Release of Information with Imperfect Memory

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Ioannis Sarafidis

Dissertation Director:Benjamin Polak

December 2001

UMI Number: 3030823

Bell & Howell Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

Release of Information with Imperfect Memory

Ioannis Sarafidis
Yale University
2001

Recent experimental evidence in psychology and neuroscience has established that the principles of similarity and repetition govern the recall process for episodic memories. I use these two principles as building blocks in developing a formal model of memory and explore its economic implications for a wide range of economic and social settings, including political campaigns, news management prior to IPOs, marketing for new products, employee-performance evaluations and public opinion formation.

In chapter 1, *What have you done for me lately? Release of Information and Strategic Manipulation of Memories* I start by deriving the memory model. I then apply the model to an economic setting, by addressing the issue of how one should time a fixed number of informative events, in order to manipulate the memories that a forgetful assessor will eventually have. I show that the spacing of events is crucial for what agents will remember and I characterize the spacing properties of optimal profiles. The theoretical results translate to normative claims that can be exploited by, among others, politicians involved in election campaigns, advertisers timing the airing of commercial spots, managers controlling the release of corporate news and employees timing their effort prior to a promotion decision.

Chapter 2 extends the model of the first chapter by relaxing the assumption that

the agent times a fixed number of events. Instead, the agent can generate events at some cost. Such an extension widens the applicability of the model and it paves the way for future empirical testing of the model. I show that the driving forces behind the two models may be different, but the following result is true in both models: favorable past events, possibly stochastic, will make the agent more eager to release more favorable events.

The dissertation concludes with chapter 3, *Revising Non-Additive Priors*, which considers the problem of updating a convex capacity upon receipt of a signal. Convex capacities arise in decision theory in an effort to model the Ellsberg paradox: the psychological finding that people are overly averse to uncertainty.

Acknowledgments

I would like to express my gratitude to my advisors Ben Polak and David Pearce. Ben taught me decision theory, spent many hours with me discussing my projects, guided me to the right direction (sometimes despite my stubbornness), showed me how to write and present my ideas, helped me control my enthusiasm and encouraged me when I suffered set backs. I thank David for useful discussions and for creating a truly perfect research environment. The Tuesday seminar he run introduced me to behavioral economics and was instrumental in getting me started with my research.

I have also benefited greatly from discussions with my colleagues Ettore Damiano, Amil Dasgupta, Jason Draho, Nick De Roos, Ricky Lam and Mario Simon. My interaction with them not only improved the material in this dissertation, but it also broadened my understanding of economics in general.

Finally, I would like to thank my mother, Lisa and all my friends for their love. Without them, no accomplishment would have any meaning.

Contents

Introduction

The underlying motivation behind this dissertation is an effort to bring insights from psychology into economics. Until very recently economists had ignored the possibility of psychological factors affecting decision making processes, because agents were assumed to be rational. However, in recent years there has been a growing interest in bounded rationality, the idea that people's ability to make rational choices is limited by cognitive limitations and psychological biases. The essays in this dissertation belong to this vein of research, whose goal is to produce more realistic models by injecting psychological insights about human behavior and cognition into economics.

In chapter 1, *What have you done for me lately? Release of Information and Strategic Manipulation of Memories,* I use assumptions based on recent experimental evidence from cognitive psychology to build a formal model of memory. I then explore the model's economic implications by addressing the issue of how one should time a sequence of informative events in order to manipulate the memories of one's forgetful assessor. I believe that this is an important problem, because in a wide range of economic interactions, agents are rewarded at some critical date on the basis of an assessment of their past performance. In many such cases, an objective criterion that summarizes past performance is not available and, as a result, assessors have to rely on their memories of past informative events, which are not perfect.

To motivate the results, assume that an incumbent senator faces re-election at some future date. Public support for the senator depends on the electorate's memories of past events pertaining to the senator, such as what side he took in a controversial dispute, or how he handled a labor union crisis. Suppose, that a few months before the election, our senator is lucky enough to get a series of positive boosts to his image from a number of recent events. Now, he has to schedule the announcement of a new popular tax plan and a public appearance that will generate a lot of positive publicity. You have just been hired as his political consultant. How do you advise him?

At first glance the problem of information release with imperfect memory seems trivial; the politician should simply release all good news close to the election date so that they are memorable. However, the experimental evidence shows that the mechanisms behind human memory are more complicated, and the problem of releasing information becomes more interesting. In particular, the experimental evidence shows that memory operates on the principles of similarity (cue dependence) and repetition (rehearsal). Loosely, cue-dependence refers to the phenomenon that current events trigger memories of similar past events, and rehearsal refers to the fact that recalling the memory of an event makes it more likely to be remembered in the future. With cue-dependance and similarity in mind, the politician now faces a trade off. On the one hand, since people forget, it is wise to time these two events as late as possible in order to make them more memorable at the time of the election. I refer to this as the recency effect. On the other hand, he wants these events to reinforce the existing good memories of past events. To maximize this rehearsal effect, the two events should be scheduled as early as possible, to reinforce the "good news", before these memories fade away. Optimal decisions are dictated by this trade off between the recency and the rehearsal effects.

I find that, as this sketch suggests, the timing of events is crucial for what the assessors remember. In quasi-rational models, it is often the case that the order of informative signals is important, since it determines how people will perceive future events. In this model, what is important is not only the order (i.e., which event comes first and which comes second), but the actual spacing (i.e., the amount of time that elapses between the two events). The shorter the amount of time between two events the more effective will the second one be in triggering the memory of the first. This logic dictates that good news, referred to as *successes,* should be bunched together in order to reinforce each other's memories and, by the same token, bad news, referred to as *failures,* should be spread apart. In addition, I show that the optimal rule for when to release information has a nice simple form. Loosely speaking, we can think of a streak of stochastic events (both good or bad) creating a "stock" of news. When this stock is above a certain threshold level, it triggers the agent to try to generate further successes. These successes reinforce fresh memories of past successes and they guarantee that past failures will *not* be reinforced.

The results can be applied to executives managing news prior to an IPO, to advisors managing political campaigns, to advertisers scheduling the airing of commercial spots prior to a product or movie release, to employees timing their efforts prior to a promotion decision, and to students timing their class contributions prior to the end of term. An application particularly close to home is the timing of publications and seminars by junior faculty prior to a tenure decision or by a graduate student prior to the job market.

The second chapter extends the model to allow for a convex cost for generating additional events and a continuous choice variable. An extension along these lines allows us to apply the model to settings where this new set of assumptions is more relevant. Consider a worker who will be evaluated at some future date on the basis of his past output performance. Each period he can increase the probability of high, as opposed to low, output level by exerting more effort. In this case, there is no reason to believe that working hard today will deplete the amount of effort that you can exert tomorrow. Instead, assuming that effort is costly is more appropriate. I find that the driving forces behind this model and that of chapter one are very different. The assumption that there is no fixed budget is crucial, because there is no longer a trade off between working hard *now or late*r. However, I find that these two different models give rise to a similar result, that the higher the stock of past news the more motivated the agent will be to deliver a new success.

The third chapter, *Revising Non-Additive Priors*, is joint work with Ricky Lam. We consider the problem of updating a convex capacity upon receipt of a signal of known additive likelihoods. Convex capacities arise in modeling the psychological finding that people are uncertainty averse, also known as the Ellsberg paradox. For motivation, consider an employer who has a subjective prior over the quality of a worker and who knows the distribution for output conditioned on each level of quality. How does she update her beliefs regarding quality upon observing some output level? If her prior is additive, this problem is trivial: Bayes' rule suffices. This involves two steps: first, calculate the probability measure on the product space, of pairs of quality and output, and then use Bayes' rule to calculate the posterior beliefs for quality.

When the employer's beliefs are represented by a capacity, calculating a measure over the product space is no longer so simple. We propose two rules: the first uses the idea of Choquet integration over identity functions and produces a non-additive measure over the product space; the second converts the initial non-additive measure to a set of additive priors, and then applies Bayes' rule to each element in this set. We show that these rules are related but are not equivalent. We argue that their non-equivalence highlights a limitation of non-additive measures. While this limitation does not matter for the representation of uncertainty-averse preferences, it results in a loss of information when beliefs have to be revised.

Chapter 1

What have you done for me Lately? Release of Information and Strategic Manipulation of Memories.

The completion of this project would be impossible without Ben Polak's constant guidance and emotional support. I would also like to thank Michael Boozer, John Geankoplos and David Pearce, for useful discussions. I have also benefited from discussions with my classmates Ricky Lam, Jason Draho, Nick de Roos, Mario Simon and Ettorre Damiano. Financial support from the Cowles Foundation in the form of an Anderson Fellowship is gratefully acknowledged.

1.1 Introduction

In a wide range of social and economic interactions, agents are rewarded at some critical date on the basis of an assessment of their past performance. In many such cases, an objective criterion that summarizes past performance is not available and, as a result, the agents that provide the assessment, referred to as the *assessors*, have to rely on their *memories* of past informative events. For example, assume that an incumbent senator faces re-election at some future date. Then, public support for the senator depends on the electorate's memories of past events pertaining to the senator, such as what side the senator took in a controversial dispute, or how he[1] handled a labor union crisis. In a labor market setting, the decision of whether to promote an associate consultant to a VP, or what his bonus should be, may depend on the memories that his superiors have over his past accomplishments. Similarly, the success that a new SUV model will enjoy when it is launched depends on the effectiveness of the advertising campaign that was used to support it.

This paper takes a bounded rationality stance, assuming that imperfect memory is an indisputable limitation of human cognition. Using assumptions motivated by research in cognitive psychology and neuroscience, I build a formal model of memory and then explore its economic consequences by addressing the issue of how one should time the sequence of informative events in order to manipulate the memories of one's

[1]Throughout, male pronouns are reserved for the agents under assessement and female ones for their assessors.

forgetful assessor. The results apply to any social or economic setting where products or people undergo periodic assessment.

Motivation

To fix ideas, return to the example of a senator facing an election. If people had perfect memory, the timing of past events would be unimportant. With imperfect memory, however, different timing profiles of past events rank differently in creating a positive image for the senator and making re-election more probable. Furthermore, assume that the senator can time some events at his discretion, for example he can choose when to schedule a public appearance on a TV show. How should he, then, time the sequence of the events which are at his discretion?

Naively one may think that the problem of releasing information to a forgetful assessor is as trivial as that: *successes* or *good news* (i.e., events whose memory increases final payoffs) should be scheduled close to the critical date of assessment so that they are more memorable and *failures* or *bad news* (i.e., events whose memory decreases payoffs) should be released as early as possible so that their memories have sufficient amount of time to decay. However, recent findings in cognitive psychology suggest that the workings of human memory are more intricate and thus the problem of releasing information becomes more complicated and more interesting than it first appears.

In particular, experimental evidence shows that memory operates on the principles of similarity and repetition. In a few words, more are to be found in the next section, similarity refers to the phenomenon that current events trigger memories of similar past events and repetition refers to the fact that recalling the memory of an event makes it more likely to be remembered in the future. In the psychology literature, the formal term for similarity is *cue dependence* and that for repetition is *rehearsal*. To see how cue dependence and rehearsal matter assume that a few months before the election date our senator is lucky enough to get a series of positive boosts to his image from a number of recent events. In addition, he has to schedule the announcement of a new popular tax plan and a public appearance that will generate a lot of positive publicity. You have just been hired as his political consultant. How do you advise him? On the one hand, since people forget it is wise to time these two events as late as possible in order to make them more memorable at the time of the election. I refer to this as the *recency effect*. On the other hand, you want these events to reinforce the existing good memories of past events. To maximize this *rehearsal effect* the two events should be scheduled as early as possible, to reinforce the existing memories of "good news", before these memories fade away. I show that optimal decisions, from the point of view of the agent (senator), are dictated by this trade off between the recency and the rehearsal effects. Describing the origins and the properties of this trade off is the core of this paper.

Results and applications

I find that the timing of events is crucial for what the assessors remember. In quasi-rational models, it is often the case that the order of informative signals is important, since it determines how people will perceive and interpret future events. See, for example, the story model of learning due to Lam (2000) and the model of confirmatory bias of Rabin and Schrag (1999). In this paper, what is important is not only the order (i.e., which event comes first and which comes second), but the actual *spacing* (i.e., the amount of time that elapses between the two events). The shorter the amount of time the more effective will the second event be in triggering the memory of the first one. This logic dictates that successes should be *bunched* together in order to reinforce each other's memories and, by the same token, failures have minimal effect when they are *spread apart* over a long period of time.

In addition, I show that the optimal rule for when to release information has a nice simple form. Loosely speaking, we can think of a streak of events (both good or bad) creating a "stock" of news. When this stock is above a certain threshold level, it triggers the agent to try to generate further successes. These successes not only reinforce fresh memories of past successes, but also guarantee that past failures will not be reinforced.

In an attempt to make the set up of the model less abstract, the paper develops the main ideas in the election context discussed above. However, the theoretical results can be applied normatively not only to advisors managing political campaigns, but to executives managing news prior to an IPO, to advertisers scheduling the airing of commercial spots prior to a movie or product launch, to employees managing their effort prior to a pending promotion decision and to students timing their class contribution prior to the end of the term. An application particular close to home is the timing of publications and seminars by junior faculty prior to a tenure decision or by a graduate student prior to the job market.

Related Literature

This paper belongs to the vein of economic literature that injects psychological insights on human behavior and cognition into existing economic models, in an attempt to improve their descriptive power. Rabin (1998) outlines the goals of the general research program. Some implications of imperfect memory for economic decisions have been explored in Dow (1991) and Hirshleifer and Welch (1997). Also, in a close predecessor to this paper, Mullainathan (1997) builds a model of memory using the same principles of rehearsal and cue-dependence[2]. In this line of work the focus is on an agent's *own* memory imperfections and on how these can explain often-observed decision making biases or empirical puzzles, such as inertia, impulsiveness, the curse of knowledge and over/under reaction in financial markets.

What differentiates my paper from previous work on memory imperfections, as well as from most of the existing literature in behavioral economics, is my investigation of the strategic considerations of bounded rationality. Such considerations arise when a fully rational agent recognizes a cognitive limitation or bias, from which *others* suffer, and he now faces the problem of how to modify his actions in order to manipulate, or correct, it for his own benefit. In this paper, for example, the senator recognizes that the electorate suffers from imperfect memory that obeys some principles and he times the sequence of events in an effort to manipulate the memories that the electorate will have at election time. From that point of view, the spirit of this paper is similar to Rabin and O'Donahue (1999). They consider the problem faced

[2]Mullainathan (1997) uses the alternative term, "associativeness".

by a principal in a moral hazard setting, who having recognized that the agent suffers from time inconsistency, has to identify the appropriate contract that will prevent the agent from procrastinating.

The ideas of the paper are organized as follows; section two summarizes the psychological evidence that supports the assumptions on which I build the model of memory. Section three is the core of the paper, where I develop the model and derive the results. Section four extends the model to advertising campaign applications. Section five concludes and discusses possible extensions. An appendix contains the mathematical proofs of the results.

1.2 Experimental Evidence on Memory

In this section, I present an array of experimental evidence to motivate and justify the assumptions, which I later use to build a formal mathematical model of memory. The following summary is largely based on Schacter (1996) and Parkin (1993). Even though, both of these sources have a strong emphasis on the work of theorists adhering to the so called "Toronto School", in this paper I use that part of their work which is fairly non-controversial and shared by most memory researchers.

Until 1960 little was known. Memory is, for the most part, a subjective experience and as a result its study fell out of the study domain of behaviorism, the dominant trend in psychology during the first half of the century. However, during the last forty years cognitive psychologists, neuroscientists and clinicians working with amnesic patients have joined forces and have begun to outline some principles that govern human memory. Two of the stylized facts that have emerged, and are relevant for our purposes, are that a) *rehearsal* of the memory of an event facilitates its subsequent recall and b) memory is a *cue dependant* process, which implies that current events can trigger past memories of similar events.

Rehearsal

Clearly, the most important determinant of memory is time. The idea that the passing of time erodes memories is indisputable, as it is supported not only by casual observation and introspection, but by controlled experimental evidence. As early as 1885, Hermann Ebbinghaus used himself as a subject and nonsense syllables as the target stimuli to study the rate of forgetting. He showed that forgetting is fast at first but it declines as time passes.[3]

However, the passing of time is not always an enemy of memory. Vast experimental and clinical evidence suggests that repeated exposure to the memory of an event, referred to as *rehearsal*, increases its memorability. In fact, with adequate rehearsal memories of distant events can be stronger than those for relative recent ones. In a

---

[3]Ebbinghaus used nonsense syllables, because, otherwise he reasoned, previous knowledge would contaminate the recall process. Critics have pointed out that such a task lacks any ecological validity and can not be fully representaive of human memory. In an effort to examine the effects of time on autobiographical memories, Crovitz asked subjects to recall a specific event from their past when presented with various words, such as *table*. As expected, few of the memories come from the distant past, with the majority of memories coming from recent periods. This procedure is widely used in clinical environments and is known as the *Crovitz technique*.

typical experiment, two groups of subjects are presented with some stimuli objects, such as pictures or words. After a short delay the first group is given a recall test. The second group is given the recall test after a longer delay, but they are guided into rehearsing the stimuli. It turns out that recall is superior for the second group, even though, more time has elapsed between the initial exposure to the stimuli and the moment of recall. In some cases, not only do distant events have a higher probability of recall than recent ones, but memorability for a single event increases over time as a result of rehearsal, a phenomenon referred to as *hypermnesia*. For example, in the above type of experiment, it has been documented that when subjects are given a series of recall tests, memory improves after every test.

The way rehearsal increases memorability is also evident in the so called "Serial Position Curve". Subjects are presented with a list of words to be remembered. When the position of the word in the list is plotted against the probability of recall, we obtain a U-shaped curve. As expected, words at the end of the list have high probabilities of recall, (*recency* effect), but quite surprisingly words in the beginning of the list have higher recall probabilities than words in the middle, (*primacy* effect). Psychologists attribute the primacy effect to the fact that words in the beginning of the list are rehearsed more than subsequent words. This hypothesis was examined by Rundus (1971) who required his subjects to rehearse words out loud. When plotting the serial position against the average number of rehearsals he obtained a downward sloping curve, which fits nice to the hypothesis that the counterintuitive primacy effect is attributed to rehearsal.

Brain physiology provides a different angle from which to explain the effects of rehearsal on memorability. Neuroscientists have established that human experience is fed into the brain by generating new synapses between neurons and by increasing the strength of already existing ones. Quite naturally then, a memory-experience that is encountered many times, due to rehearsal, should result in more and stronger synapses and, consequently, should become easier to access and recall in the future. To verify this hypothesis, Kandel et.al (1995) experimented with the sea slug Aplysia. One of its characteristics is to react to unpleasant stimulus by withdrawing its gill, with the response disappearing after a few minutes. However, as the experimenter increases the number of unpleasant stimuli, the slug withdraws its gill for longer periods of time. This suggests that repeated exposure to the stimuli makes its memory more resistant to time. More importantly, they show that this effect arises because as the number of shocks increases, an enhanced release of a neurotransmitter strengthens the connection between the neurons receiving the shock and the ones responsible for the gill withdrawal.

Finally, clinical evidence also highlights the important role of rehearsal. It is common for amnesic patients to suffer from severe loss of memory for recent events either prior (retrograde amnesia) or after (anterograde amnesia) the onset of their condition, only to have no difficulty in recovering memories from their distant past. A possible interpretation according to Schacter, is that "people talk about and think about their past experiences; the older the memory, the greater the opportunity for such post-event rehearsal."

Memory as a cue dependant process

A second theme that has emerged from memory research is that the present environment, under some conditions, functions as a cue that facilitates recall of similar past events. Watching a Brady Bunch episode, for example, can be very effective in stimulating recall of long buried childhood memories.

More formally, Tulving and Osler (1968) conducted the following four phase experiment. In phase 1 they presented subjects with a list of target words (engine) paired with a closely related cue word (black). In phase 2, subjects were given a list of words that were very closely associated with the target words of phase 1, (steam), and subjects were asked to produce words they associated with these new phase 2 words. Of course there was a high probability that subjects would generate the phase 1 words. In phase 3, subjects were asked to look at the list they produced in phase 2 and determine which words they recognize as the target phase 1 words. Finally, in phase 4, subjects were presented with the cue words of phase 1 and asked to recall the target words with which they were paired. Quite surprisingly it was common for subjects to recall words in phase 4, but they failed to recognize these same words in phase 3.[4]

---

[4]The idea that memory is a cue dependant proccess goes back to Semon in 1904. However, mainly because of his insistance to link memory to evolution his ideas were ignored. Tulving rediscovered Semon's original idea in an effort to disprove recognition generation models, according to which memory operates on two stages; first the brain *generates* a list of possible candidates, and then each candidate undergoes a recognition process. If that was the case, it would not be possible for subjects

To explain the results of their study, Tulving et. al. argued that memory is a *cue dependant process*, where current events and experiences can bring to mind past memories, even memories of distant events that appeared to be forgotten. Furthermore, Tulving and his colleague Thompson (1973) put forward the *encoding specificity principle* according to which, the way we approach and encode an event when it occurs determines what type of cue will be effective in facilitating its subsequent recall. When there is a high degree of similarity and overlap between the way we encoded an experience (or event) and the present environment, then recall will be superior.

However, similarity between the present and the past, is not always defined in a literal sense. What differentiates an effective from an ineffective cue is not always literal similarity between the cue and the target, but whether the cue is successful in reproducing the same type of emotion or mental image as the original episode. For example, someone with my greek background can be elicited to recall an episode form Homer's "Iliad" by the mention of the word "Somalia". The reason is that, as to most Westerners, the mention of the word Somalia brings to mind the well publicized scene of Somalis dragging the body of an American soldier in the streets of Mogadishu, just like Achilles dragged the body of Hector behind his chariot in Troy.

To summarize, overwhelming experimental evidence from cognitive psychology suggests that memory operates on the principles of repetition and similarity. Memories which are encountered repeatedly become increasingly accessible for future recall and current events function as cues to facilitate recall of similar past events. These two

---

to fail to recall words in phase 3 only to recall them in phase 4.

well documented stylized facts, referred to as *rehearsal* and *cue dependance,* provide the basic ingredients with which I build the formal model of the next sections.

Some additional remarks

The discussion up to this point may have given the false impression that memory imperfections are limited to our inability to remember or recognize events that occurred in the past. However, psychology research with enormous judicial and social repercussions reveals that humans suffer from a different memory imperfection; that of incorrectly "remembering" events that never occurred. In the famous "remember the time that..." study Elizabeth Loftus (1993) tried to compare childhood memories of siblings. One sibling was briefed by the psychologist and would then describe a fictitious but detailed experience, such as being lost in the mall, to the other sibling, in the form: "remember the time that you were lost at the mall, and you were found by a policeman, and our mother was so mad at you...". In numerous cases subjects agreed that they remembered these fictitious events. [5] [6]

---

[5]Some unfortunate people learn about such memory imperfections the hard way when they are falsely accused of committing some hideous crime based on the testimony of eye witnesses.Schacter (1996) discusses a nototious 1978 trial, where Fank Walus was falsely accussed of being a Nazi war criminal, based on numerous testimonies that linked him to specific events. One "witness" "remembered" that Walus killed two children and their mother, another how Walus brutalized his father.

[6]Along the same lines, there is strong evidence to suggest that the way we approach the past may have an effect on what type of memory we derive from a certain event. Consider for example the last time you were at a party. If you see yourself in the picture you have an *observer* perspective, as

Such findings may be perplexing for those who want to believe, that the brain is like a computer hard drive storing events as they occur, with memory being a flash light that enables us to search through and access the contents of the hard drive. Despite its intuitive appeal, such a metaphor, which implies a one to one correspondence between stored memories and realized events, is not correct. Rather, some psychologists theorize, memory relies heavily on *reconstruction*. In other words, memories, instead of being actual snapshots of reality, they are only fragments of the past. These fragments of the past interact with previous knowledge and elements of the present, such as our desires, biases and motives to produce the experience of a memory.

As a result, a complete model of memory should ideally focus not only on factors that affect recall of past events, but factors that could probably lead to inventing memories of events that never occurred, or memories that are only partially faithful to reality. I see the absence of these factors from the model of the next sections as a limitation. However, since psychologists have not yet identified any undisputable principles and stylized facts behind false memories formation, as they have done with the processes of encoding and recall of real memories, it is a limitation we have to

---

opposed to a *field* perspective in case you see the scene the same way you saw it when it occurred. In experiments people assume a field perspective for emotionally charged memories and an observer perspective for more objective type of episodes. More surprisingly however, when you direct people to switch from the field to the observer perspective, these same memories have a reduced emotional intensity.

live with.

1.3 The Model

To fix ideas, return to the example of the introduction, where an incumbent senator faces re-election at some future date T. During the period leading to the day of the election, T, the electorate observes a series of events pertaining to the senator's popularity, which are assumed to be of identical importance and strength, and for the moment, only beneficial to the senator's image. Such events could be for example, how he used his negotiating skills to avoid an airline strike on Thanksgiving, an announcement for a state tax cut, his instrumental role in passing a popular legislation, etc. Re-election is assumed to be more probable the more of such events the electorate remembers at date T. When memory imperfections is not an issue the timing of these events is immaterial. However, if the electorate forgets, different timing profiles rank differently in making re-election more probable. We can ask, for example, whether re-election is more probable when a certain number of events are spread apart in regular intervals, or when they are all bunched together sometime in the middle. More interestingly, if we are in a stochastic environment where events are generated by some random process, but at the same time the senator has some control on the timing of some events, how should he then, or his campaign team, time the release of those events which are at his discretion?

The first step in addressing these issues is to model explicitly the memory technology. The next subsection formalizes the problem of information release just described and develops a simple model of memory, that incorporates the findings of the experimental evidence discussed earlier. For the moment, I restrict attention to the case where all information comes in the form of *good news,* i.e. events whose memories increase final payoffs. Later on, I generalize to the case of both good and bad news.

1.3.1 The Basic Set up

Time is discrete and indexed by $t = 1, 2, 3..., T$. In each period $t \leq T$, an event e_t may be realized, in which case we say we have a *success* and we write $e_t = 1$, as opposed to $e_t = 0$ when the event is not realized, that is $e_t \in \{0, 1\}$. At the end of the last period T, we have obtained a *history* $e = \{e_1, e_2, ..., e_T\}$ of past realizations of the events e_t.

To model that recall is imperfect, let M_t^i be the memorability, or strength, of the memory $e_i = 1$ at time t. Since a realization of an event can have meaningful memorability only after it has occurred, $M_t^i = 0$ for all $i > t$, and we normalize $M_i^i = 1$. At time T, an agent (referred to as the *assessor*) who has been observing the realizations of past events remembers a period i success only with some probability, using the event $e_i = 0$ as the default memory. I refer to this probability as the *recall probability* for success i, and I assume it is some increasing function of the memorability of success i at time T, M_T^i.

By modeling the memory technology this way, we implicitly assume that the assessor can only forget successes that actually happened and can not invent memories for successes that never happened. Also, in using $e_i = 0$ as the default memory in case a period i success is forgotten, we interpret the event $e_t = 0$ as a non- event, i.e., if $e_t = 0$ then nothing happened that period. Later on, I generalize to the case where the event $e_t = 0$ is interpreted as the agent *failing* to deliver a success.

Next, I invoke the experimental results on memory, discussed in the previous section, in order to model the evolution of M_t^i over time. The fact that memories fade away with the passing of time dictates that M_t^i decreases with time. In particular, I assume that memories decay exponentially at a constant rate $(1 - \rho)$. More interestingly, however, the experimental evidence suggests that similarity and repetition increase memorability. In the context of our simple model, this translates to the fact that a successful realization today can trigger past memories of successes. As these past memories get triggered and rehearsed they become more memorable. This is incorporated into the model by assuming that a success at time t, will enhance the memorability of past successes, the enhancement being bigger for more recent successes. To formalize this, I define b_t^i to denote the enhancement (boost) to the memorability of success i by a possible success at time t.[7] The size of the increment to the memorability of an event depends on the time that has elapsed between the event itself and the later similar event that rehearses it. To capture this I assume that the rehearsal of an event from time i by a similar event at time $t+1$ is a fraction $\kappa < 1$ of the rehearsal of that event by a similar event at time t; that is $b_{t+1}^i = \kappa b_t^i$. It is also natural to assume that the increment to the memorability of a period t success

[7] As with M_t^i, $b_t^i = 0$ for all $i > t$ and we normalize $b_i^i = 1$.

from a success at $t+1$ is less than its current memorability thus $\kappa < \rho$[8]. As long as $i < t$, all this can be neatly summarized in the following pair of equations,

$$M_{t+1}^{i} = \rho M_{t}^{i} + b_{t+1}^{i} I\{e_{t+1} = 1\} \tag{1.1}$$

$$b_{t+1}^{i} = \kappa b_{t}^{i} \tag{1.2}$$

where I is the usual indicator function, i.e. $I\{e_t = 1\} = 1$ if $e_t = 1$ and 0 otherwise.

Recall that the number of successes that a forgetful assessor actually remembers is a random variable, with upper bound the actual number of successes, $\sum_{i=1} I\{e_i = 1\}$. If the probabilities of recall are proportional[9] to the memorabilities at time T, M_T^i, then the *expected* number of successes that the agent recalls at time T is proportional to the *sum of the memorabilities* at time T, denoted by $A_T = \sum_{i=1}^{T} M_T^i$. Using (1.1) and (1.2) together with the initial conditions $M_i^i = b_i^i = 1$ and summing over i we

[8]Note that at time $t+1$ the memorability of a period t success is ρ.

[9]The memorability of a success can never be greater than $\frac{1}{1-\kappa}$, for any rate of memory decay ρ. Therefore, chosing the constant of proportionality to be $(1-\kappa)$, or smaller, guarantess that the recall probabilities are well defined.

can work out the evolution of A over time[10]:

$$\begin{aligned}
\sum_{i=1}^{t+1} M_{t+1}^i &\equiv \sum_{i=1}^{t} M_{t+1}^i + I\{e_{t+1}=1\} = \sum_{i=1}^{t}[\rho M_t^i + I\{e_{t+1}=1\}b_{t+1}^i] + I\{e_{t+1}=1\} \\
&= \rho \sum_{i=1}^{t} M_t^i + I\{e_{t+1}=1\}(1+\sum_{i=1}^{t} b_{t+1}^i.) \\
&= \rho \sum_{i=1}^{t} M_t^i + I\{e_{t+1}=1\}(1+\kappa \sum_{i=1}^{t} b_t^i.) \text{ and} \\
\sum_{i=1}^{t+1} b_{t+1}^i &\equiv \sum_{i=1}^{t} b_{t+1}^i + I\{e_{t+1}=1\} = \kappa \sum_{i=1}^{t} b_t^i + I\{e_{t+1}=1\}
\end{aligned}$$

If we define $S_t = \sum_{i=1}^{t} b_t^i$, then we can rewrite the last pair of equations as

$$A_{t+1} = \rho A_t + I\{e_{t+1}=1\} \cdot (1+\kappa S_t) \qquad (1.3)$$

$$S_{t+1} = \kappa S_t + I\{e_{t+1}=1\} \qquad (1.4)$$

Casting the model in terms of this new pair of equations allows for a different interpretation. Memory decay, embedded in ρ, can be thought of as a discount rate on the *past*, rather than on the *future*. Then, one could think of A_t as a *running total* stock variable that measures the number of successes that have occurred up to time t, with past successes being "discounted" at rate $(1-\rho)$. More importantly, however, when a new success is realized the stock A_t is increased, but the incremental effect of a new success has two components. A *direct* effect (plus 1), originating simply from the fact that a new success has been realized, plus an *indirect* effect, (plus κS_t)

[10]For notational convinience it is useful to define M_t^i and b_t^i even for events $e_i = 0$, and set them at zero for all values of t.

coming from the fact that the new success triggers memories of past successes and therefore become more memorable.

Our reduced form equations (1.3) and (1.4) give a convenient way to record this indirect effect. One might think that the indirect effects of a new success would depend on the exact sequence of zeros and ones that have occurred up to that point. In fact, however, the variable S_t acts as a summary statistic for the sequence: the greater is S_t the greater the indirect effect from a new success will be. We can think of $S_t = \sum_{i=1}^{t} b_t^i$ as a "*rehearsal stock*". Just as the running total stock A_t decays at rate $(1-\rho)$, it is as if the rehearsal stock S_t "decays" at rate $(1-\kappa)$. Our earlier assumption that the rehearsal of an event from time i by a similar event at time t is less than the current memorability of that event, translates to saying that the rehearsal stock S_t decays a faster rate than the actual stock A_t $(\kappa < \rho \Leftrightarrow (1-\kappa) > (1-\rho))$.

From (1.4) it should be straightforward to see that S is a sum of powers of κ. To see the connection between the actual vector of past realizations and the summary statistic S_t, note that the n^{th} power of κ is present in S_t if and only if a success occurred in period $t-n$. The fact that each point in time we can summarize the past in a stock variable, rather than having to carry the whole vector of past realizations, is clearly an attractive feature of the model, especially given the fact that the pair (1.3) and (1.4) is derived from the structural formulation of the M_t^i and the b_t^i equations, which are, in turn, motivated directly from the psychology literature.

The rehearsal stock S_t may rehearse a memory for readers who are familiar with models of addiction, as in Becker and Murphy (1988); growth with endogenous preferences, as in Ryder and Heal (1973); or Constantinides' habit formation (1990). In this line of work, the marginal utility of consumption depends on past levels of consumption. Similarly, in the present model, the effect of a new success depends on the stock of past memories as summarized by S_t. The more vivid past memories are, i.e. the higher the rehearsal stock S_t is, the more effective a new success becomes in triggering them. However, in this memory model the rehearsal stock S_t is not imposed exogenously. Rather, it is a variable that comes out in the reduced form (A_t and S_t equations) as a consequence of assumptions on the structural model (M_t^i and b_t^i equations). Some further similarities between our memory model with the earlier literature on addiction and habit formation, will become more apparent and further discussed in a later section when I extend the present model to cover advertising applications.

Some readers may find it useful to work through the following example.

Example 1. *Let* $T = 4$ *and* $e = (1, 1, 0, 1)$. *Let us first calculate* A_4 *using (1.1) and (1.2), and then using the reduced form expressions (1.3) and (1.4).*

By definition, $M_1^1 = 1$ *and* $b_1^1 = 1$. *After one period, memorability decays to* ρ, *but at the same time the realization* $e_2 = 1$ *boosts the memory of* $e_1 = 1$ *by* $b_2^1 = \kappa$, *so that* $M_2^1 = \rho + \kappa$. *In period three memorability simply decays to give us* $M_3^1 = \rho(\rho+\kappa)$. *The success in period 4 adds a boost of* $b_4^1 = \kappa^3$, *to the memorability that has decayed for one more period for a total of* $M_4^1 = \rho^2(\rho+\kappa)+\kappa^3$. *Similarly,* $M_2^2 = 1$, $M_3^2 = \rho$, $M_4^2 = \rho^2 + \kappa^2$, *and* $M_4^4 = 1$. *Again, by definition,* $A_4 = M_4^1 + M_4^2 + M_4^4 = [\rho^2(\rho + \kappa) + \kappa^3] + [\rho^2 + \kappa^2] + 1 = \rho^3 + \rho^2(1 + \kappa) + 1 + \kappa^2 + \kappa^3$.

Now using (1.3) and (1.4) we have $A_1 = S_1 = 1$; $A_2 = \rho + (1+\kappa)$, $S_2 = 1+\kappa$; $A_3 = \rho(\rho + 1 + \kappa)$, $S_3 = \kappa(1+\kappa)$. *Finally,* $A_4 = \rho^2(\rho + 1 + \kappa) + 1 + \kappa^2(1+\kappa) = \rho^3 + \rho^2(1+\kappa) + 1 + \kappa^2 + \kappa^3$, *as it should be.*

As final remark, in the present set up, successes do not enhance past enhancements, that is there are no second order effects from rehearsal. In the example above the event $e_4 = 1$ enhances the memorability of $e_1 = 1$ by κ^3. However, the enhancement would still be κ^3 even if we had $e_2 = 0$, which would no longer reinforce the memory of $e_1 = 1$ in period $t = 2$. It is straightforward to include such second order effects by letting $S_{t+1} = \kappa S_t + I\{e_{t+1} = 1\}(1 + \kappa S_t)$. The nature of the following results remains unchanged.

1.3.2 Optimal Profiles and Properties of the Model

Suppose now that an agent (senator) can exert some influence over the process that generates the realizations of the events e_t by making a choice c_t between acting or waiting, so that $c_t \in \{Act, Wait\}$. In case he acts he secures a success in that period, whereas, if he waits, a success occurs stochastically, with probability $p < 1$. At time T, the assessor (electorate), who has been observing the realizations of the events e_t, rewards the agent (senator) on the basis of how many successes she remembers that the agent has produced.

In addition, the agent is constrained by $\sum_{i=1}^{t} I\{c_t = Act\} = B < T$, that is he can act for at most B periods. Such a constraint could refer to the fact that the agent is "sitting" on B number of news which he can release at his discretion, as a politician who can choose when to announce a tax cut and reduction in the crime rate. Alternatively, B could be a stock of effort or energy that is depleted every time the agent acts in order to increase the probability of obtaining a success. Such an interpretation may be more relevant for an employee-performance evaluation, where an employee can affect his output quality, which can be either high ($e_t = 1$) or the usual status quo ($e_t = 0$), by exerting high ($c_t = Act$) or low effort ($c_t = Wait$). What is then the optimal action profile over time?

To a risk neutral agent, the optimal profile is defined as that profile that maximizes A_T, the aggregate memorability of past successes at time T. More formally, the agent solves the following program[11];

$$
\begin{aligned}
& \max \quad A_T && (1.5)\\
s.t.\ A_{t+1} &= \rho A_t + I\{e_{t+1} = 1\} \cdot (1 + \kappa S_t) \\
S_{t+1} &= \kappa S_t + I\{e_{t+1} = 1\} \\
\Pr(e_t &= 1 | c_t = Act) = 1 \\
\Pr(e_t &= 1 | c_t = Wait) = p \\
\text{and } \sum_{i=1}^{T} I\{c_t &= Act\} = B
\end{aligned}
$$

The initial condition S_0 may be greater than zero if at the time the agents first encounters the problem there exists already a relevant history of events.

[11]In a more general fomulation we could assume that $\Pr(e_t = 1 | c_t = Act) = q$, where $p < q \leq 1$. This complicates the arguments in the proofs but the nature of the subsequent results remains unchanged.

We can cast the above problem in a more familiar setting by letting $V(S, n, \tau)$ be the expected continuation payoff, assuming that choices are made in an optimal way. There are three state variables: (S) the rehearsal stock which summarizes the past history of events, (n) the number of periods that the agent can set $c_t = Act$ before he exhausts the constraint $\sum_{i=1}^{T} I\{c_t = Act\} = B$, and ($\tau$) the number of periods left until the end. Notice that the running total stock variable A is not a state variable, since the rehearsal stock S carries all past information which is relevant for optimal decisions. Then, at time $t = 0$ we face[12]

$$V(S_0, B, T) = \max\{V_{act}(S_0, B, T), V_{wait}(S_0, B, T)\} \quad (1.6)$$

$$\text{where } V_{act}(S_0, B, T) = \rho^{T-1}(1 + \kappa S_0) + V(1 + \kappa S_0, B - 1, T - 1) \quad (1.7)$$

$$V_{wait}(S_0, B, T) = p[\rho^{T-1}(1 + \kappa S_0) + V(1 + \kappa S_0, B, T - 1)] \quad (1.8)$$
$$+(1 - p)[V(\kappa S_0, B, T - 1)]$$

Since $\kappa < \rho$[13], there exists a textittrade off between acting in any period t rather than in a subsequent period, say $t + k$. Acting in period $t + k$ is preferred from a textitrecency point of view as successes closer to period T are more likely to be remembered. However, acting in period t is preferred form a textitrehearsal point of view, as early successes are more powerful in reinforcing fresher previous memories. The following example clarifies this observation.

Example 2. *Let* $T = 2$, $B = 1$, $p = 0$ *and* $S_{T-2} = \bar{S}$. *Here, the agent has one*

---

[12] If $n = 0$, acting is not an option, so $V_{act}(S, n, \tau)$ is defined only as long as $n > 0$.

[13] For $\kappa > \rho$, the problem is trivial and it is optimal to act the last B periods.

available action with two remaining periods. We must choose between acting early in the penultimate period $T-1$, or waiting and acting in the last period, T. If he acts early, the resulting profile will be $\bar{S}\ |1,0$, where $\bar{S}\ |$ denotes the rehearsal stock entering the decision problem. If he waits the resulting profile will be $\bar{S}\ |0,1$. The respective continuation payoffs from these profile are:

$\bar{S}\ |1,0 \rightarrow \rho(1+\kappa\,\bar{S})$

$\bar{S}\ |0,1 \rightarrow (1+\kappa^2\,\bar{S})$

Clearly, the recency effect ($\rho < 1$) favors waiting and acting the last period. The rehearsal effect, however, favors acting early. An early success adds to the continuation payoff an indirect effect of κ in period $T-1$ which decays for one period, contributing a total of $\rho\kappa$, as opposed to a late success that adds only κ^2. If $\bar{S}$ is above the threshold of $\frac{1-\rho}{\kappa(\rho-\kappa)}$, then acting early is optimal.

The logic behind this example can be generalized to the multiple period, multiple action, stochastic ($p > 0$) case, yielding a simple rule for choosing optimally. Optimal choices are driven by the trade off between the recency and the rehearsal effects. When the rehearsal stock is above a certain threshold, the rehearsal effect dominates over the recency effect and acting is preferred to waiting. The critical thresholds that trigger the agent to act depend on the number of periods to go, τ, the number of available actions, n, as well as on the decay parameters ρ and κ. When there are *fewer* periods to go, early successes are *more* attractive because there is now one less period of memory decay, and as a result a *lower* rehearsal stock is needed to induce acting early. The comparative statics with respect to the other variables are less obvious and their discussion is postponed until later. All this is formalized in the following proposition, proved in the appendix

PROPOSITION 1 *(Thresholds) For every value of number of periods to go, τ, and number of actions still available, n, with $n < \tau$, there exists a threshold rehearsal stock, $H(\tau, n|\rho, \kappa)$, such that it is optimal to act if in period t and only if $S \geq H(\tau, n|\rho, \kappa)$. Moreover, the thresholds are increasing in τ; that is the fewer the periods to go, the lower the rehearsal stock need be before it triggers the agent to act.*

On the road to describing the optimal profile in greater detail it will be useful to use the following lemma.

LEMMA *(Acting increases the stock of rehearsal) For any $\kappa < 1$, $1 + \kappa S > S$.*

The proof is straightforward. Recall that S is a sum of powers of κ. Then, S is at most $\frac{1}{1-\kappa}$. Rearranging $S < \frac{1}{1-\kappa}$ yields the result. This result can also be understood in an intuitive way. Assume that at time t we have observed a certain history of events, call it e^t. A possible success in period $t+1$ will boost all successes in e^t. The sum of these boosts will be κS_t. Assume now that we observe a success in period $t+1$. Now, a possible success in $t+2$ will boost all successes in e^t *and* the $t+1$ period success. The sum of these boosts will be $1 + \kappa S_t$. To see the result in the lemma, observe that the $t+2$ period success boosts the $t+1$ period success by (weakly) more than the period $t+1$ success boosts the most recent success in e^t. Similarly, the $t+2$ period success boosts the most recent success in e^t by (weakly more) than the period $t+1$ success boosts the second most recent success in e^t and so on.

We can now give a more precise characterization to the structure of the optimal profile. It turns out that it is never optimal to act in period t, then wait for some k periods and then act again in period $t+k+1$. In the optimal profile, the periods that the agent chooses to *Act* are *bunched* together one after the other. Therefore, once the rehearsal stock passes the threshold level for some number of periods to go, τ, and number of available actions, n, we keep on acting until the supply of available actions, B, has been exhausted.

PROPOSITION 2 *(Bunching) If it is optimal to act in period t, it is also optimal to act in period $t+1$. More formally, for any level of rehearsal stock S, number of available actions n and periods to go τ, with $0 < n \leq \tau$,*

$$V(S, n+1, \tau+1) = V_{act}(S, n+1, \tau+1) \Rightarrow V(1+\kappa S, n, \tau) = V_{act}(1+\kappa S, n, \tau)$$

The formal proof is in the appendix[14], but the following example illustrates the intuition behind the bunching structure of the optimal profile. This intuition is also pivotal in understanding the forces that drive many of the results of this paper.

Example 3. *Consider the case of $T = 3, B = 2$, $p = 0$ and $S_0 > 0$. There are two available actions and three remaining periods. We can act the first two periods, the first and the last period, or the last two periods. For each case we obtain the following*

---

[14]If we assumed instead that $\Pr(e_t = 1|c_t = Act) = q > p$, we would obatin the following similar result: Assume you act in period t and you obtain a success. Then, it is optimal to act in the $t+1$ period.

profiles together with their respective continuation payoffs

$$
\begin{aligned}
e &: S_0|1,1,0 \rightarrow \rho^2(1+\kappa S_0) + \rho(1+\kappa+\kappa^2 S_0) \\
e' &: S_0|1,0,1 \rightarrow \rho^2(1+\kappa S_0) + (1+\kappa^2+\kappa^3 S_0) \\
e'' &: S_0|0,1,1 \rightarrow \rho(1+\kappa^2 S_0) + (1+\kappa+\kappa^3 S_0)
\end{aligned}
$$

I show algebraically that e', i.e., the profile where the successes are not bunched together, can not be optimal. It will be dominated by e, e'', or both.

Assume that e' is preferred to e. Then,

$$(1-\rho) > (\rho-\kappa)\kappa(1+\kappa S_0) \tag{1.9}$$

If e' is also preferred to e'', we have

$$(1-\rho) + \frac{\kappa}{\rho}(1-\kappa) < (\rho-\kappa)\kappa S_0 \tag{1.10}$$

Now using the fact that $1+\kappa S > S$, it is immediate to see that (1.9) and (1.10) can not hold simultaneously.

Now let us recover what (1.9) and (1.10) say in words. As before, there is a *recency* effect and a *rehearsal* effect. The recency effect dictates that successes be as late as possible. The rehearsal effect has two components: rehearsal of *past* successes, which requires successes to be as early as possible, and rehearsal from possible *future* successes that works in the same direction as the recency effect.

The history profile e' can be obtained from e by moving the last success from period $T-1$ to period T. Thus, comparing the histories e and e' is equivalent to asking whether it is optimal to act now rather than one period later given a stock of $1 + \kappa S_0$. Acting early is clearly a "good" move from a past rehearsal point of view (RHS of (1.9)) but not from a recency point of view (LHS). The effect of future rehearsal does not come into play here. Similarly, the profile e'' can be obtained from e' by moving the first success one period later, which is equivalent to asking whether it is optimal to act now rather than one period later when the stock of rehearsal is S_0 *and* there is a future success in the last period. Again, acting early is a "good" move from a rehearsal point of view (first term on the RHS of (1.10)), a "bad" move from the recency perspective (LHS) and a "bad" move from the point of view of future rehearsal (second term on the LHS). To complete the argument, observe that the recency effect is of the same relative magnitude in both cases, the rehearsal effect is stronger in the first case ($1 + \kappa S_0 > S_0$), and the effect from future rehearsal exists only in the second case. Therefore, if e' is preferred to e despite the large rehearsal effect, it can not be the case that e' is preferred to e''.

It may also be useful to think of the bunching result as a statement regarding the nature of the thresholds. If an agent decides to act with τ remaining periods, the next period decision involves one less period to go *and* one less action available. We know that thresholds are increasing in the number of periods to go, τ, but we do not know anything concerning the comparative statics with respect to the number of available actions, n. Nevertheless. the bunching result allows us to describe how thresholds evolve when the number of periods to go, τ, and the number of available actions, n, *together* decrease from one period to the next. To see that, notice that

$$V(S,n+1,\tau+1) = V_{act}(S,n+1,\tau+1) \Leftrightarrow S > H(\tau+1,n+1|\rho,\kappa)$$

Having acted the stock of rehearsal has now evolved to $1+\kappa S$. Similarly then

$$V(1+\kappa S,n,\tau) = V_{act}(1+\kappa S,n,\tau) \Leftrightarrow 1+\kappa S > H(\tau,n|\rho,\kappa).$$

Combining the last two inequalities yields the following corollary.

COROLLARY 1 *The thresholds satisfy* $1+\kappa H(\tau+1,n+1|\rho,\kappa) > H(\tau,n|\rho,\kappa)$.

The result refers to the rate at which thresholds change with the number of available actions and periods to go, n and τ. In particular, as the number of available actions, n, and the number of periods to go, τ, *together* decrease from one period to the next, the thresholds do not necessarily decrease. If they increase, however, they can not do so faster than a ceratin rate.

1.3.3 The non-stochastic case

Some additional results can be proved for the case where $p=0$ that is, a success is obtained if and only if the agents acts. This problem, though a special case, arises when you are given a past history of events and B successes that you can place in T periods in order to *construct* the history that maximizes A_T. To give a real world flavor, assume again that at time $t=0$ you are hired as a political consultant to an incumbent senator facing re-election in period T. She faces a past history of events pertaining to her reputation (summarized in S_0) and she has B identical, and beneficial to her reputation, events/announcements at her disposal that she can release any time before the election. Unlike before, there is no possibility of chance good news if none are released. What do you advise her? By the earlier proposition, we know that the B successes will be bunched together in consecutive periods. The following result allows us to say even more.

PROPOSITION 3 *(Either/Or) When $p = 0$, the optimal profile is to act the first B or the last B periods.*

The intuition for this result, proved in the appendix, is that rehearsal adds convexity to the model. Assume for simplicity that $B = 1$ and that acting the second rather than the first period is preferable. This means that the rehearsal effect from the initial stock S_0 is not enough to offset the recency effect. However, in deciding whether acting the second or the third period, the rehearsal effect is even smaller, since now the stock is only (κS_0), whereas the relative magnitude of the recency effect is the same as before. This means that, it will be preferable to act the third rather than the second period and so on until the end. This logic generalizes for $B > 1$.

This result simplifies significantly the task of computing explicitly the thresholds, since for every value of available actions n and periods to go τ we only need compare the continuation payoffs from acting the first n periods as opposed to the last n. Such comparisons allow us to give the thresholds $H(\tau, n|\rho, \kappa)$ in closed form, and carry out the comparative statics.

PROPOSITION 4 *(Thresholds in closed form) When $p = 0$, the critical thresholds $H(\tau, n|\rho, \kappa)$ satisfy*

a) $H(\tau, n|\rho, \kappa) = \frac{1-\rho^{\tau-n}}{(\rho^{\tau-n}-\kappa^{\tau-n})} \frac{[\frac{1-\kappa^n}{1-\kappa} - \rho\frac{\rho^n-\kappa^n}{\rho-\kappa}]\frac{1}{1-\rho}}{\kappa\frac{\rho^n-\kappa^n}{\rho-\kappa}}$

b) $H(\tau, n|\rho, \kappa)$ *is decreasing in the rate of memory decay,* ρ

c) $H(\tau + 1, n + 1|\rho, \kappa) > H(\tau, n|\rho, \kappa)$

The first fraction in (a) shows the tension between the recency and the rehearsal effect. The numerator is the benefit from acting the last n periods (recency), whereas the denominator captures the benefit from getting a higher rehearsal boost by acting the first n periods (rehearsal). The second fraction, does not depend on the number of periods to go, τ, and simply adjusts for actions still available, n, which is also the length of the string of success to be bunched. In particular, assume that we are in the beginning of period $t-n+1$, with stock κS_{t-n} and we see n consecutive successes. The numerator of the second fraction is the effect on A_t from the n consecutive successes in periods $t - n + 1$ through t, *ignoring the effect from the rehearsal* that these n successes will induce on stock S_{t-n}. This rehearsal effect is the denominator.

Thresholds are decreasing in ρ, the rate at which memory decays, because a higher ρ makes the recency effect less pronounced, thus increasing the attractiveness of acting early for every level of stock of rehearsal.

The third statement of the proposition, $H(\tau + 1, n + 1|\rho, \kappa) > H(\tau, n|\rho, \kappa)$, says that as the number of periods to go, τ, and the number of available actions, n, *together* decrease from one period to the next, the critical threshold that triggers the agent to act decreases. This condition proves the bunching result for the special case $p = 0$ and it is in fact a *stronger* statement. To see that, recall that bunching goes through as long as the critical thresholds decrease or at least do not increase fast enough[15],

[15]This was shown in the corollary, according to which, bunching requires that the critical thresh-

as n and τ together decrease from one period to the next.

The fact that $H(\tau+1, n+1|\rho, \kappa) > H(\tau, n|\rho, \kappa)$ does *not* imply that thresholds are increasing in the number of available actions, n. It turns out that the comparative statics with respect to the number of available actions, n, are ambiguous. Sceptic readers may convince themselves by verifying that $H(4, 3|0.95, 0.3) < H(4, 2|0.95, 0.3)$ (decreasing in n) and $H(3, 2|0.95, 0.3) > H(3, 1|0.95, 0.3)$ (increasing in n).

In light of the above observation, that thresholds are not increasing in the number of available actions, n, the main bunching proposition is more surprising than it may first appear. To see that, assume that with stock S, $n+1$ available actions and $\tau+1$ periods to go, it is optimal to act, that is $V(S, n+1, \tau+1) = V_{act}(S, n+1, \tau+1)$. Then, next period's decision involves a rehearsal stock of $1+\kappa S$ and τ periods to go both of which favor acting by *more* than S and $\tau+1$ favored acting in the first period. This is because thresholds are increasing in τ,the number of periods to go, and because $1+\kappa S > S$. However, having acted in the first period, the second period decision involves $n < n+1$ available actions, which *may* work in the opposite direction, decreasing the attractiveness of acting. The bunching proposition says that this last effect can not dictate choices, as it will always be offset.

---

olds satisfy $1 + \kappa H(\tau+1, n+1|\rho, \kappa) > H(\tau, n|\rho, \kappa)$.

1.3.4 Allowing for bad news

Up to this point we have been assuming that all news is good news. The nature of information conveyed by realizations of the event e_t is either beneficial ($e_t = 1$), or neutral ($e_t = 0$) to the agent under assessment. In this section, I extend the model to include *bad news,* as realizations of the event e_t that are harmful for the agent and whose effect he seeks to minimize.

Consequently, I include a third possible realization $e_t = -1$ for the events e_t, so that $e_t \in \{-1, 0, 1\}$. I refer to $e_t = -1$ realizations as *failures.* The memory technology that determines the recall probabilities for successes and failures is the same as before, with failures rehearsing only past failures and successes rehearsing only past successes. That is, the memorability for a period i failure (success) is governed by[16]

$$M^i_{t+1} = \rho M^i_t + b^i_{t+1} I\{e_{t+1} = -1\} \quad (1.11)$$

$$b^i_{t+1} = \kappa b^i_t \quad (1.12)$$

Assessment at period T depends now on both the number of successes and the number of failures that the assessor remembers at time T. Let $A^-_T = \sum_{i \in \{i|e_i=-1\}} M^i_T$ (and $A^+_T = \sum_{i \in \{i|e_i=+1\}} M^i_T$) be the some of all memorabilities of all failures (successes) up to time t.[17] Now, the reduced form version of the model consists of two pairs of

---

[16]as opposed to $I\{e_{t+1} = 1\}$ for successes

[17]Recall, that when the recall probability for a period i failure (success) is proportional to its memorability at time T, the expected number of failures (successes) that the assessor remembers at

equations; one for successes and one for failures

$$A^j_{t+1} = \rho A^j_t + I^j_{t+1} \cdot (1 + \kappa S^j_t) \tag{1.13}$$

$$S^j_{t+1} = \kappa S^j_t + I^j_{t+1} \tag{1.14}$$

where $j = -$ denotes failure and $j = +$ denotes success, so that $S^-_t = \sum_{i \in \{i | e_i = -1\}} b^i_t$ and $S^+_t = \sum_{i \in \{i | e_i = +1\}} b^i_t$ are the rehearsal stocks for failures and successes respectively.

Assuming that the agent maximizes the *net* number of successes that the assessor will remember on expectation at time T and specifying the following probabilities for the event e_t conditional on choice at time t, the maximization problem becomes

$$\max \; A^+_T - A^-_T \tag{1.15}$$

$$\begin{aligned} s.t.\; A^j_{t+1} &= \rho A^j_t + I^j_{t+1} \cdot (1 + \kappa S^j_t) \\ S^j_{t+1} &= \kappa S^j_t + I^j_{t+1} \\ \Pr(e_t &= 1 | c_t = Act) = 1 \\ \Pr(e_t &= 1 | c_t = Wait) = p_1 \\ \Pr(e_t &= -1 | c_t = Wait) = p_{-1} \\ \text{and} \sum_{i=1}^{t} I\{c_t &= Act\} = B \end{aligned}$$

where $j = +, -$, I^+_{t+1}, I^-_{t+1} are the indicator functions for $e_t = 1$and $e_t = -1$ respectively and $p_1 + p_{-1} \leq 1$.

---

time T is proportional to the sum of the memorabilities of all failures (successes) at T;

Clearly, the model including only good news is a special case (when $p_{-1} = 0$) of this more general model. I chose to discuss the special case first in an attempt to bring out the main ideas in the least complicated setting possible. In this subsection, we want to explore how the addition of failure events changes the nature of the existing results. The intuition developed in the context of the simpler, success-only, model suggests that optimal decisions should somehow depend on the rehearsal stocks, S^+ and S^-. One might think that since the objective function subtracts failures from successes, optimal decisions should depend on the difference between the two rehearsal stocks. This turns out to be false and the two rehearsal stocks, S^+ and S^-, enter optimal decisions in a more surprising fashion.

As in the model including only good news, optimal decisions are dictated by the trade off between the recency and the rehearsal effects. When the rehearsal stock of successes, S^+, is above some threshold level, the rehearsal effect dominates and it is optimal for the agent to act and generate a success. However, generating a period t success has now two functions; first it reinforces past successes and second it guarantees that past failures *will not be reinforced* by a period t failure. As a result, the agent is also triggered to act when the rehearsal stock for failures, S^-, is above some threshold level. Therefore, acting is an optimal choice when either the rehearsal stock for successes or the rehearsal stock for failures is above some threshold levels. We can therefore state the following result,

PROPOSITION 5 *(Thresholds for successes-failures) For every value of periods to go, τ, and any number of available actions, n, it is optimal to act if either of the following* *two statements holds*

a) Fixing the rehearsal stock for failures, S^-, the rehearsal stock for successes is above some threshold level $H^+(\tau, n, S^-|\rho, \kappa)$

b) Fixing the rehearsal stock for successes, S^+, the rehearsal stock for failures is above some threshold level $H^-(\tau, n, S^+|\rho, \kappa)$

Moreover, the critical thresholds $H^+(\tau, n, S^-|\rho, \kappa)$ and $H^-(\tau, n, S^+|\rho, \kappa)$ are both increasing in the number of periods to go, τ, that is the fewer the periods to go, the less the respective critical threshold need be before it triggers the agent to act.

Since one agent's successes can be someone else's failures, this result has important implications for the behavior of two, or more, agents who are competing for assessment in front of a common assessor. To see that, return to our motivating example and assume that our senator has received a substantial series of positive boosts to his image. Then, *both* our senator *and* his opponent will be triggered to act and generate further successes. Their motivation, however, is different. Our senator wants to generate more successes to reinforce the fresh memories of his success (his rehearsal stock for successes, S^+, is high), whereas the opponent wants to generate successes in an effort to guarantee that the fresh memories of his failures will not be reinforced by future failures (his rehearsal stock for failures, S^-, is high).

A special case of the model with both successes and failures is when $p_1 + p_{-1} = 1$. That is, all news are either good or bad: there are no null events. This case is identical to what we would have got in our original model with $e_t \in \{0, 1\}$, if we had interpreted the events $e_t = 0$ as the agent failing to deliver a success and with these failures subject to rehearsal. In this case, the benefit from generating a success is completely symmetric with the benefit from avoiding a failure. As a result, optimal decisions depend on the *sum* of the two rehearsal stocks, S^+ and S^-. Notice that this is not to say that we subtract the failure rehearsal stock from the success stock. Rather, the incentive to act to avoid rehearsing failures is *added* to that aimed at rehearsing successes.

PROPOSITION 6 *(Single threshold) When* $p_1 + p_{-1} = 1$, *for every value of periods to go,* τ, *and any number of available actions,* n, *there exists a single threshold* $H^{+,-}(\tau, n|\rho, \kappa)$, *so that it is optimal to act if and only if* $(S^+ + S^-) > H^{+,-}(\tau, n|\rho, \kappa)$.

The surprising implication of this result is, that if we return to our original set up with $e_t \in \{0, 1\}$ and reinterpret the events $e_t = 0$ as failures subject to rehearsal, optimal decisions will still depend on a *single* threshold level, as before.

Finally, the bunching structure of the optimal profile goes through, with the intuition being the same as before. For completeness I state this as

COROLLARY 2 *In the general case where* $e_t \in \{-1, 0, 1\}$, *if it is optimal to act in period* t, *then it will be optimal to act in period* $t+1$, *as long as the constraint* $\sum_{i=1}^{t} I\{c_t = Act\} = B$ *permits to do so.*

1.3.5 Discussion of results and applications

In an effort to keep the set up of the model as concrete as possible, I have been addressing the model in the context of a senator running for re-election. Nevertheless, the general framework can be applied to almost any setting where products or people undergo periodic assessment. The following is a list of potential applications, in addition to the election campaign that we have already discussed. They cover a wide range of settings, from finance to marketing, and from evaluations at the workplace to opinion formation on moral and public policy issues.

- executives managing the release of news prior to an IPO
- advertisers timing the airing of commercial spots prior to a new product/movie release
- marketing executives choosing when to launch new products, following a recent streak of good or bad news (Bridgestone-Firestone)
- employees timing their effort prior to a pending promotion/bonus decision
- evaluation of forecasters based on their past predictions
- public opinion formation; the NRA and the gun control lobbies scheduling their actions following a streak of Columbine type shooting incidents
- students timing their class contribution prior to the end of term
- junior faculty timing their publications prior to a tenure decision
- Pop artists choosing when to release a new hit album, following a streak of hits or flops

Keeping these applications in mind, the theoretical results can now be translated to the following normative claims.

Bunch together your successes-spread apart your failures. In models of quasi-bayesian and quasi-rational decision making the *order* of informative signals matters. Rabin and Schrag (1999) argue that this is a consequence of a confirmatory bias, whereas Lam (2000) shows that early signals may lock the decision maker into a scenario-story, which affects the way new information is interpreted. In our model, what is important is not only the order (i.e., which events comes first and which comes second), but the actual *spacing* (i.e., the amount of time that elapses between informative events), since it determines the effectiveness of rehearsal. In particular, rehearsal is most effective when all events are bunched together in consecutive periods. This implies that agents evaluated on the basis of their past performance should seek to time in close proximity the events that enhance their image and, by the opposite logic, spread apart those events that are harmful to their image. This implies that a stock market forecaster is more likely to achieve a "guru " status in the mind of investors, when he is lucky enough to make successful predictions in consecutive months or years, rather than when his successes come in intervals. Similarly, a student will create a more memorable good impression to his teacher if he manages to make intelligent comments in a short interval of time.

The same logic can also be applied for public opinion formation on public and moral issues, such as abortion and gun control. Assume, for example, that the a new gun control legislation has been brought to congress. Public opinion is influenced by various gun related crimes that make headlines, such as the Columbine high-school shooting incident. The timing of these events is crucial for what people remember and consequently for whether they will choose to support the new bill. The bunching result predicts that public support will reach its peak, when these random incidents are recent, as dictated by common sense, but also when they are clustered together in a short period of time, because this profile is most efficient in reinforcing the memorabilities of each of these events. By the opposite logic, the cause of the NRA will be least affected if these events are spread apart by big intervals of time.

Finally, advertisers seem to be aware of this bunching intuition, when they time the airing of commercial spots *back to back*, or when they place the same billboard as on *consecutive* billboards along a highway or a railway track.

Random events trigger information releases - Act immediately following a streak of random events. The model implies a simple rule for releasing information. We can think of random events as creating a rehearsal stock. When enough random events push the rehearsal stock above some threshold level, the agent finds it optimal to act and release information, in an effort to reinforce the past good memories (and avoid reinforcement of the past bad memories), summarized in the rehearsal stock. Indeed, this happens in our model despite the fact that there is a finite number of actions to be used up, which implies that early actions come at the cost of not being able to act later on. Moreover, once a random success pushes the rehearsal stock above the threshold, the agent is triggered to act and he *continues* doing so until his stock of available actions has been depleted.

This intuition can be manipulated by an employee timing his effort in the workplace prior to a promotion decision. Our model implies that he can enhance his future reputation by working hard the periods immediately following a streak of random successes (periods with above average performance). Similarly, a student might start offering frequent comments in a seminar class immediately following a couple lucky guesses of his on hard questions.

Stochastic events, good or bad alike, trigger actions. Related to the above, is the idea that random events induce information releases regardless of their nature. Return to the gun-control legislation example. Assume that a year from now congress will take a vote on the controversial bill, and a series of high school shooting accidents have recently occurred. The recommendation that follows from our discussion of the model including both good and bad news is that the recent accidents will trigger *both* the gun control lobbies *and* the NRA to use their resources to try to generate events that promote their respective causes. This way, the gun control lobbies can reinforce the memories of the accidents and the NRA can prevent these same memories from being reinforced.

1.4 Advertising Campaigns

Advertising is an obvious domain where memory imperfections play an important role. For motivation, assume that the senator of our motivating example has raised B dollars for his election campaign. How should he spend it over time in order to obtain the maximum effect at the date of the election. Firms that market new products and studios releasing new movies also face the same problem, the critical date being the day the new product, or movie, is launched.

In the model developed in the previous section, each period the agent makes a binary choice $c_t \in \{Act, Wait\}$, and the assessor remembers, with some probability, individual events $e_t \in \{-1, 0, 1\}$. Advertising is a continuous choice variable. This section is an attempt to combine some of the insights from the model of section three together with the existing literature in advertising in order to incorporate memory imperfections in advertising applications.

Goodwill and optimal advertising expenditures

In the existing literature, dating back to Nerlove and Arrow (1962), advertising is modeled by defining a ***stock of goodwill*** that summarizes the past effect of advertising on demand. Then, the underlying structure becomes an optimal control problem with the stock of goodwill as the state variable and the flow of advertising expenditure as the control. For example, Arrow and Nerlove state the advertising expenditure problem of a firm as

$$\max_{a_t} \int_0^{\infty} e^{-rt}[\pi(G_t) - c(a_t)]dt \tag{1.16}$$

$$s.t. \overset{\bullet}{G_t} = a_t - \delta G_t \tag{1.17}$$

with G_t as the stock of goodwill, a_t the advertising expenditure flow, δ a decay parameter, r the rate of interest, and $c(\cdot)$ and $\pi(\cdot)$ as the cost and revenue functions respectively. Gould (1970) derives solutions for a wide range of combinations of $\pi(\cdot)$ and $c(\cdot)$ functions and Sethi (1977) provides an excellent summary for the use of optimal control in extending the original Nerlove-Arrow formulation.

Memory appears even in the early Nerlove-Arrow formulation through the decay parameter δ, but the framework is not rich enough to incorporate the principles of repetition and similarity. Marketing textbooks devote whole chapters in discussing their importance for effective advertising (see Sissors and Bumba, chapters 6-9) and advertisers seem to realize their importance when scheduling two identical spots to be aired back to back in the hope that the second will reinforce the memory of the first.

Following the existing literature, I adopt the concept of the stock of goodwill, G_t, which I assume it increases with advertising expenditure, a_t. Memory imperfections are incorporated in two ways: First memory decay implies that the stock of goodwill depreciates at some rate δ. Second, in order to model rehearsal and cue dependence, I assume that the marginal effect of advertising depends on the past stock of goodwill. The higher the stock of goodwill, the more memorable past advertising is. Therefore, current advertising will be more effective in triggering these memories and the marginal effect on goodwill will be greater. These restrictions imply that

$$\dot{G}_t = u(a_t, G_t) - \delta G_t \tag{1.18}$$

where $u_a > 0$, $u_{aa} < 0$, $u_G > 0$, $u_{GG} < 0$ and $u_{aG.} > 0$.

Equation (1.18) can be thought of as an analog of (1.3). Now, a senator with a budget B for his election campaign, maximizes goodwill at time T, $G(T)$, subject to

(1.18) and the isoperimetric constraint implied by the fixed budget B. That is,

$$\max_{a_t} \; G(T) \tag{1.19}$$

$$s.t. \; \overset{\bullet}{G_t} = u(a_t, G_t) - \delta G_t \tag{1.20}$$

$$\int_0^T a_t dt = B \tag{1.21}$$

with $G_0 > 0$ given.

The euler equation can be reduced to (see appendix)

$$\frac{u_{aa}}{u_a} \overset{\bullet}{a} + \frac{u_{aG}}{u_a} \overset{\bullet}{G} - u_G + \delta = 0 \tag{1.22}$$

To extract more intuition from the euler equation, consider the discrete time analog of (1.41). Letting $\rho \equiv 1 - \delta$, replacing integrals with summation signs and (1.18) with $G_{t+1} = \rho G_t + u(a_{t+1}, \rho G_t)$ we get (see appendix)

$$u_a(a_{t+1}, \rho G_t) = \rho u_a(a_t, \rho G_{t-1})[1 + u_G(a_{t+1}, \rho G_t)] \tag{1.23}$$

with the usual "marginal cost equals marginal benefit" interpretation. Assume, that we are on the optimal path and we contemplate shifting one unit of expenditure (a) from period t to period $t+1$. The marginal benefit from increased expenditure in period t+1 is $u_a(a_{t+1}, \rho G_t)$, the LHS of (1.23). The marginal cost has two components. First, there is the direct effect of decreased consumption in period t, which equals $\rho u_a(a_t, \rho G_{t-1})$, the first term in the RHS of the euler equation. Second, there is an indirect effect: less advertising in period t results to a lower (by $u_a(a_t, \rho G_{t-1})$) stock of rehearsal at time $t+1$ and therefore lowers (by $\rho u_G(a_{t+1}, \rho G_t)$) the marginal benefit from expenditure in period $t+1$, since there is less to be rehearsed. The total[18] $\rho u_a(a_t, \rho G_{t-1}) u_G(a_{t+1}, \rho G_t)$ is the second term in the expansion of the RHS of (1.23).

A closed form solution for the optimal advertising expenditure profile exists for the special case where $u(a, G)$ belongs to the Cobb-Douglas family, i.e. $u(a, G) = Ka^\beta G^\gamma$, with $\beta + \gamma \leq 1$. The parameter β captures the concavity of $u(a, G)$ with respect to a and the parameter γ measures the importance of rehearsal. Substituting (1.18) in the euler equation and carrying out the algebra, we get

$$\begin{aligned} \frac{\dot{a}}{a} &= \frac{\delta(1-\gamma)}{(1-\beta)} \Rightarrow \\ a(t) &= a(0) \exp\{\frac{\delta(1-\gamma)}{(1-\beta)} t\} \end{aligned} \tag{1.24}$$

where $a(0) = \frac{B\frac{\delta(1-\gamma)}{(1-\beta)}}{\exp\{\frac{\delta(1-\gamma)}{(1-\beta)}T\}-1}$, which can be obtained using the isoperimetric constraint.

In the optimal path, advertising expenditures increase over time *exponentially* at rate $\frac{\delta(1-\gamma)}{(1-\beta)}$. There are three forces at work here: a) memory decay, embodied in the parameter δ, that dictates expenditures to be timed as late as possible (recency effect), b) degree of concavity of the u function with respect to a, captured in the parameter β, which implies expenditure smoothing over time and c) rehearsal, through the parameter γ, that complements (b). One can easily verify, that the rate at which expenditures grow as we approach the election date T is *lower* when the environment is *more* rehearsal sensitive (high γ), when $u(a, G)$ is *more* concave in a (low β) and when memory decay is *less* pronounced (low δ).

---

[18] Here we invoke the chain rule: $\frac{d}{da_t}(u(a_{t+1}, \rho G_t(a_t))) = u_G(a_{t+1}, \rho G_t)\rho\frac{d}{da_t}(u(a_t, \rho G_{t-1})) = u_G(a_{t+1}, \rho G_t)\rho u_a(a_t, \rho G_{t-1})$

Therefore, rehearsal *induces smoothing* in excess of that implied by concavity alone. When too little is spent early, future expenditures have little to rehearse. In practical terms, this implies that advertising campaigns should start early on. Advertising is most effective when there is a substantial stock of goodwill that can be reinforced. Early spending creates that stock of goodwill and increases the effectiveness of future advertising expenditures. The following quote from a political campaign manual, Guzzetta (1987), bears witness to the empirical validity of this theoretical prediction for political campaigns. The same logic applies to any advertising campaign for a new brand or a new movie.

Even when the product is not necessarily an "offensive" one, it needs to be stated over and over in order to penetrate your consciousness. Regrettably, politicians are in an even more difficult position than most new products because generally they have a high "negative" rating in the minds of the public. This is the primary reason why campaigns have to start so early: only repetition over a long period of time can overcome this natural barrier in the minds of the voters.

In parallel to the model of the previous section, one might expect a rehearsal effect that competes with the recency effect and according to which, given some existing stock of goodwill, it is optimal to act early when the stock is still memorable, rather than later when the past stock has decayed and been forgotten. Such an effect is not present here. Consider the choice between spending one unit today when the stock is G_t, rather than tomorrow with stock ρG_t. By spending it today we get ρG_t^γ, as opposed to $\rho^\gamma G_t^\gamma$ by spending it tomorrow. Since $\gamma, \rho < 1$ it is always optimal to wait and spend tomorrow. This occurs because the stock of goodwill (stock of rehearsal in the previous terminology) does not decay fast enough to offset the recency effect from memory decay. In the model of section three, a trade off between the recency and the rehearsal effects was achieved because the structural form of the model (M_t^i and b_t^i equations) implied that the stock of rehearsal decays at a rate faster than memory decay. To get such a rehearsal effect here, we have to either modify the functional form for $u(\cdot, \cdot)$, or directly introduce a rehearsal stock that decays at a rate faster than δ.

Relation to addiction and habit formation

In the advertising model just presented, the effect of past advertising was summarized in a single variable, the stock of goodwill, which affected the marginal benefit from current advertising. This mathematical structure can also be found in models of addiction and habit formation. For example, when working with preferences that exhibit habit formation past consumption levels are assumed to be summarized in a single state variable, which in turn affects the marginal utility from current consumption. A seminal paper along these lines is the model of rational addiction due to Becker and Murphy (1989). To capture the effect of addiction, they assume that the instantaneous utility function is a non separable function of current consumption c_t, and of past levels of consumption, summarized in the stock S_t;that is utility at time t is $u(c_t, S_t)$ with $u_{cS} \neq 0$. The stock S_t evolves similarly to the stock of goodwill in the Nerlove-Arrow model. Then, the consumer solves[19]

---

[19]In fact, this is a "stripped down" version of the Becker and Murphy original set up.

$$\max_{a_t} \int_0^\infty u(c_{t,}, S_t)e^{-rt}dt \tag{1.25}$$

$$s.t. \; \overset{\bullet}{S}_t = c_t - \delta S_t \tag{1.26}$$

$$\int_0^T e^{-rt}[p_t c_t]dt \leq L_0 + \int_0^T e^{-rt} w_t(S_t)dt \tag{1.27}$$

with δ as the rate of depreciation, r the interest rate, p_t the price of the consumption good, $w(\cdot)$ the earnings function and L_0 the initial value of assets.

The mathematical structure of the rational addiction model and the advertising model is so similar that one could even use the Becker and Murphy set up word for word, and after renaming variables ($c_t \rightarrow$advertising expenditure, $S_t \rightarrow$ stock of goodwill, $u(\cdot,\cdot) \rightarrow$profit function, etc.) use it in order to address the advertising problem of a firm over a horizon, when people forget at rate δ, but current advertising expenditures can also reinforce past memories.

1.5 Conclusion

Imperfect memory is an undeniable fact of human cognition with profound social and economic implications. Despite the current interest in psychology and economics, with the exception of Mullainathan (1997), memory limitations have not found their way into economic literature, even though they are discussed, but not modeled, in marketing textbooks. It is true, that memory is a complex process and memory researchers have not fully understood its functions. However, some stylized facts that govern the recall of episodic memories have emerged, and allow us to make a first step. The purpose of this paper was twofold: a) to use these stylized facts (cue dependance and rehearsal in particular), as building blocks in developing a simple model of memory, and b) to apply the formal mathematical model to an economic setting, by addressing the issue of how informative events be released to a forgetful agent, who at some future date T is expected to make an assessment based on her memories of these events. I showed that the *spacing* of events is an important determinant of what memories people have and I described the spacing properties of optimal profiles for releasing information. The general framework and the theoretical results can be applied to wide range of settings including, political campaigns, advertising or marketing of new products, evaluations at the workplace and public opinion formation.

I conclude by briefly discussing some possible extensions to the existing framework.

Many assessors: diversification. Throughout this paper, I assume that the agent is assessed by only one agent (or all assessors use the same history of events). However, there are many interesting applications, where an agent is assessed by two or more different assessors, each using a different history of events. Once again, assume that our politician maybe interested in gaining the support of two distinct, but equally important, lobbies. Furthermore, assume that the event $e_t = 1$ now means that "the politician took an action that favors the objectives of the lobby at time t". Different lobbies have different agendas and as a result they use a different history of events to form their assessments. Should the politician concentrate his effort and energy in satisfying one of the two lobbies or should he *diversify* trying to grant favors evenly? Concavity in payoffs favors diversification, but memory imperfections dictate the opposite. Less diversification allows for greater benefits from rehearsal, as favors granted to one lobby rehearse the memories of past favors granted to the same lobby.

Many agents: competing for assessment. In other settings two or more agents are competing in front of a common assessor. In a presidential race, for example, two candidates struggle to create the best impression for themselves in the eyes of the public. Or in the gun control example the gun control lobbies competes against the NRA. It would be of interest to extend the present model to include the game theoretic interaction that emerges when each candidate, or lobby, releases informative events worrying not only for his past record, but for the past record of his opponent as well.

1.6 References

[1] BECKER, G. S. AND K. M. MURPHY (1988), "A theory of rational addiction", *Journal of Political Economy,* 96 (4), 675-700.

[2] CONSTANTINIDES, G. M. (1990), "Habit formation: A resolution of the equity premium puzzle", *Journal of Political Economy,* 98 (3), 519-543.

[3] DOW, J. (1991), "Search decisions with limited memory", *Review of Economic Studies,* 58, 1-14.

[4] GOULD, J. P. (1970), "Diffusion processes and optimal advertising policy". In E. S. Phelps (ed.) *Microeconomic foundations of employment and inflation theory,* 338-368.

[5] GUZZETTA S. J. (1987), *The campaign manual,* (2nd ed.), Alexandria, VA:Political Publishing Company.

[6] HIRSHLEIFER, D. AND I. WELCH (2001), "An economic approach to the psychology of change: Amnesia, inertia and impulsiveness", working paper.

[7] KANDEL, E. R., SCHWARTZ, J. H. AND T. M. JESSEL (1995), *Essentials of neural science and behavior.* Norwalk, CT:Appleton & Lange.

[8] LAM, R. (1993), "Learning through stories", mimeo, Yale University.

[9] LOFTUS, E. F. (1993), "The reality of repressed memories", *American Psychologist,* 48, 518-537.

[10] MULLAINATHAN, S. (1997), "A memory based model of bounded rationality", mimeo, MIT.

[11] NERLOVE, M. AND K. J. ARROW (1962), "Optimal advertising policy under dynamic conditions" *Economica (*May*),* 124-142.

[12] PARKIN A. J. (1993), *Memory: Phenomena, experiment and theory.* Oxford:Blackwell.

[13] RABIN, M. (1998), "Psychology and Economics", *Journal of Economic Literature,* 36, 11-46.

[14] RABIN, M. AND T. O'DONAHUE, (1999), "Incentives for Procrastinators", *Quarterly Journal of Economics,* 114 (3), 769-816.

[15] RABIN, M. AND ,J. SCHRAG (1999), "First Impressions Matter: A Model of Confirmatory Bias", *Quarterly Journal of Economics,* 114 (1), 37-82.

[16] RUNDUS, D. (1971), "Analysis of rehearsal processes in free recall", *Journal of Experimental Psychology,* 89, 63-77.

[17] RYDER, H. E., JR. AND G. M. HEAL (1973), "Optimum growth with intertemporally dependent preferences", *Review of Economic Studies,* 40, 1-33.

[18] SCHACTER, D. L. (1996), *Searching for memory.* New York: BasicBooks.

[19] SETHI, P. S. (1977), "Dynamic Optimal Control Models in Advertising: A Survey", SIAM Review, 19 (4), 685-725.

[20] SISSORS, J. Z. AND L. BUMBA (1994), *Advertising media planning,* (4th ed.), Lincolnwood, IL:NTC Business Books.

[21] TULVING, E. AND S. OSLER (1968), "Effectiveness of retrieval cues in memory for words", *Journal of Experimental Psychology,* 77, 593-601.

[22] TULVING, E. AND D. M. THOMPSON (1973), "Encoding specificity and retrieval processes in episodic memory", *Psychological Review,* 80, 353-373.

1.7 Appendix

1.7.1 Notation

1. Let $\bar{V}_{Wait}$ $(S, n, \tau | S', n', \tau')$ be the continuation payoff starting from the state vector (S, n, τ) when one follows the following rule: wait the first period and then take actions that *would have been optimal* if you had started out with the state vector (S', n', τ') rather than (S, n, τ). Also, let $\bar{V}_{Act}$ $(S, n, \tau | S', n', \tau')$ be the continuation payoff starting from the state vector (S, n, τ) when you follow the rule: a) act the first period, then wait, b) if $n > 1$, when you act again you have to keep on acting until you run out of available actions, that is until $n = 0$, c) in deciding when to act for the second time take actions that *would have been optimal* if you had started out with state vector (S', n', τ') instead.

2. Let $Z(i)$ be the set of all profiles of zeros and ones of length i, that is $Z(i) = \{z \in R^i | z_i = 1 \text{ or } z_i = 0\}$. Also let $|z| = i$ be the order (length) of z and $|z^1| = \sum_{i=1}^{i} z_j$ be the number of ones in z. For example, letting $z = (1, 0, 1, 1, 0, 1)$, we have $|z| = 6, |z^1| = 4$.

3. Assume that the state vector is (S, n, τ). Let $\Omega(S, n, \tau)$ be the set of all profiles of zeros and ones after which the first action would be optimal. (Then, $\Omega(S, n, \tau) = \emptyset \Rightarrow V(S, n, \tau) = V_{Act}(S, n, \tau)$). For example, let $n = 1$, $\tau = 3$ and the rehearsal stock is S, satisfying $S < H(3, 1|\rho, \kappa)$, but $1 + \kappa S > H(2, 1|\rho, \kappa)$. Then, $\Omega(S, n, \tau) = \{(1), (0, 1), (0, 0)\}$. Also, for every $\omega \in \Omega$, let $|\omega|$ be the order of the profile, and $|\omega^1|$ be the number of ones in ω.

4. Let S be the rehearsal stock with τ periods to go. Assume that in the τ periods to come we see a profile $e \in R^{\tau}$. Let $CP(e, S)$ be the continuation payoff from a history profile e starting with initial rehearsal stock S. For example, let $e : S|1, 1, 0, 1, 0$. Then, $CP(e, S) = \rho^4(1 + \kappa S) + \rho^3(1 + \kappa + \kappa^2 S) + \rho(1 + \kappa^2 + \kappa^3 + \kappa^4 S)$.

5. Decompose $CP(e, S)$ into the following two terms; a) $D(e)$: the continuation payoff that profile e would induce even if $S = 0$, and b) $RE(e)$: the total rehearsal effect that the successes in profile e will induce on the past rehearsal stock S, as a proportion of S. Returning to the previous example of $e : S|1, 1, 0, 1, 0$ we have; $RE(e) = \rho^4\kappa + \rho^3\kappa^2 + \rho\kappa^4$ and $D(e) = \rho^4 + \rho^3(1 + \kappa) + \rho(1 + \kappa^2 + \kappa^3)$. Then, $CP(e, S) = D(e) + S \cdot RE(e)$

6. Let e be a profile of zeros and ones. A success immediately after this profile e will rehearse all successes in e. Call the sum of all these rehearsal boosts $F(e)$. For example, with $e = (1, 1, 0, 1, 0)$, $F(e) = \kappa^2 + \kappa^4 + \kappa^5$.

1.7.2 Bunching

PROPOSITION If it is optimal to act in period t, it is also optimal to act in period $t + 1$. More formally, for any S, n and τ, with $0 < n \leq \tau$,

$$V(S, n + 1, \tau + 1) = V_{act}(S, n + 1, \tau + 1) \Rightarrow V(1 + \kappa S, n, \tau) = V_{act}(1 + \kappa S, n, \tau)$$

PROOF

I prove the contrapositive statement that

$$V(1+\kappa S,n,\tau) = V_{Wait}(1+\kappa S,n,\tau) \Rightarrow V(S,n+1,\tau+1) = V_{Wait}(S,n+1,\tau+1) \tag{1.28}$$

The proof is by induction on n.

Let $n = 1$ be the base case for our induction.

Step 1: Show that with $n = 1$

$$\begin{aligned} V_{Wait}(1+\kappa S,n,\tau) &> V_{Act}(1+\kappa S,n,\tau) \Rightarrow \\ \bar{V}_{Wait}\,(S,n,\tau|1+\kappa S,n,\tau) &> \bar{V}_{Act}\,(S,n,\tau|S,n,\tau+1) \end{aligned} \tag{1.29}$$

Let the state vector be $(1+\kappa S,n,\tau)$. Then, the instructions in $V_{Wait}(1+\kappa S,n,\tau)$ and $V_{Act}(1+\kappa S,n,\tau)$ give rise to the profiles e^1 and e^2 respectively

$$\begin{aligned} e^1 &: (1+\kappa S)|\omega,1,z \\ e^2 &: (1+\kappa S)|1,\omega,z \end{aligned}$$

with probability $p_{z\omega} \equiv p^{|\omega^1|+|z^1|}(1-p)^{\tau-1-|\omega^1|-|z^1|}$, where $\omega \in \Omega(1+\kappa S,n,\tau)$, and $z \in Z(\tau-1-|\omega|)$. Similarly $\bar{V}_{Wait}\,(S,n,\tau|1+\kappa S,n,\tau)$ and $\bar{V}_{Act}\,(S,n,\tau|S,n,\tau+1)$ give rise to profiles e^3 and e^4 with the same probability.

$$\begin{aligned} e^3 &: S|\omega,1,z \\ e^4 &: S|1,\omega,z \end{aligned}$$

The continuation payoff from profiles e^1 and e^3 is the same with the exception that the successes in the profile $(\omega,1,z)$ rehearse $1+\kappa S$ in e^1 rather than S, as in e^3. Similarly, the continuation payoff from profiles e^2 and e^4 is the same with the exception that the successes in the profile $(1, \omega, z)$ rehearse $1 + \kappa S$ in e^2 rather than S, as in e^4. Then the difference in the continuation payoffs between e^1 and e^3 is $CP(e^1, 1+\kappa S) - CP(e^3, S) = RE((\omega, 1, z))(1 + \kappa S - S)$. Similarly, $CP(e^2, 1 + \kappa S) - CP(e^4, S) = RE((1, \omega, z))(1 + \kappa S - S)$. The statement in (1.29) follows if[20]

$CP(e^1, 1 + \kappa S) - CP(e^2, 1 + \kappa S) < CP(e^3, S) - CP(e^4, S) \Leftrightarrow$

$CP(e^1, 1 + \kappa S) - CP(e^3, S) < CP(e^2, 1 + \kappa S) - CP(e^4, S) \Leftrightarrow RE((\omega, 1, z)) < RE((1, \omega, z))$

The desired statement $RE((\omega, 1, z)) < RE((1, \omega, z))$ holds because the rehearsal effect from the earliest success in $(1, \omega, z)$ is weakly bigger than the earliest success in $(\omega, 1, z)$, the effect from the second earliest success in $(1, \omega, z)$ is weakly bigger than the second earliest success in $(\omega, 1, z)$ and so on. This holds for all $\omega \in \Omega$ and for all $z \in Z$.

Step 2: Show that (with $n = 1$)

$$\begin{aligned} \bar{V}_{Wait}\left(S, n, \tau | 1 + \kappa S, n, \tau\right) &> \bar{V}_{Act}\left(S, n, \tau | S, n, \tau + 1\right) \Rightarrow \\ \bar{V}_{Wait}\left(S, n, \tau + 1 | 1 + \kappa S, n, \tau\right) &> \bar{V}_{Act}\left(S, n, \tau + 1 | S, n, \tau + 1\right) \end{aligned} \tag{1.30}$$

The instructions implied by $\bar{V}_{Wait}\left(S, n, \tau | 1 + \kappa S, n, \tau\right)$ and $\bar{V}_{Act}\left(S, n, \tau | S, n, \tau + 1\right)$

---

[20]To see this, it is important to note that the profiles e^1, e^2, e^3 and e^4 are all obtained with the same probability, $p_{z\omega}$.

give rise to profiles e^1 and e^2 with probability $p_{z\omega} \equiv p^{|z^1|+|\omega^1|}(1-p)^{\tau-1-|z^1|-|\omega^1|}$

$$e^1 \;:\; S|\omega, 1, z$$

$$e^2 \;:\; S|1, \omega, z$$

where $\omega \in \Omega(1+\kappa S, n, \tau)$ and $z \in Z(\tau - 1 - |\omega|)$ as before. Also, $\bar{V}_{Wait}$ $(S, n, \tau + 1|1 + \kappa S, n, \tau)$ and $\bar{V}_{Act}$ $(S, n, \tau+1|S, n, \tau+1)$ give rise to profiles e^3 and e^4.

$$e^3 \;:\; S|\omega, 1, z, ?$$

$$e^4 \;:\; S|1, \omega, z, ?$$

The last element can be either a one or a zero. It is zero with probability $1-p$ and one with probability p. The "if" statement of (1.30) is equivalent to $\sum_{z\in Z}\sum_{\omega\in\Omega} p_{z\omega}[CP(e^1, S) - CP(e^2, S)] > 0$. Similarly, the "then" statement of (1.30) is equivalent to $\sum_{z\in Z}\sum_{\omega\in\Omega} p_{z\omega}[CP(e^3, S) - CP(e^4, S)] > 0 \Leftrightarrow (1-p)\sum_{z\in Z}\sum_{\omega\in\Omega} p_{z\omega}[\rho CP(e^1, S) - \rho CP(e^2, S)] + p\sum_{z\in Z}\sum_{\omega\in\Omega} p_{z\omega}[\rho CP(e^1, S) + 1 + F(\omega, 1, z) - \rho CP(e^2, S) - 1 - F(1, \omega, z)] > 0 \Leftrightarrow \rho\sum_{z\in Z}\sum_{\omega\in\Omega} p_{z\omega}[CP(e^1, S) - CP(e^2, S)] + p\sum_{z\in Z}\sum_{\omega\in\Omega} p_{z\omega}[F(\omega, 1, z) - F(1, \omega, z)] > 0$

The first of the last two summation terms is positive by assumption. The second summation term is also positive because for every z and every ω, $F(\omega, 1, z) > F(1, \omega, z)$.This is because the first term in $F(\omega, 1, z)$ is weakly greater than the first term of $F(1, \omega, z)$ and similarly for the second, third, and all other terms. (The first term is the one coming from the boost to the memorability of the earliest success in the respective profiles $(\omega, 1, z)$ and $(1, \omega, z)$.)

Step 3: Show that (with $n = 1$)

$$\begin{aligned}
\bar{V}_{Wait}\left(S,n,\tau+1|1+\kappa S,n,\tau\right) &> \bar{V}_{Act}\left(S,n,\tau+1|S,n,\tau+1\right) \Rightarrow \\
\exists V^{*}(S,n+1,\tau+1|1+\kappa S,n,\tau)s.t.\ \overset{*}{V} &> V_{Act}(S,n+1,\tau+1) \qquad (1.31)
\end{aligned}$$

where $V^{*}(S,n,\tau|S^{'},n^{'},\tau^{'})$ is the continuation payoff generated by some set of instructions with τ periods to go, n available actions and rehearsal stock S. In particular let that set of instructions be: a) act in consecutive periods and b) in deciding which period to start acting pretend that the state vector you started out with were $(S^{'},n^{'},\tau^{'})$ instead and act optimally to that. The first two expressions in (1.31) give rise to profiles

$$\begin{aligned}
e^{1} &: S|\omega,1,?,z \\
e^{2} &: S|1,\omega,?,z
\end{aligned}$$

where $z \in Z(\tau - |\omega| - 1)$ and $\omega \in \Omega(1+\kappa S,n,\tau)$. The last two expressions in (1.31) give rise, with the same probability, to

$$\begin{aligned}
e^{3} &: S|\omega,1,1,z \\
e^{4} &: S|1,\omega,1,z
\end{aligned}$$

With probability p, the uncertain term "?" in e^{1} and e^{2} is a one and we have $CP(e^{1},S) - CP(e^{3},S) = CP(e^{2},S) - CP(e^{4},S) = 0$. With probability $(1-p)$, the uncertain term is a zero, and we have $CP(e^{1},S) - CP(e^{3},S) = -\rho^{|z|}(1+F(\omega,1))$ and $CP(e^{2},S) - CP(e^{4},S) = -(1+\rho^{|z|}F(1,\omega))$.The result in (1.31) holds as long $CP(e^{1},S) - CP(e^{3},S) < CP(e^{2},S) - CP(e^{4},S) \Leftrightarrow F(\omega,1) > F(1,\omega)$

This is true, as explained in step 2.

We have therefore established the contrapositive statement in (1.28) for $n = 1$, which means that if the number of available actions is two, in the optimal profile they are bunched together.

Now we can proceed using the induction hypothesis that the result holds for all n up to k. Again we need a series of steps

Step 0: Show that for $n = k$

$$\begin{aligned} V_{Wait}(1+\kappa S,n,\tau) &> V_{Act}(1+\kappa S,n,\tau) \Rightarrow \\ V_{Wait}(1+\kappa S,n,\tau) &> \bar{V}_{Act}\ (1+\kappa S,n,\tau|S,n,\tau+1) \end{aligned} \tag{1.32}$$

The desired result follows by the definition of $V_{Wait}(1+\kappa S,n,\tau)$ which is defined as the optimal set of instructions.

The rest of the proof mimics the proof of the base case $n = 1$. We need to conduct the following three steps

Step1$'$: Show that

$$\begin{aligned} V_{Wait}(1+\kappa S,n,\tau) &> \bar{V}_{Act}\ (1+\kappa S,n,\tau|S,n,\tau+1) \Rightarrow \\ \bar{V}_{Wait}\ (S,n,\tau|1+\kappa S,n,\tau) &> \bar{V}_{Act}\ (S,n,\tau|S,n,\tau+1) \end{aligned} \tag{1.33}$$

Step 2$'$. Show that

$$\begin{aligned} \bar{V}_{Wait}\ (S,n,\tau|1+\kappa S,n,\tau) &> \bar{V}_{Act}\ (S,n,\tau|S,n,\tau+1) \Rightarrow \\ \bar{V}_{Wait}\ (S,n,\tau+1|1+\kappa S,n,\tau) &> \bar{V}_{Act}\ (S,n,\tau+1|S,n,\tau+1) \end{aligned} \tag{1.34}$$

Step $3'$. Show that

$$\bar{V}_{Wait}(S,n,\tau+1|1+\kappa S,n,\tau) > \bar{V}_{Act}(S,n,\tau+1|S,n,\tau+1) \Rightarrow \quad (1.35)$$
$$V^*(S,n+1,\tau+1|1+\kappa S,n,\tau) > V_{Act}(S,n+1,\tau+1)$$

where $V^*(S,n,\tau|S',n',\tau')$ is the same as in step 3 of the base case, $n=1$.

The proof of each step is identical to that of the base case $n = 1$. I prove in detail only step $1'$ in order to show its similarity to step 1 and the importance of the induction hypothesis.

The first two set of instructions in (1.33) give rise with probability $p_{z\omega} \equiv p^{|z^1|+|\omega^1|}(1-p)^{\tau-n-|z^1|-|\omega^1|}$ to

$$e^1 \; : \; 1+\kappa S|\omega,1,1,1,..,1,z$$
$$e^2 \; : \; 1+\kappa S|1,\omega,1,1,...,1,z$$

where $\omega \in \Omega(1+\kappa S,n,\tau)$ and $z \in Z(\tau-n-|\omega|)$. The induction hypothesis is crucial, because it imposes the restriction that the n successes will be bunched together in profile e^1. The last two sets of instructions in (1.33) give rise to

$$e^3 \; : \; S|\omega,1,1,1,...,1,z$$
$$e^4 \; : \; S|1,\omega,1,1,...,1,z$$

The desired result follows once you note that $RE(\omega,1,1,1,...,1,z) < RE(1,\omega,1,1,...,1,z)$. (See step 1 for the base case, $n=1$.)

1.7.3 Thresholds

PROPOSITION For every value of number of periods to go, τ, and number of actions still available, n, with $n < \tau$,there exists a threshold rehearsal stock, $H(\tau, n|\rho, \kappa)$, such that it is optimal to act if in period t and only if $S \geq H(\tau, n|\rho, \kappa)$. Moreover, the thresholds are increasing in τ; that is the fewer the periods to go, the lower the rehearsal stock need be before it triggers the agent to act.

PROOF

For the first part of the proposition it suffices to show that for any rehearsal stock $S' < S$,

$$\begin{aligned} V_{Wait}(S, n, \tau) &> V_{Act}(S, n, \tau) \Rightarrow \\ \bar{V}_{Wait}(S', n, \tau|S, n, \tau) &> V_{act}(S', n, \tau) \end{aligned} \tag{1.36}$$

This implies that, fixing τ, the number of periods to go, and n, the number of available actions, if it is optimal to wait with a rehearsal stock S, it is also optimal to wait with a lower stock.

I use the same logic as in the proof of the bunching result. The first two statements in (1.36) give rise to the following profiles with probability $p_{z\omega} \equiv p^{|z^1|+|\omega^1|}(1-p)^{\tau-n-|z^1|-|\omega^1|}$

$$\begin{aligned} e^1 &: S|\omega, 1, 1, ..., 1, z \\ e^2 &: S|1, 1, ..., 1, \omega, z \end{aligned}$$

where $\omega \in \Omega(S, n, \tau)$ and $z \in Z(\tau - |\omega| - n)$. We know that profiles e^1 and e^2 have the property that the n successes are together by the bunching result. Similarly the last two statements in (1.36) give rise to

$$e^3 \; : \; S'|\omega, 1, 1, ..., 1, z$$

$$e^4 \; : \; S'|1, 1, ..., 1, \omega, z$$

with the same probability, $p_{z\omega}$. The profile e^1 is almost identical to profile e^3, the only difference being that all future successes rehearse the stock S rather than S'. Therefore their difference, in terms of final payoff, is $(S - S')RE(\omega, 1, 1, ..., 1, z)$. Similarly, the difference between profiles e^2 and e^4 is $(S - S')RE(1, 1, ..., 1, \omega, z)$. The desired result in (1.36) holds because $RE(1, 1, ..., 1, \omega, z) > RE(\omega, 1, 1, ..., 1, z)$, for all ω and z. (This is as in step 1 of the bunching result).

Now, I show that the thresholds are increasing in τ, the number of periods to go. It suffices to show that

$$\begin{aligned} V_{Wait}(S, n, \tau) \; &> \; V_{Act}(S, n, \tau) \Rightarrow \\ \bar{V}_{Wait}\,(S, n, \tau + 1|S, n, \tau) \; &> \; V_{act}(S, n, \tau + 1) \end{aligned} \tag{1.37}$$

This means that if it is optimal to wait with τ periods to go, $(S < H(\tau, n|\rho, \kappa))$, then it would also be optimal to wait when we have the same rehearsal stock and $\tau + 1$ periods to go $(S < H(\tau + 1, n|\rho, \kappa))$.

Once again, I use the same logic as in the proof of the bunching result. The first two statements in (1.37) give rise to the following profiles with probability $p_{z\omega} \equiv p^{|z^1|+|\omega^1|}(1-p)^{\tau-n-|z^1|-|\omega^1|}$

$$e^1 \; : \; S|\omega, 1, 1, ..., 1, z$$

$$e^2 \; : \; S|1, 1, ..., 1, \omega, z$$

where $\omega \in \Omega(S, n, \tau)$ and $z \in Z(\tau - |\omega| - n)$. We know that profiles e^1 and e^2 have the property that the n successes are together by the bunching result. Similarly the last two statements in (1.37) give rise to

$$e^3 \; : \; S|\omega, 1, 1, ..., 1, z, ?$$

$$e^4 \; : \; S|1, 1, ..., 1, \omega, z, ?$$

with the same probability, $p_{z\omega}$. The last element can be one or zero with probabilities p and $1-p$. The rest of the proof is *identical* to step $2'$ of the bunching result. (Note that step $2'$ leads to the same profiles e^1, e^2, e^3 and e^4.)

1.7.4 The non-stochastic case

Let A_t and S_t be the running total and rehearsal stocks respectively at time t. Assume that starting at time $t+1$ we have a sequence of n successes. Let 1^n denote a profile of n consecutive successes. At time $t+n$ the running total stock A_{t+n} equals

$A_{t+n} = \rho^{t+n} A_t + CP(1^n, S_t) = \rho^{t+n} A_t + D(1^n) + RE(1^n) \cdot S_t$

One can easily verify that

$D(1^n) = \rho^{n-1} + \rho^{n-2}(1+\kappa) + \rho^{n-3}(1+\kappa+\kappa^2) + ... + (1+\kappa+\kappa^2+...\kappa^{n-1}) =$

$\frac{1}{1-\rho}[\frac{1-\kappa^n}{1-\kappa} - \rho\frac{\rho^n-\kappa^n}{\rho-\kappa}]$

$RE(1^n) = \kappa(\rho^{n-1} + \rho^{n-2}\kappa + \rho^{n-1}\kappa^2 + ... + \kappa^{n-1}) = \kappa\frac{\rho^n-\kappa^n}{\rho-\kappa}$

We can, now, prove the results.

PROPOSITION (Either/Or) When $p = 0$, the optimal profile is to act the first B or the last B periods.

PROOF

Let $p = 0$, $S_0 = S$, $T > B$ and $1 \leq t \leq T - B - 1$. We know by the bunching result that it is optimal to act in consecutive periods.

Assume that acting in periods $t + 1$ through period $t + B$ is preferred to acting in periods t through $t + B - 1$. Call the resulting history profiles e' and e respectively. I will show that preferring e' to e, i.e. $CP(e', S) > CP(e, S)$, implies that acting in period $t + 2$ through $t + B + 1$ results to a profile, e'', that is preferred to e', i.e. $CP(e'', S) > CP(e', S)$

$e : S|0, 0, 0, 1, 1, ..., 1, 0, 0, 0, 0$

$e' : S|0, 0, 0, 0, 1, 1, ..., 1, 0, 0, 0$

$e'' : S|0, 0, 0, 0, 0, 1, 1, ..., 1, 0, 0$

$CP(e', S) > CP(e, S) \Rightarrow \rho^{T-(t+B)}[D(1^B)+RE(1^B)\kappa^t S] > \rho^{T-(t+B-1)}[D(1^B)+RE(1^B)\kappa^{t-1} S]$

$\Leftrightarrow D(1^B)[1 - \rho] > RE(1^B)\kappa^{t-1}(\rho - \kappa)S \Leftrightarrow S < \frac{1-\rho}{\kappa^{t-1}(\rho-\kappa)} \frac{D(1^B)}{RE(1^B)}$

Then, profile e'' is preferred to profile e' if and only if

$CP(e', S) > CP(e, S) \Leftrightarrow \rho^{T-(t+B+1)}[D(1^B) + RE(1^B)\kappa^{t+1} S] > \rho^{T-(t+B)}[D(1^B) + RE(1^B)\kappa^t S]$

$\Leftrightarrow \frac{1-\rho}{\kappa^t(\rho-\kappa)} \frac{D(1^B)}{RE(1^B)} > S$

Since $\kappa < 1$, $\frac{1-\rho}{\kappa^{t-1}(\rho-\kappa)} \frac{D(1^B)}{RE(1^B)} > S \Rightarrow \frac{1-\rho}{\kappa^t(\rho-\kappa)} \frac{D(1^B)}{RE(1^B)} > S$. This proves the result.

PROPOSITION When $p = 0$, the critical thresholds $H(\tau, n|\rho, \kappa)$ satisfy:

a) $H(\tau, n|\rho, \kappa) = \frac{1-\rho^{\tau-n}}{(\rho^{\tau-n}-\kappa^{\tau-n})} \frac{[\frac{1-\kappa^n}{1-\kappa} - \rho\frac{\rho^n-\kappa^n}{\rho-\kappa}]\frac{1}{1-\rho}}{\kappa\frac{\rho^n-\kappa^n}{\rho-\kappa}}$

b) $H(\tau, n|\rho, \kappa)$ is decreasing in ρ

c) $H(\tau + 1, n + 1|\rho, \kappa) > H(\tau, n|\rho, \kappa)$

PROOF

Part (a): Let $p = 0$, $S_0 = S$ and $T > B$. We know that it is optimal to act in consecutive periods starting in period 1 (profile e), or in period $T-B+1$ (profile e'). The respective continuation payoffs are

$e : S|1, 1, ..., 1, 0, 0, 0, 0, 0 \rightarrow \rho^{T-B}[D(1^B) + S \cdot RE(1^B)]$

$e' : S|0, 0, 0, 0, 0, 1, 1, ..., 1 \rightarrow [D(1^B) + \kappa^{T-B} S \cdot RE(1^B)]$

Therefore, we act the first B periods if $S > \frac{1-\rho^{T-B}}{(\rho^{T-B}-\kappa^{T-B})} \frac{D(1^B)}{RE(1^B)}$. This implies that $\frac{1-\rho^{T-B}}{(\rho^{T-B}-\kappa^{T-B})} \frac{D(1^B)}{RE(1^B)}$ is the threshold when we have T periods to go, $\tau = T$ and B available actions, $n = B$. Substituting $D(1^B)$ and $RE(1^B)$ from above we obtain part (a) of the proposition.

Part (b): Rewrite $H(\tau, n|\rho, \kappa)$ as $\{\frac{1-\rho^{\tau-n}}{\kappa(\rho^{\tau-n}-\kappa^{\tau-n})} \frac{1}{1-\rho}\}\{\frac{[\frac{1-\kappa^n}{1-\kappa} - \rho\frac{\rho^n-\kappa^n}{\rho-\kappa}]}{\frac{\rho^n-\kappa^n}{\rho-\kappa}}\}$. Notice that

$\frac{\rho^n-\kappa^n}{\rho-\kappa} = \rho^{n-1} + \rho^{n-2}\kappa + \rho^{n-1}\kappa^2 + ... + \kappa^{n-1}$

which is increasing in ρ. We can also rewrite $D(1^n)$ as

$\frac{1-\kappa^n}{1-\kappa} - \rho\frac{\rho^n-\kappa^n}{\rho-\kappa} = 1 + \kappa + \kappa^2 + ... + \kappa^{n-1} - \rho^n - \rho^{n-1}\kappa - ... - \kappa^{n-1}\rho$

which is decreasing in ρ. Therefore, the expression $\frac{[\frac{1-\kappa^n}{1-\kappa} - \rho\frac{\rho^n-\kappa^n}{\rho-\kappa}]}{\frac{\rho^n-\kappa^n}{\rho-\kappa}}$ is decreasing in ρ. Now I show that the expression $\frac{1-\rho^{\tau-n}}{\kappa(\rho^{\tau-n}-\kappa^{\tau-n})} \frac{1}{1-\rho}$ is also decreasing in ρ. Forgetting the κ in the denominator and replacing $x = \tau - n$, we can rewrite it as $\frac{1+\rho+\rho^2+...+\rho^{x-1}}{\rho^x-\kappa^x}$.

The derivative w.r.t. ρ is negative as long as the following is also negative

$[1 + 2\rho + 3\rho^2 + ... + (x-1)\rho^{x-2}][\rho^x - \kappa^x] - (1 + \rho + \rho^2 + ... + \rho^{x-1})x\rho^{x-1} =$

$[1 + 2\rho + 3\rho^2 + ... + (x-1)\rho^{x-2}][\rho^x - \kappa^x] - (1 + \rho + \rho^2 + ... + \rho^{x-1})x\rho^{x-1} -$

$x\rho^{x-1} - x\rho^x - x\rho^{x+1} - ... - x\rho^{2x-2} =$

$-x\rho^{x-1} - \rho^x(x-1) - \rho^{x+1}(x-2) - ... - \rho^{2x-2}(x-(x-1)) - \kappa^x[1+2\rho+3\rho^2+...+(x-1)\rho^{x-2}]$

which is indeed negative, since all terms are negative.

Part (c): $H(\tau+1, n+1|\rho, \kappa) > H(\tau, n|\rho, \kappa)$ if and only if $\frac{[\frac{1-\kappa^{n+1}}{1-\kappa} - \rho\frac{\rho^{n+1}-\kappa^{n+1}}{\rho-\kappa}]}{\rho^{n+1}-\kappa^{n+1}} > \frac{[\frac{1-\kappa^{n}}{1-\kappa} - \rho\frac{\rho^{n}-\kappa^{n}}{\rho-\kappa}]}{\rho^{n}-\kappa^{n}} \Leftrightarrow (1-\kappa^{n+1})(\rho^n - \kappa^n) - (1-\kappa^n)(\rho^{n+1} - \kappa^{n+1}) > 0 \Leftrightarrow (\rho^n - \kappa^n) - (\rho^{n+1} - \kappa^{n+1}) + \rho^n\kappa^n(\rho - \kappa) > 0 \Leftrightarrow \rho^{n-1}(1-\rho) + \rho^{n-2}\kappa(1-\rho^2) + \ldots + \kappa^{n-1}(1-\rho^n) > 0$ which is indeed true.

The propositions for the model including bad news, i.e., $e \in \{-1, 0, 1\}$ are proved using identical techniques. Detailed proofs are available upon request from the author.

1.7.5 Advertising campaigns

In order to formulate the problem in a more familiar mathematical setting, we can integrate forward (1.18). Then,

$\dot{G}_t = u(a_t, G_t) - \delta G_t \Leftrightarrow e^{\delta t}\{\dot{G}_t + \delta G_t\} = e^{\delta t}u(a_t, G_t) \Leftrightarrow \frac{d}{dt}(e^{\delta t}G_t) = e^{\delta t}u(a_t, G_t) \Leftrightarrow$

$e^{\delta t}G_t|_0^T = \int_0^T e^{\delta t}u(a_t, G_t)dt \Leftrightarrow G_T e^{\delta T} - G_0 = \int_0^T e^{\delta t}u(a_t, G_t)dt \Leftrightarrow$

$G_T = G_0 e^{-\delta T} + e^{-\delta T}\int_0^T e^{\delta t}u(a_t, G_t)dt$

Therefore maximizing $G(T)$ subject to (1.18) and (1.21) is identical to maximizing $\int_0^T e^{\delta t}u(a_t, G_t)dt$ under the same set of constraints. This is a standard problem in the calculus of variations (or optimal control theory) with the following first order conditions:

$$\partial a \; : \; u_a[1 + \psi_t] = \phi_t \tag{1.38}$$

$$\partial G \; : \; \dot{\psi} + u_G[1 + \psi_t] = 0 \tag{1.39}$$

with $\psi_t = \lambda_t e^{-\delta t}$, $\phi_t = \mu e^{-\delta t}$, where λ_t and μ are the lagrange multipliers corresponding to the $\overset{\bullet}{G}$ and the isoperimetric constraints respectively. Differentiating (1.38) w.r.t. time we obtain:

$$[u_{aa}\overset{\bullet}{a}+u_{aG}\overset{\bullet}{G}][1+\psi_t]+u_a\overset{\bullet}{\psi}=\overset{\bullet}{\phi} \tag{1.40}$$

Dividing by $u_a[1+\psi_t]=\phi_t$, using (1.39) and the fact that $\frac{\overset{\bullet}{\phi}}{\phi}=-\delta$, we obtain the euler equation as a function of the parameter ρ and the derivatives of $u(\cdot,\cdot)$.

$$[u_{aa}\overset{\bullet}{a}+u_{aG}\overset{\bullet}{G}]\frac{1}{u_a}+\frac{\overset{\bullet}{\psi}}{[1+\psi]}=\frac{\overset{\bullet}{\phi}}{\phi}\Leftrightarrow\frac{u_{aa}}{u_a}\overset{\bullet}{a}+\frac{u_{aG}}{u_a}\overset{\bullet}{G}-u_G+\delta=0 \tag{1.41}$$

The discrete time version of the problem leads to the following Lagrangian

$$L=\sum_{t=1}^{T}\rho^{T-t}u(a_t,\rho G_{t-1})-\lambda_t\sum_{t=1}^{T}[G_t-\rho G_{t-1}-u(a_t,\rho G_{t-1})]-\mu\sum_{t=1}^{T}a_t$$

where $\rho\equiv(1-\delta)$. The first order conditions are

$$\partial a_t\ :\ \rho^{T-t}u_a(a_t,\rho G_{t-1})+\lambda_t u_a(a_t,\rho G_{t-1})=\mu$$

$$\partial G_t\ :\ \rho^{T-t}u_G(a_{t+1},\rho G_t)-\lambda_t+\lambda_{t+1}\rho+\lambda_{t+1}\rho u_G(a_{t+1},\rho G_t)=0$$

Substituting and carrying out the algebra yields the euler equation

$$u_a(a_{t+1},\rho G_t)=\rho u_a(a_t,\rho G_{t-1})[1+u_G(a_{t+1},\rho G_t)]$$

Chapter 2

Release of Information and Strategic Manipulation of Memories: the convex cost case

2.1 Introduction

In chapter one, I built a simple model of memory, which I then used to examine the implications of imperfect memory for the problem of releasing information. In particular, I addressed the issue of how one should time a sequence of informative events in order to manipulate the memories that his forgetful assessor will have at some critical date. There are two important features in this model. First, the choice variable for the agent who controls the release of information is binary: in each period t you either time an event, or you do not. Second, the total amount of information to be released is predetermined, as there is a fixed number of events to be scheduled. I will refer to this as the fixed-budget binary choice model, or FBBC.

Here, I extend the model to allow for a continuous, rather than binary, choice variable and, more importantly, a convex cost for generating events, rather than a fixed "budget" of events. I believe that an extension along these lines is important for two reasons. First, it widens the applicability of the model to economic settings where this set of assumptions is more appropriate. Consider, for example, an employee who is evaluated at some critical future date on the basis of his past output performance. In each period he has to decide how much effort to exert. There is no reason to believe that working hard today decreases the amount of effort that he can spare in the future. Rather, it is more appropriate to assume that additional effort can always be generated by incurring some cost, psychic or physical. Second, extending the model in this direction paves the way to an empirical investigation of the model. For example, one could use data from actual advertising campaigns to see how companies spend advertising money prior to the launch of a new product. Once again, the decision variable, i.e. how much to spend on advertising in each period, is continuous and the company can always advertise more if it is willing to pay the additional cost.

Summary and Results

Recall from the detailed discussion in chapter one that there is a vast array of experimental evidence to support the idea that memory operates on the principles of similarity and repetition. This implies that current events have the direct effect of creating new memories and the indirect effect of reinforcing the memories of similar past events. We can think of past events as if they are creating a rehearsal stock, S. I show that the indirect effect of a current event depends only on the level of this rehearsal stock S, rather than the exact history profile of past events.

Then, we can cast the problem of information release as a stochastic dynamic optimization problem, where the agent tries to maximize the memorability of his past *successes*, i.e. events that are beneficial to his reputation, subject to the memory technology of his assessor. The state variable is exactly the rehearsal stock, S. In each period the agent controls the probability of a success by exerting effort, but unlike the model of chapter one there is no isoperimetric constraint on how much total effort the agent can exert. Rather, the cost of effort is given by some convex cost function.

The problem can be solved by the means of backward induction. A closed form solution is not available and I therefore solve the model numerically, by a grid search. I also prove some qualitative results that hold under the general case. In particular, I show that effort is an increasing function of the rehearsal stock, S. That is, the more past successes the agent has obtained, the more motivated he will be to produce new successes in the future. This is an important result because it demonstrates how stochastic events can generate action by the agents. Consider for example two identical employees who are evaluated at some future date T on the basis of their past output levels. At first, both start with equal rehearsal stocks and consequently they choose equal effort levels. The luckier of the two will enjoy higher output and in the next period he will have a higher rehearsal stock that he can reinforce. He will therefore choose a higher effort level than his unlucky colleague. Now, as he puts more effort it becomes more likely that he will be the lucky one in this second period and accumulate an even greater rehearsal stock, compared to his colleague's. He will therefore choose more effort in the subsequent period and so on. In an ironic way, if someone with perfect memory were monitoring the performance and the effort of the two employees, she would identify the "lucky" one as a better employee, failing, perhaps, to recognize that motivation depends crucially in early stochastic events that are beyond anyone's control.

A similar result was obtained in the context of the FBBC model of chapter one. There, I showed that the agent will be triggered to act and try to produce a success in a given period if the rehearsal stock in that period is above some threshold level. Nevertheless, these two similar results are shown to be attributed to entirely different forces. The assumption of convex cost of effort versus a fixed budget is crucial, altering decisively the driving forces between the two models.

In an effort to make this paper "read on its own", some repetition between the material in this chapter and chapter one is inevitable. In fact, I start by briefly summarizing the memory technology as described, in more detail, in chapter one. In section three I set up the model, solve it and prove some qualitative results. I conclude in section four by briefly discussing how the model lends itself to future empirical work.

2.2 A Brief Summary of the Memory Technology

The material presented in this section is a stripped down discussion of the memory technology as presented in chapter one. It can be bypassed by readers acquainted with it, but even they should know, by now, that rehearsal of past information can help solidify memories.

The model is cast in discrete time indexed by $t = 1, 2, 3.., T$. In each period $t \leq T$, an event may be realized, in which case we write $e_t = 1$, as opposed to $e_t = 0$, in case the event does not realize, that is $e_t \in \{0, 1\}$. A realization of the event in some period t, i.e. $e_t = 1$, will be referred to as a *success*.

To model imperfect recall, I define M_t^i to be the memorability, or strength, of the memory $e_i = 1$ at time t. Since a realization of an event can have meaningful memorability only after it has occurred, $M_t^i = 0$ for all $i > t$, and we normalize $M_i^i = 1$. At time T, an agent (referred to as the *assessor*) who has been observing the realizations of past events remembers a period i success only with some probability, using the event $e_i = 0$ as the default memory. I refer to this probability as the *recall probability* for success i, and I assume it is some increasing function of the memorability of success i at time T, M_T^i.

To model the evolution of M_t^i over time, I invoke the experimental results on memory. First, memories fade away with the passing of time, which dictates that M_t^i decreases with time. In particular, I assume that memories decay exponentially at a constant rate $(1-\rho)$. Second, to incorporate the properties of rehearsal and similarity, I assume that a success at time t will enhance the memorability of past successes, the enhancement being bigger for more recent successes.

I formalize this by defining b_t^i to denote the enhancement (boost) to the memorability of success i by a possible success at time t.[1] I let the rehearsal of a period i success from a success at time $t+1$ be a fraction $\kappa < 1$ of the rehearsal that this period i success would enjoy from a success at time t instead, that is $b_{t+1}^i = \kappa b_t^i$. It is also natural to assume that the increment to the memorability of a period t success from a success at $t+1$ is less than its current memorability, thus $\kappa < \rho$.[2] As long as $i < t$, we can neatly summarize this information in the following pair of equations:

$$M_{t+1}^i = \rho M_t^i + b_{t+1}^i I\{e_{t+1} = 1\} \tag{2.1}$$

$$b_{t+1}^i = \kappa b_t^i \tag{2.2}$$

where I is the usual indicator function, i.e. $I\{e_t = 1\} = 1$ if $e_t = 1$ and 0 otherwise.

---

[1]As with M_t^i, $b_t^i = 0$ for all $i > t$ and we normalize $b_i^i = 1$.

[2]Note that at time $t+1$ the memorability of a period t success is ρ.

Recall that the number of successes that a forgetful assessor actually remembers is a random variable, with upper bound the actual number of successes, $\sum_{i=1} I\{e_i = 1\}$. If the probabilities of recall at time T are proportional to the memorabilities of each past success at time T, M_T^i, then the *expected* number of successes that the agent recalls at time T is proportional to the *sum of the memorabilities* at time T, denoted by $A_T = \sum_{i=1}^{T} M_T^i$. Using (1.1) and (1.2) together with the initial conditions $M_i^i = b_i^i = 1$, we can sum over i to show that A_t evolves as follows:

$$A_{t+1} = \rho A_t + I\{e_{t+1} = 1\} \cdot (1 + \kappa S_t) \tag{2.3}$$

$$S_{t+1} = \kappa S_t + I\{e_{t+1} = 1\} \tag{2.4}$$

where $S_t = \sum_{i=1}^{t} b_t^i$.

This new pair of equations allows for a different interpretation of the model. Memory decay, embedded in ρ, can be thought of as a discount rate on the *past*, rather than on the *future*. Then, one could think of A_t as a *running total* stock variable that measures successes that have occurred up to time t, each one "discounted" at rate $(1-\rho)$. In each period that a new success is realized the stock A_t is increased, but the incremental effect of a new success has two distinct components. A *direct* effect (plus 1), originating simply from the fact that a new success, and thus a new memory, has been realized, plus an *indirect* effect, (plus κS_t) coming from the fact that the new success triggers memories of past successes and therefore become more memorable.

Our reduced form equations (2.3) and (2.4) give a convenient way to record this indirect effect. One might think that the indirect effects of a new success would depend on the exact sequence of zeros and ones that have occurred up to that point. In fact, however, the variable S_t acts as a summary statistic for the sequence: the greater is S_t the greater the indirect effect from a new success will be. We can think of $S_t = \sum_{i=1}^{t} b_t^i$ as a "*rehearsal stock*". Just as the running total stock A_t decays at rate $(1-\rho)$, it is as if the rehearsal stock S_t "decays" at rate $(1-\kappa)$. Our earlier assumption that the rehearsal of an event from time i by a similar event at time t is less than the current memorability of that event, translates to saying that the rehearsal stock S_t decays a faster rate than the actual stock A_t $(\kappa < \rho \Leftrightarrow (1-\kappa) > (1-\rho))$.

I believe that being able to summarize, at each point in time, the past in a stock variable, rather than having to carry the whole vector of past realizations, is an attractive feature of the model. This proves to be a great simplification that allows us to set up the problem of information release as a dynamic optimization problem, where the rehearsal stock, S_t, is one of the state variables.

2.3 The Model with Convex Cost of Effort

To motivate the model, consider an employee who will be evaluated by his boss, at some known future date T, on the basis of his past performance. We can think of this evaluation at time T as an end-of-year promotion, or bonus, decision. In each period the employee's performance can be either plain and ordinary, making no impression to the boss, or it may be well above average and hence noticeable. In the latter case, we say that the employee had a *success*. These successes are the events that the boss is called on to remember at time T, her faulty memory being governed by the technology described in the previous section. The more successes she remembers the higher the bonus to the employee, or the higher the chance for a promotion.

Each period the employee chooses effort, how hard to work. Let c_t denote the choice variable at time t and assume it lies in a closed interval, that is $c_t \in [\underline{c}, \bar{c}]$. Effort comes with some increasing *convex* cost function $q(c)$, but it increases the probability of a success for that period, through some increasing *concave* function $p(c)$, with $p(\underline{c}) = 0$, $p(\bar{c}) = 1$.

A risk neutral employee will choose his effort so as to maximize the expected number of successes that his assessor will remember at time T, which is A_T, minus the total cost of effort. The maximization is of course subject to the assessor's memory technology, as summarized in the pair of equations (2.3) and (2.4). Therefore, the employee solves

$$
\begin{aligned}
& \max \;\; A_T - \sum_{i=1}^{T} q(c_t) && (2.5)\\
s.t.\; A_{t+1} &= \rho A_t + I\{e_{t+1} = 1\} \cdot (1 + \kappa S_t) \\
S_{t+1} &= \kappa S_t + I\{e_{t+1} = 1\} \\
\Pr(e_t &= 1 | c_t) = p(c_t)
\end{aligned}
$$

Comparing this to the maximization problem in chapter one we can spot two differences. First, the decision variable is no longer binary, $c_t \in \{\underline{c}, \bar{c}\}$,[3] rather it

[3]In the jargon of the old model, $c_t \in \{Act, Wait\}$.

is continuous in the interval $[\underline{c}, \bar{c}]$. More importantly, the employee of this model can generate more effort each period by incurring some direct psychic cost. In the previous model, there was no direct cost of effort, but today's effort reduced the amount of effort that the agent could spare in subsequent periods. This may be the situation faced by a politician trying to decide when to time a predetermined number of events, say, announce a popular tax cut and schedule a TV appearance that will generate positive publicity. Or an advertiser deciding when to schedule the airing of a limited number commercial spots. However, in the employee-evaluation example the assumption of convex cost of effort fits better.

2.3.1 The Bellman equation

We now transform the stochastic dynamic maximization problem of (2.5) to the usual Bellman equation formulation. Assume that we are in period t, carrying over some running total A_t and a rehearsal stock S_t. Recall that a success in period t has three effects. First, it produces a new memory, adding one unit to the running total. Second, it reinforces past memories, adding the term κS_t to the running total for a combined effect of $(1 + \kappa S_t)$. Third, this new memory increases the rehearsal stock for the future.

These are the effects that a period t success has on the running total A_{t+1} and on the rehearsal stock S_t, but it increases the agent's payoff only through the effect it has on A_T, the running total at the critical time T. Since the running total decays at the rate $(1 - \rho)$, the first two effects increase A_T, the running total at time T, by

$\rho^{T-t}(1+\kappa S_t)$. The third effect is accounted for by "updating" the rehearsal stock S_t to S_{t+1} through equation (2.4). Therefore, even though the agent receives her payoff at time T, it is as if she receives a *flow* payoff each period there is a success, which is discounted at rate ρ.[4] As a result, A_t, the running total at time t, can be dropped, leaving us with two state variables: the rehearsal stock S and the number of periods to go, denoted by τ[5]. Then, denoting by $V(S,\tau)$ the continuation payoff starting with a rehearsal stock S and with τ periods to go, we have

$$\max_c V(S,\tau) = p(c)[\rho^{\tau-1}(1+\kappa S) + V(1+\kappa S, \tau-1)] + [1-p(c)]V(\kappa S, \tau-1) - q(c) \tag{2.6}$$

Since the problem terminates at time T, it can be solved by backward induction. With only one period to go, the agent simply solves

$$\max_c V(S,1) = p(c)(1+\kappa S) - q(c) \tag{2.7}$$

If we assume specific functional forms for $p(\cdot)$ and $q(\cdot)$ we can solve for the optimal choice for c as a function of the rehearsal stock, S, which we can then plug into (2.7) to obtain $V(S,1)$ as a function of the rehearsal stock, S, and the two decay parameters, ρ and κ. For example, let $p(c) = c^\beta$ and $q(c) = c^\gamma$, with $0 \le \beta \le 1$ and $\gamma > 1$. Then,

---

[4]Recall that here we discount the past and not the future.

[5]That is $\tau = T - t + 1$

(2.7) is solved by letting[6]

$$c_1 = [\frac{\beta}{\gamma}(1+\kappa S)]^{\frac{1}{\gamma-\beta}} \tag{2.8}$$

provided that $c_1 \leq 1$ (and $c_1 = 1$ otherwise). This in turn implies (ignoring the corner solution $c_1 = 1$)

$$V(S,1) = [\frac{\beta}{\gamma}(1+\kappa S)]^{\frac{\beta}{\gamma-\beta}}(1+\kappa S) - [\frac{\beta}{\gamma}(1+\kappa S)]^{\frac{\gamma}{\gamma-\beta}} = [\frac{\beta}{\gamma}(1+\kappa S)]^{\frac{\gamma}{\gamma-\beta}}(\frac{\beta}{\gamma}-1) \tag{2.9}$$

Now, with two periods to go the agents solves

$$\max_c p(c)[\rho(1+\kappa S) + V(1+\kappa S, 1) - V(\kappa S, 1)] - q(c) + V(\kappa S, 1) \tag{2.10}$$

The optimal choice is[7].

$$c_2 = \{\frac{\beta}{\gamma}[\rho(1+\kappa S) + V(1+\kappa S, 1) - V(\kappa S, 1)\}^{\frac{1}{\gamma-\beta}} \tag{2.11}$$

Using (2.9) to substitute for $V(1+\kappa S, 1)$ and $V(\kappa S, 1)$ one can obtain the optimal choice with two periods, c_2, only as a function of the rehearsal stock, S, and the parameters ρ, κ, β and γ. This, in turn, will give us $V(S,2)$, the value function with two periods to go, only as a function of the rehearsal stock S and the parameters, which we can use to compute the optimal choice with three periods to go, c_3, and so on. However, one can see that the expressions for the optimal choice and the value

[6]The subscript i on the choice variable, c_i, denotes that this is the optimal choice with i periods to go.

[7]provided that $c_2 \leq 1$ and $c_2 = 1$ otherwise

function quickly become very complicated as we increase the number of periods to go, and a closed form solution is not easily obtained. Of course, there may be other functional forms for the probability mapping $p(\cdot)$ and the cost function $q(\cdot)$ that are more convenient. Since I could not find any I proceed with numerical simulations.

2.3.2 Simulation results

In this section I simulate the model using the functional forms from the earlier example, i.e. $p(c) = c^\beta$ and $q(c) = c^\gamma$. Also, for the rest of the discussion I assume the following parameter values: $\rho = 0.99, \kappa = 0.35, \beta = 1, \gamma = 2$. To simulate the model numerically, we need to discretize the state space. I let the rehearsal stock S take the values $\{0, d, 2d, 3d, .., \frac{1}{1-\kappa}\}$[8], where d is some small increment, say 0.001. Time is discrete by assumption. By discretizing the state space, we are able to compute the continuation payoff and the optimal effort level, with τ periods to go, for *each possible* value of the rehearsal stock S, instead of having to express them as (complicated) functions of the rehearsal stock S.

With τ periods to go, knowing the exact numerical value of the continuation payoff for each possible value of the rehearsal stock S, allows us to compute the optimal effort level for each possible value of the stock and $\tau + 1$ periods to go. This enables us to compute the value that the continuation payoff takes, for each possible value of the stock, and with $\tau + 1$ periods to go. And so on, inductively.

---

[8]Since the rehearsal stock S is a sum of powers of κ, it is at most $\frac{1}{1-\kappa}$.

I then simulate the model to obtain the expected effort level for each period, starting with zero initial rehearsal stock and $T = 10$ periods to go.

Figure 1

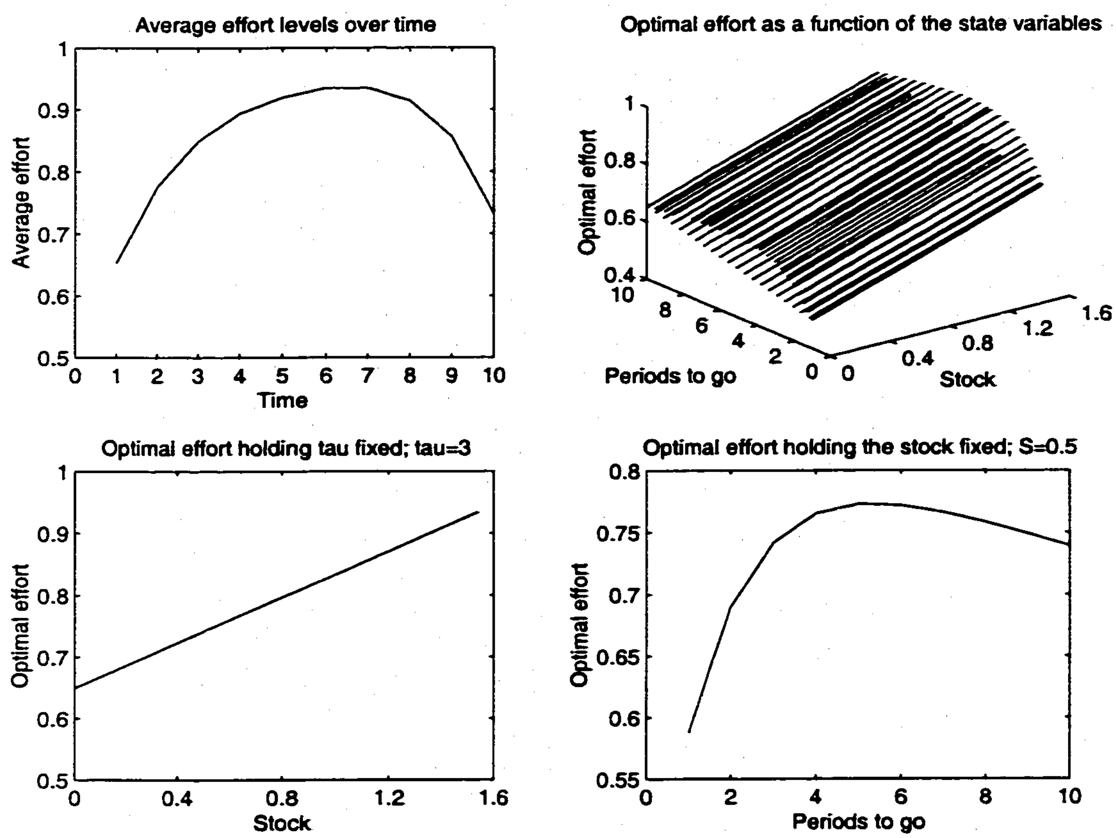

The results are shown in the upper left portion of figure one. Average effort takes this inverted (and tilted to the right) U-shape over time, which is typical and holds for the whole range of possible parameter values. The agent initially increases his effort level over time. After a certain point in time, when there are not many periods to go, he starts decreasing his effort. In other words, he works early on to create a high rehearsal stock. This makes his future actions more effective, allowing him to slack off and reap the benefits of his past hard work.

To give some intuition for this result, I plot in the upper right portion the optimal effort level as a function of the state variables, periods to go and the rehearsal stock. For a clearer picture look at the lower left portion, which plots the optimal effort level against all possible values of the rehearsal stock S, keeping the number of periods to go *fixed*, at $\tau = 3$. We should extract the fact that effort is increasing in the rehearsal stock[9]. This result is discussed at great length in the next subsection. Similarly, the lower right portion plots the optimal effort level against the number of periods to go, keeping the rehearsal stock fixed, at $S = 0.5$. We see that effort initially[10] increases as the number of periods to go decreases. Recent successes are more memorable, therefore the agent should work harder as time goes by. However, there is a second effect that decreases effort as the number of periods to go decreases. Working harder leads to a higher probability of success and therefore a higher rehearsal stock for the future. The fewer the periods to go, the fewer the number of periods that the agent will be able to exploit this higher stock, and as a result the less of an incentive she has to work hard. For example, with two periods to go, a higher probability of success gives you a higher chance for grabbing the flow payoff of $\rho(1 + \kappa S)$, *and* a higher probability of making your future rehearsal stock $(1 + \kappa S)$ rather than κS. With only one period to go, you only care about the flow payoff. The rehearsal stock that you will carry over for the future is not important since the model ends.

---

[9]The linearity is a consequence of the fact that we have chosen $\beta = 1$ and $\gamma = 2$. It does not hold in general. See equation (2.8).

[10]"Initially" here means reading the graph from the end, when we have many periods to go.

2.3.3 A general result

In the numerical simulations effort appeared to be increasing in S, the rehearsal stock. Next, I show that this result is a general feature in the model, and holds for arbitrary choices for the parameters (ρ and κ), for any concave probability mapping $p(\cdot)$ and any convex cost function $q(\cdot)$. Focus on the first order condition, obtained by the Bellman equation in (2.6) given by

$$p'(c_\tau)[\rho^{\tau-1}(1+\kappa S)+V(1+\kappa S,\tau-1)-V(\kappa S,\tau-1)]=q'(c_\tau) \tag{2.12}$$

Applying the implicit function theorem we can differentiate both sides with respect to the rehearsal stock, S, to obtain, after some algebraic manipulation,

$$\frac{\partial c^*(S,\tau)}{\partial S}=\frac{p'(c^*_\tau)\kappa[\rho^{\tau-1}+V_1(1+\kappa S,\tau-1)-V_1(\kappa S,\tau-1)]}{\{q''(c^*_\tau)-p''(c^*_\tau)[\rho^{\tau-1}(1+\kappa S)+V(1+\kappa S,\tau-1)-V(\kappa S,\tau-1)]\}} \tag{2.13}$$

where c^* is the optimal choice that solves (2.12), and $V_1(\cdot,\cdot)$ denotes the partial derivative with respect to the first argument. Since the cost function $q(\cdot)$ is convex and the probability mapping $p(\cdot)$ is concave the denominator is positive. The numerator will be positive provided that the continuation payoff is convex with respect to its first argument, i.e. $V_{11}(\cdot,\cdot)>0$. To compute this second derivative, we differentiate the Bellman equation in (2.6) with respect to the rehearsal stock S, twice. First we get[11]

$$V_1(S,\tau)=p(c_\tau)\kappa[\rho^{\tau-1}+V_1(1+\kappa S,\tau-1)-V_1(\kappa S,\tau-1)]+V_1(\kappa S,\tau-1) \tag{2.14}$$

[11] Recall, that we can ignore the change in the choice variable c_τ due to the envelope theorem.

Differentiating once again with respect to the rehearsal stock, S,

$$\begin{aligned} V_{11}(S,\tau) &= p'(c_\tau)\frac{\partial c(S,\tau)}{\partial S}\kappa[\rho^{\tau-1}+V_1(1+\kappa S,\tau-1)-V_1(\kappa S,\tau-1)] \quad (2.15) \\ &\quad +p(c_\tau)\kappa V_{11}(1+\kappa S,\tau-1)+[1-p(c_\tau)]\kappa V_{11}(\kappa S,\tau-1) \end{aligned}$$

We now have the machinery to prove the following result.

PROPOSITION 1 *For all periods to go, τ, effort is increasing*[12] *in the rehearsal stock, that is $\frac{\partial c(S,\tau)}{\partial S} \geq 0$. Morover, the continuation payoff is convex with respect to its first argument, that is $V_{11}(\cdot,\cdot) > 0$.*

Proof. The proof is by induction on the number of periods to go, τ. It is straightforward to verify that $\frac{\partial c^*(S,1)}{\partial S} \geq 0$ (with equality when we have the corner solution $c^*(S,1) = 0)$) and $V_{11}(S,1) > 0$. Assume that for some $\tau > 1$, $V_{11}(S,\tau) > 0$. Then, by (2.13) we can be assured that $\frac{\partial c^*(S,\tau+1)}{\partial S} \geq 0$ (with equality for the corner solution $c^*(S,\tau+1) = 1$). This in turn implies, by (2.15), that $V_{11}(S,\tau+1) > 0$. ■

This result has the important implication that stochastic events that increase an agent's rehearsal stock can motivate the agent to exert more effort. More importantly, stochastic events will force two identical agents to choose different effort levels, thus inducing their actions to diverge in the future as well. To see that, assume that two identical employees will be evaluated at some future period T. They start with the same rehearsal stock, and as a result they choose the same effort level for the first period. Assume that one gets lucky and obtains a success, whereas the second is

[12]It is weakly increasing only for the special case of a corner solution, $c^*(S,\tau) = 1$

unlucky and does not get a success. In the second period, the lucky employee will have a higher rehearsal stock and will choose a higher effort level. This means that a success in the second period is more likely for him than for his unlucky colleague who chose a lower effort level. Therefore, it is even more likely the first one will choose to work harder in the third period and so on.

This situation is depicted in figure two where I plot the optimal effort level over time for the "most lucky" employee, who always obtained a success, and the "most unlucky" employee who never obtained a success. In period one both choose the same effort level. In period two the lucky one has a higher stock to rehearse and he has an incentive to work harder. In period three the difference between the rehearsal stocks of the two employees is even greater and their respective actions diverge even further, and so on. Notice, that in producing the plot in figure two we *assumed* that the lucky employee will always produce a success and the unlucky one will never produce one. Nevertheless, this assumption is "reinforced" by the actions of the agents, since the lucky employee works harder and has indeed a higher chance of obtaining a success.

It is also ironic that an assessor with *perfect* memory monitoring the two employees would conclude that the "lucky" employee has worked more and deserves a higher payoff, perhaps not realizing that the difference in their performance is attributed to early stochastic events, beyond the employees' control.

A similar result was obtained in the context of the FBBC model of chapter one. In particular, it was shown that the agent will be triggerred to act early on when the rehearsal stock is above some threshold. In other words, the greater the rehearsal

Figure 2

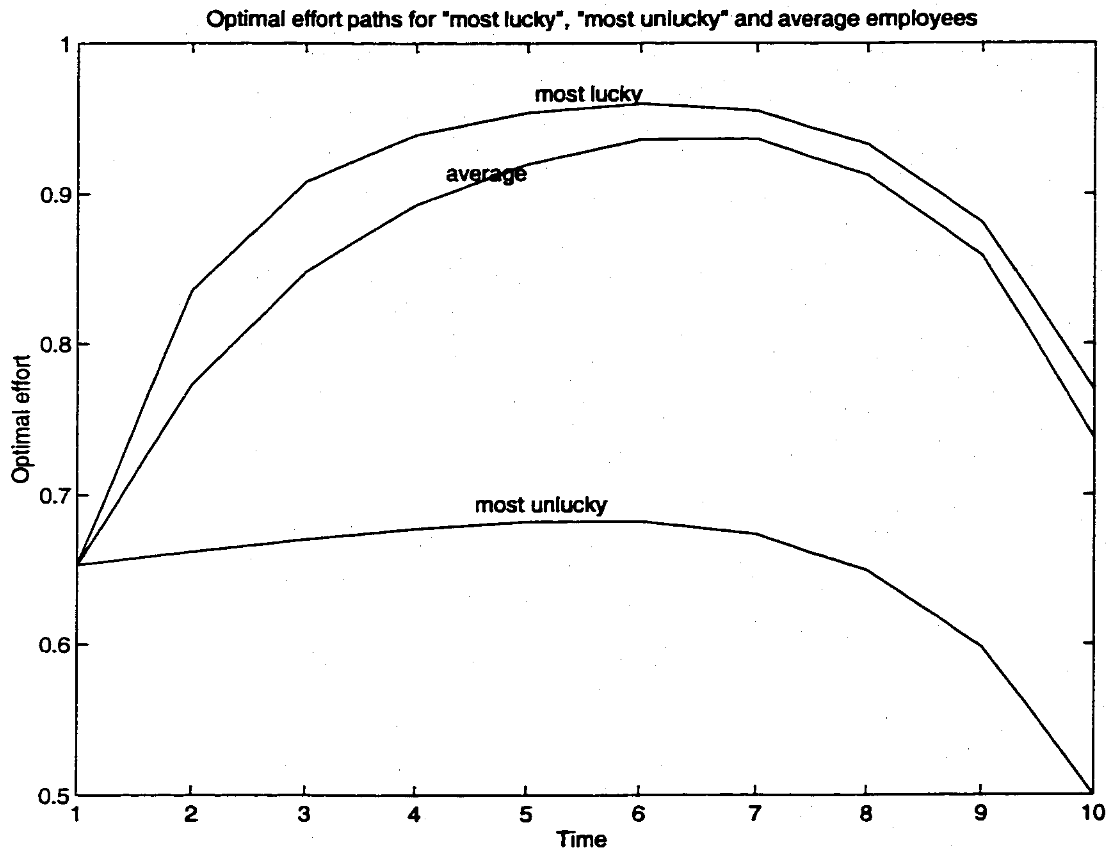

stock, the greater the incentive to deliver a success now, rather than later on. Nevertheless this result, seemingly related to proposition 1, is attributed to different forces. A quick way to see this is to notice that the result of the FBBC model rested crucially in the assumption $\rho > \kappa$, that the rehearsal stock decays faster than memory. This assumption, however, was not used in the proof of proposition 1.

For a better explanation, recall that in the FBBC model the agent had a fixed budget of effort, facing the question of *when* it is optimal to spare some effort and increase the probability of a success. On the one hand, sparing some effort early on is desirable because a potential success will reinforce existing past successes, whose memories are decaying very fast. This was termed the rehearsal effect. On the other hand, waiting to spare the limited amount of effort later on produces successes that are themselves more memorable, termed the recency effect. Optimal actions are dictated by a trade off between the recency and the rehearsal effect. The greater the rehearsal stock, the greater the rehearsal effect, and hence the greater the incentive to work hard. In the present model this trade off between recency and rehearsal is not present since additional effort can always be generated, albeit at some cost. Instead the agent equates the marginal benefit of effort to the marginal cost of effort, as in (2.12). The higher the rehearsal stock, the greater the marginal benefit of effort, and this is for two reasons. First a higher stock results in a higher flow payoff which is received only if there is a success, and second, more rehearsal stock increases the difference between the continuation payoff if a success occurs, $V(1+\kappa S,\tau)$, and the continuation payoff if a success does not occur, $V(\kappa S,\tau)$.

2.4 Testable Implications?

In this chapter I extended the model of chapter one to include a continuous choice variable and convex cost for generating events. I do not wish to make the claim that one model is better than the other, or that one is a special case of the other. Rather, I see the two models as complementing each other, each being applicable in settings where its assumptions are more relevant. However, I do wish to claim that the continuous choice convex cost model lends itself more easily to future empirical work, which I briefly outline here.

The idea that memory relies heavily on similarity and repetition means that current events not only create new memories, but that they rehearse similar past memories as well. This fact has the important implication that the present decisions of an agent manipulating her assessors' memories will depend on past events. In particular, the main result of this chapter was that the more memorable past favorable events are, the more motivated will the agent be to work hard, or spend more.

Political campaigns provide a real world setting where one could potentially test this hypothesis. Politicians spend big sums of money in advertisng campaigns in an effort to create favorable impressions for themselves, and to a lesser extend, unfavorable impressions for their opponents. Assume that we had access to the expenditure profile over time for various political campaign teams. We could then use opinion polls at regular intervals as a proxy to summarize the effectiveness of past advertising on the public's perceptions, just as the rehearsal stock S summarizes the past history of events in the abstract model. The hypothesis of the model would be verified if we could detect a tendency to spend more money right afer the publication of a favorable poll.

Chapter 3

Revising Non-Additive Priors

With Ricky Lam

Many thanks to Ben Polak for teaching us decision theory, and for many fruitful discussions. Ettore Damiano, David Pearce and Mario Simon made many helpful comments during the course of this project.

3.1 Introduction

In a wide range of dynamic economic situations with incomplete information, agents are required to update their initial beliefs upon the receipt of some informative signal or message. For example, an employer may update her prior on the quality of a worker after observing the worker's output. Or the manager of a potential entrant in an industry may update his prior of being fought, after observing the actions of the incumbent firm to previous entrants.

In the belief revision process, five different probability measures may be involved. In order to clarify this, and to illustrate the questions addressed by this paper, we consider the following concrete example. An employer has just hired an employee. The worker can either be one who exerts high effort (θ_H) or one who exerts low effort (θ_L). It is assumed that the employer does not observe the worker's type but she does observe the worker's output, which again can be either high (y_H) or low (y_L). Define the set of possible types and output levels to be $\Theta = \{\theta_H, \theta_L\}$ and $Y = \{y_H, y_L\}$, respectively. The state space is all possible combinations of the worker's type and the output produced; we denote this by: $S = \Theta \times Y = \{(\theta_H, y_H)$, (θ_H, y_L), (θ_L, y_H),

$(\theta_L, y_L)\}$. The five probability measures are:

$\nu : 2^{\Theta} \to [0,1]$	The unconditional prior over types Θ
$\mu : 2^{Y} \to [0,1]$	The unconditional measure over signals Y
$\sigma : 2^{S} \to [0,1]$	The joint measure over the product space $S = \Theta \times Y$
$\mu(. \mid \theta) : 2^{Y} \to [0,1]$	Conditional likelihood over signals Y, given a type θ in Θ
$\nu(. \mid y) : 2^{\Theta} \to [0,1]$	Posterior over types Θ, given a signal y in Y

If these measures are all additive, the information contained in them can be summarized in three equivalent ways: (a) by σ, the joint probability over the product space; (b) by the prior for types ν, together with the set of likelihoods $\{\mu(. \mid \theta)\}_{\theta \in \Theta}$; and (c) by the unconditional measure over signals μ, together with the set of posteriors over types $\{\nu(. \mid y)\}_{y \in Y}$. Bayes's theorem allows us to move among these representations.

From the point of view of economic applications, agents typically possess information in the form of (b). It is natural to assume that the employer has initial beliefs over the quality of the worker, and that her knowledge of the production process implies knowledge about the distribution over output, conditioned on each of the worker's types.

Now, imagine that the employer's knowledge is indeed in the form of (b), and that a low output (y_L) is realized. How does she update her beliefs on the employee's type? The updating problem is trivial. First, the employer transforms the representation in (b) to that in (a) by a simple rearrangement of Bayes's rule. For all $\theta \times y$ in $\Theta \times Y$,

$$\sigma[(\theta, y)] = \nu(\theta) \cdot \mu(y \mid \theta) \tag{3.1}$$

Having obtained the joint beliefs over the product space, another application of Bayes's rule produces the posterior probability that the employee is type θ:

$$\nu(\theta \mid y_L) = \sigma[(\theta, y_L) \mid \{(\theta_H, y_L), (\theta_L, y_L)\}] = \frac{\sigma[(\theta, y_L)]}{\sigma[(\theta_L, y_L)] + \sigma[(\theta_H, y_L)]} \tag{3.2}$$

This is essentially moving from representing the information using (a) to representing it by (c).

In this updating framework, the prior measure ν is subjective while the likelihood distributions $\{\mu(. \mid \theta)\}_{\theta \in \Theta}$ are often objective.[1] Recent work in decision theory has sought to represent the subjective beliefs of uncertainty-averse agents in the form of a non-additive measure. Schmeidler (1989) and Gilboa (1987) show that if the decision maker's preferences satisfy certain axioms that are consistent with uncertainty aversion, then they choose as if they are maximizing Choquet expected utility. That is, preferences can be represented by a utility function which requires an expectation with respect to a non-additive measure.

If our hypothetical employer possesses such a prior, the three ways of representing information discussed above are no longer equivalent: revising her beliefs over types

[1]In the traditional view among economists, the subjective probability measure should be thought of as arising from some representation of the decision maker's preferences. In particular, Savage (1954) and Anscombe and Aumann (1963) outline the axioms for an expected-utility representation. Bayes's rule can also be justified through preferences. (See Myerson 1991.)

in the light of signals is no longer so obvious. This non-equivalence arises because Bayes's theorem does not hold for non-additive measures. One may think that the Dempster–Shafer rule for calculating conditional capacities—which we will describe subsequently—can be used in place of Bayes's rule. This is partly justified by the work of Gilboa and Schmeidler (1993), who show that a particular form of "pessimism" in preferences leads to the rule as an updating device for non-additive measures.[2]

By analogy to the additive case, one may attempt to use the Dempster-Shafer rule to construct joint beliefs σ and then condition on the relevant partition of the product space to obtain posterior beliefs over types $\nu(. \mid y)$. Unfortunately, the first stage of this procedure fails. In general, unique beliefs over the state space S cannot be obtained from the Dempster-Shafer rule alone. This has important implications because in many economic applications, such as the example here, beliefs over the state space S are not given in the specification of the problem. Although the Dempster-Shafer rule can be used to calculate posteriors once the joint measure is known, our maintained assumption is that information is presented to the economist in the form of (b).

We propose two rules for defining a measure over the space S. Under the first proposal, the value of a set in 2^S is given by the iterated expectation of the corresponding indicator function. Expectation is first taken with respect to $\{\mu(. \mid \theta)\}_{\theta \in \Theta}$,

[2]The assumption is that when conditioning preferences on a particular event, the agent assumes that the best possible outcome obtains in the impossible states. See Gilboa and Schmeidler for details.

and then with respect to ν. We refer to this procedure as the *Choquet-indicator rule*. With additive probability measures, this is of course the correct thing to do because the expectation of an indicator function over a set is the probability of that set. When beliefs are non-additive, we show that this rule still has desirable properties. In the second approach, we recognize that the perception of uncertainty embodied in ν can be equivalently represented by a set of additive measures, which we denote by P. Each of these distributions over the type space Θ can be taken in turn and used to construct a probability over the state space S. We refer to this rule as the *multiple-priors rule*. It produces a set of distributions, denoted by Q.

The two rules are closely related, but not equivalent. This non-equivalence arises because non-additive measures are unable to capture certain restrictions on the relative likelihood of events. While this does not matter for the representation of uncertainty-averse beliefs, it results in a loss of information when beliefs have to be revised.

The updating problem considered in this paper is in fact closely related to a theoretical question which has received some attention in the literature. When an individual has non-additive beliefs, whether the objects of choice are Anscombe-Aumann "*horse-lotteries*" (functions from states to *lotteries* over consequences) or Savage *acts* (functions from states to consequences) affects her preference for randomization. (Eichberger and Kelsey 1996) This in turn has implications for the desirability of mixed strategies in games with uncertainty-averse players.

To see the relationship between this literature and our paper, note that the signal processing example has two stages of randomness. The first relates to *uncertainty* about which element of Θ corresponds to reality (no objective probabilities are available), and the second relates to *risk* about which signal from Y will be received (objective probabilities). This problem can thus be placed within the Anscombe-Aumann model, where the objects of choice are precisely such two-stage lotteries. Within this framework, the non-additive measure over types ν should not be viewed as primitive, but rather arising from the representation of some preference ordering $\succsim^{AA}$.[3]

Obtaining beliefs over the state space can now be rephrased in terms of preferences. We will show how the binary relation $\succsim^{AA}$ over two-stage horse-lotteries *induces* an ordering over one-stage acts in the Savage framework. Denote this induced relation by $\succsim^{SV}$. Finding a measure over the product space S is then equivalent to finding a Choquet expected utility representation for the Savage preferences $\succsim^{SV}$.

Based on the non-equivalence of the Choquet-indicator rule and the multiple-priors rule, we argue that the difference between Anscombe-Aumann decision making and the Savage framework arises, not from inherent differences between one and two-stage lotteries, but from the inability of non-additive priors to model uncertainty as precisely as multiple priors.

The rest of the paper is organized as follows. Section 2 provides the notation and outlines some existing results. Section 3 presents the theoretical framework for our updating problem. In section 4, we introduce the two rules for constructing beliefs

[3]The superscript refers to the Anscombe-Aumann setting.

and discuss the relationship between them. Section 5 makes the case that, at least within dynamic updating problems, a multiple-priors representation of uncertainty is more appropriate. A summary, together with some conclusions, are to be found in section 6.

3.2 Notation and Preliminaries

Let Θ be a finite set of *types* and Y denote the set of *signals*. From the specification of the problem, we have a convex *capacity* ν over Θ.

Definition (Capacity). *A* capacity *or* non-additive measure *over* Θ *is a function* $\nu : 2^{\Theta} \to [0,1]$ *satisfying the following:*

(i) $\nu(\emptyset) = 0$, $\nu(\Theta) = 1$

(ii) *For* $A_1, A_2 \subseteq \Theta$, $A_1 \subseteq A_2 \Rightarrow \nu(A_1) \leq \nu(A_2)$

If (ii) holds, ν *is* monotone.

We say that ν *is* convex, *or* supermodular, *if in addition, the following holds:*

(iii) $\nu(A_1 \cup A_2) \geqslant \nu(A_1) + \nu(A_2) - \nu(A_1 \cap A_2)$, *for all* $A_1, A_2 \in 2^{\Theta}$

It is superadditive, *if (iii) holds for disjoint* A_1 *and* A_2.

For each type $\theta \in \Theta$, there is an additive probability distribution over the set of signals Y which may be received. These lotteries represent objective risk and we denote them by $\{\mu(\cdot \mid \theta)\}_{\theta \in \Theta}$.

The most popular scheme for updating a convex capacity is the Dempster-Shafer rule. For additive measures, this rule corresponds to Bayes's rule.

Definition (Dempster-Shafer). *The* Dempster-Shafer update *of a convex capacity* ν *conditioned on event is* $A \subseteq \Theta$ *is defined by the following expression. For all* $A_1 \subseteq A$,

$$\nu(A_1 \mid A) = \frac{\nu(A_1 \cup A^c) - \nu(A^c)}{1 - \nu(A^c)} \tag{3.3}$$

To see how our problem relates to the literature on decision theory, we need to introduce preferences. Let $\succsim^{AA}$ represent a preference ordering over *horse-lotteries*. A horse-lottery in our notation is simply a mapping from Θ onto the set of probability distributions over consequences, $h : \Theta \to \Delta C$, where C is the set of consequences. Denote the set of horse-lotteries by H.

Throughout, we assume that preferences $\succsim^{AA}$ are primitive and that they satisfy the Schmeidler (1989) axioms for representation as a Choquet expected utility function so that for all h and h' :

$$h \succsim^{AA} h' \text{ if and only if } \int U \circ h \, d\nu \geqslant \int U \circ h' \, d\nu \tag{3.4}$$

where U is a von-Neumann-Morgenstern linear utility function with a Bernoulli utility function, $u : C \to \mathbb{R}^+$, over consequences. The capacity over types ν is obtained from this representation. Calculating utility involves a two-stage expectation. In the von-Neumann-Morgenstern utility, the expectation is with respect to the lotteries over consequences. Expectation is then carried out using the Choquet integral which is defined as follows:

Definition (Choquet Integral). *Let* $g : \Theta \to \mathbb{R}$ *be a random variable. The* Choquet integral *of* g *with respect to the capacity* ν *is defined as:*

$$\int g \, d\nu = g_1 \nu(A_1) + \sum_{i=2}^{n} g_i \left[\nu \left(\cup_{j=1}^{i} A_j \right) - \nu \left(\cup_{j=1}^{i-1} A_j \right) \right] \tag{3.5}$$

where g_i *is the* i^{th} *highest consequence under* g *and* $A_i \in 2^{\Theta}$ *is the event in which the consequence* g_i *occurs.*

Because a preference ordering which admits a Choquet expected utility representation can always be represented as a maxmin expected utility, we know from Gilboa and Schmeidler (1989) that there exists a closed convex set P of additive probability measures on Θ, such that for all h and h' :

$$h \succsim^{AA} h' \text{ if and only if } \min_{p \in P} \int U \circ h \, dp \geqslant \min_{p \in P} \int U \circ h' \, dp \tag{3.6}$$

Moreover, the set of multiple priors P is the *core* of ν. We will abuse notation by referring to p as both a measure and a vector.

Definition (Core). *The* core *of a non-additive measure* ν*, denoted by* $core(\nu)$ *is defined, as in the cooperative theory for transferable-utility games, by:*

$$core(\nu) = \left\{ p = (p_1, ..., p_{|\Theta|}) \in \Delta^{|\Theta|-1} | \sum_{i \in A} p_i \geqslant \nu(A), \text{ for all } A \subseteq \Theta \right\} \tag{3.7}$$

For the purpose of the signaling problem, the set of probability distributions, $\{\mu(. \mid \theta)\}_{\theta \in \Theta}$, are objective and fixed. Therefore, to place our problem within the Anscombe-Aumann decision setting, we have to restrict the set of horse-lotteries to those in which the second-stage risk is given by some element of $\{\mu(. \mid \theta)\}_{\theta \in \Theta}$. The consequences attached to these probabilities can differ between horse-lotteries. We denote this set of restricted horse-lotteries by $H_\mu \subseteq H$. To illustrate in the context of our motivating example, consider the following capacity and likelihoods:

$$\begin{aligned} \nu(\theta_H) = \tfrac{1}{4},\ \nu(\theta_L) = \tfrac{1}{4},\ \text{and}\ \nu(\{\theta_H, \theta_L\}) = 1 \\ \mu(y_H \mid \theta_H) = \tfrac{4}{5},\ \mu(y_L \mid \theta_H) = \tfrac{1}{5} \\ \mu(y_H \mid \theta_L) = 0,\ \mu(y_L \mid \theta_L) = 1 \end{aligned} \tag{3.8}$$

The employer's prior over the worker's type is characterized by ambiguity and results in a non-additive measure. The production process yields a high output (y_H) with probability $\frac{4}{5}$ when the worker is a high-effort type (θ_H). It yields low output (y_L) with probability 1 if the worker is of the low-effort type (θ_L). With these numbers, elements of H_μ take the form of the following pair of lotteries: $h(\theta_H) = \langle \frac{4}{5}, c_1;\ \frac{1}{5}, c_2 \rangle$; $h(\theta_L) = \langle 0, c_3;\ 1, c_4 \rangle$ where $c_i \in C$, for all $i \in \{1, 2, 3, 4\}$.

In the next section, it is necessary to compare preference orderings under $\succsim^{AA}$ with those under Savage preferences $\succsim^{SV}$. Call the product space $S = \Theta \times Y$ the set of *states*. Savage preferences are defined over *acts*, which are mappings from states to consequences, $f : S \to C$. Denote the set of acts by F. To facilitate comparison between $\succsim^{AA}$ and $\succsim^{SV}$, we need this additional definition:

Definition (Induced Act). *Write $Y = \{y_1, y_2, ..., y_n\}$ and consider lotteries over consequences with the dimension of the support equal to the cardinality of Y; that is, for all θ in Θ, $h(\theta) = \langle \mu(y_1 \mid \theta), c_{\theta,y_1}; ...; \mu(y_n \mid \theta), c_{\theta,y_n} \rangle$. The act over states induced by the horse-lottery h in H_μ is a mapping, $f^h : S \to C$, defined as:*

$$f^h(\theta \times y) = c_{\theta,y} \tag{3.9}$$

Notice that induced acts do not depend on the probabilities which are part of the specification of horse-lotteries. In the Savage setting, the risk contained in lotteries over consequences is modeled explicitly as part of the description of the state. Continuing with the example above, the horse-lottery—$h(\theta_H) = \langle \frac{4}{5}, c_1;\ \frac{1}{5}, c_2 \rangle$, $h(\theta_L) = \langle 0, c_3;\ 1, c_4 \rangle$—in the Anscombe-Aumann framework induces the following Savage act: $f^h = (c_1, c_2, c_3, c_4)$. Figure 1 illustrates this example.

Figure 1

Type		Signal	State	Likelihood	Consequence
		$\nearrow$	$s_1 = (\theta_H, y_H)$	$\mu(y_H \mid \theta_H) = \frac{4}{5}$	c_1
$\nearrow$	θ_H	$\rightarrow$	$s_2 = (\theta_H, y_L)$	$\mu(y_L \mid \theta_H) = \frac{1}{5}$	c_2
$\searrow$	θ_L	$\rightarrow$	$s_3 = (\theta_L, y_H)$	$\mu(y_H \mid \theta_L) = 0$	c_3
		$\searrow$	$s_4 = (\theta_L, y_L)$	$\mu(y_L \mid \theta_L) = 1$	c_4

Axiomatizations for both Choquet and maxmin expected utility exist in the Savage setting. For Choquet expected utility, see Gilboa (1987) and Sarin and Wakker (1992). Casadesus-Masanell *et al.* (1998) axiomatize the maxmin expected utility representation. One final piece of notation. We denote the capacity and the set of multiple priors over the state space, $S = \Theta \times Y$, by σ and Q, respectively.

3.3 Theoretical Framework

As we pointed out in the introduction, updating the capacity over types upon the receipt of a signal requires the construction of beliefs on the product space $S =$

$\Theta \times Y$. How should this—in general, non-additive—measure be constructed? What desiderata should σ possess?

By placing our problem within the framework of preferences, we obtain a very natural property that σ should satisfy. Assume that the capacity ν over Θ is the result of a representation of the primitive ordering $\succsim^{AA}$ over horse lotteries in H_μ. Based on this preference ordering, we can *define* a relation $\succsim^{SV}$, over acts, according to the following. For all $h, h' \in H_\mu$,

$$h \succsim^{AA} h' \Leftrightarrow f^h \succsim^{SV} f^{h'} \qquad (3.10)$$

Having done so, constructing beliefs on the state space S amounts to finding a measure σ that represents $\succsim^{SV}$. That is, we want σ to satisfy the following utility representation:

$$f^h \succsim^{SV} f^{h'} \Leftrightarrow \int_S u\left[f^h(s)\right]\, d\sigma(s) \geq \int_S u\left[f^{h'}(s)\right]\, d\sigma(s) \qquad (3.11)$$

We can obtain another perspective by re-stating the requirement in (3.10) as that of finding a σ such that the expected utility representation of the two preference orderings are equivalent. For all h in H_μ we want,

$$\int_\Theta U\left[h(\theta)\right]\, d\nu(\theta) = \int_S u\left[f^h(s)\right]\, d\sigma(s) \qquad (3.12)$$

The left-hand-side contains the utility function which represents $\succsim^{AA}$; the right-hand-side contains the representation of $\succsim^{SV}$. This equation can be written more explicitly as:

$$\int_\Theta \int_Y u\left[f^h(\theta \times y)\right]\, d\mu(y \mid \theta)\, d\nu(\theta) = \int_{\Theta\times Y} u\left[f^h(\theta \times y)\right]\, d\sigma(\theta \times y) \qquad (3.13)$$

Equation (3.13) allows us to restate the problem. The aim is to find a measure, σ, on the product space, $S = \Theta \times Y$, for which part of Fubini's theorem holds. Fubini's theorem states that the order of the iterated integrals with respect to two marginal measures do not matter and that both are equal to integration with respect to the product measure. Condition (3.13) requires only that integration over Y, *then* Θ, be equivalent to integration with respect to the product measure.

Sarin and Wakker (1992) were the first to observe that, in general, it is not possible to find a capacity σ on S which satisfies (3.13). This can be illustrated in the context of our example. Consider the following three horse-lotteries, h_1, h_2, h_3:

$$\begin{aligned} h_1(\theta_H) &= \left\langle \tfrac{4}{5}, 1;\ \tfrac{1}{5}, 0 \right\rangle, \quad h_1(\theta_L) = \langle 0,0;\ 1,0 \rangle \\ h_2(\theta_H) &= \left\langle \tfrac{4}{5}, 1;\ \tfrac{1}{5}, 0 \right\rangle, \quad h_2(\theta_L) = \langle 0,0;\ 1,1 \rangle \\ h_3(\theta_H) &= \left\langle \tfrac{4}{5}, 2;\ \tfrac{1}{5}, 0 \right\rangle; \quad h_3(\theta_L) = \langle 0,0;\ 1,1 \rangle \end{aligned} \tag{3.14}$$

Recall that each lottery is of the form $h(\theta) = \langle \mu(y_1 \mid \theta), c_1; \mu(y_2 \mid \theta), c_2 \rangle$. These are all elements of H_μ because the probabilities are identical and given by $\{\mu(\cdot \mid \theta)\}_{\theta \in \Theta}$. Respectively, they induce the following acts in the Savage formulation:

$$\begin{aligned} f^{h_1} &= (1,0,0,0) \\ f^{h_2} &= (1,0,0,1) \\ f^{h_3} &= (2,0,0,1) \end{aligned} \tag{3.15}$$

Without loss of generality, we assume that consequences are in utils, or that the Bernoulli utility function is given by $u(c) = c$. To satisfy equation (3.13) for f^{h_1} and

f^{h_2}, it can then be verified that we need:

$$\begin{aligned} \sigma\left[(\theta_H, y_H)\right] &= \tfrac{4}{20} \\ \sigma\left[\{(\theta_H, y_H), (\theta_L, y_L)\}\right] &= \tfrac{17}{20} \end{aligned} \tag{3.16}$$

However, with these values, the Choquet expected utility of the horse lottery h_3 in the Anscombe-Aumann framework is given by $\int_X \int_Y u\left(f^{h_3}\right) d\mu d\nu = \frac{23}{20}$ while the Choquet expected utility of the corresponding induced act f^{h_3} in the Savage framework is given by $\int_S u\left(f^{h_3}\right) d\sigma = \frac{21}{20}$.

This difference between the two frameworks does have important implications. For example, Eichberger and Kelsey (1996) show that in the Anscombe-Aumann framework (represented by the left-hand-side of equation 3.13), uncertainty-averse agents exhibit a preference for randomization, but they do not necessarily do so when the objects of choice are Savage acts (the representation of the right-hand-side of 3.13). We have shown that the difference also matters when agents are revising non-additive priors upon the receipt of some signal.

3.4 Obtaining Beliefs Over the State Space

Having established that it is impossible to obtain a capacity σ which satisfies the desideratum of (3.13), we now consider some weaker desirable properties which we may want a capacity over the state space to satisfy.

One obvious feature which we would like our rule to possess is that it should correspond to Bayes's rule in the special case of additive distributions. In order to ensure this, rectangular sets formed by partitioning S according to some element of Θ—that is, sets of the form $\theta \times B$, where $\theta \in \Theta$ and $B \in 2^Y$—must have measure given by:

$$\sigma(\theta \times B) = \nu(\theta) \cdot \mu(B \mid \theta) \tag{3.17}$$

This is of course just Bayes's rule when ν is additive. We will refer to (3.17) as the *multiplicative property*. However, this still leaves the measure of many subsets in S, including many rectangles, unspecified.

From the Dempster-Schafer rule, we have the following:

$$\mu(B \mid \theta) = \frac{\sigma\left(\theta^C \cup B\right) - \sigma\left(\theta^C\right)}{1 - \sigma\left(\theta^C\right)} \tag{3.18}$$

Rearranging,

$$\begin{aligned} \sigma\left(\theta^C \cup B\right) &= \mu(B \mid \theta)\left[1 - \sigma\left(\theta^C\right)\right] + \sigma\left(\theta^C\right) \\ &= \mu(B \mid \theta)\left[1 - \nu\left(\theta^C\right)\right] + \nu\left(\theta^C\right) \end{aligned} \tag{3.19}$$

The right-hand-side of (3.19) is given by the specification of the problem. Thus using the Dempster-Shafer rule we can obtain the value of the joint capacity on sets of the form $\theta^C \cup B$, where $B \in 2^Y$ and $\theta \in \Theta$. We say that a capacity σ satisfies the *Dempster-Shafer property* if it obeys (3.19).

Even if we impose the multiplicative property of (3.17), as well as the Dempster-Shafer property of (3.19), we do not obtain a unique capacity. In a similar spirit to Hendon *et al.* (1991), we can characterize the *set* of capacities that we do obtain by some limits if we require that σ be monotone. Take for example the set $E =$

$\{(\theta_H, y_L), (\theta_L, y_L)\}$ from figure 1. Although equations (3.17) and (3.19) do not provide a unique value for its measure, we can derive the following bound using a simple set inclusion argument:

$$\max\{\sigma(\theta_H, y_L), \sigma(\theta_L, y_L)\} \leqslant \sigma(E) \leqslant \min\{\sigma(\theta_H \cup Y), \sigma(\theta_L \cup Y)\} \tag{3.20}$$

These bounds can be calculated using the multiplicative property and the Dempster-Shafer property.

3.4.1 The Choquet-Indicator Rule

We now propose a rule for obtaining a *unique* capacity over the state space which satisfies the multiplicative and Dempster-Shafer properties. The definition is as follows.

Definition (Choquet-Indicator Rule). *A capacity σ on S with marginal ν over Θ and a set of likelihood distributions $\{\mu(\cdot \mid \theta)\}_{\theta \in \Theta}$ over Y, is said to be generated by the* Choquet-indicator rule *if for every $E \in 2^{\Theta \times Y}$:*

$$\sigma(E) = \int_{\Theta} \int_{Y} 1_E \, d\mu(y \mid \theta) \, d\nu(\theta) \tag{3.21}$$

where 1_E is an indicator function over E.

Result. *Let σ be a capacity on $\Theta \times Y$ calculated using the Choquet-indicator rule. Then σ satisfies the multiplicative and Dempster-Shafer properties in (3.17) and (3.19), respectively.*

Remark. *By construction, over the set of acts, $\{f \in F \mid u[f(s)] \in \{0,1\}$ for all s in $S\}$, the Choquet-indicator rule satisfies (3.13), our original desideratum.*

The result is easy to verify. The remark says that, if one restricts attention to acts which take on only two consequences, the Choquet-indicator (CI) rule maintains the equivalence between the Anscombe-Aumann and the Savage frameworks. It is constructed to do so. A comparison of equations (3.13) and (3.21) makes this obvious.

To illustrate how the CI rule works, consider the set $\{(\theta_H, y_H), (\theta_L, y_L)\}$ from our example above. Naively, one could make the following calculation:

$$\begin{aligned} \sigma\left[\{(\theta_H, y_H), (\theta_L, y_L)\}\right] &= \nu(\theta_H) \cdot \mu(y_H \mid \theta_H) + \nu(\theta_L) \cdot \mu(y_L \mid \theta_L) \\ &= \frac{1}{4} \times \frac{4}{5} + \frac{1}{4} \times 1 = \frac{9}{20} \end{aligned} \tag{3.22}$$

This calculation assigns a probability of $\frac{1}{4}$ to each of the two types, θ_H and θ_L. It ignores the fact that, with the residual probability of $\frac{1}{2}$, either θ_H or θ_L will necessarily occur. The CI rule corrects for this in the most "pessimistic" way. It assigns the residual probability to that outcome which would produce the lowest Choquet expectation, in this case (θ_H, y_H), as $\mu(y_H \mid \theta_H) = \frac{4}{5} < \mu(y_L \mid \theta_L) = 1$.

Despite being intuitive, and despite satisfying the multiplicative property and the Dempster-Shafer property—both of which seem desirable—the CI rule does not imply an equivalence between the Anscombe-Aumann and Savage frameworks. As we pointed out in the previous section, no rule which generates a capacity can. If we are willing to leave the non-additive framework, and allow uncertainty-averse beliefs to be represented by a set of multiple *additive* priors, can we do better? In the next subsection, we present a rule for calculating multiple priors over states which yields an equivalence result between the one and two-stage frameworks for decision making.

Before this is done, we discuss the relationship between the Choquet–indicator rule and the work of Ghirardato (1997). Ghirardato considered a situation where two non-additive marginal measures are known. He asked what conditions are necessary to obtain a capacity σ on the product space which satisfies the Fubini theorem. He showed that the theorem will hold if one restricts the set of acts to those which are *slice comonotonic* and imposes on σ a strengthening of independence, which he termed the *Fubini property*[4].

Definition (Fubini Property). *A function $g : \Theta \times Y \to \mathbb{R}$ is slice comonotonic if for every $\theta, \theta' \in \Theta$, $g(\theta, \cdot)$ and $g(\theta', \cdot)$ are comonotonic, and if for every $y, y' \in Y$, $g(\cdot, y)$ and $g(\cdot, y')$ are comonotonic. Define a* comonotonic set *as one in which the indicator function over that set is* slice-comonotonic. *Now, a capacity σ is said to satisfy the* Fubini property *with respect to marginals, ν over Θ and μ over Y, if the following equation holds for every comonotonic set $E \in 2^{\Theta \times Y}$:*

$$\sigma(E) = \int_{\Theta} \int_{Y} 1_E \, d\mu(y) \, d\nu(\theta) \tag{3.23}$$

where 1_E is the indicator function over E.

Our Choquet-indicator (CI) rule can be viewed as a strengthening of the Fubini property; it imposes that (3.23) hold for *all* sets $E \in 2^S$. (The CI rule also differs from equation (3.23) in that one of the measures in the integral is a conditional one.)

Ghirardato's definition does not require that equation (3.23) hold for all sets

---

[4]This is not to be confused with the Fubini theorem. The Fubini property is a characteristic of capacities.

because, when applied to all elements in 2^S, the capacity generated by (3.23) does not satisfy the "iterated integration" part of the Fubini theorem. It is not clear whether one should define $\sigma(E)$ as $\int_\Theta \int_Y 1_E d\mu d\nu$ or $\int_Y \int_\Theta 1_E d\nu d\mu$. For comonotonic sets, these two integrals are equivalent.

In our updating framework, the order of integration is clear so this part of Fubini's theorem is not a desirable restriction. By strengthening the Fubini property, we are able to obtain a *unique* capacity over the product space; there are multiple capacities over the product space S which satisfy the Fubini property. Ghirardato requires the additional assumption of convexity to obtain uniqueness. In general, our CI rule does not produce a convex capacity. However, we argue that this is not a weakness since convexity restricts the kind of uncertainty which one can model. This point will become clearer when we define the alternative way to obtain beliefs in the next subsection.

3.4.2 The Multiple-Priors Rule

For an agent with preferences $\succsim^{AA}$ over lotteries $h \in H$ who satisfy the axioms for a Choquet expected utility representation with a convex capacity ν over Θ, Gilboa and Schmeidler (1989) showed that the agent is behaviorally equivalent to one with a maxmin expected utility; that is, one who maximizes $\min_{p\in P} \int_\Theta U \circ h \, dp$, where $P = core(\nu)$. In the other direction, assume that an agent possesses a maxmin expected utility representation with a set of multiple priors P. The agent is identical to one who maximizes Choquet expected utility with a capacity ν, defined by $\nu(A) =$

$\min_{p \in P} p(A)$, if and only if ν is convex and $core(\nu) = P$.

In light of these results, we can convert the capacity ν, which is convex by assumption, to the corresponding set of multiple priors $P = core(\nu)$. Is it then possible to obtain a set of *additive* beliefs Q over the product space $S = \Theta \times Y$ so as to obtain equivalence between the one and two-stage formulations? More formally, we require Q to satisfy:

$$\min_{p \in P} \int_{\Theta} U\left[h(\theta)\right] \; dp(\theta) = \min_{q \in Q} \int_{S} u\left[f^{h}(s)\right] \; dq(s) \tag{3.24}$$

where again U is a von-Neumann-Morgenstern utility function with Bernoulli utility u. This is simply the multiple priors analogue to equation (3.12). The left-hand-side is the maxmin expected utility representation of $\succsim^{AA}$; the right-hand-side is the representation of $\succsim^{S}$. Consider the following rule for calculating a set of distributions Q over the state space using the priors over types, $P = core(\nu)$, and the likelihoods, $\{\mu(. \mid \theta)\}_{\theta \in \Theta}$.

$$Q = \left\{q = (q_1, ..., q_{|S|}) \in \Delta^{|S|-1} \mid q_s = p(\theta) \times \mu(y \mid \theta) \text{ for all } p \in P\right\} \tag{3.25}$$

where $s = \theta \times y$. We refer to this as the multiple-priors (MP) rule. It simply takes each prior in P in turn and applies a rearrangement of Bayes's rule. In the case of our motivating example, we obtain:

$$Q \;\; = \;\; \left\{q \in \Delta^3 \mid q_1 \in \left[\tfrac{4}{20}, \tfrac{12}{20}\right], \; q_2 \in \left[\tfrac{1}{20}, \tfrac{3}{20}\right], \; q_3 = 0, \; q_4 \in \left[\tfrac{5}{20}, \tfrac{15}{20}\right], \; q_1 = 4q_2\right\} \tag{3.26}$$

Figure 2 illustrates the projection of this set onto the (q_1, q_2) space. We can completely represent the set in two dimensions because $q_3 = 0$ and q is on the simplex.

These two restrictions reduce the degrees of freedom to two.

Figure 2

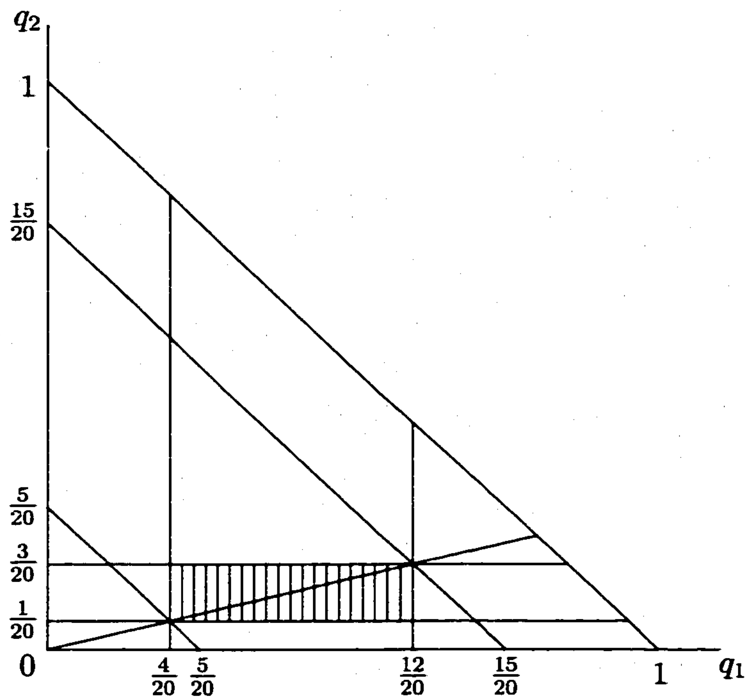

Proposition (Multiple-Priors Rule). *The multiple-priors rule satisfies the Anscombe-Aumann and Savage equivalence. That is, it satisfies equation (3.24) for all* $h \in H_\mu$.

We omit the proof, since the result is a direct consequence of the fact that the Fubini theorem holds with additive priors. This result, though simple, is somewhat surprising given that it cannot be obtained in terms of a capacity on S. Intuition will be provided in the next subsection.

Remark. *The set of additive measures* Q *generated from the multiple-priors rule cannot in general be expressed as the core of any capacity.*

The figure above is an example of this remark. From the definition of equation (3.7), we can see that a set can only be expressed as the core of some capacity if it can be defined by a system of linear inequalities of the form $\sum_{i \in A} p_i \geqslant \nu(A)$. Geometrically, the set must have sides which are parallel to the sides of the simplex. In figure 2, the set of distributions, Q, is represented by the line segment connecting the points $\left(\frac{4}{20}, \frac{1}{20}, 0, \frac{15}{20}\right)$ and $\left(\frac{12}{20}, \frac{3}{20}, 0, \frac{5}{20}\right)$. Since this line is not parallel to any of the sides of the triangle, it cannot be expressed as the core of any capacity.

3.4.3 The Relationship Between the Two Rules

The two procedures for updating beliefs over types can be summarized as follows. We begin with a convex, non-additive prior ν over the set of types Θ. One can think of these beliefs as deriving from some Choquet expected utility representation of the agent's preferences in an Anscombe-Aumann setting. To obtain posterior beliefs after the receipt of some signal from Y, we need to first define beliefs over the product space, $S = \Theta \times Y$. One way to do this is to use a rule based on the Choquet integration of indicator acts. This yields a measure σ which is in general non-additive, reflecting the transfer of uncertainty and uncertainty aversion over the types to uncertainty and uncertainty aversion over states. The measure σ satisfies the multiplicative property and the Dempster-Shafer property. It does not, however, represent an ordering over Savage acts which is equivalent to the Anscombe-Aumann preference ordering.

An alternative approach is to convert the non-additive measure ν to an equivalent set of multiple priors, P. Multiple probability distributions over states can then be obtained using Bayes's rule. The resulting set is labeled Q.

Figure 3

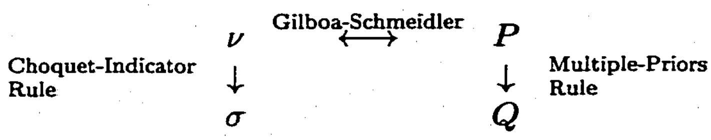

The main advantage of the multiple-priors approach is that it ensures an equivalence between the one and two-stage frameworks.

If one wanted to remain within the non-additive framework, then an obvious question is: which capacity comes "closest" to representing the uncertainty over states embodied in the set Q? It turns out that the capacity which does so is indeed the one calculated from the CI rule. This idea can best be described graphically using our example. Figure 2 shows that, among all sets of distributions which can be expressed as the core of some capacity, the shaded rectangle is the *smallest* one which contains Q. One can verify that, if we define σ using the CI rule, then this rectangle is precisely the set $core(\sigma)$. The theorem and corollary below formalize this.

Theorem (CI and MP Rules). *Assume that ν on Θ is convex. Let σ be the capacity on S defined by the Choquet-indicator rule and let Q be the set of multiple additive measures on S derived from the multiple-priors rule. Then,*

$$\sigma(E) = \min_{q \in Q} q(E) \tag{3.27}$$

for all $E \in S$.

Proof. From Schmeidler (1989), proposition (x), we have:

$$\int_{\Theta} U[h(\theta)]\, d\nu(\theta) = \min_{p \in P} \int_{\Theta} U[h(\theta)]\, dp(\theta) \tag{3.28}$$

for any act h where ν is convex. From our rules, $\sigma(E) = \int_{\Theta} 1_E d\nu(\theta)$ and $\min_{q \in Q} q(E) = \min_{p \in P} \int_{\Theta} 1_E dp(\theta)$. The result follows, by using the act 1_E for h in (3.28). ■

Corollary. *Let σ be the capacity on S defined under the Choquet-indicator rule and let Q be the set on multiple additive measures on S derived from the multiple-priors rule. Then,*

$$Q \subseteq core(\sigma) \tag{3.29}$$

Proof. Let $q \in Q$. Assume, contra-hypothesis, that $q \notin core(\sigma)$. Then, there exists $E \in S$, such that $q(E) < \sigma(E)$. But this contradicts the theorem above. ■

As an aside, the CI rule produces a capacity σ which is not necessarily convex. This is easy to verify in the example above. Despite this, we have the following remark.

Remark. *Because Q is always non-empty, the corollary implies that σ has a core which is non-empty.*[5]

The theorem of this section, together with its corollary, provides a formal justification for the Choquet-indicator rule. Given that the multiple-priors rule satisfies the

[5]In a recent working paper, Ghirardato and Marinacci (1998) argue that, within the Savage framework, ambiguity aversion corresponds to nonemptiness of the core, a property strictly weaker than convexity. In light of this result, the CI rule maintains the initial uncertainty aversion even though the capacity σ on S is not convex.

requirement of (3.24), we argue that the agent's beliefs over the product space should in fact be given by this rule. The Choquet-indicator rule can then be justified on the grounds that it produces that capacity which comes the closest to the "correct" beliefs.

Before concluding this section, we return to the original motivation for this paper and construct, for our example, the posteriors on types, Θ, after the observation of a low output, y_L. Applying the Dempster-Schafer rule to the capacity σ, generated by the CI rule, we obtain the following posterior:

$$\begin{aligned} \nu(\theta_H \mid y_L) &= \tfrac{1}{16} \\ \nu(\theta_L \mid y_L) &= \tfrac{13}{16} \\ \nu(\{\theta_H, \theta_L\} \mid y_L) &= 1 \end{aligned} \tag{3.30}$$

This measure has the following core:

$$core\left[\nu(. \mid y_L)\right] = \left\{p = (p_1, p_2) \in \Delta^1 \mid p_1 \in \left[\tfrac{1}{16}, \tfrac{3}{16}\right]\right\} \tag{3.31}$$

We now wish to construct posterior beliefs over types using the set of additive distributions on the product space calculated from the multiple-priors rule. Gilboa and Schmeidler (1993) show that, for decision makers who can be represented both by Choquet expected utility and by maxmin expected utility, the Dempster-Shafer rule on capacities coincides with the combination of maximum likelihood and Bayes's rule applied to the set of multiple priors. Therefore, to enable comparison between the Choquet indicator rule and the multiple-priors rule, we apply maximum likelihood to the set Q and then use Bayes's rule, element-by-element, to obtain posterior beliefs over types. This gives the following unique additive posterior:

$$\{p = (p_1, p_2) \in \Delta^1 \mid p_1 = \tfrac{1}{16}\} \tag{3.32}$$

Comparing equations (3.31) and (3.32), it is clear that the CI rule yields posterior beliefs which contain greater uncertainty than those obtained from the multiple-priors rule.

3.5 The Argument for Multiple Priors

From the work of Gilboa and Schmeidler (1989), we know that the multiple-priors framework is more general than that of convex capacities. Any convex capacity can be represented as a set of multiple priors, whereas the converse is not true. What is surprising about the updating example is that, even though we begin with beliefs on types which can be represented equivalently by a capacity ν or by a set of multiple priors P, as soon as we introduce signals and attempt to construct beliefs on the product space, the two frameworks diverge.

We believe that the updating problem considered in this paper highlights the importance of the additional generality of multiple priors. A comparison of the core of σ from the CI rule, with Q from the MP rule, makes this point. In the example:

$$core(\sigma) = \left\{q \in \Delta^3 \mid q_1 \in \left[\tfrac{4}{20}, \tfrac{12}{20}\right],\ q_2 \in \left[\tfrac{1}{20}, \tfrac{3}{20}\right],\ q_3 = 0,\ q_4 \in \left[\tfrac{5}{20}, \tfrac{15}{20}\right]\right\} \tag{3.33}$$

$$Q = \left\{q \in \Delta^3 \mid q_1 \in \left[\tfrac{4}{20}, \tfrac{12}{20}\right],\ q_2 \in \left[\tfrac{1}{20}, \tfrac{3}{20}\right],\ q_3 = 0,\ q_4 \in \left[\tfrac{5}{20}, \tfrac{15}{20}\right],\ q_1 = 4q_2\right\} \tag{3.34}$$

These two sets are identical except for the equation $q_1 = 4q_2$ in (3.34). *Capacities are unable to capture restrictions on the relative likelihood of some events.* The capacity σ is unable to restrict the probability of the first state, (θ_H, y_H), to be four times that of the second, (θ_H, y_L). Clearly, given that the risk associated with the signal is objective, no matter what the agent's original beliefs over types, the likelihood ratio between these two states should remain 4 to 1. By trying to use capacities to capture beliefs on the product space, the agent loses some of the information contained in the signal and attributes to the problem greater uncertainty than is in fact present. In turn, this leads to greater uncertainty in the posterior, as the sets in (3.31) and (3.32) demonstrate.

Moreover, the inability of capacities to capture relative likelihoods is the reason for the non-equivalence between the Anscombe-Aumann and the Savage frameworks. (Recall that in the multiple-priors setting, the one and two-stage frameworks are equivalent.) As a result, Eichberger and Kelsey's (1996) claim that the Savage framework is more appropriate for modeling uncertainty aversion is not justified. There is really no inherent difference between one and two-stage lotteries as objects of choice. Differences arise from a limitation of capacities.

3.6 Conclusion

In many dynamic economic situations, beliefs over the relevant state space are not given by the specification of the problem. With additive measures, Bayes's rule usually suffices to define a unique distribution over states. However, with ambiguous beliefs represented by a non-additive measure, unique beliefs over the state space cannot be obtained from the Dempster-Shafer rule alone.

We argued that obtaining beliefs over the state space is closely related to the issue of whether one can move from the Anscombe-Aumann to the Savage setting while maintaining the "same" preference ordering. This is impossible when capacities are used to model uncertainty-aversion. However, using multiple additive distributions, this equivalence between the two frameworks is possible. We then set out to find the capacity which comes closest to the beliefs obtained using multiple-priors. Such a capacity can be constructed by taking Choquet expectations of appropriate indicator functions.

Finally, we showed that the updating problems studied in our paper highlight the advantage of multiple priors relative to non-additive measures. Capacities are unable to place restrictions on the relative likelihood of events. This is a severe limitation in dynamic problems where such ratios arise naturally from the updating of beliefs.

3.7 References

[1] ANSCOMBE, F. AND R. AUMANN(1963), "A Definition of Subjective Probability", *Annals of Mathematical Statistics.* 34, 199-205.

[2] CASADESUS-MASANELL, R., P. KLIBANOFF, AND E. OZDENOREN(1998), "Maxmin Expected Utility Over Savage Acts With a Set of Priors", mimeo., Northwestern University.

[3] EICHBERGER, J. AND D. KELSEY (1996), "Uncertainty Aversion and Preference for Randomization", *Journal of Economic Theory* 71, 31-43.

[4] GHIRARDATO, P. (1997), "On Independence for Non-Additive Measures With a Fubini Theorem", *Journal of Economic Theory* 73, 261-291.

[5] GHIRARDATO, P. AND M. MARINACCI (1998), "Ambiguity Made Precise: A Comparative Foundation", mimeo..

[6] GILBOA, I. (1987), "Expected Utility Theory With Purely Subjective Non-Additive Probabilities", *Journal of Mathematical Economics* 16, 65-88.

[7] GILBOA, I. AND D. SCHMEIDLER(1989), "Maxmin Expected Utility With a Non-Unique Prior", *Journal of Mathematical Economics* 18, 141-153.

[8] GILBOA, I. AND D. SCHMEIDLER(1993), "Updating Ambiguous Beliefs", *Journal of Economic Theory* 59, 33-49.

[9] HENDON, E., H. JACOBSEN, B. SLOTH AND T. TRANÆS(1991), "The Product of Capacities and Lower Probabilities", mimeo., University of Copenhagen.

[10] MUKERJI, S.(1997), "Understanding the Non-Additive Probability Decision Model", *Economic Theory* 9, 23-46.

[11] MYERSON, R.(1991), *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, Massachusetts.

[12] SARIN, R. AND P. WAKKER(1992), "A Simple Axiomatization of Non-Additive Expected Utility", *Econometrica* 60, 1255-1272.

[13] SAVAGE, L.(1954), *The Foundations of Statistics*, Wiley, New York.

[14] SCHMEIDLER, D.(1989), "Subjective Probability and Expected Utility Without Additivity", *Econometrica* 57, 571-587.