immediate payoffs and "rebate coupons" of varying values and deadlines to enable a precise identification of consumers' memory beliefs and true parameters (as in Ericson's experimental design). Another interesting treatment would be to allow consumers to choose their deadline length, as in Ariely and Wertenbroch (2002).[35] Field evidence on rebates and free trials may also shed greater light on the equilibrium effects of fallible memory that may be obscured in experimental settings, as memory aids and other commitment-type devices may be effectively used to limit welfare losses in "real-world" settings. For instance, consumer advocacy groups are now touting the use of disposable, "virtual credit cards" – numbers that are issued for a single-use only, and then invalidated – for the purposes of avoiding unwanted takeup with negative option goods.[36]

From a theory standpoint, there are potentially interesting extensions of our model that we have not yet examined. In order to compare and contrast our findings to the present-biased preferences literature, we have thus far focused on examining onerous tasks – that is, ones with upfront costs and deferred benefits. We have followed the literature in assuming an environment with stochastic costs and some fixed benefit to task execution. However, our framework is more generally applicable to environments with time-varying costs and benefits (whether stochastic or deterministic). In such environments, it may be interesting to examine the bounds on losses (relative to the first-best) for different types of agents.

Another important theoretical consideration concerns potential learning about prospective memory. This learning may occur across tasks (e.g., if I forget to submit one rebate, I submit my next rebate right away), or more relevantly to our current framework, within tasks. In particular, it may be that NFs update their beliefs such that after a string of recall failures, they revise beliefs downward and correspondingly increase their threshold costs. For instance, in everyday life, often times if we have forgotten to perform a task for a long while and it suddenly pops into

[34] See link to Visa's notice to customers referenced in footnote 28.

[35] This proposed study would be different in that the task – completing a rebate – would have greater applicability to our PM model. Their task – completing a class project – may have too many natural reminders to be relevant in our context.

[36] This service is offered by most large credit card companies, as well as payment services such as Paypal.

our head, we may think “better do it now before I forget again.” However, if, by luck of the draw, the same task continually comes to our attention, we are much more likely to think we can defer the task and bank on future recall.

Chapter 2: Exploiting Consumer Forgetting: Camouflage Pricing with Negative-Option Subscriptions

2.1. Introduction

Chapter 1 developed a model of how fallible prospective memory and overconfident beliefs about memory may affect consumer behavior. One could imagine a real-world consumer solving a multi-task generalization of this model – that is, each period, there is some probability that each consumer choice (e.g., whether to buy a new stereo) or task (e.g., mailing in a rebate coupon) is recalled, and once it is in the consumer's attention, she can choose whether to complete it or not.

As the single-task model generalizes to a multi-task setting, the top-of-mind dynamic for consumers becomes increasingly complex – the question is no longer *if* a task is recalled, but *which* tasks are recalled, and why. To the extent that consumers can endogenize the probability of recall across tasks, they may allocate their limited attention span as they would any other scarce resource – that is, by equating the net marginal benefit across all possible uses of their attention. Gabaix, Laibson, Moloche, and Weinberg (2006) find that, in an experimental setting, subjects adopt an attention allocation strategy closely approximating perfect rationality.[37]

However, while consumers may seek to allocate their limited memory capacity as efficiently as possible across tasks, tactics employed by firms to exploit limited memory may confound these efforts.[38] In certain cases, firms may attempt to increase sales by *reminding* consumers of certain choices at specific periods of time, perhaps cueing them to make purchases they would not otherwise make (indeed, this is one interpretation of the value of advertising). On the other hand, there may be cases where firms may wish to *obscure* certain choices from consumers in order to maximize profits (two examples, the submission of rebates and cancellation of free trial offers, are discussed in Section 1.7).

This study provides field evidence of this latter scenario from the market for recurrently billed subscriptions and memberships. There is growing evidence that consumers continue paying for such negative-option plans long after they have stopped using them. DellaVigna and Malmendier (2006) find evidence of this pattern with respect to health club memberships. In particular, they find that, on average, among health club members being billed on a recurrent monthly basis, "2.31 full months elapse between the last attendance and contract termination with associated membership payments of $187..[and] this lag is at least four months for 20 percent of the users." While DellaVigna and Malmendier attribute this behavior to naïve present-biased preferences,[39] , an alternative story is that these firms are exploiting forgetting by consumers. For the vast majority of goods, the consumption decision is extremely salient at the time of purchase. For instance, an individual in a grocery store is well aware of her decision to consume a certain good

[37] The experiment is about attention allocation and not memory recall in particular, but may be applicable to the prospective task framework if one considers it a related form of allocating limited cognitive bandwidth.
[38] See Ellison (2006) for a survey of bounded rationality in the Industrial Organization field, including a review of the related psychology and economics literature that focuses on how "consumers depart from rationality in some particular way….then explore[s] how a monopolist would exploit such a bias."
[39] Such members, they reason, may be overconfident that their "future selves" will use the membership to exercise, even as their "current selves" procrastinate.

at the point of purchase – she makes that choice in selecting the item off the shelf, and yet again when the item is rung up on the register. In contrast, firms selling subscriptions or memberships may implement recurrent billing methods to capitalize on the *low* salience of the consumption decision at the time of renewal. In particular, by the time a subscription or membership is up for renewal, consumers who are not making regular use of the service may have forgotten that they are paying for it at all. Thus, firms benefit from the low salience of automatic renewals as some portion of subscribers will not *choose to renew*, but *forget to cancel*, leading to a higher overall retention rate.[40]

The simplest field test of this memory hypothesis – that certain consumer choices are subject to memory recall, and that firms can increase profits by manipulating consumers' awareness of these choices – would be for a firm to randomly vary its consumers' awareness of a particular choice, and monitor the effects of this treatment. In the context of subscription goods, this might mean sending members reminders of impending renewals, and seeing if these reminders increase cancellation rates.

Instead, this field study exploits a billing strategy recently employed by membership-based web sites, in which the treatment condition actually *lowers* consumers' awareness of the renewal choice. Membership-based web sites generally offer subscriptions on a recurrent billing basis, charging customers' credit cards on a 30-day cycle. Some such web sites have reported that switching from charging "standard" amounts, such as $10 or $9.95, to "camouflage" amounts such as $9.87 or $10.06, has led to a substantial decline in cancellation rates.[41] The leading explanation for this result is that credit card statements serve as a reminder of the subscription purchase (and more importantly, a reminder to cancel the subscription if it is not being used), and consumers are much more likely to notice, hence react to, a standard amount, since these tend to stand out more on a credit card bill. Conversely, camouflage amounts blend in with other similar charges such as grocery, gas, or taxed clothing expenses, and are thus much less likely to get noticed by consumers. Randomly quoting prospective subscribers standard or camouflage membership fees then serves as the desired manipulation of consumer awareness of the renewal decision – it varies the size of the positive shock to the memory process from reviewing a credit card bill, since the standard amounts are much more likely to draw a subscriber's attention, hence remind her of her subscription (and to cancel if appropriate).

The partnering website implemented this camouflage pricing test for a 40-day period in the spring of 2004. During that period, visitors to the site who clicked on a membership sign-up link were randomly allocated to one of several pricing groups – some were charged standard amounts, others camouflage amounts. Over the next 4 months, US customers randomly assigned a camouflage price had an 8-10% lower cancellation rate than US customers assigned a standard price. This translates to a 3 percentage-point lower monthly hazard rate (probability of canceling conditional on being active at the start of a month).[42] Furthermore, non-US customers – all of whom are effectively charged camouflage amounts due to currency conversion on their credit

[40] The fact that prospective customers are often offered discounts or bonus gifts for enrolling in automatic recurrent billing plans is further suggestive evidence of their value to firms.
[41] The president of the website in this study heard this claim made by other Internet executives at various e-commerce conferences.
[42] The average monthly hazard rate for the 4-month period was roughly 42%.

card statements – also had an 8-10% lower cancellation rate than US customers in standard pricing groups. Additionally, within the non-US customer segment, there is no significant difference in cancellation rates between subscribers allocated to standard price and camouflage price groups, suggesting that selection at signup is not driving the US-customer results. The lower hazard rate from camouflage pricing translates into a 10% increase in revenue per subscriber over the first year after signup.

Furthermore, the U-shaped pattern of hazard rates over the 30-day billing cycle suggests that forgetting plays a key role in cancellation outcomes. The spike in cancellations at the end of the cycle may be explained by a number of factors, including procrastination, stochastic cancellation costs, or retaining the option value of renewal while learning the value of the subscription. However, none of these explanations predict a cancellation spike at the beginning of the billing cycle. On the other hand, if some customers who would otherwise cancel simply forget about their subscriptions, an early-cycle cancellation spike might be due to a reminder effect. Seeing the new charge for the ensuing month reminds forgetters of their subscriptions, and may prompt them to cancel immediately. While they may still have some incentive to defer cancellation for the reasons listed above, these subscribers may simply opt to cancel the subscription before they forget again.

2.2.1 Description of the Website

The website employs a subscription-based offering, with 6 types of memberships available to customers.[43] In the first half of 2004, the website averaged over 150,000 members. Memberships are primarily offered on a monthly or annual basis (occasional promotions offer quarterly or semi-annual memberships). The website periodically runs special offers for free trial periods, which vary in length from three days to several weeks, but the majority of members begin paying membership dues immediately upon sign-up.

Roughly 90% of subscribers are enrolled in one of the two main membership packages. The most popular package, the standard offering, costs $19.95 per month, or $95.40/year. The second most popular package, the "premium" offering, costs $29.95 per month, or $155.40/year.[44]

The website features a free section which is accessible to non-members. New memberships are generated when site visitors who are browsing the free section click on either an ad banner or pop-up window promoting the various club memberships. The visitors are then taken to a sign-up page where they can choose between the various membership options (e.g., monthly vs. annual), and enter their personal and payment information. At the time of the test, the web site's only accepted form of payment was credit card.

---

[43] In accordance with the Non-Disclosure Agreement signed for this project, the name and specific industry of the website will not be disclosed.

[44] The annual memberships are advertised as costing $12.95/month and $7.95/month, to make salient the substantial savings over the month-to-month subscriptions.

2.2.2 Description of the Experiment

The experiment was designed to test three primary effects of pricing on initial take-up and subsequent renewal:

1) A threshold effect (i.e., pricing at $20 vs. $19.95);
2) A camouflage pricing effect (i.e., a $19.83 or $20.13 charge being less noticeable on a credit card bill than a $20 charge, leading to lower cancellation rates); and
3) A switch-price effect (i.e., any difference between offering and subsequently charging $20 for a membership, and offering $20, but charging a lesser amount such as $19.83)

The experiment was run from late February until the end of March of 2004, on both of the two main packages offered by the website. During that month, visitors clicking on a membership offer ad (for either the standard or premium package) were randomly sorted into one of five groups. Based on their group assignment, prospective members were taken to a sign-up page stipulating their group's offered monthly price, and, if they signed up, were subsequently billed the assigned charge for their group.[45] The offer price and actual charge for the five groups for each membership package are listed in Table 2.1.

Table 2.1: Monthly Offered Prices and Actual Charges by Group

Membership	Group 1	Group 2	Group 3	Group 4	Group 5
Standard					
Offer Price	$19.83	$19.95	$20.00	$20.00	$20.13
Actual Charge	$19.83	$19.95	$20.00	$19.83	$20.13
Premium					
Offer Price	$29.78	$29.78	$29.95	$30.00	$30.14
Actual Charge	$29.67	$29.78	$29.95	$30.00	$30.14

The attrition of all members who subscribed during the test period was tracked over the next five months. The company provided data on cancellations through mid-July 2004.

2.3 Data Description

Table 2.2 lists summary statistics from the test, including offer page views, signups, takeup rates, and monthly attrition by membership type. The attrition figures above are only for non-censored observations. While no observations are censored in the first three months, some observations (i.e., subscribers who signed up towards the end of the test period at the end of March) are censored in the 4th month. Thus, the month 4 hazard and retention rates are based only on data

[45] The annual price option mentioned earlier ($155.40 for the premium membership and $95.40 for the standard membership) was omitted during the test month.

from subscribers who were retained through the first 3 billing cycles, and whose 4th billing cycle ended before the mid-July endpoint for the cancellation data.

Table 2.2: Takeup and Attrition by Membership Type

Takeup	Standard	Premium	Total
Offer Page Views	2,633,506	1,149,885	3,783,391
US Signups	3,986	1,374	5,360
Non-US Signups	403	138	541
Total Signups	4,389	1,512	5,901
Overall Takeup Rate (%)	0.166	0.131	0.156
Attrition (Hazard %/Retention%)	Standard	Premium	Total
Month 1	56/44	68/32	59/41
Month 2	28/31	28/22	28/29
Month 3	16/26	20/18	17/24
Month 4	13/23	13/16	13/21

Potential takeup selection issues arising from the test design are addressed in Section 2.5.3.

Table 2.3 lists the month-by-month cancellations for each pricing group among US signups (see Tables 2.4 in Appendix B for non-US signups). The cancellation rates for the camouflage price groups do appear to be consistently lower than for the standard price groups. This pattern can also be seen in the month-by-month hazards by price group shown in Figures 2.1A and 2.1B (see Appendix B).

2.4 Identification Strategy

To get a more precise estimate of the camouflage price effect, I fit a logistic regression model[46] of the form:

$$log[P_{igmt}/(1 - P_{igmt})] = \alpha_t + \beta X_{igmt} \qquad (2.1)$$

where P_{igmt} is the probability that subscriber *i*, in price group *g* of membership type *m*, in month *t* of membership, cancels during the month; α_t is a set of constants to capture the baseline time trend in hazard rates; and X_{igmt} is a set of observed characteristics of subscriber *i*.

[46] See Allison (1982), Ashenfelter and Card (2002), and Efron (1988) for discussions and examples of using logistic regression models to model hazard probabilities.

Table 2.3: Monthly Cancellations by Pricing Group, US Signups

Standard			Month 1		Month 2		Month 3		Month 4	
Group	Offer	Actual	Start	Cancels	Start	Cancels	Start	Cancels	Start	Cancels
1	19.83	19.83	797	441	356	95	261	47	145	17
				55%		27%		18%		12%
2	19.95	19.95	781	444	337	106	231	36	125	23
				57%		31%		16%		18%
3	20	19.83	781	436	345	105	240	42	121	18
				56%		30%		18%		15%
4	20	20	802	475	327	107	220	39	109	16
				59%		33%		18%		15%
5	20.13	20.13	782	440	338	90	243	39	143	14
				56%		27%		16%		10%
Total			3943	2236	1703	503	1195	203	643	88
				57%		30%		17%		14%
Camouflage Pricing Groups (1,3,5)			2360	1317	1039	290	744	128	409	49
				56%		28%		17%		12%

Premium			Month 1		Month 2		Month 3		Month 4	
Group	Offer	Actual	Start	Cancels	Start	Cancels	Start	Cancels	Start	Cancels
1	29.78	29.67	267	175	92	26	66	13	35	9
				66%		28%		20%		26%
2	29.78	29.78	272	184	88	29	59	7	35	1
				68%		33%		12%		3%
3	29.95	29.95	278	203	75	26	49	14	24	5
				73%		35%		29%		21%
4	30	30	253	178	75	22	53	12	23	1
				70%		29%		23%		4%
5	30.14	30.14	278	192	86	22	64	11	33	2
				69%		26%		17%		6%
Total			1348	932	416	125	291	57	150	18
				69%		30%		20%		12%
Camouflage Pricing Groups (1,2,5)			817	551	266	84	189	31	103	12
				67%		32%		16%		12%

Table 2.3 Notes: The difference between the active subscriptions at the end of month 3 (month 3 starts less month 3 cancels) and the start of month 4 is the number of censored observations in month 4. I.e., for standard price group 1, there were 261 subscription at the start of month 3, 47 month 3 cancellations, and 69 subscriptions that were active to start month 4 but were censored, and 261 – 47 – 69 = 145.

In addition, I analyze the data at a daily level. While the monthly-level specification above most naturally corresponds to effective outcomes (given the 30-day billing cycle), the daily level analysis has two primary advantages. First, it allows a more efficient use of the censored data – while the monthly analysis discards observations from subscribers in mid-billing-cycle as of the cancellation data cutoff in mid-July, the daily analysis includes all available data up until the cutoff day. Secondly, the daily analysis permits a closer inspection of cancellation rates over the course of the billing cycle, and sheds some light on subscriber preferences and their subsequent motivation for canceling. The daily level specification is:

$$log[P_{igmdt}/(1 - P_{igmdt})] = \delta_d + \alpha_t + \beta X_{igmdt} \qquad (2.2)$$

where P_{igmdt} is the probability that subscriber i, in price group g of membership type m, in cycle-day d of month t of membership, cancels on that day; δ_d is a set of constants to capture the baseline time trend in hazard rates within a billing cycle; α_t is a set of constants to capture the baseline time trend in hazard rates across billing cycles; and X_{igmdt} is a set of observed characteristics of subscriber i.

2.5 Results

2.5.1 Monthly Level

Table 2.5 shows the results of specification *(2.1)* for all signups, US signups, and non-US signups during the test period. Below-threshold and switch pricing do not have a statistically significant effect on cancellation rates (columns 1-3). However, camouflage pricing lowers the odds of canceling in a given month by 10.7 percent, or a mean percentage-point drop in the cancellation rate of 2.7 (column 1). For US-signups camouflage pricing reduces the cancellation odds by 12.4 percent, or a mean percentage-point drop of 3.2 (column 2). Interaction terms between the camouflage pricing and monthly dummy variables are statistically insignificant (results not shown), suggesting that the effect is constant over time.

For non-US signups, the camouflage pricing group coefficient is positive but statistically insignificant (column 3). However, this non-US result is consistent with the camouflage pricing hypothesis, since, after currency conversion, all prices show up as camouflage prices on non-US credit card statements.[47] In fact, the camouflage pricing effect among US sign-ups is statistically indistinguishable from the difference in cancellation rates between US signups in standard pricing groups and non-US signups (Table 2.6, Appendix B). The coefficients on a camouflage price group * US-signup interaction term, and a non-US sign-up term, are both statistically significant and close in magnitude (column 1). Furthermore, an alternate model imposing the restriction that the two effects are equal (column 2) cannot be rejected – the chi-square test statistic is 0.15 with one degree of freedom. However, while this evidence is certainly suggestive that the lower cancellation rates among non-US signups is due to effective camouflage pricing from currency conversions, it could also be due to other, unobserved differences between US and non-US members.

[47] This non-result for non-US signups should also alleviate concerns that the effect for US signups is driven by selection.

Table 2.5: Estimates for Logistic Hazard Model of Membership Cancellation, Monthly Level

Explanatory variable	All Members (1)	US Members (2)	Non-US Members (3)
Camouflage Price Group	-0.139*** (0.047)	-0.158*** (0.049)	0.030 (0. 156)
Below-Threshold Price Group	0.001 (0.047)	-0.006 (0.013)	0.048 (0.157)
Switched Price Group	0.065 (0.062)	0.054 (0.065)	0.181 (0.208)
Standard Membership	-0.353*** (0.040)	-0.365*** (0.052)	-0.265* (0.162)
Month 2	-1.27*** (0.05)	-1.24*** (0.05)	-1.52*** (0.18)
Month 3	-1.91*** (0.06)	-1.93*** (0.07)	-1.68*** (0.20)
Month 4	-2.22*** (0.10)	-2.25*** (0.10)	-2.49*** (0.36)
Observations (subscriber-months)	10696	9689	1007
Mean (s.d.) dependent variable	0.427 (0.494)	0.429 (0.495)	0.408 (0.491)
Coefficient on Camouflage Price Group Converted to Odds Ratio	0.869*** (0.041)	0.853*** (0.042)	1.031 (0.161)
Implied Mean Percentage-Point Change in Cancellation Rate from Camouflage Pricing	-3.04*** (1.14)	-3.87*** (1.20)	0.72 (3.76)

Table 2.5 Notes: Standard errors in parentheses. Significantly different than zero at 99 (***), 95 (**), and 90 (*) percent confidence.

To calculate the increased revenue associated with a 3% reduction in the monthly hazard rate, I estimate the change in expected duration per subscriber, using the average rates for subscribers in 2003 as a baseline (see Figure 2.2, Appendix B). While average duration for the first 12 months after signup was 3.2 months for standard members and 2.6 months for premium members in 2003, a 3% reduction in the monthly hazard rate would increase these durations to 3.6 months and 3.0 months, respectively – both increases of roughly 14%. Assuming the camouflage effect only lowers the hazard in the first 4 months after signup (the length of time observed in the field test), the durations would increase to 3.5 months and 2.9 months, or by roughly 9%.

2.5.2 Daily Level

Table 2.6 (Appendix B) shows the results of specification *(2.1)* for all signups, US signups, and non-US signups during the test period. The coefficients on the camouflage price group indicator variable are slightly smaller in magnitude to the monthly level coefficients, though they remain significant at the 90% level and 95% level for all signups and US signups, respectively.

Figure 2.3a shows the estimates and confidence intervals of the billing cycle time-trend coefficients, δ_d from specification *(2.2)*, converted to odds ratios. Cancellation rates are significantly higher in the first 5 and last 5 days of the billing cycle than during the middle 20 days.

Figure 2.3a: Odds Ratio Coefficients on Cycle-Day Indicators from Daily Analysis Logistic Regression

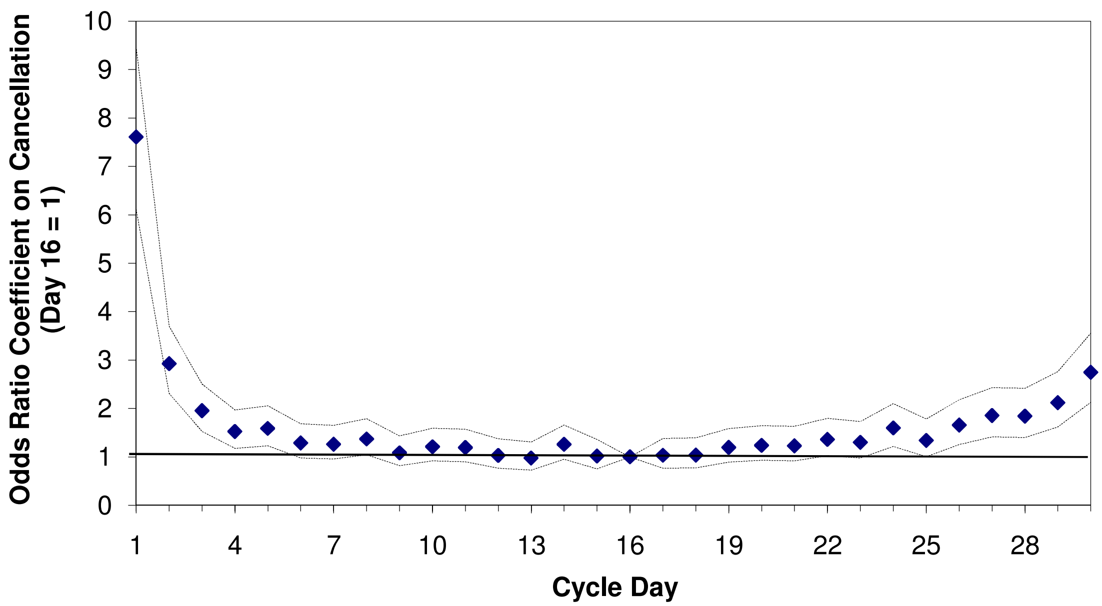

The spike in cancellations at the end of the cycle may be explained by a number of factors. The most obvious explanation is that subscribers defer the cancellation decision until late in the cycle to retain the option value of renewal while learning the value of the subscription. However, even if subscribers know at the beginning of a cycle that their preference is to cancel the subscription, they may still defer the actual task of canceling until later in the cycle. For example, if subscribers are present-biased, they may procrastinate to defer incurring the "task cost" associated with canceling until the latest possible date. Alternatively, if this task cost is stochastically determined, even time-consistent agents will, all other things equal, be more likely to cancel later in the cycle. The logic is that they will only cancel early in the cycle when it is extremely easy/cheap to do so, since they know they will have many more chances to draw a low task cost; late in the cycle, when fewer draws remain, subscribers are more likely to cancel even for high-cost draws (i.e., even on days when they are busy or where getting to a computer is inconvenient).

Importantly, however, none of these explanations predict a cancellation spike at the beginning of the billing cycle. Assuming subscriber forgetfulness, on the other hand, would generate this early cancellation spike in the form of a reminder effect. Suppose that subscribers who stop using the site, with some probability, forget about their subscriptions altogether. Seeing the new charge for the ensuing month reminds these forgetters of their subscriptions, and may prompt them to cancel immediately. While they may still have some incentive to defer cancellation for the reasons listed above, these subscribers may simply opt to cancel the subscription before they forget again.

There is an alternative explanation for early cancellations in the billing cycle. It may be that some subscribers sign up to use a particular feature or access particular content on the site, and cancel immediately thereafter. If subscribers' have this "one-time access" motivation for signing up, it may be least costly for them to cancel immediately after this use, since they will already be on the site. Note, however, that this explanation likely applies to only the first month. Figure 2.3b (Appendix B) shows the odds ratios estimates and confidence intervals of the billing cycle time-trend coefficients, δ_d from specification *(2.2)*, when also including interaction terms between the cycle-day indicators and the first-billing-cycle indicator (thus isolating the cancellation pattern over the billing cycle *after* the first month of subscription). While the early-cycle spike is slightly dampened, it is still large in magnitude and statistically significant. Interestingly, the late-cycle spike is substantially curtailed after the first month. This may be interpreted as evidence in favor of the option-value explanation for the late-cycle spike in cancellations. If subscribers have more or less learned the value of the subscription to them within the first 30 days of subscribing, there would be a smaller option-value incentive to defer canceling until the end of ensuing business cycles.

Thus, the pattern of daily hazard rates over the billing cycle – in particular, the spike in cancellations in the first 5 days – provides evidence consistent with the proposed explanation of the camouflage price effect. It suggests that the "reminder" subscribers receive in the form of the charge on their credit card bill does indeed meaningfully impact the probability of canceling. To the extent that firms can manipulate (i.e., reduce) the strength of this reminder in the form of camouflage pricing, they may be able to improve retention.

This explanation implicitly assumes that subscribers are more likely to see the renewal charge in the first 5 days of the new billing cycle. However, this is certainly not an obvious assumption. For instance, if subscribers' all review their credit card charges on through online statements, and if subscribers' credit card billing cycles are independent of their subscription billing cycles – a reasonable assumption – then the probability of actually seeing a credit card statement with the subscription charge would be constant over the course of the billing cycle. One possible explanation is that some subscribers monitor their credit card statements online at regular intervals. If this is the case, then at least this portion of subscribers will be more likely to see the new charge early in the billing cycle. While this may not be a reasonable assumption to make about the credit-card monitoring behavior of the population at-large, it is more plausible with a sample of individuals who purchase online subscriptions.

2.5.3 Selection at Takeup

A primary concern given the experimental design is potential selection problems. Non-random selection could arise in a number of ways. First, prospective members could become aware of the randomly assigned price points, and repeatedly click through to the membership sign-up page until they receive the lowest offer. Secondly, even if offer page views are random, take-up may still be nonrandom. For instance, it is possible that certain types of individuals would be suspicious of a camouflage charge, and thus would not sign-up, whereas the same individual, if allocated to a group with a standard charge, would have signed up.[48]

Table 2.8 shows the page views, signups, and take-up rates for each price group in the test. A chi-square test cannot reject the null hypothesis that the page views for each membership type were allocated randomly across groups.[49] Furthermore, neither below-threshold nor camouflage pricing has a statistically significant impact on takeup rates (results not shown).

Table 2.8: Offer Page Views, Signups, and Takeup Rates by Group

Price Group	Offer	Actual	Views	Signups	Takeup (x100)
Standard					
1	$19.83	$19.83	525,536	877	0.1669
2	$19.95	$19.95	526,378	854	0.1622
3	$20.00	$20.00	527,140	974	0.1848
4	$20.00	$19.83	526,576	866	0.1645
5	$20.13	$20.13	527,876	818	0.1550
Premium					
1	$29.78	$29.67	230,757	294	0.1274
2	$29.78	$29.78	229,485	295	0.1286
3	$29.95	$29.95	229,342	335	0.1461
4	$30.00	$30.00	230,053	277	0.1204
5	$30.14	$30.14	230,248	312	0.1355

2.6 Conclusion

This field experiment was designed to test if a firm, through manipulating consumer forgetting, could meaningfully alter behavior (and effective demand) to their benefit. The result that the camouflage pricing treatment is associated with a 3% reduction in the monthly hazard rate

[48] Of course, "sophisticated" consumers may not merely be suspicious of camouflage prices, but rather may be aware that such a charge may inefficiently decrease their probability of canceling at the appropriate time.

[49] More specifically, testing H_0 that p_i, the probability that an offer clickthrough directs to price group i, is equal to 0.2 for all price groups (i = 1…5) for both the premium and standard memberships.

suggests that this pricing tactic is one such manipulation. Of course, the design of this test itself is evidence that firms are aware of consumers' fallible memory and the potential to manipulate their attentiveness, and that they strategize on how to best exploit it.

A natural question arising from this result concerns the equilibrium effects of a broader adoption of this pricing tactic. In this test, camouflage pricing was implemented by just one website, for a limited period of time. If other subscription-based firms became convinced of its effectiveness and implemented similar pricing tactics, would consumers recognize the costs of their inattention and adjust their behavior accordingly? While in Chapter 1, we assumed that the probability of recalling a consumer task is exogenous, there may very well be an endogenous component to recall. In particular, in a multi-task environment, individuals surely are more likely to recall more important tasks, perhaps by implementing reminders to selectively improve recall probability. If consumers recognized the increased cost of overlooking camouflage prices as more firms implemented them, perhaps they would respond by reading their credit card bills more carefully, or by setting up monthly reminders to evaluate the value of continuing the subscription. Ultimately, the evidence from this test cannot speak definitively on the equilibrium effects. However, the similar magnitude of the camouflage pricing effect among non-US members – who, because of currency conversion, are operating in an environment where all US-dollar-denominated subscription products are effectively camouflage-priced – may suggest that consumers are slow to adjust.

Another important issue is whether consumer inattention in cases such as this field test is simply an example of bounded rationality, or if it is influenced by some bias. As discussed in Chapter 1, it may be that consumers are not just forgetful, but also overconfident in their future memory recall – a bias which would lead to even lower task completion rates than would result from limited memory or attention alone. Unfortunately, this evidence on camouflage pricing cannot discern between boundedly-rational and overconfident forgetting. Given the stakes involved with these memberships, and assuming that individuals have a limited capacity to recall tasks, it may, in fact, be inefficient to recall the membership and the option to cancel with probability one each period. The camouflage pricing result merely suggests that the recall probability is less than one, and that firms, through camouflage pricing, may be able to further reduce this probability. One potential extension of this study would be an empirical design that identifies the effects of such attention manipulation by firms on not just behavior, but welfare as well.

Chapter 3: Do Probability Weights in the Field Match Predictions from the Lab? Evidence from Online Trading Card Auctions

3.1 Introduction

While the psychology and economics literature features extensive work on experimental calibrations of risk models,[50] there is an increasing emphasis on measuring its risk parameters in the field.[51] Recent field studies – ranging in context from key household decisions such as purchasing insurance (Barseghyan, Molinari, O'Donoghue, and Teitelbaum 2010) to risk-seeking activities such as racetrack betting (Snowberg and Wolfers 2010) – have found empirical evidence consistent with prospect theory, and nonlinear probability weighting in particular. Figure 3.1 illustrates a characteristic prospect theory probability weighting function (Kahneman and Tversky 1979), wherein small probabilities are overweighted and medium and large probabilities are underweighted, capturing both tastes (e.g., the certainty effect) and cognitive errors (e.g., the difficulty of distinguishing between small probabilities).

Figure 3.1: Characteristic Probability Weighting Function Estimated in the Laboratory

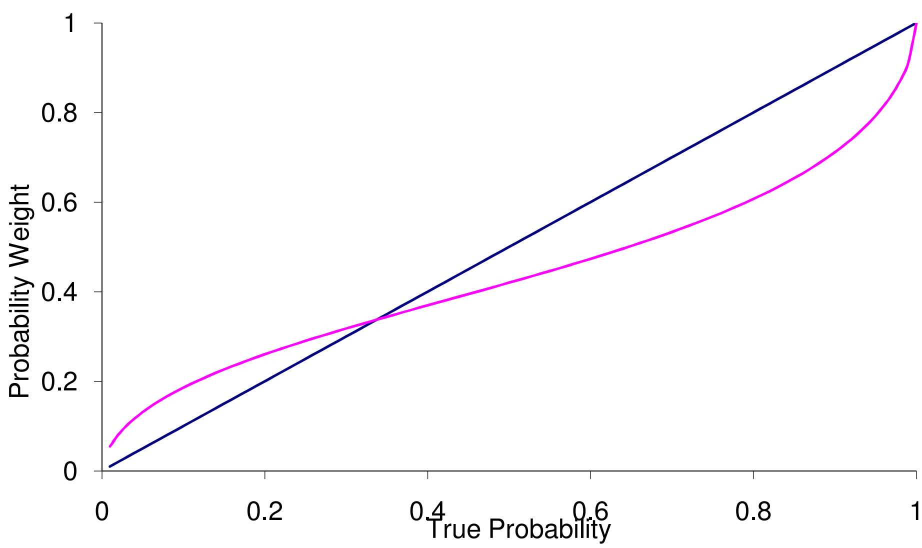

Figure 3.1 Notes: The most common parameterizations for the function have been a one-parameter version, where $\pi(p) = p^{\gamma}/[p^{\gamma} + (1\text{-}p)^{\gamma}]^{1/\gamma}$ (Tversky and Kahneman 1992), and a two-parameter form $\pi(p) = \delta p^{\gamma}/[\delta p^{\gamma} + (1\text{-}p)^{\gamma}]^{1/\gamma}$ first proposed by Lattimore, Baker and Witte (1992). This figure is generated using the one-parameter form with $\gamma = 0.61$ as estimated over gains by Tversky and Kahneman (1992).

---

[50] For instance, see prospect theory calibrations in Tversky and Kahneman (1992), Abdellaoui (2000), Abdellaoui, Vossmann, and Weber (2004), Bleichrodt and Pinto (2000), Camerer and Ho (1994), Tversky and Fox (1995), and Gonzalez and Wu (1996) and (1999).

[51] The applicability of experimental results to "real-world" decision-making under risk has long been a subject of debate in the field. Over 50 years ago, Edwards (1954) commented on some early experimental results suggesting that people distorted probabilities, that "the results of such pilot experiments too often are picked up and written into the literature without adequate warning about the conditions under which they are performed and the consequent limitations on the significance of the results."

This study adds to that field evidence by exploiting a "natural experiment" from the trading card industry in which a particular card experienced a sudden jump in value. The card was a highly sought-after collectible of Tom Brady, the star quarterback of the National Football League's New England Patriots, and the price jump coincided with the Patriots winning the Super Bowl in February 2004. The Brady card was randomly inserted into sealed boxes at the rate of 1 per 60 boxes.[52] Collectors buying these randomized, unopened boxes were essentially buying a fairly complex gamble that included, among other things, the possibility of receiving the Brady card. Thus, in abstract terms, the experiment measured how the market price for a lottery ticket (i.e., an unopened box) with fixed, known odds of winning adjusted when the value of the "grand prize" (the Brady card) increased. I track auction sales of the particular card and sealed boxes for 10 days before and 10 days after the sudden jump, and provide evidence that other card prices in the set were not fluctuating over the same time frame, so that changes in the box price must have been driven only by changes in the price of the one particular card. I find that, while the probability of obtaining the Brady card was 1 in 60 (or 0.017), the probability weight (or "decision weight") that reconciles the price movements is roughly 0.27. The full probability weight range predicted by experimental calibrations (0.03 – 0.10) can be rejected with 99% confidence.

While acknowledging the identification strategy's limitation of measuring just one point on the probability weighting function, the result is an example of consumer behavior in the field diverging from experimental estimates of risk preferences. I discuss two potential interpretations for the discrepancy. The first is that the salience the Tom Brady rookie card – particularly in light of the heavy media coverage Brady received around and after the Super Bowl – may have led collectors to overly weight the "jackpot" outcome of pulling the Brady card in pricing the sealed box gamble. The notion that the salience of outcomes affects choice under risk has been discussed in the literature, both in terms of its psychological underpinnings (Loewenstein, Weber, Hsee, and Welch 2001) and as a theoretical construct (Bordalo, Gennaioli, and Shleifer 2010). This salience explanation is also linked to the empirical literature on how selective consumer attention impacts market outcomes (Chetty, Looney, and Kroft 2009, Finkelstein 2009, Hossain and Morgan 2006, Lee and Malmendier forthcoming, Lacetera, Pope, and Sydnor 2010, and Chapters 1 and 2 of this dissertation). That consumers may overly value risky prospects by having their attention directed towards particularly salient, high-payoff outcomes is a potential extension of the finding that the salience of add-on costs or product attributes affects consumer demand for goods.

A second interpretation of the result is that trading card collecting, and box-buying in particular, is a gambling and thrill-seeking activity, and that estimated preferences from such activities cannot be broadly generalized to all choice under risk. While the use of price changes rather than the static box and card prices potentially "differences out" gambling utility, the disparity between the estimated decision weight and previous experimental estimates may reflect that collectors' behavior is driven by factors other than their general risk preferences – such as neurological stimuli that are associated with compulsion and addiction. Understanding what motivates gambling is certainly an important economics question in itself – Barberis (2010) cites the

[52] The random insertion of rare cards such as this particular one is not an uncommon practice in the trading card industry, as I discuss in Section 3.2.

American Gaming Association's estimate that gambling revenues exceeded $85 billion in 2007 – but ultimately, it may be a separate issue from calibrating the risk preferences that more generally govern economic behavior.

3.2. Background

3.2.1 Overview of The Trading Card Industry

Trading cards are a $1.5 billion-a-year industry in the United States, with sports trading cards accounting for approximately one-fifth of that total.[53] A considerable portion of industry sales is generated by the secondary market, which includes trading card stores, conventions, and Internet sales. Online sports trading card auction sales have been estimated at over $1 billion annually, with the vast majority on eBay.[54] Trading card sets have been produced for numerous sports (including basketball, football, soccer, ice hockey, tennis, golf, and even wrestling) and on a global scale (trading cards have been produced for sports leagues based in North America, Europe, Asia, and South America). Non-sports trading cards, including sets produced for strategy games (e.g., Magic the Gathering, Pokemon, Yu-Gi-Oh) and for TV shows and movies (e.g., Lord of the Rings) account for another $5 billion in annual sales internationally.[55]

Trading cards are generally sold on a pack or box basis. Each pack contains a "random" selection of 4-15 cards (depending on the particular set), and boxes usually contain 12-36 packs. For many years, the random selection of cards found in any given pack was out of a full set of 300-700 cards, and each card had the identical production run (usually in the tens or hundreds of thousands). A relatively recent phenomenon in the industry is the production of rare, or short-printed "insert" or "chase" cards. This practice began in the early 1990s as industry revenues declined, competition intensified (with the entry of a number of new producers in the 1980's), and trading card companies searched for new ways to differentiate their products. Initially, chase cards were seeded at ratios of one every several packs. But as card companies continued to attempt to "one-up" each other and develop the most appealing products, the quality of chase cards was improved and insertion odds grew longer. Early chase cards were often distinguished only by an alternate design or special foil stamping, but more recent issues include player autographs, serial numbering, or small pieces of game-used equipment or uniforms embedded into the card. In fact, many sets produced in the past year contained inserts with print runs of just a single card. These "1 of 1" editions often sell for thousands (if not tens of thousands) of dollars in online auctions.

Whereas pack prices prior to the development of chase cards never rose above $1, today's high-end products often carry pack prices of over $100. In 2004, The New York Times ran a feature

[53] See MacMillan 2007. Estimates of industry size are complicated by the difficulty of tracking and assessing secondary market activity, and because the three major sports card manufacturers – Topps, Upper Deck, and Panini America – are privately owned. Prior to being privatized in 2007, Topps reported trading card sales of $294 million in 2006.
[54] At any given point, there are roughly 300,000-500,000 active sports trading card auctions on eBay, with current bids ranging from 1 cent to several thousand dollars. There are an additional 4 million lots for sale at fixed prices.
[55] "It's All in the Yu-Gi-Oh Cards," Northeast Times, March 19, 2003.

on the industry's first $500-per-pack product, Upper Deck Exquisite.[56] On July 22, 2004, a card from this product sold for $150,100 on eBay. While the pre-chase-card era was based on set collecting, contemporary collectors are often motivated to buy packs or boxes by the hopes of "hitting the jackpot" and pulling a card worth thousands of dollars. Indeed, it is not unusual to see individuals opening packs or boxes at trading card stores, and throwing most (if not all) of the non-chase cards they receive away.

3.2.2 Description of The "Experiment"

On February 1, 2004, the National Football League's New England Patriots defeated the Carolina Panthers to capture their second Super Bowl championship in three years. Cards of the Patriots' quarterback, Tom Brady, instantly became the most sought-after items in the trading card hobby.

Brady's most popular card – and the one generating the most market activity during the Patriots' Super Bowl run – is part of the 2000 SP Authentic football set. The product was released in the December 2000 by the Upper Deck Company, a privately-owned card producer based in Carlsbad, California. Boxes containing a random assortment of 120 cards (5 cards per pack, 24 packs per box) were issued with an initial suggested retail price of $119.76 ($4.99 per pack).

On average, each 2000 SP Authentic football box yields 110 base cards, 8 inserts, 1 of 60 rookie cards (each serially numbered out of 1250), and 1 autograph card. Base cards range in book value from $0.30-$2.50; insert cards from $0.60-$6; and autograph cards from $10-$80.[57] Prior to the 2004 Super Bowl, the 60 rookie cards ranged in book value from $10 to $200, with Brady's being the most valuable.

3.3. Data Description

eBay sales of 2000 SP Authentic boxes, and all 2000 SP Authentic rookie cards, were tracked for the 10 days before and after the Super Bowl (i.e., from January 23-February 11). In all there were 78 box sales and 111 (including 46 Tom Brady) rookie card sales. Prices for base cards, inserts, and autograph cards are assumed to be stable before and after the Super Bowl and thus were not tracked.[58]

---

[56] "Who'll Trade Two for a Freddy Adu?," New York Times, July 25, 2004.

[57] Two special versions of autograph cards were also randomly inserted into boxes – "buy backs" with print runs of 3 to 620 were inserted into one of every 3 boxes, and "gold" versions with print runs ranging from 2 to 92 were inserted into one of every 20 boxes. These rarer autograph cards generally sell at a premium of 2 to 3 times above the regular autograph cards. Trading cards generally sell for 25-75% of book value in online auctions. Book values are compiled from reported transactions at trading card shops and shows across the country.

[58] This is a plausible assumption for a number of reasons. First, other than an increase in the book value of the Tom Brady card from $200 to $300, there were no changes in the book value of any base cards, inserts, or autographs in the Beckett Football Collector monthly price guides issued immediately before (February 2004) and after (March 2004) the Super Bowl. Beckett Football Collector is widely considered the industry's most reliable price guide for football cards. Second, none of the Patriots' other star players had valuable or sought-after cards in the 2000 SP Authentic set. Third, rookie card prices are generally more volatile than other card prices – for most base cards and inserts (including autographs), book values stabilize within a few months of their release, whereas rookie cards fluctuate considerably based on recent performance.

Figure 3.2 shows the eBay realized sale prices of the 2000 SP Authentic Tom Brady rookie card during the studied period. The time frame is broken up into 3 periods – pre-Super Bowl (January 22, 2004 until 8 PM Pacific Standard Time on February 1), the night of the Super Bowl (February 1, 8 PM PST-midnight PST), and post-Super Bowl (February 2-11). There were 46 total sales of the card during this time frame (22 pre-SB, 4 night of, 20 post-SB). The average sale price increased from $174.07 pre-SB to $244.74 post-SB, a change of over $70.

Figure 3.2: eBay Sales of 2000 SP Authentic Tom Brady Rookie Card, 1/23/04-2/11/04

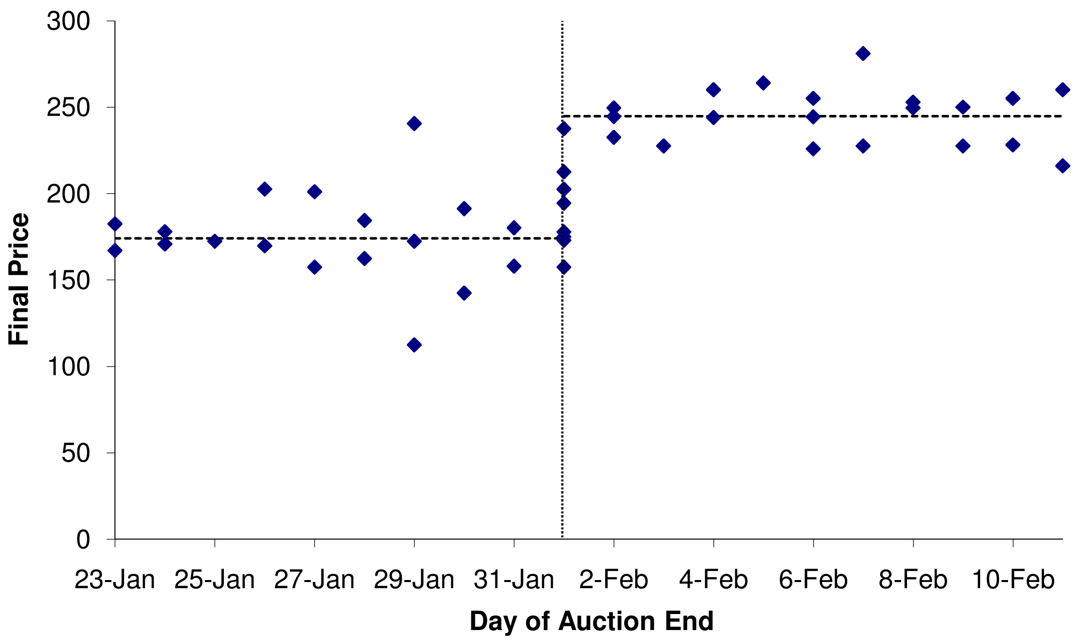

Figure 3.2 Notes: Sales of ungraded version of card only (46 total). The trendlines represent the average final price pre-Super Bowl ($174.07) and post-Super Bowl ($244.74).

Figure 3.3 (see Appendix C) shows the time pattern of 2000 SP Authentic box sales during the same period. There were 78 sales in all – 35 pre-SB, 6 night of, and 37 post-SB. The average sale price for boxes rose from $84.85 to $109.14 from the pre-SB period to the post-SB period, or an increase of nearly $25.

Figure 3.4 (Appendix C) shows the sales of all 2000 SP Authentic rookie cards other than Tom Brady in the period. There were 65 sales, 38 pre-SB and 27 post-SB. The average realized sale price/book value ratio in these auctions was nearly identical in the pre-SB (0.42) and post-SB periods (0.43), supporting the claim that the increase in box prices was generated solely by the surge in the value of the Tom Brady card.

3.4 Identification Strategy

I employ a basic OLS regression framework to get more precise estimates of the pre-SB to post-SB price changes for both the Brady rookie card and unopened boxes. The regression specification is as follows:

$$EndBid_i = \alpha + \gamma * NightOf_i + \lambda * PostSB_i + X_i'\beta, + \varepsilon_i \quad (3.1)$$

where $NightOf_i$ and $PostSB_i$ are indicator variables for auction *i* ending in those periods, and X_i is a vector of auction-level characteristics that previous studies have found to be significant determinants of final auction price.[59]

I then use the regression results to back-out the probability weight that would reconcile the Tom Brady rookie card and unopened box price movements. Importantly, I assume a linear value function in the gains domain – that is, I assume that the box price increases by exactly the "modified expected resale value" change of the box, or the expected resale value change based on probability weights rather than true probabilities. Note, this assumption means that the resulting probability weight estimate represents a lower bound for the true, market-level probability weight. Alternatively, assuming diminishing returns of the value function in the gains domain would imply that individuals would raise their bids on unopened boxes by less than the "modified expected resale value" change of the box. Thus, the probability weight assigned to receiving the Tom Brady rookie card would have to be even greater than the linear-value-function estimate to justify the increase in the box price.[60]

Assuming risk-neutrality in gains, and maintaining the claim that all other card prices in the set were stable during this period, it must be that:

$$\Delta \text{ Box Price} = \pi(\text{True probability of pulling Tom Brady card}) * \Delta \text{ Tom Brady Card Price}, \quad (3.2)$$

or

$$\lambda_{BOX} = \pi(1/60) * \lambda_{BRADY} \Rightarrow \hat{\pi}(1/60) = \hat{\lambda}_{BOX} / \hat{\lambda}_{BRADY}, \quad (3.3)$$

where $\pi(.)$ is the probability weighting function and $\hat{\lambda}_{BRADY}$ and $\hat{\lambda}_{BOX}$ are the estimates of the post-SB coefficients for Brady card sales and unopened box sales, respectively, from the equation *(3.1)* regressions.

---

[59] These characteristics include the timing of the auction end (time of day and day of week), seller feedback, length of auction (in days), "effective reserve" (starting bid + shipping and handling fee), and shipping and handling fee. Papers that have investigated the impact of these characteristics include Kauffman and Wood (2004), Lucking-Reiley, Bryan, Prasad, and Reeves (2000), and Hossein and Morgan (2006).

[60] Abdellaoui, Vossmann, and Weber (2004), in a similar vein, find a less skewed probability weighting function when they assume a linear value function than when they assume a power value function (i.e., $v(x) = x^{\alpha}$)

While this basic procedure results in a point estimate for the probability weight implied by the Tom Brady card and unopened box price movements, it may also be of interest to calculate a confidence interval for the estimate, and test the null hypotheses of the estimated probability weight equaling 1) the true probability or 2) the weight predicted by experimental calibrations (Table 3.1 summarizes estimates of *π(1/60)* from previous experimental studies, which range from 0.03 to 0.10). By the Delta Method:

$$Var(\hat{\lambda}_{BOX} / \hat{\lambda}_{BRADY}) \approx D'VD \qquad (3.4)$$

where

$$D' = \left(\partial(\hat{\lambda}_{BOX} / \hat{\lambda}_{BRADY}) / \partial(\hat{\lambda}_{BOX}) \qquad \partial(\hat{\lambda}_{BOX} / \hat{\lambda}_{BRADY}) / \partial(\hat{\lambda}_{BRADY})\right)$$
$$= \left(-\hat{\lambda}_{BOX} / (\hat{\lambda}_{BRADY})^2 \qquad 1/\hat{\lambda}_{BRADY}\right), \text{ and}$$
$$V = \begin{pmatrix} Var(\hat{\lambda}_{BRADY}) & Cov(\hat{\lambda}_{BRADY}, \hat{\lambda}_{BOX}) \\ Cov(\hat{\lambda}_{BRADY}, \hat{\lambda}_{BOX}) & Var(\hat{\lambda}_{BOX}) \end{pmatrix}$$

To use the Delta Method in this specification requires some assumption regarding $Cov(\hat{\lambda}_{BRADY}, \hat{\lambda}_{BOX})$. Given that the two post-Super Bowl price-jumps are estimated from different datasets (corresponding to the auction sales of the respective items), an estimate of the covariance is not directly calculable. One option is to simply assume a 0 covariance, but this is clearly problematic given the nature of the data.

Table 3.1: Estimates of *π(1/60)* from Previous Experimental Studies

Study	γ	δ	*π(1/60)*
One-Parameter			
Tversky and Kahneman (1992)	0.61		0.073
Camerer and Ho (1994)	0.56		0.086
Wu and Gonzalez (1996)	0.71		0.051
Two-Parameter	0.60		
Tversky and Fox (1995)	0.69	0.77	0.043
Gonzalez and Wu (1999)	0.44	0.77	0.098
Abdellaoui (2000)	0.60	0.65	0.052
Bleichrodt and Pinto (2000)	0.55	0.81	0.074
Abdellaoui, Vossmann, and Weber (2004)	0.83	0.97	0.032

Table 3.1 Notes: In studies where the probability weighting function was estimated separately for gains and losses, the figure for gains is used.

An alternative is a two-stage procedure to transform the two datasets into equivalent units of observation and directly calculate the covariance. In particular, the datasets can be collapsed to daily observations, in the following manner:

1) Re-run the equation *(3.1)* regressions, now including an unrestricted set of 20 dummies for each day in the studied period. Normalizing for the morning of the Super Bowl (Feb. 1), the 20 dummies will be for Jan 23-31, the night of Feb 1, and Feb 2-11. Label the coefficients on these as $d_{BOX,t}$ and $d_{BRADY,t}$, with $t = 1\ldots21$ and $d_{BOX,10} = d_{BRADY,10} = 0$;
2) For both sets of coefficients, regress:

$$d_t = \alpha' + \gamma'^* I_{(t=11)} + \lambda'^* I_{(t>11)} + e_t. \qquad (3.5)$$

Now $Cov(\hat{\lambda}'_{BRADY}, \hat{\lambda}'_{BOX}) = Var(\hat{\lambda}'_{BRADY}) * Cov(e_{BRADY}, e_{BOX}) / Var(e_{BRADY})$, and the confidence interval for the probability weight estimate can be constructed accordingly.

3.5 Results

3.5.1 First-Stage

The first-stage regression results are shown in Tables 3.2 and 3.3. The standard errors are corrected for heteroskedasticity and serial correlation using the Newey-West estimator with 3 lags.[61] The post-Super Bowl price jumps for the Tom Brady card and unopened boxes are roughly \$70 and \$19, respectively. Other than the indicator variables for night of SB and post-SB, none of the other independent variables have a significant effect (at 95% confidence) on the final price of the Tom Brady rookie card auctions (Table 3.2, regressions 2-5). For unopened boxes, having "Tom Brady" in the searchable auction title increases realized prices by roughly \$30, and sealed boxes command a \$12 premium over unsealed boxes (Table 3.3, regression 1). However, these two factors do not have a differential impact in the post-SB period versus the pre-SB period (Table 3.3, regression 2). No other auction characteristics are significant at the 95% level (Table 3.3, regressions 3-6) for unopened boxes.

3.5.2 Second-Stage

The new independent variable for the second-stage – the coefficients on the 20 daily dummies – are taken from Table 3.2, regression 6 (for Brady card sales) and Table 3.3, regression 7 (for unopened box sales). Figure 3.5 (Appendix C) shows a plot of both sets of coefficient values. The second-stage regressions yield similar post-Super Bowl price jump estimates to those from the first-stage (Table 3.4 in Appendix C, regressions 1 and 2). Again, the standard errors are calculated using the Newey-West estimator with 3 lags.[62]

---

[61] Breusch-Pagan and Durbin-Watson tests on both sets of results provide evidence of heteroskedasticity and negative autocorrelation (results not shown).
[62] While the first-stage data with observations on the auction level is negatively autocorrelated, Durbin-Watson tests of the daily-level data reject the null hypothesis of autocorrelation. Breusch-Pagan tests also suggest that heteroskedasticity in the daily-level data is less severe (results not shown).

Table 3.2: Determinants of Final Price in 2000 SP Authentic Tom Brady Rookie Card Auctions on eBay, 1/22/04-2/11/04

	Dependent variable: Final Auction Price					
Explanatory variable	(1)	(2)	(3)	(4)	(5)	(6)
Auction Ended in Post-SB Period	70.6*** (3.6)	72.5*** (4.8)	69.3*** (4.9)	72.5*** (3.4)	70.2*** (3.8)	71.3*** (4.8)
Auction Ended Night of SB	37.6*** (7.9)	36.3*** (8.3)	35.5*** (9.3)	35.7*** (8.1)	38.6*** (7.5)	40.9*** (11.7)
Auction Ended Thursday or Friday		-0.4 (5.5)				
Auction Ended Saturday or Sunday		0.7 (4.7)				
Auction Ended between noon and 6 PM PST		0.0 (8.9)				
Auction Ended between 6 PM and midnight PST		4.2 (9.5)				
Seller Feedback > 100			10.4 (11.0)			
Seller Feedback > 1000			-0.9 (7.5)			
5-Day Listing			-0.8 (5.9)			
3-Day Listing			-3.5 (9.8)			
1-Day Listing			6.5 (7.1)			
Effective Reserve (Starting Bid + Shipping Cost) > $20				-6.9 (4.4)		
Shipping Cost					-1.0 (1.1)	
Daily Fixed Effects	No	No	No	No	No	Yes
R^2	0.74	0.74	0.75	0.74	0.74	0.76
Mean (s.d.) dependent variable	208.0 (39.6)	208.0 (39.6)	208.0 (39.6)	208.0 (39.6)	208.0 (39.6)	208.0 (39.6)
95% confidence interval on Post-SB Indicator	[63.2, 78.0]	[62.8,82.2]	[59.3,79.2]	[65.7, 79.3]	[62.6, 77.8]	[61.7, 80.9]

Notes for Table 3.2: All regressions are OLS with N = 46. Heteroskedasticity and autocorrelation consistent standard errors are calculated using the Newey-West estimator with 3 lags (in parentheses). Significantly different than zero at 99 (***), 95 (**), and 90 (*) percent confidence.

Table 3.3 Determinants of Final Price in 2000 SP Authentic Unopened Box Auctions on eBay, 1/22/04-2/11/04

	Dependent variable: Final Auction Price						
Explanatory variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Auction Ended in Post-SB Period	19.3*** (2.6)	27.0*** (7.4)	21.5*** (3.2)	18.8*** (3.1)	19.1*** (2.4)	19.2*** (2.4)	21.1*** (6.5)
Auction Ended Night of SB	6.6*** (2.0)	5.9** (2.3)	10.6** (3.2)	7.3** (3.0)	6.7** (2.7)	6.5*** (2.1)	1.9 (3.8)
"Tom Brady" in Item Title	30.2*** (5.9)	28.5*** (3.4)	31.8*** (6.3)	30.1*** (7.8)	30.0*** (5.4)	30.2*** (5.1)	28.3*** (7.1)
Sealed Box	12.2*** (3.9)	17.5*** (4.2)	10.5** (4.1)	11.3** (5.1)	11.9** (4.1)	11.9** (4.2)	13.5*** (4.5)
Post-SB * "Tom Brady" in Item Title		2.8 (7.3)					
Post-SB * Sealed Box		-9.4 (7.0)					
Auction Ended Thursday or Friday			0.1 (4.1)				
Auction Ended Saturday or Sunday			-0.2 (4.3)				
Auction Ended between 6 AM and noon PST			8.5 (10.1)				
Auction Ended between noon and 6 PM PST			12.9 (10.5)				
Auction Ended between 6 PM and midnight PST			5.4 (10.2)				
5-Day Listing				6.9* (3.9)			
3-Day Listing				-1.7 (2.9)			
1-Day Listing				-2.4 (3.7)			
Effective Reserve (Starting Bid + Shipping Cost) > $10					0.5 (2.7)		
Shipping Cost						-0.1 (0.8)	
Daily Fixed Effects	No	No	No	No	No	No	Yes
Seller Feedback Effects	No	No	No	Yes	No	No	No
R^2	0.73	0.73	0.75	0.75	0.73	0.73	0.81
Mean (s.d.) dependent variable	96.9 (19.5)	96.9 (19.5)	96.9 (19.5)	96.9 (19.5)	96.9 (19.5)	96.9 (19.5)	96.9 (19.5)
95% confidence interval on Post-SB Indicator	[14.1, 24.4]	[12.2, 41.8]	[15.1, 27.9]	[12.6, 25.0]	[14.3, 23.9]	[14.4, 24.0]	[8.1, 34.1]

Notes for Table 3.3: All regressions are OLS with N = 78. Heteroskedasticity and autocorrelation consistent standard errors are calculated using the Newey-West estimator with 3 lags (in parentheses). Significantly different than zero at 99 (***), 95 (**), and 90 (*) percent confidence.

3.5.3. Probability Weight Point-Estimate and Confidence Interval

Table 3.5 (Appendix C) summarizes the point-estimate and confidence interval calculations using the one-stage and two-stage methodologies described in Section 3.4. Both methodologies yield similar results – the one-stage point estimate is .27 with a 95% confidence interval of [.19, .37] and the two-stage point estimate is .28 with a 95% confidence interval of [.19, .35]. The null hypothesis that the implied probability weight equals the true probability can be rejected with 99% confidence, as can the null hypothesis that the implied probability weight is equal to anything in the range predicted by the experimental calibrations in Table 3.1 (that is, a probability weight of between .03 and 10).

3.6. Conclusion

3.6.1 Discussion of Results

Salience and Attention

Taylor and Thompson (1982) describe salience as "the phenomenon that when one's attention is differentially directed to one portion of the environment rather than to others, the information contained in that portion will receive disproportionate weighting in subsequent judgments." To the extent that Tom Brady and his collectibles were particularly salient to collectors around the time of the Patriots' 2004 Super Bowl win, sealed-box purchasers' "disproportionate weighting" of the Tom Brady card outcome may provide a potential explanation for the result. This explanation is consistent with the "risk-as-feelings" hypothesis (Loewenstein, Weber, Hsee, and Welch 2001) that behavior in the face of risky prospects is influenced not only by an agent's cognitive evaluation of outcomes and probabilities, but also by affective factors, including the "vividness" of outcomes.[63] As examples, LWHW cite the Johnson, Hershey, Meszaros, and Kunreuther (1993) result that people were willing to pay more for airline travel insurance covering death from "terrorist acts" than death from "all possible causes,"[64] and the Browne and Hoyt (2000) finding that knowing someone who has been in a flood or earthquake, or being in one oneself, greatly increases the likelihood of purchasing insurance (even after controlling for subjective expectations). That Tom Brady's collectibles were top of mind for collectors during the month of the Super Bowl may have made the possibility of pulling his rookie card particularly "vivid," and influenced valuations of the sealed-box gamble accordingly. In contrast, the best-case outcomes in experimental settings are generally small monetary payoffs, none of which elicit particularly strong mental imagery or emotional reactions from subjects[65] – this may

[63] Other affective factors discussed by the authors include the immediacy of the risk and background mood. The authors argue that behavior that appears to manifest cognitive errors in probability assessments may not be cognitive errors at all, but rather, be driven by the parallel, and often competing, influence of emotions incited in a separate part of the brain from the area believed to be responsible for "rational" thought.

[64] LWHW note that while the latter form of coverage clearly subsumes the former, death from terrorist acts is a "highly imaginable event," while the more general wording of universal coverage "does not spontaneously bring fear-provoking images to mind."

[65] Interestingly, this does not suggest that the effect of "vividness" is always muted in lab settings. For instance, Rottenstriech and Hsee (1999) find that subjects exhibit systematically different sensitivities to probability variations when outcomes are more emotional-visceral (e.g. painful electric shocks) than when they are relatively pallid (e.g., small cash payoffs).

lead subjects to rely more on their cognitive evaluation of probabilities rather than the vividness of outcomes. Thus, the gap between this field result and previous experimental calibrations may be explained by the relative vividness of favorable outcomes – the possibility of pulling a Tom Brady card from a pack may have evoked stronger mental imagery for a card collector than the prospect of a monetarily-equivalent $200 payoff would for an experimental subject.

This "vividness" hypothesis is consistent with the substantial premium received by sealed box auctions that included "Tom Brady" in their searchable eBay item description – over $30, or roughly 30% above the average market price (importantly from an identification standpoint, there was no significant difference in this premium before and after the Super Bowl).[66] In a related result, Lee and Malmendier (forthcoming) find that eBay auction prices exceed fixed prices for the same item (displayed on the same webpage as the auction listings) in 40-50 percent of observed cases. The authors' leading explanation for this finding is bidders' "limited attention towards the fixed price," just as bidders in this field test appeared to exhibit limited attention towards identical sealed boxes that did not list "Tom Brady" in the item description. The salience/attention explanation for this field result is also linked to other empirical studies on limited attention, which have found that demand for goods can vary with consumer attentiveness to their full cost (Chetty, Looney, and Kroft 2009, Finkelstein 2009, and Hossain and Morgan 2006) and the relative salience of their attributes (Lacetera, Pope and Sydnor 2010). Just as consumer demand for goods can be influenced by their attentiveness to price and product attributes, their demand for risky prospects may be influenced by their relative attentiveness to the various outcomes and payoffs. Trading card manufacturers seem well aware of their incentive to focus collectors' attention on the best potential outcomes – their advertisements generally picture several cards from their product, but almost always disproportionately feature tougher-to-find, more valuable insert cards.

This salience/attention interpretation is also related to the model of choice under risk proposed by Bordalo, Gennaioli, and Shleifer (2010), wherein decision weights are endogenized as a function of payoffs by assuming that states are more salient as the distance between payoffs increases. However, a straightforward application of their model in its most general form would not predict this field result. Their model treats salience as an ordinal rather than cardinal property, and as such the Brady card outcome would be no more salient after the price appreciation, as it would remain in its position as the highest payoff (and thus most salient) outcome. The authors note that they did not want to impose the cardinal properties of any particular functional form on salience in the model, although this limits its predictions in scenarios like this field experiment where payoffs change but their ranking remains the same. In contrast to their model, it is possible that the dynamic price setting of this experiment directed collectors' attention even more disproportionately towards the Brady card. As discussed in Section 3.3, the Brady card was the one card in the set appreciating in value during the time of the experiment, and as Kahneman (2003) notes, "changes and differences are more accessible to a decision maker than absolute values." This type of "delta salience" may have applicability in other lottery contexts as well – for instance, it would suggest that state-run lotteries may generate

[66] This premium does not appear to be driven for the obvious reason that boxes with Tom Brady in the item title are seen by more potential buyers – controlling for list days and end date, there is no significant difference in the page views generated by boxes with and without "Tom Brady" in the item title.

more cumulative revenue by implementing a "growing pot" format over a series of single-round drawings.

Box-Buying as Gambling

The salience/attention explanation assumes that collectors were applying their standard risk preferences to their sealed-box valuations, but were additionally influenced by other affective factors. However, to the extent that box-buying is a form of gambling – supported by the Schaefer and Aasved (1997) finding that nearly two-thirds of surveyed attendees at a sports card convention could be classified as "pathological gamblers"[67] – the "standard risk preferences" assumption may not be appropriate. Barberis (2010) and Wolfers and Snowberg (2010) are recent studies that used prospect theory to explain the motivation for casino gambling and the favorite-longshot bias in horse race betting, respectively, but preferences estimated from such obvious risk-seeking activities may not be broadly applicable to other contexts of decision under risk. For instance, it may be that participants in such activities derive "gambling utility" that is not captured by standard risk parameters.[68] The identification strategy of this field experiment may partially "difference out" this factor – since gambling utility is derived both before and after the appreciation of the Brady card – but it may not completely abstract from it, thus limiting the generalizeability of the result.

The psychology literature has found increasing evidence of addiction as a neurological disorder, which may further limit the credibility of calibrating preferences from an addictive activity and generalizing them to a broader realm of economic activity. Dagher and Robbins (2009) survey the research on the effects of Parkinson's Disease medication – this research has found that the drug-induced release of dopamine in the brain impairs decision-making and reinforces addictive behaviors such as smoking or pathological gambling. While compulsive behavior such as gambling has been previously modeled in the literature in a formal decision-making framework (most notably the rational addiction model in Becker and Murphy 1988), the psychology evidence on addiction – that behavior such as gambling may be triggered by chemical stimuli that impair the brain's functional capabilities – suggests that its study in economics should perhaps be self-contained, and that preferences elicited from such activities (such as the decision weight calculated in this field experiment) should be generalized with caution. Ultimately, this field result may not be so much an estimate of a probability weight, but a calibration of the potential error from generalizing risk preferences derived from compulsive forms of behavior.

Finally, given the heterogeneity of risk preferences in the population at large, the gap between this field result and past experimental estimates may simply be an issue of sampling. Even if we assume card collectors are not pathological gamblers, they have selected into engaging in a risky activity, and even within this risk-seeking group, it is perhaps the most extreme risk-lovers that are bidding up the unopened box auction prices to levels that imply such a high decision weight on the "find-a-Tom-Brady" outcome. In contrast, experiments generally calibrate probability weights based on the risk preferences of the median subject.

[67] The authors conducted a modified version of the South Oaks Gambling Screen on 209 subjects.
[68] See Diecidue, Schmidt, and Wakker (2004) for a review of the "utility of gambling" literature, and a discussion of the difficulties of modeling the phenomenon in an analytically tractable manner.

Equilibrium Effects

One may wonder how price movements generated by such nonlinear probability weights could be sustained in a market setting. That is, why could an arbitrager not buy up individual rookie cards and offer "random rookie card" lots for sale, taking advantage of the high probability weight buyers seem to be attributing to finding a Tom Brady card? One possibility is that, to some extent, opening factory sealed boxes is an irreversible process. Buyers are leery that sellers may deliberately omit the Tom Brady card from any "re-packaged" product, and it is difficult for any individual buyer to monitor sellers and ensure that, across a large number of sales, the card allocations statistically resemble random allotments. Sellers' inability to establish sufficient credibility to sell such "randomized" products may prevent the pricing dynamic from being arbitraged away. That doesn't stop enterprising eBay sellers from trying – a quick glance through the trading card auctions and lots on the site reveals numerous listings for bulk lots with the carrot that one in every x lots will include a special chase card, or that one particular lot will include a highly sought after card such as the 2000 SP Authentic Tom Brady rookie card. However, these lots are generally not popular items, perhaps because of the credibility issue described above.

3.6.2 Potential Extensions

A clear limitation of the study is that it is an estimate of a single point on the probability weighting function, based on limited data (46 observations for the Tom Brady card, 78 observations for the unopened boxes) from a one-time event. There are two possible extensions of this study that could use trading card prices to further calibrate probability weights.

The first is to identify more event studies like this one and replicate the analysis for the associated cards and boxes. A series of such replications would allow for an estimation of the probability weighting function rather than just a single point.

A second extension would be to analyze what trading card prices in a "static" setting imply about probability weighting. This study has used a "dynamic," or "differences" approach, focusing on the change in the market price of a certain card and box due to a one-time event. However, even before this event, buying the 2000 SP Authentic football box (or any other box of sports cards for that matter) was already a risky prospect, with many different outcomes at different probabilities. Using static price data for single cards and boxes, the odds of pulling the various single cards, and some parameterization of the probability weighting function, it would be possible to estimate best-fit parameter values and hence the curvature of the function. One hypothesis is that salience has some effect on decision-making in such stable scenarios, but less than in the dynamic case – this hypothesis would predict that the estimated probability weighting function would be less skewed than in a dynamic analysis, but more skewed than in experimental calibrations. Finding that consumers are more likely to diverge from rational assessments of risk in dynamic, rather than static settings would have interesting implications for the markets for risky goods.

References

Abdellaoui, Mohammed. "Parameter-free elicitation of utility and probability weighting functions," *Management Science* 46 (2000), 1497-1512.

_____, Frank Vossman, and Martin Webber. "Choice-based Elicitation and Decomposition of Decision Weights for Gains and Losses under Uncertainty," mimeo, 2004.

Allison, Paul D. "Discrete-Time Methods for the Analysis of Event Histories," *Social Methodology* 13 (1982), 61-98.

Ariely, D. and K. Wertenbroch. "Procrastination, Deadlines, and Performance: Self-Control by Precommitment," *Psychological Science* 13 (2002), 219-224.

Angner, Erik and George Loewenstein. "Behavioral Economics," in Dov Gobbay, Paul Thagard, and John Woods eds., *Handbook of the Philosophy of Science* (Amsterdam: Elsevier, forthcoming)

Ashenfelter, Orley and David Card. "Did the Elimination of Mandatory Retirement Affect Faculty Retirement?" *American Economic Review* 92 (2002), 1-24.

Baddeley, Alan. *Human Memory: Theory and Practice* (East Sussex: Psychology Press, 1997).

Barberis, Nicholas. "A Model of Casino Gambling," mimeo, Yale University, 2010.

Benartzi, Shlomo and Richard Thaler. "Naïve Diversification Strategies in Defined Contribution Savings Plans," *American Economic Review* 91 (2001), 79-98.

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum. "The Nature of Risk Preferences: Evidence from Insurance Choices," mimeo, 2010.

Becker, Gary S. and Kevin M. Murphy. "A Theory of Rational Addiction," *Journal of Political Economy* (96) 1988, 675-700.

Bernheim, Douglas and Raphael Thomadsen. "Memory and Anticipation," mimeo, Stanford University, 2003.

Bleichrodt, Han and Jose Luis Pinto. "A Parameter-Free Elicitation of the Probability Weighting Function in Medical Decision Analysis," *Management Science* 46 (2000), 1485-1496.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. "Salience Theory of Choice Under Risk," mimeo, Harvard University, 2010.

Brewer, N., A. Keast amd A. Rishworth. "Improving the Confidence-Accuracy Relation in Eyewitness Identification: Evidence from Correlation and Calibration," *Journal of Experimental Psychology: Applied* 8 (2002), 44-56.

Browne, Mark J., and Hoyt, Robert E. “The Demand for Flood Insurance: Empirical Evidence,” *Journal of Risk and Uncertainty* 20 (2000), 271-289.

Camerer, Colin F. and Teck-Hua Ho. “Violations of the Betweenness Axiom and Nonlinearity in Probability,” *Journal of Risk and Uncertainty* 8 (1994), 167-196.

_____, George Loewenstein, and Martin Weber. “The Curse of Knowledge in Economic Settings: An Experimental Analysis,” *Journal of Political Economy* 97 (1989), 1232-1254.

Chen, Y., S.Moorthy, and J.Zhang. “Price Discrimination: Price Discrimination After the Purchase: A Note on Rebates as State-Dependent Discounts,” *Management Science* 51 (2005), 1131-1140.

Chetty, Raj, Adam Looney, and Kory Kroft. “Salience and Taxation: Theory and Evidence,” *American Economic Review* 99 (2009), 1145-77.

Choi, James, David Laibson, Brigitte Madrian, and Andrew Metrick. “Saving for Retirement on the Path of Least Resistance,” in Edward J. McCaffrey and Joel Slemrod, eds., *Behavioral Public Finance: Toward a New Agenda* (New York: Russell Sage Foundation, 2006).

Chuang, T. “Rebates: The Discount that Makes People See Red,” *Orange County Register*, November 2003.

Cockburn, J. and P.T. Smith. "Effects of Age and Intelligence on Everyday Memory Tasks," in Michael M. Gruneberg, Peter E. Morris, and Robert N. Sykes, eds., *Practical Aspects of Memory: Current Research and Issues, Volume 1: Memory in Everyday Life* (Chichester: John Wiley & Sons, 1988).

Cohen, Gillian. “Memory and Ageing,” in Graham M. Davies and Robert H. Logie, eds., *Memory in Everyday Life* (Amsterdam: North-Holland, 1993).

Cook, Philip J. and Robert H. Frank. *The Winner-Take-All Society* (New York: Penguin Books, 1996).

Dagher, Alain and Trevor W. Robbins. "Personality, Addiction, Dopamine: Insights from Parkinson’s Disease," *Neuron* 61 (2009), 502-510.

DellaVigna, Stefano. “Psychology and Economics: Evidence from the Field,” *Journal of Economic Literature* 47 (2009), 315-372.

_____ and Ulrike Malmendier. “Paying Not to Go to the Gym,” *American Economic Review* 96 (2006), 694-719.

Diecidue, Enrico, Ulrich Schmidt, and Peter P. Wakker. "The Utility of Gambling Reconsidered," *The Journal of Risk and Uncertainty* 29 (2004), 241-259.

Ebbinghaus, Hermann. *Über das Gedchtnis. Untersuchungen zur experimentellen Psychologie* (Leipzig: Duncker & Humblot, 1885). English edition is *Memory: A Contribution to Experimental Psychology* (New York: Teachers College, Columbia University, 1913).

Edwards, Matthew A. "The Rebate Rip-Off: New York's Legislative Responses to Common Consumer Rebate Complaints," *Pace Law Review* 29 (2009), 471-478.

Edwards, Ward. "The Theory of Decision-Making." *Psychological Bulletin* 51 (1954), 380-417.

Efron, Bradley. "Logistic Regression, Survival Analysis, and the Kaplan-Meier Curve," *Journal of the American Statistical Association* 83 (1988), 414-425.

Ellis, J.A. "Memory for Future Intentions: Investigating Pulses and Steps," in Michael M. Gruneberg, Peter E. Morris, and Robert N. Sykes, *Practical Aspects of Memory: Current Research and Issues, Volume 1: Memory in Everyday Life* (Chichester: John Wiley & Sons, 1988).

_____ and I. Nimmo-Smith. "Recollecting Naturally-Occurring Intentions: A Study of Cognitive and Affective Factors," *Memory* 1 (1993), 107-126.

Ellison, Glen. "Bounded Rationality in Industrial Organization," mimeo, Massachusetts Institute of Technology, 2006.

Ericson, Keith. "On Memory and Procrastination," mimeo, Harvard University, 2010.

_____. "Forgetting We Forget: Overconfidence and Prospective Memory," *Journal of the European Economic Association*, forthcoming.

Finkelstein, Amy. "EZ-Tax: Tax Salience and Tax Rates," *Quarterly Journal of Economics* 124 (2009), 969-1010.

Fischhoff, B. "Hindsight / Foresight: The Effect of Outcome Knowledge on Judgment Under Uncertainty," *Journal of Experimental Psychology: Human Perception and Performance*, 1 (1975), 288-299.

Gabaix, Xavier, David Laibson, Guillermo Moloche and Stephen Weinberg. "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model," *American Economic Review* 96 (2006), 1043-1068.

Gates, S.J. and D.K. Colburn. "Lowering Appointment Failures in a Neighborhood Health Center," *Medical Care* 14 (1976), 263-267.

Gonzalez, Richard and George Wu. "On the Shape of the Probability Weighting Function." *Cognitive Psychology* 38 (1999), 129-166.

Grow, Brian and Rishi Chatwal. "The Great Rebate Runaround," *Business Week*, December 5 2005, p. 34-37.

Harris, John E. "Remembering to do things: A forgotten topic." in Jeff E. Harris and Peter E. Morris, eds., *Everyday Memory, Actions and Absent-Mindedness* (London: Academic Press, 1984).

Hossein, Tanjim and John Morgan. "…Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay," *B.E. Journals in Economic Analysis and Policy: Advances in Economic Analysis and Policy* 6 (2006), 1-27.

Johnson, Eric J., John Hershey, Jacqueline Meszaros, and Howard Kunreuther. "Framing, Probability Distortions, and Insurance Decision," *Journal of Risk and Uncertainty* 7 (1993), 25-51

Kahn, Debra. "Will 'Smart' Electric Meters Lead to Smarter Consumers?" *New York Times*, December 11, 2009.

Kahneman, Daniel. "Maps of Bounded Rationality: Psychology for Behavioral Economics." *American Economic Review* 93 (2003), 1449-1476.

________________ and Amos Tversky. "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica* 47 (1979), 263-291.

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman. "Getting to the Top of Mind: How Reminders Increase Saving," NBER Working Paper, 2010.

Kauffman, Robert J. and Charles A. Wood. "Doing Their Bidding: An Empirical Examination of Factors That Affect a Buyer's Utility in Internet Auctions," forthcoming in *Information Technology and Management*, 2004.

Kearney, Melissa S., Peter Tufano, Jonathan Guryan, and Erik Hurst. "Making Savers Winners: An Overview of Prize-Linked Savings Product," NBER Working Paper, 2010.

Lacetera, Nicola, Devin Pope, and Justin Sydnor. "Heuristic Thinking and Limited Attention in the Car Market," mimeo 2010.

Lattimore, Pamela, Joanna R. Baker, and A. Dryden White. "The Influence of Probability on Risky Choice." *Journal of Economic Behavior and Organization* 17 (1992), 377-400.

Lee, Young Han and Ulrike Malmendier. "The Bidder's Curse," *American Economic Review*, forthcoming.

Leonesio, R.J., and T.O. Nelson. "Do Different Metamemory Judgments Tap the Same Underlying Aspects of Memory?" *Journal of Experimental Psychology: Learning, Memory & Cognition*, 16 (1990), 464-470.

Levy. R. and V. Claravall. "Differential Effects of a Phone Reminder on Appointment Keeping for Patients with Long and Short Between-Visit Intervals," *Medical Care* 15 (1977), 435-438.

_____ and Geoffrey R. Loftus. "Compliance and memory," in Jeff E. Harris and Peter E. Morris, eds., *Everyday Memory, Actions and Absent-Mindedness* (London: Academic Press, 1984).

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. "Projection Bias in Predicting Future Utility," *Quarterly Journal of Economics* 118 (2003), 1209-1248.

_____, Elke U. Weber, Christopher K. Hsee, and Ned Welch. "Risk as Feelings," *Psychological Bulletin* 127 (2001), 267-286.

Loftus, E. F. *Eyewitness Testimony* (Cambridge: Harvard University Press, 1979).

Lucking-Reiley, David, Doug Bryan, Naghi Prasad and Daniel Reeves. "Pennies from eBay: the Determinants of Price in Online Auctions," mimeo.

MacMillan, Douglas. "Can Topps Still Play Ball?" *Businessweek*, April 19, 2007.

Madrian, Brigitte and Shea, Dennis. "The Power of Suggestion: Intertia in 401(k) Participation and Savings Behavior," *Quarterly Journal of Economics* 116 (2001), 1149-1187.

McDaniel, Mark A. and Giles O. Einstein. *Prospective Memory: an Overview and Synthesis of an Emerging Field* (Thousand Oaks: Sage Publications, 2007).

Meacham, J.A. and B. Leiman. "Remembering to Perform Future Actions," in U. Neisser, ed., *Memory Observed: Remembering in Natural Contexts* (San Francisco: W.H. Freeman and Company, 1982).

_____, and J. Singer. "Incentive Effects in Prospective Remembering." *The Journal of Psychology* 97 (1977): 191-197.

Morris, Peter E. "The Validity of Subjective Reports on Memory," in Jeff E. Harris and Peter E. Morris, eds., *Everyday Memory, Actions and Absent-Mindedness* (London: Academic Press, 1984).

____________. "Prospective Memory: Remembering to Do Things," in Michael Gruneberg and Peter Morris, eds., *Aspects of Memory* (New York: Routledge, 1992).

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

O'Donoghue, Ted, and Matthew Rabin. "Incentives for Procrastinators," *Quarterly Journal of Economics* 114 (1999a), 769-816.

_________________________________. "Doing it Now or Later," *American Economic Review* 89 (1999b), 103-124.

Orne, M.T. "Hypnosis, Motivation, and the Ecological Validity of the Psychological Experiment," in W.J. Arnold and M.M. Page, eds., *Nebraska Symposium on Motivation* Vol. 18 (Lincoln: University of Nebraska Press, 1970).

Palmer, Kimberly. "Why Shoppers Love to Hate Rebates," *US News and World Report*, January 18 2008.

Parker, J.W., E.K. Taylor, R.S. Barnett, and S. Martens. "Rating Scale Content: The Relationships Between Supervisory- and Self-Rating." *Personnel Psychology* 12 (1959), 49-63.

Piccione, M. And A. Rubinstein. "On the Interpretation of Decision Problems with Imperfect Recall," *Games and Economic Behavior* 20 (1997), 3-24.

Rottenstreich, Yuval and Christopher K. Hsee. "Money, Kisses, and Electric Shocks: On the Affective Psychology of Risk," *Psychological Science* 12 (2001), 185-190.

Rubin, David C., and Amy Wenzel. "One hundred years of forgetting: A quantitative description of retention." *Psychological Review* 103 (1996), 734-760.

Schaefer, James M. and Mikal J. Aasved. "Gambling in Sports Card Collecting: An Empirical Study," mimeo, 1997.

Shafir, Eldar and Amos Tversky. "Choice Under Conflict: The Dynamics of Deferred Decision," *Psychological Science* 3 (1992), 358-361.

Shepard, D.S. and T.A. Moseley. "Mailed Versus Telephoned Appointment Reminders to Reduce Broken Appointments in a Hospital Outpatient Department," *Medical Care* 14 (1976) 268-273.

Silk, Timothy. "Examining Purchase and Non-Redemption of Mail-In Rebates: The Impact of Offer Variables on Consumers' Subjective and Objective Probability of Redeeming," Unpublished dissertation, University of Florida, 2004.

_____ and Chris Janiszewski. "Managing Mail-In Rebate Promotions," mimeo, University of South Carlina, 2003.

Snowberg, Erik and Justin Wolfers. "Explaining the Favorite-Long Shot Bias: Is it Risk-Love or Misperceptions?" *Journal of Political Economy* 118 (2010), 723-746.

Spencer, J. "Rejeced! Rebates Get Harder to Collect," *Wall Street Journal*, June 2002, p. D2.

Svenson, O. “Are We All Less Risky and More Skillful Than Our Fellow Drivers?” *Acta Psychology* 47 (1981), 143-148.

Taylor, Shelley, and Suzanne Thompson. "Stalking the elusive vividness effect." *Psychological Review*, 89 (1982), 155-181.

Thaler, Richard H. “From Homo Economicus to Homo Sapiens,” *Journal of Economic Perspectives* 14 (2000), 133-141.

Tugend, Alina. “A Growing Anger Over Unpaid Rebates,” *New York Times*, March 4 2006, p. C5.

Tversky, Amos and Craig R. Fox. “Weighing Risk and Uncertainty.” *Psychological Review* 102 (1995), 269-283.

_____ and Daniel Kahneman. “Advances in Prospect Theory: Cumulative Representation of Uncertainty.” *Journal of Risk and Uncertainty* 5 (1992), 297-323.

Vesonder, G.T., and J.F. Voss. “On the Ability to Predict One's Own Responses While Learning,” *Journal of Memory and Language* 24 (1985), 363-376.

Whitaker, Barbara. “A Web Offer Too Good To Be True? Read the Fine Print,” *New York Times*, December 26, 2004.

Wilkins, A.J. "A Failure to Demonstrate the Effects of the "Retention Interval" in Prospective Memory," mimeo, 1976.

_____ and Alan D. Baddeley. “Remembering to Recall in Everyday Life: An Approach to Absentmindedness," in Michael M. Gruneberg, Peter E. Morris, and Robert N. Sykes, *Practical Aspects of Memory* (London: Academic Press, 1978).

Willis, Gerri. “‘Free trial’ or Internet scam?" *CNNMoney.com*, December 17, 2009.

Wilson, Andrea. “Bounded Memory and Biases in Information Processing,” mimeo, Princeton University, 2003.

Wilson, B., J. Cockburn, A. Baddeley, and R. Hiorns. "The Development and Validation of a Test Battery for Detecting and Monitoring Everyday Memory Problems." *Journal of Clinical and Experimental Neuropsychology*, 11 (1989), 855-870.

Wu, George and Richard Gonzalez. “Curvature of the Probability Weighting Function.” *Management Science* 42 (1996), 1676-1690.

_____, Jiao Zhang, and Richard Gonzalez. “Decision Under Risk,” mimeo, 2003.

Wylie, R.C. *The Self-Concept* vol. 2. (Lincoln: The University of Nebraska Press, 1979).

Appendix

Appendix A: Proofs of Chapter 1 Propositions

Proof. [Proof of Proposition 1] We will show that $s_t^*(1,1)$, the PR's optimal s_t, is increasing in t. Since $h_t = E(m_t)F(s_t)$ and F is increasing in s_t, this will establish the claim.

Note that a PR is defined by $m = (1,1,\cdots,1)$ with probability 1. We show that the continuation value to a PR of deferring in time t, V_t, is decreasing in t:

$$\begin{aligned} V_t &= \sum_{t'=t+1}^{T} \sum_{m \in 2^T} \mu_t(m) \pi_{t,t'}(\cdot \mid m) m_{t'} F(s_{t'}) \left[b - E\left(c_{t'} \mid c_{t'} \le s_{t'}\right)\right] \\ &= \sum_{t'=t+1}^{T} \pi_{t,t'}(\cdot \mid (1,\cdots,1)) F(s_{t'}) \left[b - E\left(c_{t'} \mid c_{t'} \le s_{t'}\right)\right] \\ &= \pi_{t,t+1}(\cdot \mid (1,\cdots,1)) F(s_{t+1}) \left[b - E\left(c_{t+1} \mid c_{t+1} \le s_{t+1}\right)\right] + V_{t+1}, \end{aligned}$$

Since the first term is always non-negative we see that $V_t > V_{t+1}$ as claimed.

Then since $s_t^* = b - V_t$, s_t^* is increasing in t. ■

Proof. [Proof of Proposition 2] First, we argue that the hazard at a given number of periods away from the deadline converges to zero as the length of the deadline grows. In our notation,

$$\lim_{T \to \infty} h_{T-j} = 0, \forall j, p, \theta$$

The argument is as follows: consider an infinite population of identical SFs, and examine those that have not completed by the deadline. Note that the fraction of these individuals who complete in period T is by definition the hazard at time T. Some fraction will have $m_T = 1$, and those individuals complete at rate $F(b)$, since $s_T^* = b$. However, the rest with $m_T = 0$ will complete at rate 0. Rerun this same experiment for a larger T. The fraction with $m_t = 0$ will increase, and in fact will converge to 1 as T grows. This same argument holds for $m_T - j$, for each fixed j, as T grows. This argument establishes the claim.

We can prove this formally by noting that with strictly positive probability, an SF forgets and never remembers. We establish this in the following two lemmas by noting that $\{N_t\}$ forms a Markov Chain, and showing that 0 is a transient state:

Lemma 11 *The process returns to zero almost surely, i.e.* $P(N_{t+k} = 0 \text{ for some } k \geq 1 \mid N_t = 0) = 1$, *if and only if* $(1-p_0)(1-p_1)...(1-p_n) \to 0$ *as* $n \to \infty$.

Proof. Let R be the (random) time of the first return to zero, conditional on the process beginning at zero. Note that

$$P(N_{t+k} = 0 \text{ for some } k \geq 1 \mid N_t = 0) = 1$$

if and only if

$$P(R = \infty) = 0.$$

That is, the process returns to zero almost surely if and only if R is finite almost surely. Next, observe that

$$\{R > n\} \downarrow \{R = \infty\}, \text{i.e.}$$
$$\{R > 1\} \supset \{R > 2\} \supset \cdots \bigcup_{n=1}^{\infty} \{R > n\} = \{R = \infty\}.$$

So, by the downward continuity of the measure P we have

$$P(R > n) \downarrow P(R = \infty).$$

Thus

$$P(R = \infty) = 0$$

if and only if

$$P(R > n) \to 0.$$

The result follows from the fact that

$$P(R > n) = (1 - p_0)(1 - p_1)...(1 - p_{n-1}).$$

■

Lemma 12 *$P(N_{t+k} = 0$ for some $k \geq 1 \mid N_t = 0) < 1$. That is, there is a strictly positive probability that the individual forgets forever.*

Proof. By Proposition 1, we need only show that

$$(1 - p_0)(1 - p_1)...(1 - p_n) = (1 - p)(1 - p\theta)\cdots(1 - p\theta^n) \not\to 0$$
$$\text{as} n \to \infty.$$

This is equivalent to showing that

$$\log[(1 - p)(1 - p\theta)...(1 - p\theta^n)] \not\to -\infty \text{ as } n \to \infty,$$

which is in turn equivalent to showing that

$$\sum_{k=0}^{\infty} \log(1-p\theta^k)$$

converges. Note that

$$\log(1-p\theta^k) = -p\theta^k + O(\theta^{2k}) \text{as} \theta^{2k} \to 0.$$

Thus

$$\sum_{k=0}^{\infty} \log(1-p\theta^k) = \sum_{k=0}^{\infty} [-p\theta^k + O(\theta^{2k})],$$

and this series converges if and only if $\theta < 1$, which is true by assumption. ■

Next, we claim that h_1 for a SF is bounded below for all T. This will establish the result. For fix j, and suppose $\underline{h}$ is the lower bound on h_1. By the above, for a sufficiently large T we can get $h_{T-j} < \underline{h}$. But then $h_{T-j} < h_1$, which shows the hazard must decrease at some point.

To show that h_1 is bounded below, we show that s_1^* is bounded below, since $h_1 = F(s_1^*)$. The argument is as follows: suppose h_1 were not bounded below. Then as T grows, $h_1 \to 0$. This implies that in the infinite-horizon problem the expected completion time is infinite; the intuition is that the SF waits around forever for an arbitrarily low cost draw. But this cannot be optimal, because in following this strategy the SF will at some point forget and never remember, with probability one. For we have argued above that 0 is a transient state of N_t, and so each period of remembrance carries a constant, positive probability that the SF will forget in the next period and never remember again. So the probability of such an eternal forgetting occurs with probability 1 in the infinite horizon problem. This implies that the expected utility from such a strategy is zero, which cannot be optimal. Therefore h_1 must be bounded below. ■

Proof. [Proof of Proposition 3] The statement is that $Q(s^*, \hat{p}, \hat{\theta})$ decreases in $\hat{p}$ and $\hat{\theta}$ for $\hat{p} \le p, \hat{\theta} \le \theta$.

First we will show that $s_t(p,\theta)$ is decreasing in p and θ, by induction on t. Recall that $s_T = b$, and observe that $s_{T-1} = b - E_{T-1}(V_T \mid m_{T-1} = 1)$. Then note that

$$\begin{aligned} E_{T-1}(V_T \mid m_{T-1} = 1) &= E_{T-1}(\max(m_T(b-c_t), V_{T+1} \mid m_{T-1} = 1)) \\ &= E_{T-1}(m_T(b-c_T) \mid m_{T-1} = 1) \end{aligned}$$

$$= p(b - E_{T-1}(c_T))$$

which is clearly increasing in p. Next suppose that s_t is decreasing in p, θ. This implies that $E_t(V_{t+1} \mid m_t = 1)$ is increasing in p, θ. Recall that $s_{t-1} = b - E_{t-1}(V_t \mid m_{t-1} = 1)$ and $V_t = \max(m_t(b - c_t), V_{t+1})$, so

$$\begin{aligned} E_{t-1}(v_t \mid m_{t-1} = 1) &= E_{t-1}(\max(m_t(b - c_t), V_{t+1} \mid m_{t-1} = 1)) \\ &= \max(p(b - Et - 1(c_t)), E_{t-1}(V_{t+1} \mid m_{t-1} = 1)) \\ &= \max(p(b - Et - 1(c_t)), pE_t(V_{t+1} \mid m_t = 1, m_{t-1} = 1) + \\ &(1 - p)E_t(V_{t+1} \mid m_t = 0, m_{t-1} = 1)) \end{aligned}$$

The first term in the max in clearly increasing in p. Consider the second term, $pE_t(V_{t+1} \mid m_t = 1, m_{t-1} = 1) + (1 - p)E_t(V_{t+1} \mid m_t = 0, m_{t-1} = 1))$. The first term is increasing in p, θ by the induction hypothesis. The second term is as well, since we can continue writing V_{t+1} in terms of V_{t+2} and so forth, using the induction hypothesis each time to deal with terms of the form $E_{t-j}(V_{t-j+1} \mid m_{t-j} = 1)$. The only term we will need to deal with is of the form $E_{T-1}(V_T \mid m_{T-1} = 0, m_{T-2} = 0, \cdots, m_{t-1} = 1)$, which is clearly increasing in p, θ. Finally, note that $E_t(V_{t+1} \mid m_t = 1, m_{t-1} = 1) > E_t(V_{t+1} \mid m_t = 0, m_{t-1} = 1)$, so increasing its weight, namely p, will increase the entire quantity. The same argument applies for the weights that expand out of the second term.

Therefore $s_t(p, \theta)$ is decreasing in p and θ. Since an overconfident agent follows the exact same strategy as a sophisticated forgetter with the same $\hat{p}, \hat{\theta}$, and since we have just established that her reservation costs will be uniformly lower than an agent with the same true memory parameters but less overconfidence, we have established that the probability of task completion, $Q(s^*, \hat{p}, \hat{\theta})$, is decreasing in $\hat{p}$ and $\hat{\theta}$. ■

Proof. [Proof of Proposition 4] The statement is that $E_0(V_0(s^*, \hat{p}, \hat{\theta}))$ decreases in $\hat{p}$ and $\hat{\theta}$ for $\hat{p} \leq p, \hat{\theta} \leq \theta$.

From the previous proof, we have shown that as $(\hat{p}, \hat{\theta})$ go to $(1,1)$, the strategy of the NF converges monotonically away from that of the PR and towards that of the PR. We claim (without proof) that as the NF's strategy moves monotonically away from the PR's strategy as $(\hat{p}, \hat{\theta})$ go to $(1,1)$, that the utility loss increases monotonically as well. Put another way, the farther the strategy is from the optimal, the greater will be the utility loss, and we have established that the distance from the NF's strategy to the optimal PR's strategy increases in the degree of naïvete. ■

Proof. [Proof of Proposition 5] This proposition becomes obvious once one observes that the structure of the last T periods of a $T+\delta$-deadline problem is identical to the structure of a T-deadline problem. That is, $s^*_{T+\delta,t+\delta} = s^*_{T,t} \text{for} t \in \{1,2,\cdots,T\}$. Thus the probability of completing in periods $1+\delta, 2+\delta, \cdots, T+\delta$ with a $T+\delta$-deadline, conditional on having not completed until $1+\delta$, is the same as the probability of completing in the T-deadline case. Since the probability of completing in periods $1,2,\cdots,\delta$ is non-negative, it is clear that $Q_{T+\delta}(s^*_{T+\delta}) > Q_T(s^*_T)$.

As for the utility from the task, a similar argument holds – since the optimal strategy in $1+\delta, 2+\delta, \cdots, T+\delta$ for the $T+\delta$ problem is the same as in the original $1,2,\cdots,T$ problem, the utility from those periods conditional on arriving is the same as the total utility of the original problem. Then setting $s_1, s_2, \cdots, s_\delta$ equal to b in the longer deadline case will result in the same utility for the two problems, and by letting them be arbitrarily close but strictly below b we can add on strictly positive utility. ■

Proof. [Proof of Proposition 6] The argument is similar to the utility portion of Proposition 5. However, while we before relied on the fact that the PR's optimal strategy is identical in the ending periods of an extended deadline strategy, we will now only rely on the fact that an SF *could* use the same strategy in an extended deadline setting.

Specifically, consider extending from a T deadline to a $T+\delta$ deadline. An SF could use s^*_T for the first T periods of her $s_{T+\delta}$ strategy. This would result in an ex ante utility obtained during those periods equal to the total ex ante utility from a T deadline problem. The addition of additional ex ante utility in the final δ periods would give the utility nod to the $T+\delta$ problem. The fact that the SF will not in general actually use this replicated strategy does matter; all we argue is that the replicated strategy presents a lower bound on utility. ■

Proof. [Proof of Proposition 7] We begin with notation. Let $Q_{T,t}$ be the ex ante probability that the first task completion occurs in period t of a T-deadline problem. Then $Q_T = \sum_{t=1}^{T} Q_{T,t}$. Note that

$$
\begin{aligned}
Q_{T,1} &= P(m_1 = 1)P(c_1 < s_{T,1}) \\
&= pF(s^*_{T,1}), \\
Q_{T,2} &= P(m_2 = 1)P(c_1 < s_{T,1})(1 - Q_{T,1}) \\
&= [P(m_2 = 1 \mid N_2 = 0)P(N_2 = 0) + P(m_2 = 1 \mid N_2 = 1)P(N_2 = 1)]F(s^*_{T,2})(1 - Q_{T,1}), \\
&= \left[p \frac{1 - F(s_{T,1})}{1 - pF(s_{T,1})} p + p\theta \frac{1}{1 - pF(s_{T,1})}(1-p) \right] F(s^*_{T,2})(1 - Q_{T,1}), \text{and in general}
\end{aligned}
$$

$$Q_{T,t} = P(m_t = 1)P(c_t < s^*_{T,t})\prod_{j=1}^{t-1}(1-Q_{T,j})$$
$$= \left[\sum_{k=0}^{t-1} P(m_t = 1 \mid N_t = k)P(N_t = k)\right] F(s^*_{T,t})\prod_{j=1}^{t-1}(1-Q_{T,j})$$

Recall that $P(m_t = 1 \mid N_t = k) = \theta^k p$. Also note that $P(N_t = 0)$ is of the form $\sum_{i=0}^{t-1}\sum_{j=0}^{t-2} a_{ij} p^i \theta^j$ where $a_{i0} > 0$ if and only if $i = t-1$. In fact we have $a_{t-1,0} = 1$, since the only term of the double sum contributing to $N_t = 0$ involving p^{t-1} and θ^0 corresponds to the case where $m_1 = 1, m_2 = 1, \cdots, m_{t-1} = 1$. Thus

$$\lim_{\theta\to 0} Q_{T,t} = \lim_{\theta\to 0}\left[\sum_{k=0}^{t-1}\theta^k pP(N_t = k)\right] F(s^*_{T,t})\prod_{j=1}^{t-1}(1-Q_{T,j})$$
$$= pF(s^*_{T,t})\lim_{\theta\to 0}\sum_{i=0}^{t-1}\sum_{j=0}^{t-2} a_{ij} p^i\theta^j \prod_{j=1}^{t-1}(1-Q_{T,j})$$
$$= pF(s^*_{T,t})\lim_{\theta\to 0}\sum_{i=0}^{t-1} a_{i0} p^i \prod_{j=1}^{t-1}(1-Q_{T,j})$$
$$= p^t F(s^*_{T,t})\lim_{\theta\to 0}\prod_{j=1}^{t-1}(1-Q_{T,j})$$

From these recursions we see that $\lim_{\theta\to 0} Q_{T,t}$ is in fact a polynomial in p and that the order of this polynomial is strictly greater than 1 except in the case where $t = 1$, and then the order is exactly 1. Add these per-period completion rates to form the total completion rate as

$$\lim_{\theta\to 0} Q_T = \lim_{\theta\to 0}\sum_{t=1}^{T} Q_{T,t}$$
$$= \sum_{t=1}^{T}\lim_{\theta\to 0} Q_{T,t}$$
$$= \sum_{t=1}^{T} p^t F(s^*_{T,t})\lim_{\theta\to 0}\prod_{j=1}^{t-1}(1-Q_{T,j})$$

and then compare Q_T to $Q_{T+\delta}$ in the limit as $\theta \to 0$:

$$\lim_{\theta\to 0}\frac{Q_T}{Q_{T+\delta}} = \frac{\sum_{t=1}^{T} p^t F(s^*_{T,t})\lim_{\theta\to 0}\prod_{j=1}^{t-1}(1-Q_{T,j})}{\sum_{t=1}^{T+\delta} p^t F(s^*_{T+\delta,t})\lim_{\theta\to 0}\prod_{j=1}^{t-1}(1-Q_{T+\delta,j})}$$

Next we take the limit of both sides as $p \to 0$, use the fact discussed above that $p^t F(s^*_{T,t}) \lim_{\theta\to 0} \prod_{j=1}^{t-1}(1-Q_{T,j})$ is a polynomial in p, and apply L'Hospital's rule to obtain

$$\lim_{p,\theta\to 0} \frac{Q_T}{Q_{T+\delta}} = \frac{F(s^*_{T,1})}{F(s^*_{T+\delta,1})}$$

In fact $s^*_{T,1} = s^*_{T+\delta,1+\delta} > s^*_{T+\delta,1}$; see the proofs for Propositions 5 and 1 for the equality and inequality, respectively. Therefore

$$\lim_{p,\theta\to 0} \frac{Q_T}{Q_{T+\delta}} > 1.$$

Then since Q_T is continuous in p and θ, we can find p', θ' such that $0 < p < p', 0 < \theta < \theta'$ implies that $\frac{Q_T(s^*_T, p, \theta)}{Q_{T+\delta}(s^*_{T+\delta,p,\theta})} > 1$, i.e. that $Q_T(s^*_T, p, \theta) > Q_{T+\delta}(s^*_{T+\delta,p,\theta})$. This establishes Part 1 of the Proposition.

For Part 2, we can use the intermediate result of Part 1 that as $p, \theta \to 0$ the fraction of task completion occurring during the first period converges to 1. Combine this with the fact that $s^*_{T,1} > s^*_{T+\delta,1}$ and we see that in the limit, the set of costs for which the agent in the T-deadline problem completes the task strictly contains those for which the $T+\delta$-deadline agent completes, and that the additional costs for which she completes add strictly positive utility. This establishes Part 2. ■

Proof. [Proof of Proposition 8] Trivial. For a utility-price of $\kappa(a^1)$ the PR can turn herself into a SF. Thus utility loss of the PR relative to the SF is bounded by this amount. ■

Proof. [Proof of Proposition 9] Trivial as well. Even when $\theta = 1$, the NF's utility from a task is at most p times the PR's utility of the same task. And as $b \to \infty$, holding the cost distribution fixed, we get that $E_0(V_1(s^*(1,1),1,1)) \to \infty$. So, $E_0(V_1(s^*(1,1),1,1)) - E_0(V_1(s^*(\hat{p},\hat{\theta}), p, \theta)) > (1-p)E_0(V_1(s^*(1,1),1,1))$ grows arbitrarily large as well. ■

Proof. [Proof of Proposition 10] The proof of this proposition is virtually identical to that of Proposition 5. Any perfectly-remembering agent, whether hyperbolic discounter or not, faces the exact same problem in the last T periods of a $T+\delta$ time frame as in a T period time frame and thus over the last T periods executes the exact same strategy with the same expected utility conditional on arriving in that period. In the δ periods prior to this replication, a hyperbolic discounter will only complete the task if the cost draw is low enough to increase overall expected utility. Thus, adding periods strictly increases the expected utility of any PR in this context, regardless of the time-discounting parameters. The increasing probability of task execution is analogous. ■

Appendix B: Figures and Tables Supplement for Chapter 2

Table 2.4: Monthly Cancellations by Pricing Group, Non-US Signups

Standard			Month 1		Month 2		Month 3		Month 4	
Group	Offer	Actual	Start	Cancels	Start	Cancels	Start	Cancels	Start	Cancels
1	19.83	19.83	80	46 58%	34	3 9%	31	3 10%	18	3 17%
2	19.95	19.95	73	44 60%	29	10 34%	19	2 11%	15	0 0%
3	20	19.83	85	55 65%	30	5 17%	25	5 20%	14	1 7%
4	20	20	85	47 55%	38	10 26%	28	6 21%	15	1 7%
5	20.13	20.13	75	40 53%	35	10 29%	25	7 28%	12	1 8%
Total			398	232 58%	166	38 23%	128	23 18%	74	6 8%
Camouflage Pricing Groups (1,3,5)			240	141 59%	99	18 18%	81	15 19%	44	5 11%

Premium			Month 1		Month 2		Month 3		Month 4	
Group	Offer	Actual	Start	Cancels	Start	Cancels	Start	Cancels	Start	Cancels
1	29.78	29.67	26	19 73%	7	2 29%	5	1 20%	2	1 50%
2	29.78	29.78	23	14 61%	9	3 33%	6	3 50%	1	0 0%
3	29.95	29.95	29	20 69%	9	4 44%	5	0 0%	2	0 0%
4	30	30	24	8 33%	16	3 19%	13	3 23%	4	1 25%
5	30.14	30.14	34	22 65%	12	2 17%	10	5 50%	4	1 25%
Total			136	83 61%	53	14 26%	39	12 31%	13	3 23%
Camouflage Pricing Groups (1,2,5)			83	55 66%	28	84 300%	21	9 43%	7	2 29%

Figure 2.1A: Monthly Hazard Rates by Pricing Group, US Signups

Standard Membership

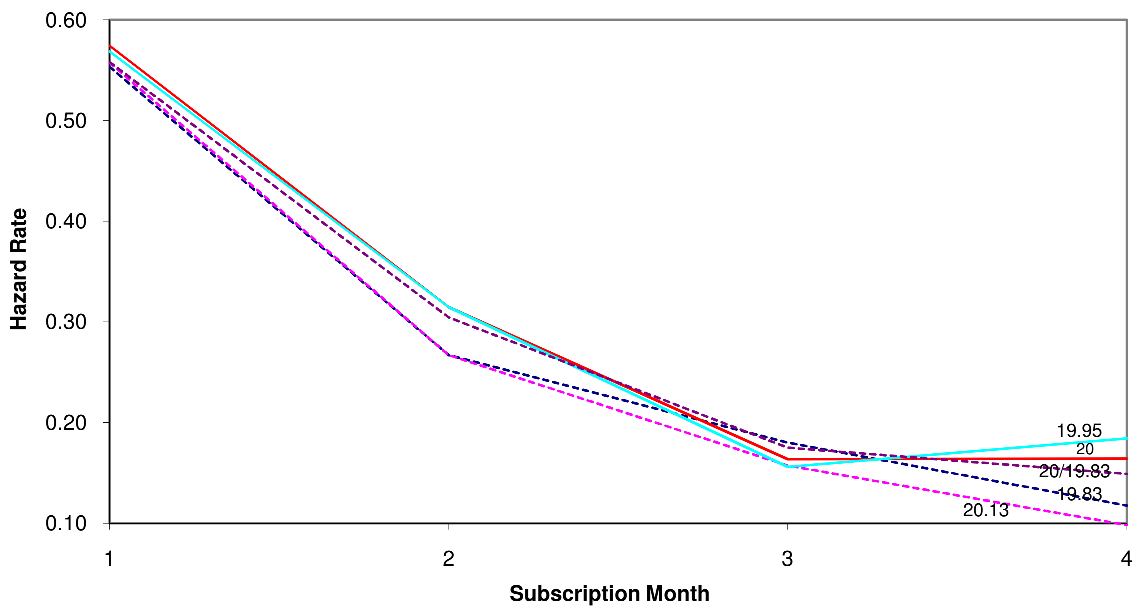

Premium Membership

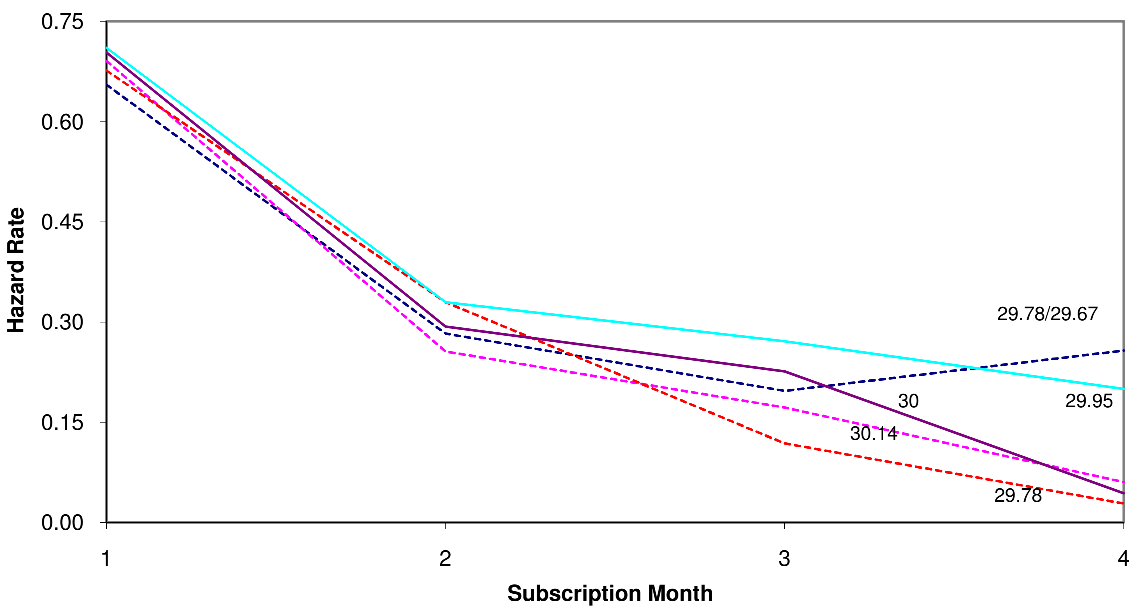

Figure 2.1B: Monthly Hazard Rates by Pricing Group, Non-US Signups

Standard Membership

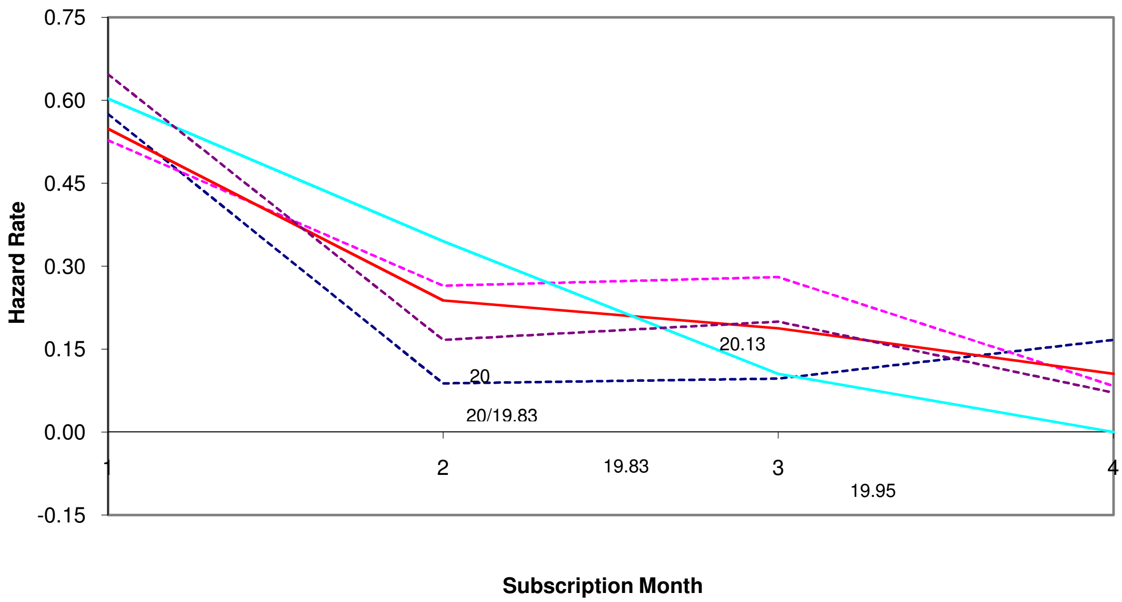

Premium Membership

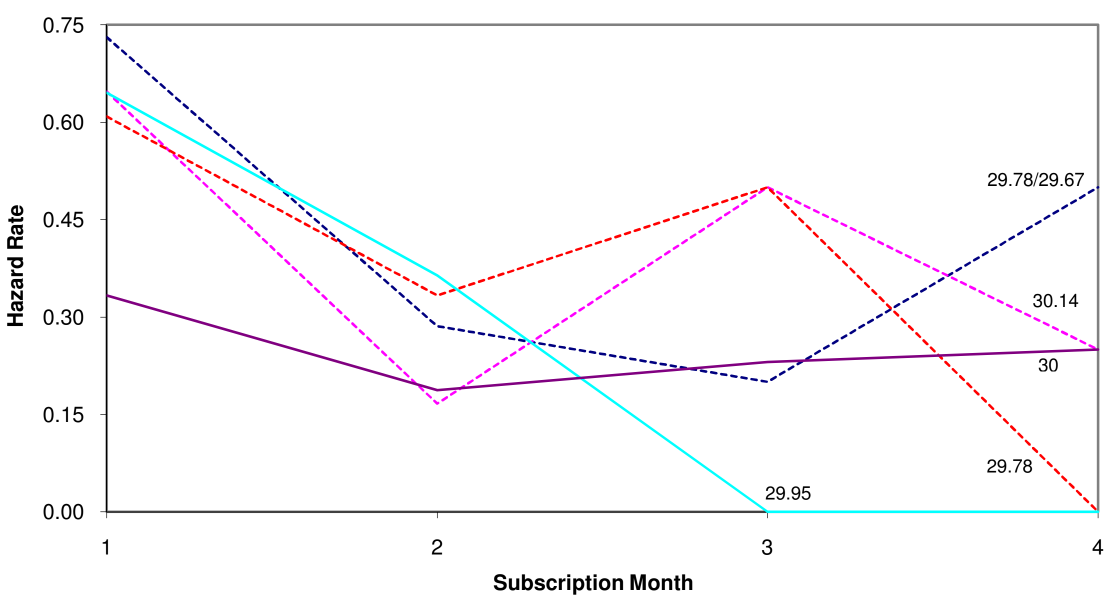

Table 2.6: Comparison of US and non-US camouflage Pricing Effect

Explanatory variable	All Members (1)	All Members (2)
Camouflage Price Group * US Sign-Up	-0.164*** (0.049)	
Non-US Sign-Up	-0.174** (0.079)	
Effective Camouflage Pricing	—	-0.166*** (0.047)
Below-Threshold Price	-0.000 (0.047)	-0.000 (0.047)
Switched Price	0.070 (0.062)	0.071 (0.061)
Standard Membership	-0.354*** (0.050)	-0.354*** (0.052)
Month 2	-1.27*** (0.05)	-1.27*** (0.05)
Month 3	-1.91*** (0.06)	-1.91*** (0.06)
Month 4	-2.22*** (0.10)	-2.27*** (0.10)
Observations (subscriber-months)	10696	10696
Mean (s.d.) dependent variable	0.427 (0.494)	0.427 (0.494)
Likelihood-Ratio Chi-Square Statistic	1761.38	1761.36

Table 2.6 Notes: Standard errors in parentheses. Significantly different than zero at 99 (***), 95 (**), and 90 (*) percent confidence.

Figure 2.2: Projected Effect of Camouflage Pricing on Retention in Year 1 of Membership

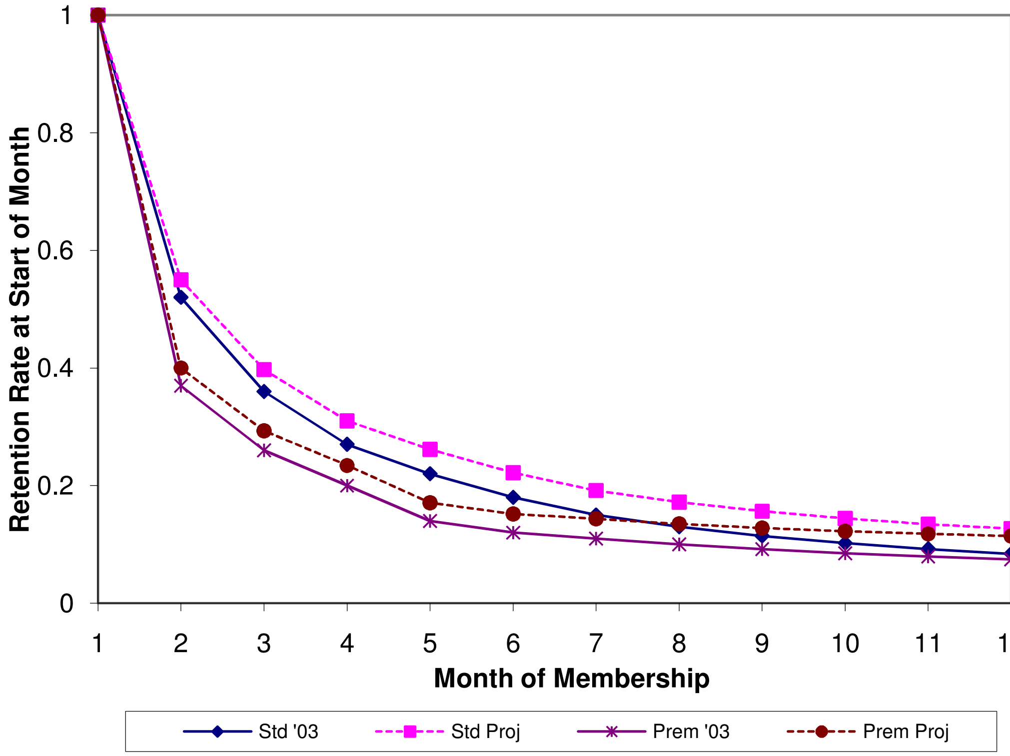

Figure 2.3b: Odds Ratio Coefficients on Cycle-Day Indicators from Daily Analysis Logistic Regression, Excluding First Billing Cycle

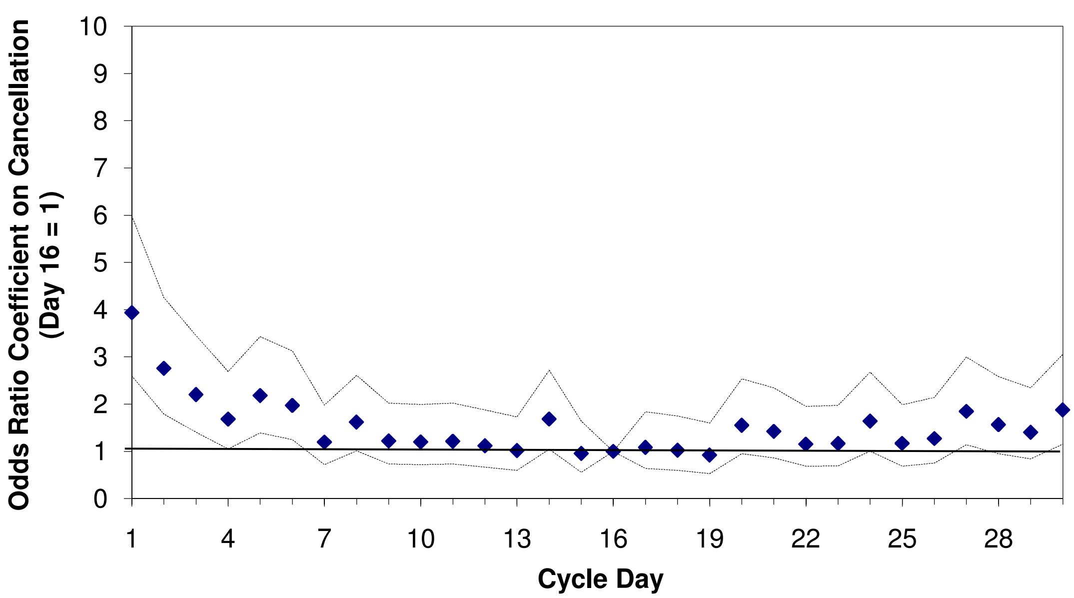

Table 2.7: Estimates for Logistic Hazard Model of Membership Cancellation, Daily Level

Explanatory variable	All Members (1)	US Members (2)	Non-US Members (3)
Camouflage Price Group	-0.082**	-0.096***	0.050
	(0.033)	(0.035)	(0.112)
Below-Threshold Price Group	0.009	0.000	0.065
	(0.033)	(0.035)	(0.112)
Switched Price Group	0.014	0.005	0.119
	(0.044)	(0.046)	(0.144)
Standard Membership	-0.322***	-0.332***	-0.253**
	(0.033)	(0.035)	(0.112)
Monthly Fixed Effects	Yes	Yes	Yes
Cycle-Day Fixed Effects	Yes	Yes	Yes
Day-of-Week Fixed Effects	Yes	Yes	Yes
Observations (subscriber-days)	259,387	234,906	24,481
Mean (s.d.) dependent variable	0.017	0.018	0.017
	(0.132)	(0.133)	(0.129)
Coefficient on Camouflage Price Group Converted to Odds Ratio	0.920**	0.907***	1.052
	(0.030)	(0.031)	(0.118)
Implied Mean Percentage-Point Change in Cancellation Rate from Camouflage Pricing	-0.13**	-0.17***	0.08
	(0.05)	(0.06)	(0.18)

Table 2.7 Notes: Standard errors in parentheses. Significantly different than zero at 99 (***), 95 (**), and 90 (*) percent confidence.

Appendix C: Figures and Tables Supplement for Chapter 3

Figure 3.3: eBay Sales of 2000 SP Authentic Boxes, 1/23/04-2/11/04

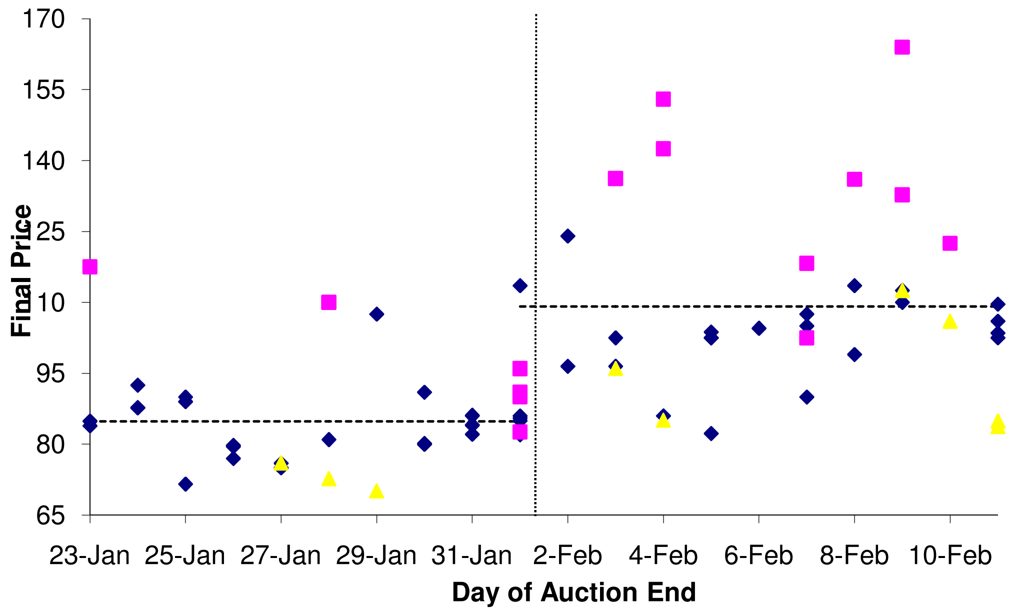

<u>Figure 3.3 Notes</u>: Diamonds represent sealed boxes w/o "Tom Brady" in item title (51 total); triangles areunsealed boxes w/o "Tom Brady" in title (10 total); and squares are sealed boxes w/ "Tom Brady" in title (17 total). The trendlines are avg for diamond series pre-Super Bowl ($85.22) and post-Super Bowl ($105.85).

Figure 3.4: eBay Sales of All Other 2000 SP Authentic Rookie Cards, 1/23/04-2/11/04

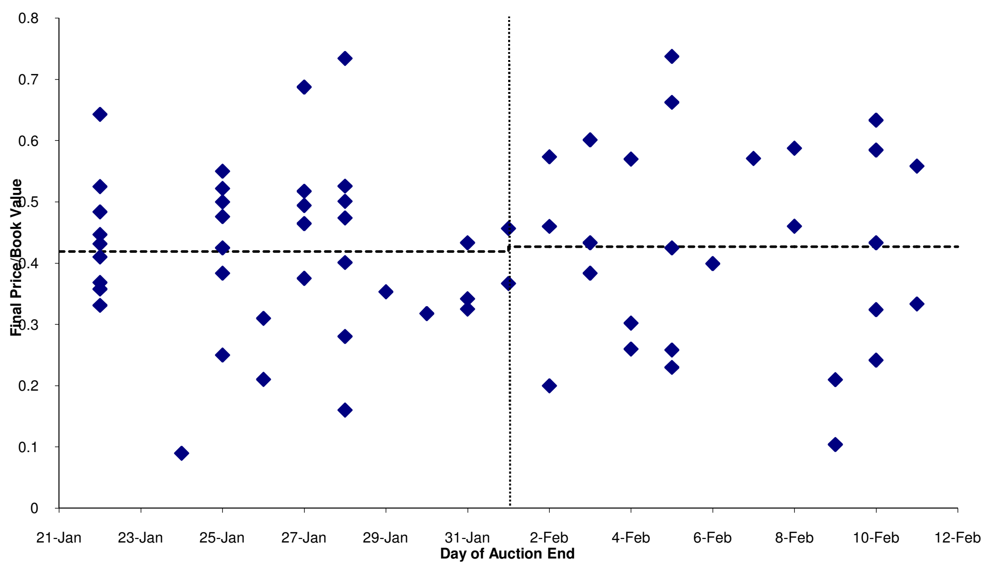

<u>Figure 4 Notes</u>: 66 total sales. The trendlines represent the average final price/book value ratio pre-Super Bowl (0.42) and post-Super Bowl (0.43).

Figure 3.5: Coefficients on Dummy Variables from First-Stage Specification for Tom Brady Rookie Card and Unopened Box Auctions on eBay, 1/22/04-2/11/04

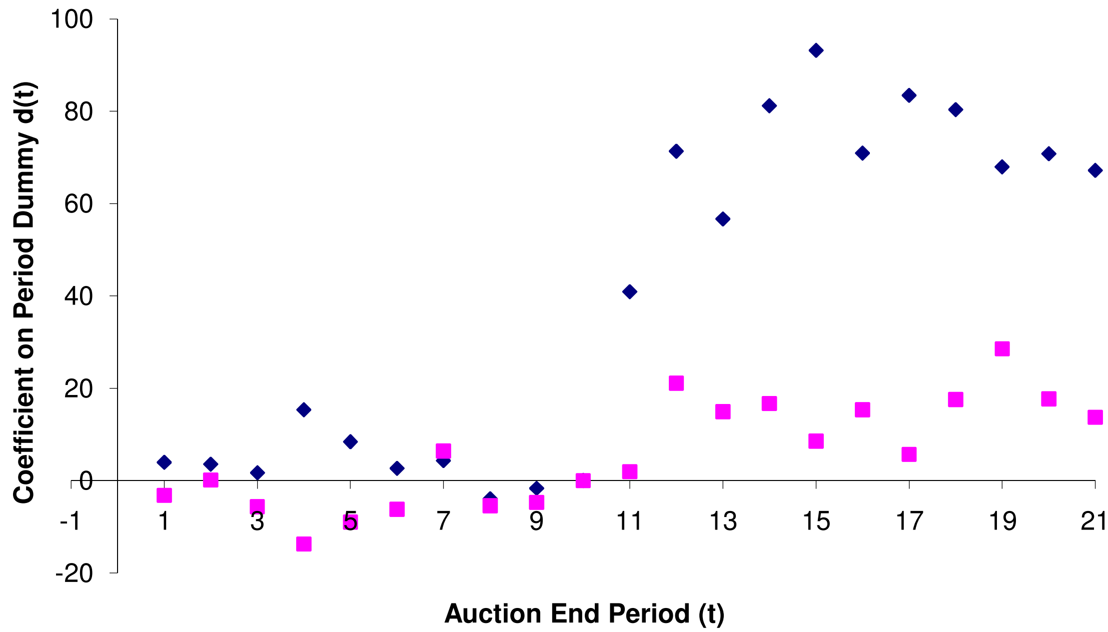

Figure 3.5 Notes: Diamonds represent dummy coefficient values for Tom Brady card (Table 3.2 regression 6), squares represent dummy coefficient values for unopened boxes (Table 3.3 regression 7).

Table 3.4: Second-Stage Regression Results for Tom Brady Rookie Card and Unopened Box Auctions on eBay, 1/22/04-2/11/04

Explanatory variable	Dependent variable: Coefficient Value on Daily Dummy Variables in Auction Sales of: Brady Card (1)	Box (2)
Auction Ended in Post-SB Period	70.8***	72.5***
	(3.6)	(4.8)
Auction Ended Night of SB	37.4***	36.3***
	(2.1)	(8.3)
R^2	0.95	0.76
Mean (s.d.) dependent variable	38.9	5.7
	(36.2)	(11.5)
95% confidence interval on Post-SB Indicator	[63.3, 78.3]	[62.5,82.5]

Notes for Table 3.4: All regressions are OLS with N = 21. Heteroskedasticity and autocorrelation consistent standard errors are calculated using the Newey-West estimator with 3 lags (in parentheses). Significantly different than zero at 99 (***), 95 (**), and 90 (*) percent confidence.

Table 3.5: Summary of π(1/60) Point Estimate and Confidence Interval

Study	One-Stage (1)	Two-Stage (2)
Estimated post-SB jump in price of 2000 SP Authentic Tom Brady card	70.6^{***} (3.6)	70.8^{***} (3.6)
Estimated post-SB jump in price of 2000 SP Authentic unopened box	19.3^{***} (2.6)	20.1^{***} (2.3)
Assumed covariance in post-SB price jumps of Brady card and unopened box	0	-4.5
$\hat{\pi}$ *(1/60)*	0.27	0.28
Est. Var($\hat{\pi}$ *(1/60))*	0.0015	0.0018
95% confidence interval on $\hat{\pi}$ *(1/60)*	[0.19,0.35]	[0.19,0.37]

Table 3.5 Notes: Row 1, column 1 estimate from Table 3.2, regression 1; row 1, column 2 estimate from Table 3.4, regression 1; row 2, column 2 estimate from Table 3.3, regression 1; row 2, column 2 estimate from Table 3.4, regression 2. Heteroskedasticity and autocorrelation consistent standard errors are calculated using the Newey-West estimator with 3 lags (in parentheses). Significantly different than zero at 99 (***), 95 (**), and 90 (*) percent confidence.